Home Equity, Mobility, and Macroeconomic Fluctuations

Vincent Sterk*

First draft: *October 2010*
This draft: *August 2012*

Abstract

How does a fall in house prices affect real activity? This paper presents a business cycle model in which a decline in house prices reduces geographical mobility, creating distortions in the labor market. This happens because homeowners face declines in their home equity levels, after which it becomes more difficult to provide the downpayment required for a new mortgage loan. Unemployed homeowners therefore turn down job offers that would require them to move. The model explains joint cyclical patterns in housing and labor market aggregates, as well as the puzzling breakdown of the U.S. Beveridge curve that occurred during 2009.

Key Words: Housing Markets, Labor Markets, Refinancing Constraints

JEL Classification: E24, E44, R21

*University College London. E-mail: v.sterk@ucl.ac.uk. I thank Jeff Campbell, Wouter den Haan, Jonas Fisher, Matteo Iacoviello, Marco del Negro, Morten Ravn, Peter Rupert, Ayşegül Şahin, Petr Sedláček, Joseph Tracy, and various seminar and conference participants for helpful comments. Part of this research was conducted at De Nederlandsche Bank and at the Federal Reserve Bank of New York. The views expressed in this paper are my own and do not necessarily reflect official positions of any policy institution. All errors are my own.

1 Introduction

The recent "Great Recession" is characterized by unusual disruptions in both housing and labor markets. In housing markets, there was a sharp fall in both prices and the number of transactions. In labor markets, there was an increase in the unemployment rate that was not only exceptional in terms of its magnitude, but also surprising given that the decline in vacancies was not as strong (Elsby, Hobijn, and Sahin (2010)). Figure 1 plots vacancies versus the unemployment rate. The figure shows that the historically strong and negative correlation between the unemployment rate and vacancies, known as the Beveridge Curve, broke down during 2009. An interpretation of this finding is that frictions in the labor market had become more severe, causing unemployed workers and firms to be matched less efficiently.

This paper develops a Dynamic and Stochastic General Equilibrium (DSGE) model in which house prices affect real activity through a *geographical mobility channel*. When house prices decline, homeowners face difficulties in moving to a new house, reducing their incentives to accept job offers that are not within commutable distances from their current homes. As a consequence, a decline in house prices causes unemployment to rise and output to fall, which in turn feeds back into house prices. The calibrated model can explain joint cyclical fluctuations in housing and labor market variables. Moreover, based on *only* house price and output data, the model generates for 2009 a flattening of the Beveridge Curve that is remarkably similar to the one observed in reality.

Joint movements in house prices and mobility arise due to the presence of a *refinancing constraint* that is new to DSGE models.[1] Standard constraints tie the amount of debt to the value of the underlying housing collateral and in each period mortgagees have to refinance their loans (see e.g. Iacoviello (2005)). The unrealistic implication is that fluctuations in house prices affect the borrowing limits of *all* homeowners during *each* period. In contrast, my constraint *only* requires those who move to a new house to refinance their loans. In this environment, households can shield their borrowing capacity from a fall in house prices by staying in their

[1] Stein (1995) considers a somewhat similar constraint and shows that it can generate positive comovement between house prices and transaction volumes in a determistic partial equilibrium model with three periods. The advantage of my constraint is that is easily embedded in a DSGE framework that can be solved using standard techniques.

current locations, avoiding the need to take out a new mortgage loan. After a fall in housing wealth, it becomes harder to provide the downpayment for a new loan, reducing incentives to move.

Fluctuations in mobility affect real activity through their effects on the process that matches unemployed workers and firms. This process is modeled following a standard version of the Diamond-Mortensen-Pissarides model, but with the addition that *some* job offers can only be accepted if the worker moves to a new location. When there are barriers to mobility due to a fall in house prices, more job offers are turned down. Thus, the economy enters a period during which, for a given level of vacancies, the unemployment rate is higher than during normal times. An interesting prediction of the model is that the mobility effects persist beyond the fall in house prices. Moreover, as the amount of leverage among households increases due to structural changes in mortgage markets, real activity becomes more sensitive to shocks, especially to those that arise in housing markets.

The model developed in this paper is the first that allows one to study the dynamic effects of house prices on mobility, as well as the spillovers to the labor market and real activity. Because outcomes in all markets are endogenous, one can also analyze feedback effects on house prices. Head and Lloyd-Ellis (2008) and Rupert and Wasmer (2011) have developed models with mobility effects to study the role of housing markets in determining long-run unemployment rates. However, their models are much less suited to study *fluctuations* in housing and labor markets, because solving stochastic versions of their models would be very challenging. In contrast, my model can be easily solved using standard methods.[2]

The proposed geographical mobility channel is consistent with joint cyclical properties of aggregate housing and labor market data. This is shown in Section 2. Section 3 describes the theoretical model. The predictions of the calibrated model are presented, compared to the aggregate data in Section 4. Section 5 explains the main mechanisms underlying the results and discusses extensions of the model. Section 6 discusses the model's predictions in relation to recent empirical findings on the importance of house-lock in the Great Recession. Section 7 concludes.

[2]Business cycle models that are related to my model include those of Iacoviello and Pavan (2009), who model borrowing constraints and infrequent housing adjustments, but not matching frictions in the labor market, and of Andres, Bosca, and Ferri (2010), who analyze a model with frictions in labor and credit markets, but without mobility effects.

2 Empirical evidence

The idea that house price declines deter geographical mobility has been supported by several micro-econometric studies, including Henley (1998), Chan (2001), Engelhardt (2003), and Ferreira, Gyourko, and Tracy (2010). The latter find that negative equity substantially reduces homeowner mobility. However, recent empirical research inspired by the turmoil in US housing and labor markets, has delivered some more contrasting findings, as reviewed by Molloy, Smith, and Wozniak (2011).[3]

A drawback of purely empirical studies on the effects of house-lock on the labor market is that the key variable of interest, the fraction of job opportunities that is foregone by homeowners because of barriers to mobility, is unobserved. My approach is to develop a structural business cycle model, and compare the predictions of this model to patterns observed in the aggregate data. I focus on the joint cyclical behavior of aggregate housing and labor market variables and consider measures of volatility and comovement that are standard in the business cycle literature. At the center of the analysis are house prices, the number of home sales, and the rate at which workers flow out of unemployment. A link between these three variables is crucial for the proposed mobility channel to be at play at the aggregate level. But of course, unconditional business cycle statistics provide only limited information. I therefore also estimate a structural Vector AutoRegressive model (VAR), which allows me to condition on shocks that arise in housing markets.

2.1 Data and methodology

The empirical analysis focuses on quarterly observations of two housing market variables and three labor market variables. The sample runs from the first quarter of 1970 until the last quarter of 2009. The housing market variables are the real house price and home sales. These data were provided by the National Association of Realtors. The house price is the median sales price of existing single-family homes, deflated by the consumer price index. Home sales are measured by the number of

[3] In particular, Schulhofer-Wohl (2011) challenges the robustness of the results of Ferreira, Gyourko, and Tracy (2010) on the basis of a different coding strategy. In a reply, Ferreira, Gyourko, and Tracy (2011) provide additional evidence to support their original finding and argue that Schulhofer-Wohl's results are due to including temporary moves and noise in the mobility measure. They also call for theoretical frameworks in which both mobility decisions and labor market conditions are endogenous.

existing single-family homes sold in a particular month.[4] The reason for analyzing home sales is that this series can be expected to be a good proxy for overall mobility among homeowners.[5]

The labor market variables are unemployment, vacancies, and the unemployment outflow hazard. Unemployment is measured by the civilian unemployment rate as released by the U.S. Department of Labor. Vacancies are measured by the Help Wanted Index, released by the Conference Board. To account for the rise in internet vacancies, I use the corrected series as constructed by Barnichon (2010) for the post 1995 period. The quarterly unemployment outflow hazard is constructed following Shimer (2007).[6]

The first part of the empirical analysis consists of a graphical investigation of the data. I consider raw data and construct their cyclical components using the Hodrick-Prescott (HP) filter. Finally, I estimate a structural VAR, which is used to condition on housing market shocks. A detailed description of the VAR used for this purpose and the identification strategy is provided in Appendix A.1.

2.2 Is the geographical mobility channel present in aggregate data?

If the geographical mobility channel is relevant, what cyclical patterns would one expect to observe? First, one would expect that during periods when house prices are low, fewer homes are sold. At the same time, there should be a fall in unemployment outflow rates during those periods. Moreover, the geographical mobility channel relies crucially on variations in the *efficiency* of labor market matching. Thus, one would expect measures of labor market efficiency to decline when house prices fall. Finally, the geographical mobility channel might cause comovements between housing and labor market variables to be particularly strong when conditioning on

[4] Monthly series on house prices and the number of home sales were converted into quarterly series by taking simple averages.

[5] The main alternative to the home sales index would be the mobility measures constructed from the Current Population Survey (CPS), that distinguish between owners and renters, but are released on a yearly frequency only. This severely limits the information content of these data, especially when constructing business cycle statistics. However, according to CPS data the mobility rate among homeowners fell thirty percent during the period 2005-2009, which is very similar to drop in the home sales during the same period.

[6] Shimer (2007) modifies the series to adjust for a structural break in January 1994, caused by a change in CPS methodology. I would like to thank Ayşegül Şahin for sharing the adjusted series.

shocks that arise in housing markets.

Raw data. Panels A and B of Figure 2 display the real house price and the number of home sales, respectively, for the period since 1970. Both variables are upward trending and display cyclical fluctuations. Two major boom-bust episodes in home sales stand out. The first boom began around 1975 and was associated with a run-up in house prices. The subsequent bust in the number of home sales started around the time Paul Volcker initiated his disinflationary monetary policy, and was accompanied by a moderate decline in house prices. The second boom seems to have started around the year 2000, and resulted in a bust that started around 2005.[7]

The unemployment outflow hazard, the unemployment rate and the vacancy index are plotted in panels C, D, and E of Figure 2. The cyclical fluctuations in these variables seem much related to those in home sales during several episodes. An exception is the period from 2001 until 2004, when home sales increased but the outflow hazard fell. The unemployment rate reached its highest level during the two housing busts discussed above, and the increase in the unemployment rate during recent years was particularly sharp. Note that during the run-up in unemployment in the last year of the sample, vacancies no longer declined. The unemployment outflow hazard, however, declined to a level that is by far the lowest in the sample.

Cyclical components. The cyclical properties of the series are displayed in Table 1, which reports the volatilities of GDP, home sales, the unemployment rate, vacancies, and the unemployment outflow hazard, relative to the volatility of the real house price. Home sales are more volatile than the real house price. The real house price, in turn, is more volatile than GDP. The unemployment rate, vacancies and the unemployment outflow hazard are somewhat more volatile than home sales.

Table 1 also displays the correlations between the above mentioned variables. Home sales are positively correlated with the real house price. Moreover, both house prices and home sales are positively correlated with output, vacancies and the outflow hazard, and negatively with the unemployment rate. These patterns are consistent with the geographical mobility channel. In Appendix A.2, I document robustness of these results using an alternative, VAR-based comovement measure.

[7] At the very end of the sample, there is a brief spike in home sales, which seems related to the Home Buyer Tax Credit program, that applied to homes purchased between January 1, 2009 and May 1, 2010.

Conditioning on housing market shocks. Figure 3 shows the dynamic responses to a joint shock in house prices and home sales of one standard deviation.[8] A negative shock leads to significant declines in house prices, home sales and the unemployment outflow hazard. Consumer prices, stock prices, industrial production and the federal funds rate also fall, while the unemployment rate increases.

More insight in the effects of housing market disturbances on the labor market matching process can be obtained by imposing a minimal degree of additional structure. Suppose that unemployed workers and firms are matched according to a function of the form $f(n_{u,t}, v_t) = \mu n_{u,t}^{\eta} v_t^{1-\eta}$, where $f(n_{u,t}, v_t)$ is the number of matches, $n_{u,t}$ is the unemployment rate, v_t is the number of vacancies, η is the elasticity of matches with respect to the unemployment rate and μ is a scaling parameter. Given a value for η and responses for $n_{u,t}$ and v_t, one can evaluate the response of the unemployment outflow hazard *implied* by the matching function, i.e. $f(n_{u,t}, v_t)/n_{u,t}$. According to Petrongolo and Pissarides (2001), plausible values for η are between 0.5 and 0.7.[9]

Figure 4 plots the responses of the actual and the implied outflow hazard for the two extreme values of η. For $\eta = 0.7$, the overall decline of the implied outflow hazard is much smaller than the decline in the actual unemployment outflow hazard. For $\eta = 0.5$, the initial declines in the two variables are similar, but the actual decline in the outflow rate is more persistent than the decline predicted by the responses of unemployment and vacancies. These results provide evidence for a reduction in matching efficiency after negative housing shocks, although the evidence is particularly convincing for high values of η.

[8] The idea behind this shock is that I want to study shocks that arise in housing markets and subsequently spill over to the real sector. I consider a joint shock for the sake of a parsimonious presentation. In Appendix A.3 it is documented that the individual responses to a house price and a home sales shock are very similar.

[9] Running a simple OLS regression of the log outflow hazard on a constant and the log of the vacancy-unemployment ration results in an estimate for η of 0.57. The value of μ is irrelevant for this exercise.

3 General equilibrium model

3.1 Main features of the model

I now turn to the theoretical model. Three main ingredients allow this model to capture the geographical mobility channel. First, agents are geographically mobile. Mobility decisions are integrated into their intertemporal optimization problems and are affected by both aggregate and individual conditions.

The second ingredient is a financial friction on the side of households, in the spirit of Kiyotaki and Moore (1997), Iacoviello (2005) and Favilukis, Ludvigson, and Van Nieuwenburgh (2010). In their models, borrowing is limited by the value of the underlying collateral and debt contracts are renewed in each period. A consequence of this modeling choice is that a decline in the price of collateral affects the debt limits of all borrowers in the economy. But in reality, borrowers who do not refinance their loans are typically not affected when house prices fall.

A key innovation of my model is that collateral requirements apply only to *new* mortgages, which are taken out at the moment when an agent moves. For existing mortgages, debt is simply limited not to exceed the amount of the previous period.[10] Precisely this feature generates a decline in mobility when house prices fall. Consider for example a fall in house prices that is so large that borrowers' home equity levels shrink to zero. Without any wealth left, it becomes nearly impossible for agents to provide the downpayment required for a new mortgage loan, even a small one. However, when the agent decides not to move there is no renewed downpayment requirement, so the agent can sustain her current level of debt without problems. Consequently, moving is very unattractive in this situation.[11]

The final main ingredient of the model is a friction in the labor market. As in Pissarides (2000), unemployed agents search for vacancies and occasionally receive a job offer. But a fraction of those job offers can *only* be accepted if the agent moves,

[10] As a consequence, my constraint does not allow non-movers to increase debt after house prices increase. Allowing for home equity loans would introduce a nonlinearity that creates severe difficulties when solving the model. Note, however, that the key aspect of the proposed geographical mobility channel is that agents can protect their debt limits from *declines* in house prices by not moving.

[11] My model abstracts from mortgage default. But note that homeowners who default are likely to have difficulties in getting a new mortgage for an even longer period of time.

as commuting would be infeasible. When moving is sufficiently unattractive, e.g. because of difficulties in obtaining a new mortgage loan, the job offer is rejected and the agent continues searching for job offers.

3.2 Model description

The model economy is populated by a continuum of households of unit mass who consume housing and non-durables.[12] There are two types of households: impatient and patient households. In equilibrium, the impatient households borrow from the patient ones, but borrowing is restricted by a refinancing constraint. Each of the two representative households consists of a continuum of members, who are either employed or unemployed.[13] In each period, a certain fraction of the members moves and the household pays a fixed moving cost for each of those members. The desire to move depends on the degrees of satisfaction of members with their current locations, which are idiosyncratic and stochastic. So moving costs are only worth paying for those members who are sufficiently dissatisfied with their current locations.

Employment relationships are destroyed at an exogenous rate, after production has taken place during the period. A member whose job gets destroyed in period t can search for a job in the same period and may have a new job in period $t+1$ without becoming unproductive. If not, the member becomes unemployed in period $t+1$ and continues searching. The total number of meetings between workers and firms is determined by a standard matching function, depending on the aggregate number of job searchers and the number of vacancies. However, a fixed fraction of all meetings can *only* result in a productive relationship if the member moves. This captures the job offers from regions other than the one in which the worker resides.[14] When moving is sufficiently unattractive, the job offer is turned down and

[12] Households own the housing they consume. The possibility of renting is discussed in Subsection 5.1.

[13] This construct was introduced by Merz (1995) and was followed by others, including Gertler and Trigari (2009). According to this setup, agents are fully insured against fluctuations in consumption that arise from idiosyncratic shocks. What follows is a framework with a representative saver and a representative borrower.

[14] For reasons of simplicity, geographic locations are not explicitly modeled, although one could think of the model as one with a continuum of locations that are a priori identical to agents. But the proposed framework is consistent with two essential aspects of the geographical mobility channel, namely that (i) moving necessitates refinancing a mortgage and (ii) in some cases moving is required to accept a job offer.

the member remains unemployed.

To impatient households, there is one additional factor affecting mobility decisions, namely the effect on the borrowing capacity of the household. The fraction of debt that the household has to refinance is equal to the fraction of its members moving to a new location. After a fall in house prices, refinancing becomes unattractive, which creates a barrier to geographical mobility.

3.2.1 Impatient households

The impatient households maximize the following objective function:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left\{ \ln c_t + \alpha z_{h,t} \ln h_t + \kappa n_{u,t} + u_{lo,t} \right\}, \tag{1}$$

where β is their discount factor, c_t is non-durable consumption, h_t is the stock of housing, α is a housing preference parameter, $z_{h,t}$ is a housing preference shock, $n_{u,t}$ is the fraction of unemployed members, each generating a utility flow κ arising from time spent at home. Finally, $u_{lo,t}$ is a utility flow term that stems from the degree of satisfaction of the household members with their locations of residence, which will be specified below.

Consumption and borrowing decisions. Each period, households decide on the amount of non-durable consumption, housing and borrowing. In doing so, they are restricted by the following budget constraint:

$$c_t + p_{h,t} \left(h_t - h_{t-1} \right) + \zeta n_{m,t} + R_{t-1} d_{t-1} = \left(1 - n_{u,t} \right) y_t + d_t, \tag{2}$$

where $p_{h,t}$ is the house prices in units of non-durables, ζ is the fixed cost of moving a member, $n_{m,t}$ is the fraction of members that moves, R_t is the gross interest rate on debt to be repaid in period $t+1$, y_t is wage income per employed member, and d_t is the amount of debt.[15] So income consists of wage income, new debt and the sales value of the housing stock of the previous period, and income is spent on non-durable consumption, housing, moving costs, and servicing of old debt.

Debt is limited by a collateral constraint. In the literature on business cycle

[15] I limit the attention to loans with variable interest rates, so refinancing only involves the enforcement of a renewed collateral constraint.

models with collateral constraints, it is standard to assume that a household can borrow up to a fraction χ of the value of the house. The collateral constraint then reads $d_t \leqslant \chi p_{h,t} h_t$, with $1-\chi$ being the downpayment as a fraction of the value of the house. It is also standard to assume that all debt is refinanced in every period. Hence, a fall in house prices directly affects all collateral-constrained borrowing in the economy. In reality, however, only those households who take out a *new* mortage loan face a renewal of the downpayment requirement. Many homeowners simply roll over their existing mortgage contracts and for those households, the relevant constraint is $d_t \leqslant d_{t-1}$ (abstracting from amortizaton). Movers, however, are typically forced to take out a new mortgage loan. Accounting for these realities, I assume that the fraction debt that the households needs to refinance equals the mobility rate, which delivers the following constraint:[16]

$$d_t \leqslant n_{m,t} \chi p_{h,t} h_t + (1 - n_{m,t})\, d_{t-1}. \tag{3}$$

The important feature of this constraint is that the fraction of debt that is refinanced depends on the mobility rate, $n_{m,t}$. If all members move, that is when $n_{m,t} = 1$, all debt is refinanced and the constraint reduces to a standard collateral constraint, which has the unrealistic feature that all mortgage debt is refinanced every period. If none of the members moves, that is when $n_{m,t} = 0$, the household can borrow up the amount of the previous period, that is up to d_{t-1}.[17] For a given housing stock, h_t, a decline in the house price $p_{h,t}$ lowers $\chi p_{h,t} h_t$ relative to d_{t-1}. This makes moving less attractive. Note that in the steady state, the borrowing constraint reduces to a standard collateral constraint. Given the presence of patient households with a higher discount factor, this constraint binds in the steady state. [18]

[16] I abstract from home equity loans.

[17] Campbell and Hercowitz (2009) consider a model in which debt also evolves in a recursive way. In their model, however, the weight of old debt in the constraint depends on a fixed amortization rate. In my model, it depends on the *mobility* rate, which is a choice variable.

[18] In the simulations, one can verify that shocks are never large enough to bring the Lagrange multiplier on the borrowing constraint even close to zero. See also Iacoviello (2005) for a discussion on this issue.

The first-order conditions for the amount of housing and debt are:[19]

$$\frac{p_{h,t}}{c_t} = \frac{\alpha z_{h,t}}{h_t} + \beta E_t \frac{p_{h,t+1}}{c_{t+1}} + \lambda_{cc,t} n_{m,t} \chi p_{h,t}, \tag{4}$$

$$\frac{1}{c_t} = \beta E_t \left(\frac{R_t}{c_{t+1}} - \lambda_{cc,t+1} \left(1 - n_{m,t+1}\right) \right) + \lambda_{cc,t}. \tag{5}$$

Equation (4) is the first-order condition for the amount of housing consumed by the household. The right hand side is the *shadow value of housing*, which consists of three terms. The first term captures the direct utility gain derived from a marginal unit of housing. The second term is the utility derived from the discounted resale value of the house in the next period. The third term is proportional to the Lagrange multiplier of the borrowing constraint, $\lambda_{cc,t}$, the real house price $p_{h,t}$, and the mobility rate $n_{m,t}$. This term stems from the additional borrowing capacity that an extra unit of housing generates. If the borrowing constraint is not binding or if no member moves, this term reduces to zero. Equation (4) states that at the optimum, the shadow value of housing must be equal to the utility derived from $p_{h,t}$ marginal units of non-durables. Equation (5) is the Euler equation for debt. A binding borrowing constraint introduces a wedge in this equation. The second term within the conditional expectation represents the fact that of the new debt taken on in period t, only a fraction $n_{m,t+1}$ will be refinanced in period $t+1$. The remaining debt is rolled over to period $t+2$.

Location preferences. Geographical mobility is an essential feature of the model. Naturally, mobility decisions are affected by a number of factors. The focus of this paper is on considerations regarding employment and borrowing. However, mobility decisions also depend on more private factors, such as changes in family composition or changes in the degree of satisfaction with the neighborhood. In order for the model to generate realistic overall mobility rates, these considerations need to be taken into account.

The setup is as follows. For each individual member j, an idiosyncratic location satisfaction shock $\varepsilon_{j,t}$ is observed during period t.[20] This shock represents the

[19] The optimality conditions for the impatient households are derived in Appendix B.

[20] For simplicity, I assume that this location satisfaction shock is i.i.d. accross time and members.

private factors that affect how willing somebody is to move.[21] For members that do not move, the realization of $\varepsilon_{j,t}$ is received as a utility flow, while each mover generates a fixed utility flow ψ instead.[22] So for members with a low realization of $\varepsilon_{j,t}$, moving is relatively attractive. For those members that receive a "long-distance job offer", moving has an additional benefit, namely that it enables them to escape unemployment.

The optimal mobility decision implies a cutoff level for the location satisfaction shock. If the realization of this idiosyncratic shock is below the cutoff level, the member moves, while the member does not move if the realization is above the cutoff level. So the cutoff level represents the location satisfaction of the marginal mover, being exactly indifferent between moving and not moving. Although there may in principle be different cutoff levels for different agents, depending on individual characteristics such as wealth and labor income, the advantage of the framework with full insurance among household members is that there will be only two such values: one for members with a long-distance job offer, denoted by $\overline{\varepsilon}_{do,t}$, and one for those without such an offer, denoted by $\overline{\varepsilon}_t$. Let $F(\cdot)$ be the cumulative distribution function of the shock. Thus, $F(\overline{\varepsilon}_{do,t})$ is the mobility rate among members with a long-distance job offer and $F(\overline{\varepsilon}_t)$ is mobility rate among the members without such a job offer. It follows from this setup that the total location satisfaction utility term in equation (1) is given by

$$u_{lo,t} = n_{do,t}\left[\psi F(\overline{\varepsilon}_{do,t}) + \int_{\overline{\varepsilon}_{do,t}}^{\infty} \varepsilon dF(\varepsilon)\right] + (1 - n_{do,t})\left[\psi F(\overline{\varepsilon}_t) + \int_{\overline{\varepsilon}_t}^{\infty} \varepsilon dF(\varepsilon)\right],$$

where $n_{do,t}$ is the fraction of members with a long-distance job offer.

Mobility decisions. The two cutoff levels that determine the mobility rate are chosen optimally by the household. This decision is taken at the beginning of the period, jointly with consumption and borrowing decisions.[23] The corresponding

[21] Stein (1995) and Ngai and Tenreyro (2009) consider similar shocks.

[22] Alternatively, one could consider a model in which ψ is stochastic and ε is fixed, or a model in which both variables are stochastic. However, this would result in an observationally equivalent models under the distributional assumptions made in this paper. What truly matters is the distribution of the *difference* between $\varepsilon_{j,t}$ and ψ.

[23] When mobility decisions are taken, it is known what members have what type of job offers, because offers were received at the end of the previous period.

first-order conditions are:

$$\psi = \frac{\zeta}{c_t} + \overline{\varepsilon}_t - \lambda_{cc,t}\left(\chi p_{h,t} h_t - d_{t-1}\right), \tag{6}$$

$$\overline{\varepsilon}_{do,t} - \overline{\varepsilon}_t = \frac{y_t}{c_t} - \kappa + \left(1 - \rho_u\right) G_t. \tag{7}$$

Equation (6) is the first-order condition for $\overline{\varepsilon}_t$, the moving cutoff for members *without a long-distance job offer*. The left- and the right-hand side of this equation are, respectively, the benefits and costs of moving, for a member who is exactly indifferent between moving and staying. On the left-hand side, ψ is the utility flow that is received for being at a new location. On the right-hand side, the first term is the utility loss arising from paying the moving cost ζ. The second term, $\overline{\varepsilon}_t$, is the utility flow received when staying at the current location, which is foregone when moving. The third term arises from the effect of mobility on the borrowing capacity of the household and can be either positive or negative. If the borrowing limit on old loans exceeds the borrowing limit on new loans, that is when $\chi p_{h,t} h_t < d_{t-1}$, there is an additional cost to mobility.

Equation (7) determines the cutoff level for members *with a long-distance job offer*, $\overline{\varepsilon}_{do,t}$, relative to the cutoff for members without such a job offer. The difference between $\overline{\varepsilon}_{do,t}$ and $\overline{\varepsilon}_t$ is determined by two factors. The first factor is the difference in utility gains from the wage income of an employed member, y_t/c_t, and the utility flow from unemployment, κ. The second factor stems from a dynamic composition effect. If more members with a long-distance job offer move, more members flow into employment, and this positively affects the fraction of members that is employed in future periods. This effect is captured by G_t, which is defined in Appendix B.

Equations (6) and (7) are essential in understanding the direct mechanism through which fluctuations in house prices affect real activity. Equation (6) reveals a direct and inverse relation between house prices and mobility among members without a long-distance job offer. Ceteris paribus, a decline in the house price, $p_{h,t}$, must be offset by a decline in the cutoff, $\overline{\varepsilon}_t$. Recall that members move when their individual location satisfaction is below the cutoff. Thus, with a lower cutoff, mobility declines. Equation (7) shows that, ceteris paribus, a decline in $\overline{\varepsilon}_t$ also lowers the cutoff for members *with* a long-distance job offer, $\overline{\varepsilon}_{do,t}$. When mobility of members with long-distance job offers declines, a larger fraction of job offers is turned down,

pushing up aggregate unemployment.

Flow equations. The mobility rate among household members, $n_{m,t}$, follows from the mobility cutoffs, and is given by

$$n_{m,t} = n_{do,t} F\left(\overline{\varepsilon}_{do,t}\right) + \left(1 - n_{do,t}\right) F\left(\overline{\varepsilon}_{t}\right). \tag{8}$$

The fraction of members with a long-distance job offer, $n_{do,t}$, is determined by

$$n_{do,t} = \omega \widehat{g}_{t-1} n_{s,t}. \tag{9}$$

Here, $\widehat{g}_t$ is the probability that an unemployed member meets a firm and gets a job offer, ω is the fraction of all meetings in which a member is required to move to accept the offer, and $n_{s,t}$ is the fraction of members that is searching for a job, which is:

$$n_{s,t} = n_{u,t} + \rho_u \left(1 - n_{u,t}\right), \tag{10}$$

where ρ_u is the exogenous job destruction rate. So the group of job searchers consists of the members that are unemployed and the employed members who just became obsolete at their current job. The fraction of unemployed members is equal to the fraction of job searchers of the previous period that did not receive a job offer, or that did receive a job offer but rejected it because moving is too unattractive:

$$n_{u,t} = n_{s,t-1} \left(1 - \widehat{g}_{t-1} + \omega \widehat{g}_{t-1} \left(1 - F\left(\overline{\varepsilon}_{do,t}\right)\right)\right). \tag{11}$$

3.2.2 Patient households

Patient households are the same as impatient households, except that their discount factor, γ, is higher than the discount factor of the impatient households, β. In equilibrium, patient households therefore lend to the impatient households and as a consequence the patient households are not borrowing constrained. Also, patient households own the firms and receive their profits.

For patient households, house prices are not directly relevant for mobility decisions because they do not borrow. Let the variables of the patient households be denoted by a tilde. The first-order condition for the moving cutoff for members of

the patient household without a distant job offer, $\widetilde{\overline{\varepsilon}}_t$, is:

$$\psi = \frac{\zeta}{\widetilde{c}_t} + \widetilde{\overline{\varepsilon}}_t. \tag{12}$$

3.2.3 Labor market

Let aggregate variables be denoted by a hat, and let ν be the share of impatient households in the total population. The aggregate unemployment rate, and the aggregate number of job searchers are, respectively,

$$\widehat{n}_{u,t} = \nu n_{u,t} + (1-\nu)\,\widetilde{n}_{u,t}, \tag{13}$$

$$\widehat{n}_{s,t} = \nu n_{s,t} + (1-\nu)\,\widetilde{n}_{s,t}. \tag{14}$$

The labor market is characterized by a matching friction. The aggregate number of meetings between firms and job candidates, $\widehat{m}_t$, is a Cobb-Douglas function of the total number of job searchers and the aggregate number of vacancies, $\widehat{v}_t$:

$$\widehat{m}_t = \mu \widehat{n}_{s,t}^{\eta} \widehat{v}_t^{1-\eta}, \tag{15}$$

where η is again the elasticity parameter, and μ the scaling parameter. The probability that a job searcher meets with a firm is $\widehat{g}_t \equiv \frac{\widehat{m}_t}{\widehat{n}_{s,t}}$ and the probability of meeting with a worker for a firm with a vacancy is $\widehat{g}_{f,t} \equiv \frac{\widehat{m}_t}{\widehat{v}_t}$.

3.2.4 Firms

Firms that are matched to a worker produce $z_{a,t}$ per period, where $z_{a,t}$ follows an exogenous productivity process with a steady-state level equal to one. The wage is simply a share ξ of total revenues, that is $y_t = \xi z_{a,t}$. The firm receives the remaining share $1-\xi$.[24] Since firms are owned by the patient households, they discount future profits using the stochastic discount factor of those households. To firms, the asset

[24] I deviate from the more standard assumption that firms and workers bargain over the surplus that is created by an employment relationship. Instead, I assume that firms post wage contracts in which the worker gets a fixed fraction of the revenues. This setup makes the model tractable but also seems reasonable, given that in this model the total surplus of the match is affected by the utility derived from mobility. It does not seem very plausible that firms would be able to observe the entire surplus and engage in bargaining over it.

value of a match, V_t, is:

$$V_t = (1-\xi)\, z_{a,t} + (1-\rho_u)\, E_t \widetilde{\Lambda}_{t,t+1} V_{t+1}, \tag{16}$$

where $\widetilde{\Lambda}_{t,t+1}$ is the stochastic factor of the patient households, that is, $\widetilde{\Lambda}_{t,t+1} \equiv \gamma \frac{\widetilde{c}_t}{\widetilde{c}_{t+1}}$. Firms that search for employees pay a vacancy cost ϑ per period. Free entry of firms in the goods market implies that the vacancy cost equals the expected benefit to the firm of posting a vacancy:

$$\vartheta = \widehat{g}_{f,t} \left(1 - \omega + \omega \frac{\nu n_{s,t}}{\widehat{n}_{s,t}} F\left(\overline{\varepsilon}_{do,t+1}\right) + \omega \frac{(1-\nu)\, \widetilde{n}_{s,t}}{\widehat{n}_{s,t}} F\left(\widetilde{\overline{\varepsilon}}_{do,t+1}\right) \right) \widetilde{\Lambda}_{t,t+1} V_{t+1}. \tag{17}$$

The term between large brackets in the free-entry condition is the fraction of meetings that is unsuccessful because the worker is unwilling to move. Aggregate firm profits are given by $\widehat{\Pi}_t = (1 - \widehat{n}_{u,t})\,(1-\xi)\, z_{a,t} - \vartheta \widehat{v}_t$.

3.2.5 Exogenous processes

The housing preference shock and the productivity shock are common to all agents and evolve according to the following laws of motion:

$$\begin{aligned} \ln z_{h,t} &= \rho_h \ln z_{h,t-1} + \varepsilon_{h,t}, \\ \ln z_{a,t} &= \rho_z \ln z_{a,t-1} + \varepsilon_{a,t}, \end{aligned}$$

where $\varepsilon_{h,t}$ and $\varepsilon_{a,t}$ are i.i.d. innovations that are normally distributed, with mean zero and standard deviations σ_h and σ_a, respectively.

3.2.6 Equilibrium

The supply of the total stock of housing is fixed and normalized to one. The aggregate supply of debt is zero. The corresponding market clearing conditions are:[25]

$$\nu h_t + (1-\nu)\, \widetilde{h}_t = 1, \tag{18}$$

$$\nu d_t + (1-\nu)\, \widetilde{d}_t = 0. \tag{19}$$

[25] A formal definition of the equilibrium is provided in Appendix B.2.

3.3 Calibration

The model is calibrated to U.S. data. The frequency is monthly. Several parameters are calibrated to pin down essential steady-state properties of the model and one parameter is calibrated to match the volatility of home sales.

3.3.1 Steady-state targets

The calibration procedure targets six steady-state properties of the model. First, the aggregate unemployment rate in the steady state is five percent. Second, the steady-state aggregate mobility rate is 0.65 percent per month. This corresponds to an annual mobility rate of 7.5 percent, as measured for US homeowners using data from the Current Population Survey (CPS) for the period 2000-2005. Third, the steady-state mobility rate due to members with long-distance job offers is 0.10 percent per month. This choice is based on data from the CPS for the period 2000-2005 as well. On average, about 15 percent of the owners who had moved, indicated that the move was primarily for employment reasons. Fourth, the steady-state value of housing wealth is 140 percent of annual output.[26] Fifth, the credit-constrained households consume the same amount of housing in the steady state as the patient households.[27,28] Sixth, the probability that a vacancy is filled is 0.34 in the steady state. This implies a quarterly probability of 0.71, as in den Haan, Ramey, and Watson (2003).

3.3.2 Parameter values

The parameter values are presented in Table 2 and are discussed below.

Preferences and moving technology. The calibration of the discount factors for patient an impatient households follows Iacoviello and Neri (2009). The discount factor of the patient households, γ, is set to 0.9975, which implies a steady-state real

[26] This choice follows Iacoviello (2005) and is based on data from the Flow of Funds.

[27] This choice is supported by data from the American Household Survey 2007. For households with the very lowest downpayment ratios (up to five percent), the median home value is below the median home value for the total sample of homeowners. However, the median home value for households with a downpayment ratio between six and twenty percent is higher than the median for the total sample.

[28] Note that when the steady-state value of the housing stock owned by patient households is pinned down, the borrowing constraint determines the steady-state level of mortgage debt. The steady-state level of aggregate debt relative to aggregate income is 22.4 percent.

interest rate of about three percent per annum. The discount factor of the impatient households, β, is set to 0.9899.

The weight of housing in the utility function, α, is set differently for patient and impatient households. The values follow from the steady-state targeting procedure described above. The value for the patient and impatient households are 0.043 and 0.139, respectively. The relatively high value for the impatient households is a direct consequence of the requirement that both types of households consume the same amount of housing in the steady state.[29]

The idiosyncratic location satisfaction shock is calibrated to be normally distributed with mean zero and standard deviation σ. The standard deviation is important for the volatility of geographical mobility. A lower value of σ means that preferences for household members' current locations are less spread out. I choose the value of σ to be such that in a version of the model with only technology shocks, mobility is about 2.4 times as volatile as the real house price, corresponding to the relative volatility found in the data for home sales. This is achieved by setting $\sigma = 3.5$.

The roles of the moving cost parameter, ζ, and the utility flow from moving, ψ, are similar. This can be seen from Equation (6), the first-order condition for the mobility cutoff $\bar{\varepsilon}_t$. The parameter ζ is meant to capture physical costs of moving, such as transaction costs and fees for real estate agents. Unlike the utility flow ψ, the effect of the physical cost depends on the marginal utility of consumption, which varies in response to shocks. Following Stokey (2009), ζ is set to eight percent of the steady-state value of a unit of housing. The parameter ψ is used in the steady-state targeting procedure and its value is -7.144.

Labor market. The elasticity of matches with respect to the number of job searchers, η, is set equal to 0.6. The job destruction probability, ρ_u, is set to 0.035. These values are within the range of standard values considered in the literature and follow Gertler and Trigari (2009). The values of the scaling parameter in the matching function, μ, and the vacancy cost, ϑ, follow from the steady-state targeting procedure, and are 0.545 and 0.181, respectively. The calibration implies that

[29] Note that impatient households discount the future resale benefits from housing more heavily than patient households. This would lead them to consume much less housing in the steady state than the patient households. As a consequence, the steady state level of debt of the impatient households would be unrealistically low.

in the steady state 1.86 percent of output is devoted to vacancy costs.

The fraction of long-distance job offers, ω, is difficult to calibrate as there is no direct equivalent in the data. My strategy is to set ω to one third and check for robustness. Alternative values for ω turn out to generate very similar results.[30] The parameter ξ controls the fraction of the revenues that flows to workers in the form of wages. I assume that accounting profits of the firms are two percent, that is, $\xi = 0.98$. This choice is in line with typical calibrations of matching models, see Hornstein, Krusell, and Violante (2005).[31]

The parameter κ measures the utility flow received per unemployed worker, and is one of the parameters that are used to match the steady-state targets. Its value is -4.428. Thus, κ can be thought of as an unemployment stigma. The negative value for κ contrasts with standard models with Nash bargaining, in which the parameter is typically positive-valued. However, in those models, the main role of κ is to determine the surplus from a match, affecting the incentives for firms to post vacancies. In my model, κ is only relevant in that it affects the incentives for unemployed workers to accept long-distance job offers.[32] Therefore, a low value of κ ensures that in the steady state a realistic fraction of workers moves for employment reasons.[33]

Credit frictions. The fraction of credit-constrained households, ν, and the collateral requirement parameter, χ, are potentially very important for the dynamics of the model. In my benchmark calibration, I set ν and χ such as to represent average values over the period since 1970. In order to study the effects of structural changes in mortgage markets, I consider two alternative calibrations, denoted by "low-leverage economy" and "high-leverage economy".

Using data from the Survey of Consumer Finances (SCF), Campbell and Her-

[30] The reason is that the steady-state fraction of the workers accepting long-distance job offers is a direct target of the calibration procedure.

[31] Unlike standard models, my results are not very sensitive to this aspect of the calibration, because I do not use Nash bargaining.

[32] Note that κ only enters the first-order condition for the mobility cutoff for workers with a long-distance offer. Thus, in a version of the model in which unemployed workers would be able to accept all job offers without moving, the value of κ would be irrelevant, provided that it is low enough to ensure that unemployed workers always accept a job offer.

[33] An important reason why the value of κ has to be so low, is that in the model there are complete insurance markets, and therefore unemployed agents do not suffer from lower levels of consumption than employed agents.

cowitz (2009) document that during the period 1983-2001, the average equity stake in newly purchased homes declined from 22.6 percent to 16.4 percent.[34] In the model, $1-\chi$ is the equity stake in the steady state. In my benchmark calibration, I set χ equal 0.8. For the low- and high-leverage economy, χ is set equal to 0.75 and 0.85, respectively.

The fraction ν is meant to capture the real-world fraction of borrowing-constrained households in the *total* population (including renters). Data from the SCF show that during the period 1989-2007, the fraction of households with a mortgage or home equity loan increased from 39 percent to 46 percent. But of course, not all of these households are actually *constrained* by a borrowing limit. In my benchmark calibration, ν is set equal to 0.2.[35] In the low- and high-leverage economies, I set $\nu = 0.15$ and $\nu = 0.25$, respectively, capturing an increase in the number of households that is eligible for a mortgage.

Exogenous shock processes. The calibration of the persistence parameter of the productivity process, ρ_a, follows Kydland and Prescott (1982). The persistence parameter of the housing preference process, ρ_h, is difficult to measure directly in the data. But since there are few a priori reasons to expect that the housing preference process is either more or less persistent than the productivity process, I set ρ_h equal to ρ_a.

4 Model results

The model is solved using a first-order perturbation method and then simulated. Three types of simulations are analyzed. First, I simulate the model with random sequences of productivity shocks and calculate standard business cycle statistics, which are compared to those found in the data. Second, I discuss the dynamic responses to one-time shocks in productivity and housing preferences. Finally, I consider a series of productivity and housing preference shocks that is chosen such that the model replicates data series for output and real house prices, over the period 1970 - 2010. The main purpose of this experiment is to investigate to what extent

[34] The average is over home purchases with an equity stake of at most fifty percent.

[35] Data from the American Housing Survey (AHS) show that in 2007, 62 percent of all mortgagors had put in a downpayment of 15 percent or less at the time they purchased their home. Earlier observations are not available.

the model can explain the puzzling dynamics of unemployment and vacancies during the aftermath of the Great Recession.

4.1 Business cycle statistics

The model of this paper is very stylized in many respects. For example, the model does not feature capital investment or nominal rigidities. Although the simplicity of the model helps to highlight its essential mechanisms, it reduces the extent to which the business cycle properties of the model can be expected to match the data. Moreover, the most interesting application of the model seems to be to simulate episodes of large falls in house prices, which do not occur very frequently. Nonetheless, it seems important to know how well this simple model can explain "regular" business cycles. I therefore analyze the business cycle properties of a version of the model with only productivity shocks.

Volatilities and correlations implied by the model are displayed in Table 3. These numbers can be compared to the volatilities found in the data (Table 1). The predicted volatilities of unemployment, vacancies and the outflow hazard are very low compared to data. This is a general problem of labor market matching models, as emphasized by Shimer (2005).

For all correlations, the model predicts the correct sign. Moreover, the correlation between vacancies and the outflow hazard, and between the unemployment rate and the outflow hazard, are comparable to their data equivalents. However, there are also quantitative discrepancies between the model and the data. The correlation between the unemployment rate and vacancies is -0.70, which is less negative than the correlation coefficient of -0.90 found in the data. Most other correlations are much stronger than in the data. For example, the correlation between house prices and the unemployment rate is -0.86 in the model, whereas the correlation found in the data is -0.44.

4.2 Dynamic responses

Productivity shock. Responses to a sudden one-percent decline in productivity are displayed in Figure 5, for several calibrations. First consider the benchmark calibration. After a fall in productivity, output and house prices fall. As in standard business cycle models with search and matching frictions in the labor market, there

is a decline in vacancies and in the (average) unemployment outflow hazard, and an increase in the unemployment rate. The mobility rate falls after the decline in productivity.

To understand the effects of refinancing constraints, consider the responses for an economy in which there are no impatient households (no borrowers). In such an economy, all households have the same discount factor and there is no debt in equilibrium. Thus, refinancing constraints are irrelevant. The responses for this economy are plotted in Figure 5 as well. Without credit-constrained households, the decline in vacancies is very similar to the decline in vacancies in the benchmark model. The unemployment rate, however, does not increase as much as in the benchmark model. Thus, the presence of credit-constrained households implies a somewhat flatter Beveridge curve. This is directly related to the fact that the drop in the mobility rate is also much less pronounced than in the benchmark economy. The declines in output and house prices, however, are quite similar across the two versions of the model. This indicates that these declines are mainly driven by the direct effects of the fall in productivity.

What are the consequences of structural changes in mortgage markets for the sensitivity of the economy to productivity shocks? Figure 5 plots the responses for the "low-leverage" and the "high-leverage" economies. Financial development clearly causes the mobility rate to be more sensitive to productivity shocks. For the other variables, the responses are very similar across the two economies.

Housing preference shock. How do shocks that originate in housing markets affect the real economy? Figure 6 displays the responses to a negative housing preference shock, generating a house price decline of about one percent on impact. First consider the benchmark calibration. Consistent with the VAR evidence, the model predicts a joint decline in house prices, mobility, output, vacancies and the unemployment outflow hazard, and an increase in the unemployment rate.

Interestingly, the increase in the unemployment rate is larger and much more persistent than the decline in vacancies. For example, eight quarters after the shock, the number of vacancies has nearly returned to its steady-state level, but the unemployment rate is still well above its steady-state level.[36]

[36] In a quantitative sense, the effects on real activity are very modest for a one percent fall in house prices. For example, the maximum increase in the unemployment rate is 0.028 percentage points.

Persistence is also observed in the decline of the mobility rate, which is much more prolonged than the fall in house prices. Six years after the shock, house prices have almost fully recovered, but the mobility rate is still one percent below its steady-state level.

What are the effects of collateral constraints? Figure 6 shows that in the economy without credit-constrained households, the decline in house prices is smaller than in the benchmark economy. But more importantly, real activity variables do not respond *at all* to the shock. To see why, note that the first-order condition for housing of the patient households is given by:

$$\frac{p_{h,t}}{\widetilde{c}_t} = \frac{\alpha z_{h,t}}{\widetilde{h}_t} + \gamma E_t \frac{p_{h,t+1}}{\widetilde{c}_{t+1}}. \tag{20}$$

In the absence of credit-constrained households, this is the only model equation in which the house price enters.[37] This implies that house prices and housing preferences are irrelevant for equilibrium allocations.[38]

Now consider the low-leverage and the high-leverage economies. Figure 6 makes clear that structural innovations in mortgage markets substantially increase volatilities. In particular, the maximum responses of the unemployment and mobility rates in the high-leverage economy are much larger than in the low-leverage economy. In contrast, the differences in the response of the real house price are small. Thus, real activity has become more sensitive to fluctuations in house prices as the amount of leverage in the economy increases, but the feedback on house prices seems limited.

4.3 Great Recession experiment

This subsection discusses how well the model can explain observed dynamics of unemployment and vacancies, and in particular the recent flattening of the Beveridge curve, as highlighted in Figure 1. For this purpose, the following experiment is conducted. Using data series on the real house price and GDP for the period 1970

Thus, the model does not only generate little volatility in the unemployment rate conditional on productivity shocks, but also conditional on housing preference shocks. However, in the next subsection it will be shown that for large swings in house prices, the effects are substantial.

[37] Since the aggregate supply of housing is fixed, the house price drops out of the budget constraint.

[38] Note that one could simply remove Equation (20) from the system of equilibrium conditions, since one would loose one equation and one endogenous variable.

- 2010, I back out realizations of the innovations to the productivity and housing preference shock processes.[39] This is simply done by inverting the equilibrium laws of motions for output and the real house price.[40] The model is then simulated, using the shocks obtained from the inversion procedure.

How well can the model account for the flattening of the Beveridge curve observed in recent years? Panel A of Figure 7 plots the unemployment rate versus vacancies for the high-leverage economy. For the period 1970-2008, there is a strong negative comovement between unemployment and vacancies, although not as strong as in the data.[41] For 2009, the model predicts that the Beveridge curve becomes essentially flat, or even somewhat upward sloping, precisely as observed in reality (see Figure 1). The unemployment rate rises to about twelve percent above its long-run level, which corresponds to an increase in the unemployment rate of 0.6 percentage points. This observation reflects the fact that the model fails to generate sufficient volatility in unemployment.

What has been the role of the bust in house prices in recent years? Panel A of 7 plots the values of unemployment and vacancies in the counterfactual simulation.[42] With housing preference shocks shut off from 2005 onwards, the behavior of unemployment and vacancies is much less extreme. In particular, the unemployment rate only increases to seven percent above its trend level. Thus, the effects of shocks arising in housing markets are substantial.

The simulation results for the low-leverage economy are plotted in Panel B of Figure 7. The model generates a correlation between the unemployment rate and vacancies of -0.91, which is very close to the correlation coefficient of -0.90 observed in the data. Thus, for the overall sample, this version of the model may be more appropriate than its high-leverage counterpart. At the same time, the low-leverage version is less successful in explaining the sharp increase in the unemployment rate

[39] Data were logged and linearly detrended. Quarterly GDP data were converted into monthly data by linear interpolation.

[40] I assume that at the beginning of the sample, just prior to 1970, the economy was at the steady state. The results are not very sensitive with respect to this assumption, especially given that my focus is on the realizations in recent years.

[41] The correlation between the unemployment rate and vacancies is -0.83, whereas it is -0.9 in the data.

[42] Appendix B.3 presents the counterfactual evolutions of house prices and output, and shows that without preference shocks, house prices decline only mildly during last 5 years of the sample.

during the Great Recession. Moreover, the flattening of the Beveridge curve during 2009 is not as pronounced as in the high-leverage economy.

4.3.1 Sticky wages

A clear failure of the model is its inability to generate realistically large volatilities of the unemployment rate and vacancies. In this subsection, I apply the fix proposed by Hall (2005) by considering an alternative version of the model with a constant wage. In particular, I redo the Great Recession experiment, setting the wage equal to its steady-state value in the benchmark model.[43]

Figure 8 plots the Beveridge Curve predicted by the model under the assumption of sticky wages. Again, the model predicts a flattening of the Beveridge curve in 2009, which is particularly pronounced in the high-leverage economy. Comparing panel A of this figure to the data plotted in Figure 1 makes clear that this version of the model predicts an increase in the unemployment rate that is even somewhat larger than the one observed in the data. The contribution of the housing preference shocks to the total increase in the unemployment rate is about eleven percent. But recall that even with only productivity shocks house prices would fall, so this number does not capture the full effect of the fall in house prices on unemployment.

5 Understanding the interactions between housing and labor markets

In the presence of refinancing constraints, outcomes in housing markets are not just passively determined by economic conditions, but have an active role in shaping real economic outcomes. The purpose of this subsection is to better understand this interaction. I focus on housing preference shocks, because this allows me to study the roles of credit frictions and mobility in isolation. I will answer three main questions that arise naturally from the analysis of dynamic responses in Subsection 4.2. First, why are the mobility effects propagated over time? Second, why do vacancies fall following a negative housing preference shock? And finally, what determines feedback effects of fluctuations in real activity on house prices? Subsection 5.1 discusses extensions of the model with options that would allow agents to move

[43] In the simulation, the probability of finding a job and of filling a vacancy is always between zero and one, and neither firms or employees have incentives to break up a match.

without refinancing their mortgages.

5.0.2 Sources of propagation

An interesting model prediction is that the responses of mobility and unemployment to a housing preference shock are more persistent than the response of house prices (Figure 6). But what drives this propagation?

Recall that Equation (6), the first-order condition for the moving cutoff for impatient households, is essential in determining fluctuations in the mobility rate. This equation relates the cutoff, $\overline{\varepsilon}_t$, directly to the real house price, $p_{h,t}$, but also to the stock of housing owned by impatient households, h_t, and their debt level, d_{t-1}.[44] Panel A of Figure 9 plots the responses of h_t and d_{t-1} to a negative housing preference shock, and shows that the reduction in h_t is larger than the fall in d_{t-1}. More importantly, the gap widens during the first four years, indicating that impatient households gradually become poorer, and thus less capable of providing downpayments.

To obtain a clearer view on the roles of $p_{h,t}$, h_t and d_{t-1} in driving propagation effects, consider the response of the moving cutoff $\overline{\varepsilon}_t$, which is plotted in Panel B of Figure 9.[45] Note that this response can be reconstructed from the responses of the other variables in Equation (6). As a mechanical exercise, I reconstruct the response of $\overline{\varepsilon}_t$, but consider two variations. First, I keep all other variables in Equation (6) except for $p_{h,t}$ at their steady-state levels. This reconstructed response is also plotted in Panel B of Figure 9. Initially, the fall in the cutoff is about as large as observed in the original response. However, the reconstructed response is much less persistent than the original one and it is also not hump-shaped. So clearly, propagation effects are not driven by the real house price itself. Second, I repeat the exercise, but no keep all variables except for $p_{h,t}$, h_t and d_{t-1} at their steady-state levels. This reconstructed IRF is quite similar to the original one and in particular, it shows a persistent and hump-shaped decline of the cutoff. Thus, the fall in housing assets owned by the impatient households is essential in driving the propagation effects.

Why do impatient households sell housing stock *in equilibrium*, following a nega-

[44] Equation (6) also relates the cutoff to the marginal utility of consumption, $\frac{1}{c_t}$, and the Lagrange multiplier on the borrowing constraint, $\lambda_{cc,t}$. But fluctuations in these variables are of secondary importance in terms of understanding the main sources of propagation, as will be discussed below.

[45] The shape of this response is very similar to one of the overall mobility rate.

tive preference shock?[46] Note that Equation (4), the first-order condition for housing of the impatient households, can be rewritten as:

$$\frac{p_{h,t}}{c_t} = E_t \sum_{k=0}^{\infty} \beta^k \frac{\alpha z_{h,t+k}}{h_{t+k}} + E_t \sum_{k=0}^{\infty} \beta^k \lambda_{cc,t+k} n_{m,t+k} \chi p_{h,t+k}, \tag{21}$$

and recall that the right-hand side of this equation is the shadow value of housing. The first term on the right-hand side is the present value of all future utility delivered by a marginal unit of housing. The second present value arises from the fact that housing serves as collateral for loans. Figure 10 plots the response of the shadow value of housing for both types of households and shows that there is a larger fall for impatient households.

To understand why incentives to buy housing decrease more for those who are credit-constrained, note the following. For the patient households, the second present value in the first-order condition for housing always equals zero, because for them the borrowing constraint never binds (i.e., $\widetilde{\lambda}_{cc,t} = 0$ in each period t). For impatient households, however, there is a decline in the second present value following a negative housing preference shock, because the house prices, $p_{h,t}$, and the mobility rate, $n_{m,t}$, both decline.[47] Figure 10 also plots the response of the shadow value for impatient households, reconstructed from Equation (21), but with the second present value is kept at its steady-state level. Without the effects of the refinancing constraint, the decline in the shadow value is much more similar for patient and impatient households.

The intuition behind the propagation effects is that as house prices decline, a unit of housing provides fewer collateral services, creating disincentives for credit-constrained households to hold housing stock. Moreover, these households expect to move less, and thus rely more on old debt. For those reasons, they gradually decumulate housing assets. But this means that by the time house prices have recovered, they are poorer than before the shock. The patient households, who purchased additional housing stock, have become richer.[48]

[46] Recall that the aggregate stock of housing is fixed, so any reduction in housing stock owned by the impatient households must be absorbed by the patient households.

[47] The Lagrange multiplier on the refinancing constraint, $\lambda_{cc,t}$, increases following the shock, but this effect does not dominate.

[48] In the model, low mobility does not deter trade in housing stock. In reality, it is difficult for

5.0.3 The role of vacancies

Why is there a fall in vacancies following a negative housing preference shock? This decline is related to the fact that the patient households purchase housing stock from the impatient households. In order to finance these purchases, the patient households cut back on non-durable consumption expenditures, implying a decrease in their stochastic discount factor, $\widetilde{\Lambda}_{t,t+1}$. Since the firms are owned by the patient households, a decrease in $\widetilde{\Lambda}_{t,t+1}$ implies that the benefits from posting a vacancy are more heavily discounted. Therefore, vacancy posting decreases. This effect becomes clear when observing the free-entry condition, Equation (17). The intuition is that in order to take advantage of the increase in returns on housing, patient households reduce their investments in vacancies.

5.0.4 Feedback on house prices

An advantage of the DSGE framework adopted in this paper is that one cannot only study the effects of house prices on real activity, but also feedback effects. However, it turns out that the response of the house price is quite similar across simulations with small and large mobility effects (Figure 6). To understand why feedback effects are quantitatively limited, note that the first-order condition for the patient households, Equation (20), needs to be satisfied in any of the model versions. This equation makes clear that all possible feedback effects of real activity on house prices must operate through the marginal utility of non-durable consumption of the patient households. Whereas impatient households suffer income losses when job offer rejection increases, patient households are not directly affected.[49]

5.1 Additional options: renting, default and portable mortgages

A limitation of the model is that it does not explicitly feature the option to either rent a home or default on a mortgage. Possibly, the use of these options is increased

non-movers to decrease their housing stock beyond cutting back on maintanance. But in reality there do exist other assets in which non-movers can trade much more flexibly. For example, households can take on credit card debt or decumulate savings that were initially intended for the downpayment needed for a future home purchase.

[49] An increase in job offer rejection affects patient households via a decline in firm profitability, but this effect is small.

after a bust in house prices, which may dampen the adverse effects on labor market outcomes. Exercising the rental or default option, however, has long-lasting costs for credit-constrained homeowners. Defaulting on a loan hampers the ability to borrow for a prolonged period of time. Becoming a renter has disadvantages as well, since renters typically cannot fully customize their homes to accommodate their individual tastes. They also may forego tax advantages and even intrinsic felicity from owning a home. Moreover, the decision to exit homeownership is difficult to reverse in situations in which a household has lost a substantial fraction of its wealth due to a fall in house prices. Haurin and Rosenthal (2004) estimate a duration model of homeownership using data from the national longitudinal survey of youth. They report that a lower level of home equity actually *reduces* the probability of exiting homeownership to become a renter.

Introducing rental and default options, as well as the associated costs, would strongly reduce the model's tractability and challenge computational feasibility. But it is possible to obtain a quantitative measure of the size of the costs that would prevent agents from exercising an option that allows them to move without renewing the mortgage downpayment. For this purpose, one can extend the model as follows.

Suppose mortgage contracts include a portability option that allows borrowers to transfer their *current* mortgages to a *new* home, without a renewed downpayment. Exercising the portability option, however, creates a monthly utility loss, denoted by $\overline{c}$, which is suffered over a period equal to the average duration between two moves (160 months).[50] When exercising the option, mobility is no longer impaired by the credit constraint. But in what situations would agents actually prefer the portability option over refinancing? Equation 6 makes clear that this is the case when the utility loss associated with refinancing, $-\lambda_{cc,t}\left(\chi p_{h,t}h_t - d_{t-1}\right)$, exceeds the present value of the utility loss created by refinancing, $\sum_{t=0}^{159}\beta^t\overline{c}$.

In a model simulation, one can compute the *minimum* value of the utility loss $\overline{c}$, that would fully prevent agents from exercising the portability option. This is done by simply calculating the maximum value of the term $-\lambda_{cc,t}\left(\chi p_{h,t}h_t - d_{t-1}\right) / \sum_{t=0}^{159}\beta^t$. Figure 11 plots the response to a housing preference shock that reduces house prices by twenty percent. It turns out that a cost of $\overline{c} > 0.026$ is enough to prevent agents

[50] In models with homeowners and renters, the disadvantage of being a renter is typically modeled as a period-period utility loss. See for example Head and Lloyd-Ellis (2008).

to exercise the portability option along the simulated path. Given the logarithmic utility specification, this cost is equivalent to a loss of at most 2.6 percent of consumption. Hence, even under large house price volatility moderate costs can hinder credit-constrained homeowners from exercising options that would allow them move without facing a new mortgage downpayment.

6 Model predictions and recent empirical evidence

This section discusses recent empirical research on the importance of house-lock for mobility patterns and labor market outcomes, in particular during the Great Recession. Sasser Modestino and Dennett (2012) show that states with higher shares of homeowners with negative equity experience lower rates of outmigration, which is consistent with the geographical mobility channel proposed in this paper. However, based on a back-of-the-envelope calculation, they conclude that the observed changes in interstate migration flows cannot account for a large fraction of the increase in the aggregate rate of unemployment rate during the Great Recession. In the light of the model discussed in the previous sections, this point is not surprising. First, most of the increase in the unemployment rate can be directly related to a fall in vacancies. But what has created most concerns among academics and policy makers is an *additional* increase in unemployment that does not seem related to a lack of vacancies. Second, interstate migration accounts for only a modest fraction of overall mobility. The upper part of Table 4 displays the number of movers as reported by Current Population Survey of 2010-2011 and shows that about 67 percent of all moves are within the same state. Moreover, 85 percent of all moves can be classified as short-distance moves.[51] What about job-related moves? The bottom part of Table 4 shows that 64 percent of all moves that are reported as being primarily for employment reasons can be classified as short-distance. More than half of these moves are actually within the same county.

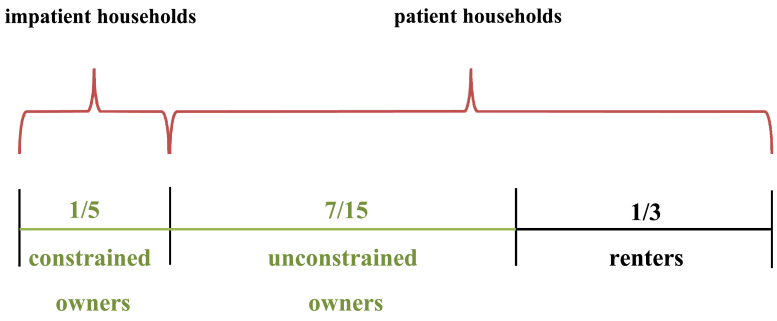
Other recent research has focused on comparing the labor market performance of owners and renters during the Great Recession.[52] Of course, renters are not an ideal control group to assess the importance of barriers to mobility among homeowners,

[51] Here, short distance moves are defined as those moves that are either within the same county, or between counties but covering a distance less than 200 miles.

[52] See Valletta (2010).

since renters and owners differ in many aspects such as their average age, education and income.[53] Moreover, many homeowners are not actually credit-constrained and may therefore not become less mobile after a fall in house prices. Given differences in the characteristics of owners and renters, it does not seem implausible that a typical renter was affected more by the reduction in job opportunities during the Great Recession than a typical *unconstrained* homeowner.

How can we compare the predictions of the model to empirical evidence on owners and renters? The model makes a distinction between those who are collateral-constrained and those who are not, rather than distinguishing explicitly between renters and owners. But as discussed in the section on the calibration of the model, one can relate the agents in the model to owners and renters in the data. The figure below depicts this relation.



The impatient households in the model represent the credit-constrained homeowners in the data. In the benchmark calibration, the size of this group is set to one fifth of the population. The other eighty percent of the population, the patient households, crucially have in common that the collateral constraint is not relevant to them. Hence, the data equivalent of the group of patient households consists of both renters and unconstrained homeowners. Using the fact that the homeownership rate in the U.S. is about two thirds, one can construct weights to compute the average job finding rate among a group labeled as "renters" and group labeled as "owners". The weights imply that all collateral-constrained households are labeled as owners, but most owners are not actually credit-constrained.

[53] According to the American Housing Survey 2005, the median income of households living in renter-occupied units was \$24.651, while among owners the median income was \$53.953.

The left panel of Figure 12 plots the responses of the job finding rates among owners and renters to a negative housing preference shock. The average response of the job finding rate is very similar among the groups, like in the data. The panel also plots the difference of the two responses, which has a useful interpretation. Suppose that due to differences in individual characteristics, renters suffer an additional decline in the job finding rate after the shock. Hence, the job finding rate of renters decreases relative to *unconstrained* homeowners.[54] The difference between the two responses thus measures the amount of heterogeneity that would render the *average* responses of the average job finding rates within the groups of owners and renters entirely indistinguishable. Figure 12 makes clear that this is only a small amount, given that the maximum difference is less than seven percent of the maximum decline in the job finding rate among renters.

The exercise makes clear that caution is required when comparing owners and renters. Despite the fact that even without introducing any additional heterogeneity, the model predicts only minor differences in the change of average job finding rates among owners and renters during a housing bust. But at the same time, a housing preference shock does create a substantial shift of the Beveridge curve, as can be seen from the responses in Figure 6.

The average responses of owners and renters mask more pronounced underlying differences in the job finding rates of credit-constrained households and unconstrained households. This is shown in the right panel of Figure 12, which plots the responses for the patient and impatient households. But the panel also makes clear that credit-constrained households are not alone in facing additional difficulties in finding a job during a housing bust. Although the patient households are not directly hindered by collateral constraints, they are affected in equilibrium as well. In particular, they suffer from the fall in vacancies and the congestion in the labor market created by credit-constrained households.

7 Concluding remarks

Both the empirical and the theoretical evidence presented in this paper support the idea that geographical mobility can be an important channel through which changes

[54] Whether the job finding rate of renters also declines relative to collateral-constrained owners depends on the magnitude of the renter-specific fall in the job finding rate.

in house prices spill over to the real economy. The model that was developed captures the essence of the mobility channel, but has been kept relatively simple in order to retain transparency and tractability. But it would be worthwhile to enrich the way in which housing and labor markets are modeled.

An interesting extension would be to introduce search frictions in the housing market. In the model, the housing stock is essentially traded on a spot market. Ngai and Tenreyro (2009) show that a model with search frictions in the housing market can generate joint (seasonal) movements in house prices and transaction volumes. A final simplification of my model is that it avoids wealth heterogeneity. This has the benefit of simplicity, but since housing wealth and financial wealth have been shown to be important drivers of fluctuations in mobility, dynamics are potentially even richer in a model with wealth heterogeneity.

References

Andres, J., R. Bosca, and J. Ferri (2010): "Labor Market Search, Housing Prices and Borrowing Constraints," mimeo.

Barnichon, R. (2010): "Building a Composite Help-Wanted Index," *Economics Letters*, 109, 175–178.

Campbell, J. R., and Z. Hercowitz (2009): "Welfare Implications of the Transition to High Household Debt," *Journal of Monetary Economics*, 56(1), 1–16.

Chan, S. (2001): "Spatial Lock-In: Do Falling House Prices Constrain Residential Mobility?," *Journal of Urban Economics*, 49(3), 567–586.

den Haan, W. J., G. Ramey, and J. Watson (2003): "Liquidity Flows and Fragility of Business Enterprises," *Journal of Monetary Economics*, 50, 1215–1241.

Elsby, M., B. Hobijn, and A. Sahin (2010): "The Labor Market in the Great Recession," *Brookings Papers on Economic Activity*, spring.

Engelhardt, G. (2003): "Nominal Loss Aversion, Housing Equity Constraints, and Household Mobility: Evidence from the United States," *Journal of Urban Economics*, 53(1), 171–195.

Favilukis, J., S. Ludvigson, and S. Van Nieuwenburgh (2010): "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium," *Working Paper*.

Ferreira, F., J. Gyourko, and J. Tracy (2010): "Housing Busts and Household Mobility," *Journal of Urban Economics*, 68, 34–45.

——— (2011): "Housing Busts and Household Mobility: An Update," *Working Paper*.

Gertler, M., and A. Trigari (2009): "Unemployment Fluctuations with Staggered Nash Wage Bargaining," *Journal of Political Economy*, 117(1), 38–86.

Hall, R. (2005): "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, 95(1), 50–65.

Haurin, D., and S. Rosenthal (2004): "The Sustainability of Homeownership: Factors Affecting the Duration of Homeownership and Rental Spells," *Report for the U.S. Department of Housing and Urban Development, Office of Policy Development and Research*.

Head, A., and H. Lloyd-Ellis (2008): "Housing Liquidity, Mobility, and the Labour Market," *Queen's Economics Department Working paper*.

Henley, A. (1998): "Residential Mobility, Housing Equity and the Labour Market," *Economic Journal*, 108.

Hornstein, A., P. Krusell, and G. L. Violante (2005): "Unemployment and Vacancy Fluctuations in the Matching Model: Inspecting the Mechanism," *Federal Reserve Bank of Richmond Economic Quarterly*, 91, 19–51.

Iacoviello, M. (2005): "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle," *American Economic Review*, 95(3), 739–764.

Iacoviello, M., and S. Neri (2009): "Housing Market Spillovers: Evidence from an Estimated DSGE Model," *American Economic Journal: Macroeconomics*, forthcoming.

Iacoviello, M., and M. Pavan (2009): "Housing and Debt over the Life Cycle and over the Business Cycle," unpublished manuscript, Boston College, Boston.

Kiyotaki, N., and J. Moore (1997): "Credit Cycles," *Journal of Political Economy*, 105, 211–248.

Kydland, F., and E. C. Prescott (1982): "Time to Build and Aggregate Fluctuations," *Econometrica*, 50, 1345–1370.

Merz, M. (1995): "Search in the Labor Market and the Real Business Cycle," *Journal of Monetary Economics*, 34, 269–300.

Molloy, R., C. Smith, and A. Wozniak (2011): "Internal Migration in the United States," *Working Paper*.

Ngai, R., and S. Tenreyro (2009): "Hot and Cold Seasons in the Housing Market," *CEP Working paper*.

Petrongolo, B., and C. Pissarides (2001): "Looking Into the Black Box: A Survey of the Matching Function," *Journal of Economic Literature*, 39, 390–431.

Pissarides, C. A. (2000): *Equilibrium Unemployment Theory, Second Edition*. Cambridge, MA: MIT Press.

Ravn, M. O., and H. Uhlig (2002): "On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations," *Review of Economics and Statistics*, 84, 371–375.

Rupert, P., and E. Wasmer (2011): "Housing and the Labor Market: Time to Move and Aggregate Unemployment," *Journal of Monetary Economics*, forthcoming.

Sasser Modestino, A., and J. Dennett (2012): "Are Americans Locked into Their Houses? The Impact of Housing Markert Conditions on State-to-State Migration," .

Schulhofer-Wohl, S. (2011): "Negative Equity Does Not Reduce Homeowners' Mobility," *NBER Working Paper No. 16701*.

SHIMER, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95, 25–49.

——— (2007): "Reassessing the Ins and Outs of Unemployment," *NBER working paper*.

STEIN, J. (1995): "Prices and Trading Volume in the Housing Market: A Model with Down-Payment Effects," *Quarterly Journal of Economics*, 110, 379–406.

STOKEY, N. (2009): "Moving Costs, Nondurable Consumption, and Portfolio Choice," *Journal of Economic Theory*, 144.

VALLETTA, R. (2010): "House Lock and Structural Unemployment," *Federal Reserve Bank of San Francisco Working Paper*.

Tables and Figures

Table 1: Standard deviations and correlations: data.

	house pr.	GDP	home sales	unemp.rate	vacancies	outfl.haz.
Standard deviations						
relative to house price	1	0.368	2.445	2.675	3.008	2.524
Correlations						
	house pr.	GDP	home sales	unemp.rate	vacancies	outfl.haz.
gdp	0.554	1				
home sales	0.552	0.608	1			
unemploment rate	−0.437	−0.858	−0.331	1		
vacancies	0.400	0.819	0.472	−0.895	1	
unemp. outflow hazard	0.464	0.861	0.360	−0.966	0.888	1

Notes: Data are quarterly. Following Shimer (2005), variables are logged and HP-detrended with smoothing parameter value 10^5.

Table 2: Parameter theoretical model (benchmark calibration).

parameter	description	value	source/target
β	discount factor impatient h.h.	0.9899	Iacoviello and Neri (2009)
γ	discount factor patient h.h.	0.9975	Iacoviello and Neri (2009)
α^{imp}	housing pref. impatient h.h.	0.139	steady state
α^{pat}	housing pref. patient h.h.	0.043	steady state
ψ	utility from new location	-7.144	steady state
σ	stdev. location preference shock	3.5	volatility mobility
ζ	moving cost	1.6	Stokey (2009)
ρ_u	rate of job destruction	0.035	Gertler and Trigari (2009)
η	elasticity parameter matching function	0.6	Petrongolo and Pissarides (2001)
μ	level parameter matching function	0.545	steady state
ϑ	vacancy cost	0.181	steady state
ω	fraction of long-distance job offers	1/3	no source, check for robustness
ξ	wage rule parameter	0.98	2% accounting profits
κ	utility from unemployment	-4.428	steady state
ν	share impatient h.h.	0.2	AHS / SCF data
χ	collateral requirement	0.8	Campbell and Hercowitz (2009)
ρ_h	autocorr. housing pref. process	0.983	same as tech. process
ρ_a	autocorr. technology process	0.983	Kydland and Prescott (1982)

Table 3: Standard deviations and correlations. Model with only productivity shocks.

	house pr.	output	mob.rate	unemp.rate	vacancies	outfl.haz.
Standard deviations						
relative to house price	1	0.980	2.362	0.824	0.8958	0.5649
	(0)	(0.002)	(0.033)	(0.016)	(0.023)	(0.002)
Correlations						
	house pr.	output	mob.rate	unemp.rate	vacancies	outfl.haz.
output	1.000	1				
	(0.001)	(0)	-	-	-	-
mobility rate	0.993	0.994	1			
	(0.029)	(0.027)	(0)	-	-	-
unemploment rate	−0.859	−0.858	−0.868	1		
	(0.026)	(0.026)	(0.026)	(0)	-	-
vacancies	0.967	0.968	0.955	−0.700	1	
	(0.003)	(0.003)	(0.004)	(0.044)	(0)	-
outflow hazard	0.916	0.915	0.919	−0.957	0.801	1
	(0.016)	(0.017)	(0.016)	(0.008)	(0.030)	(0)

Notes: The business cycle statistics are averages across 10000 simulations. Standard deviations over these simulations are displayed between brackets. Each simulation has a monthly frequency, has length 1480 and starts from the steady state. For each simulation, the first 1000 timer periods were discarded so that 40 years of data remained. Variables were logged and HP-detrended with smoothing parameter value $81 \cdot 10^5$. This value corresponds to the one used by Shimer (2005) for quarterly data, but is adjusted for the frequency using the factor recommended by Ravn and Uhlig (2002)

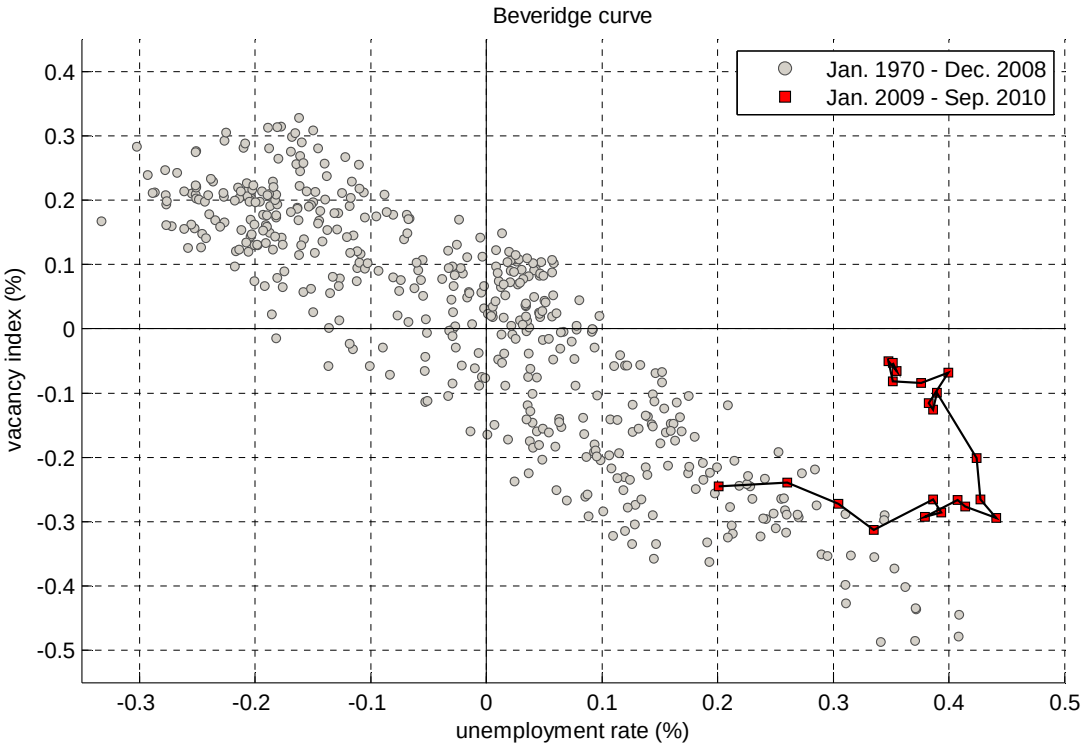
Table 4: Number of movers.

A) All movers	thousands	percentage
Total	35.075	100.0
same county	23.325	66.5
short distance	29.728	84.8

B) Movers for employment reasons	thousands	percentage
Total	6481	100.0
same county	2293	35.4
short distance	4133	63.8

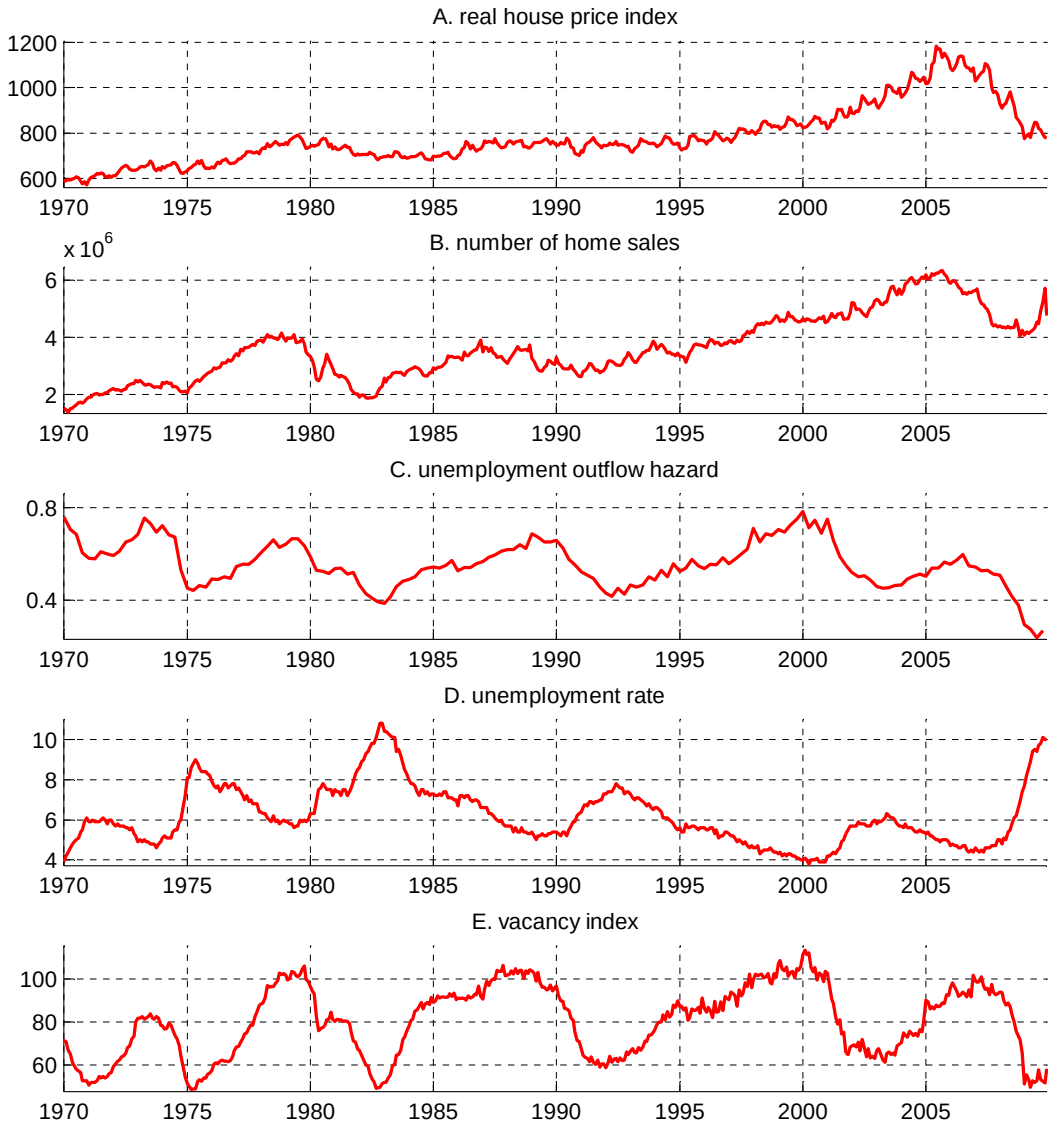
Notes: Data from the Current Population Survey 2010-2011. Short distance moves are defined as either intracounty moves or intercounty moves up to a distance of 200 miles.

Figure 1: Beveridge curve.



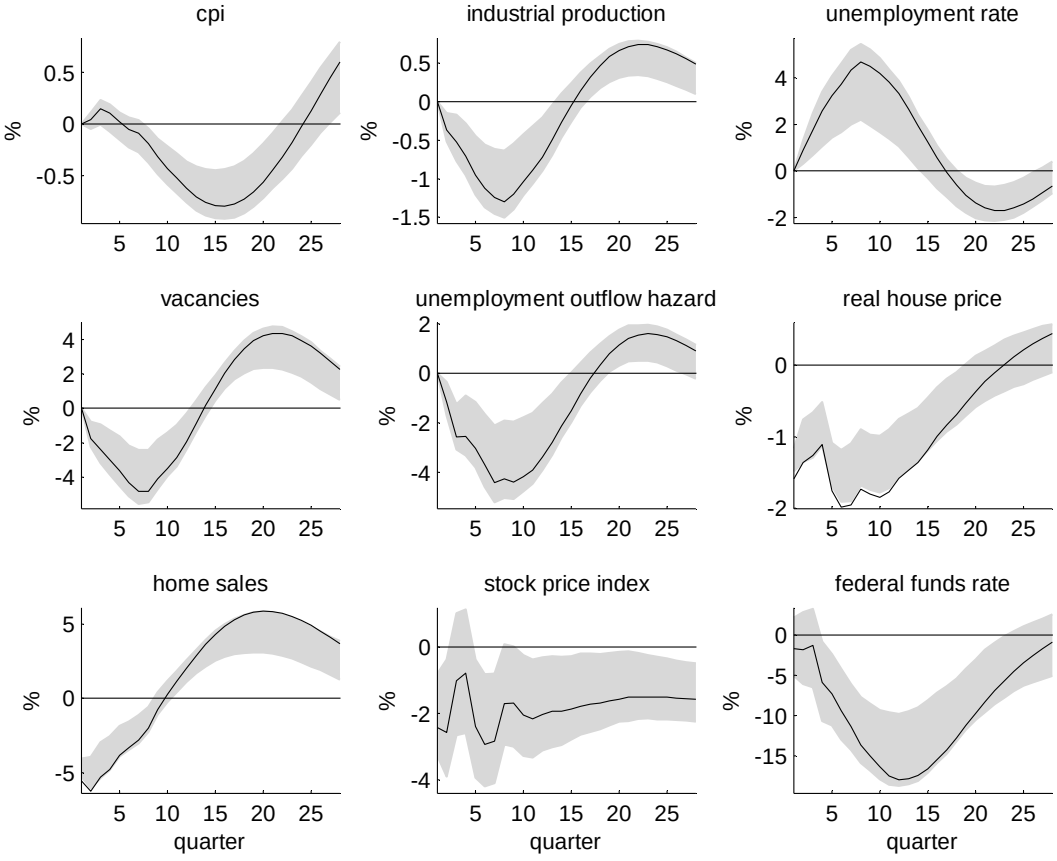
Notes: Log deviations from trend. Data are monthly and cover the period from January 1970 until December 2009. Variables were logged and HP-detrended with smoothing parameter value $81 \cdot 10^5$. This value corresponds to the one used by Shimer (2005) for quarterly data, but is adjusted for the frequency using the factor recommended by Ravn and Uhlig (2002) Source unemployment rate: U.S. Department of Labor. Source vacancy index: Barnichon (2010).

Figure 2: Raw data series.



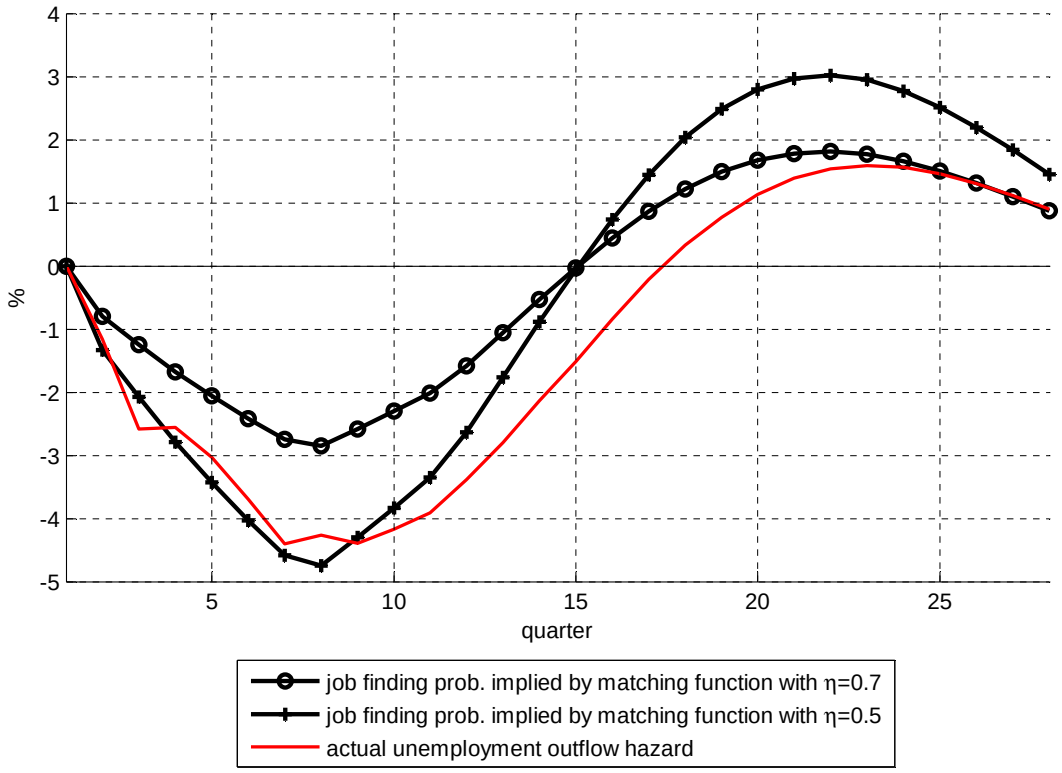
Notes: Quarterly obervations. Monthly series for house prices, home sales, the unemployment rate and the vacancy index were converted into quarterly series by simple averaging. Sources: see main text.

Figure 3: Structural VAR: housing market shock.



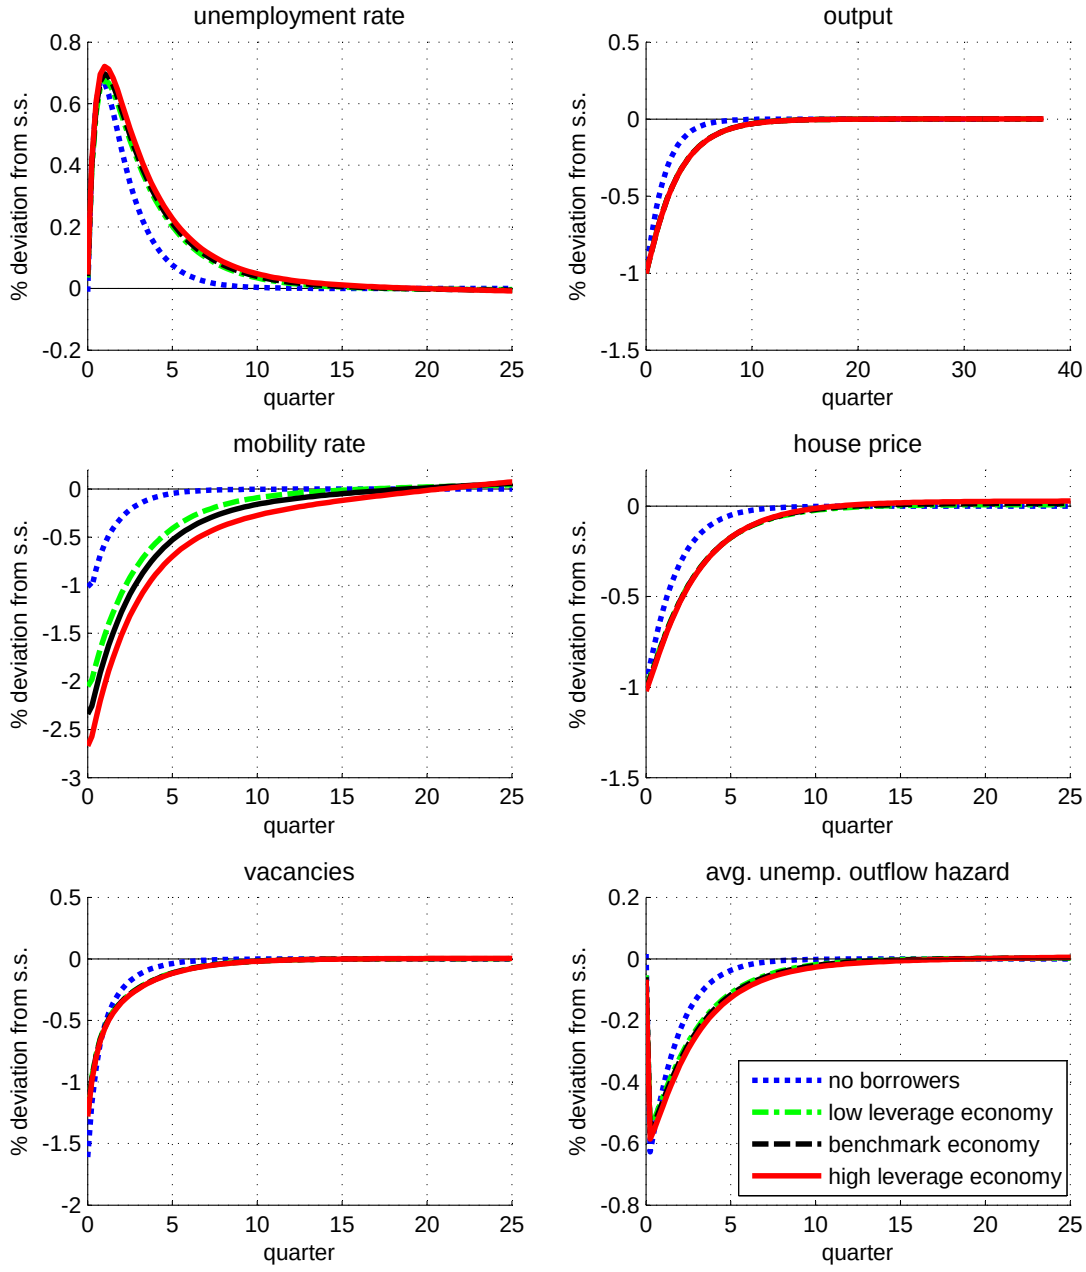
Notes: Responses to a joint and negative one standard deviation shock in real house prices and homesales. Grey areas are 90% confidence intervals, which are obtained by a bootstrap procedure.

Figure 4: Structural VAR: housing market shock. Unemployment ouflow hazard implied by the matching function versus the actual outflow hazard.



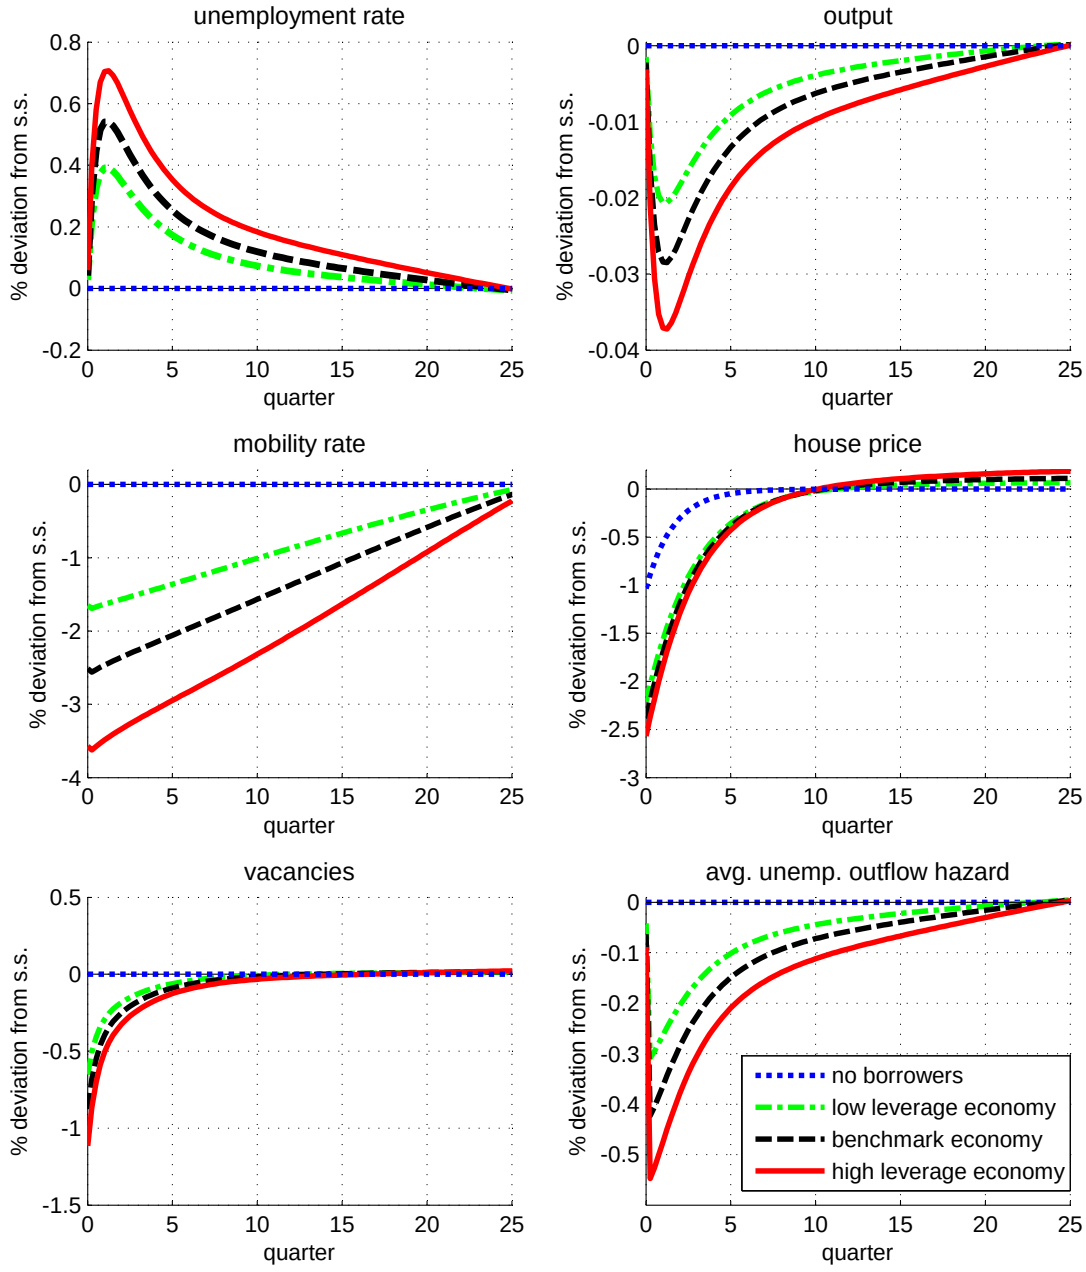
Notes: Responses to a joint and negative one standard deviation shock in real house prices and homesales. VAR responses are plotted for the unemployment outflow hazard and the job finding probability implied by a matching function of the form $f(n_{u,t}, v_t) = \mu n_{u,t}^{\eta} v_t^{1-\eta}$, for various values of η.

Figure 5: Responses to a productivity shock in the theoretical model.



Notes: Responses to a 1% decline in productivity for the high-leverage economy ($\nu = 0.25, \chi = 0.85$), the benchmark economy ($\nu = 0.2, \chi = 0.8$), the low-leverage economy ($\nu = 0.15, \chi = 0.75$), and the economy without borrowers ($\nu = 0$).

Figure 6: Responses to a housing preference shock in the theoretical model.



Notes: Responses to a housing preference shock for the high-leverage economy ($\nu = 0.25, \chi = 0.85$), the benchmark economy ($\nu = 0.2, \chi = 0.8$), the low leverage-economy ($\nu = 0.15, \chi = 0.75$), and the economy without borrowers ($\nu = 0$).

Figure 7: Great recession experiment: Beveridge curves predicted by the model.

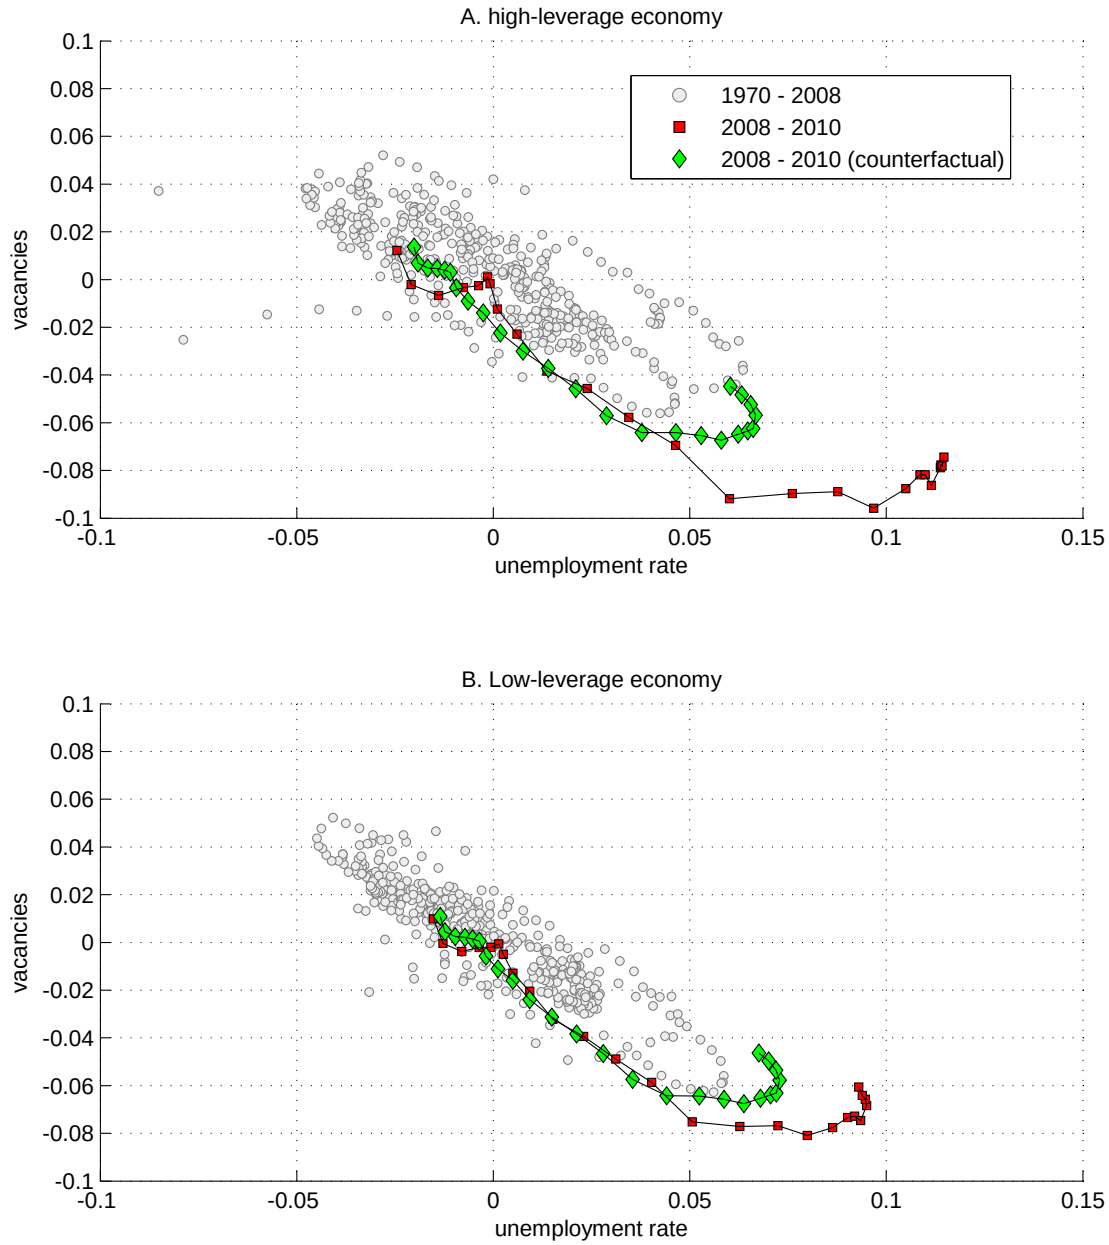
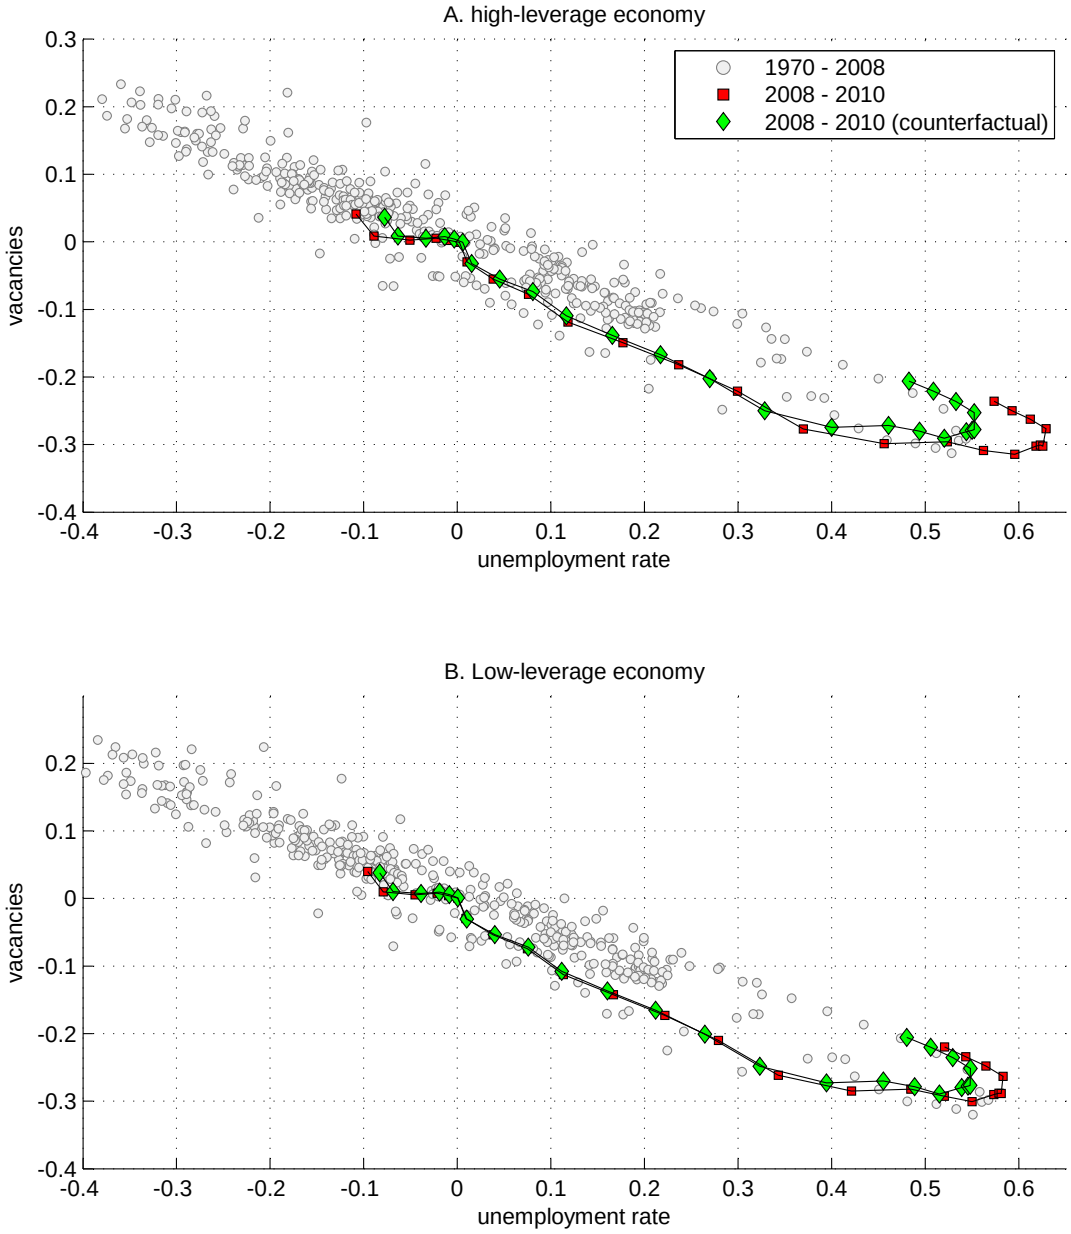
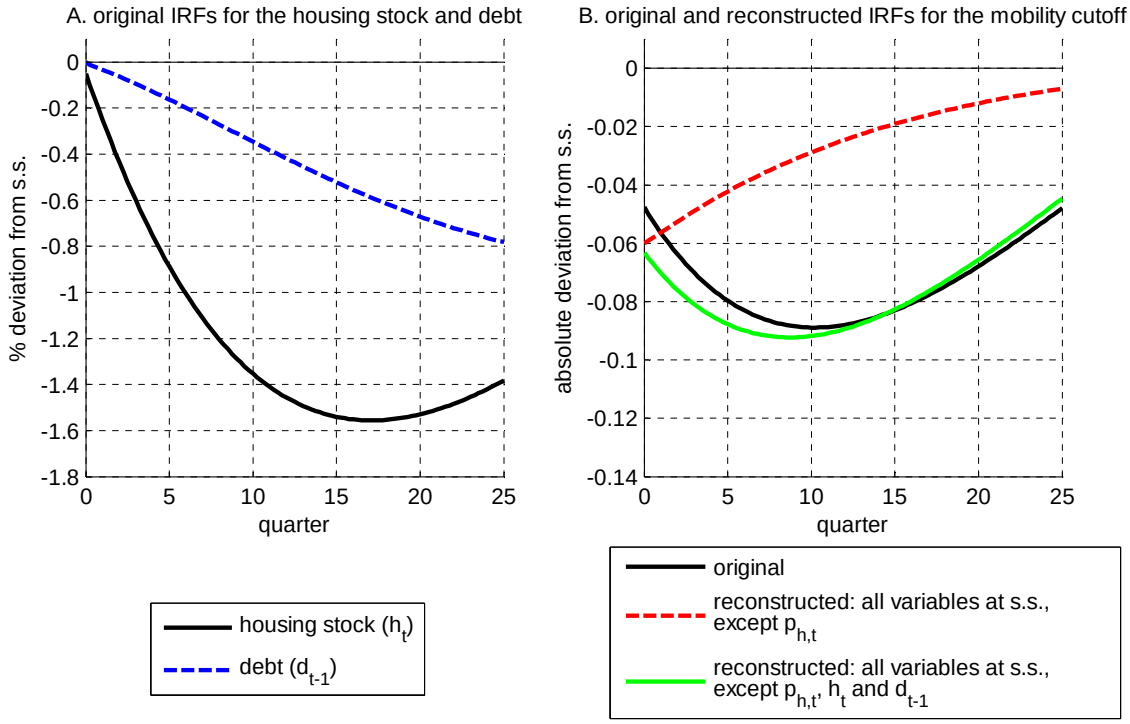


Figure 8: Great recession experiment: Sticky wage model.



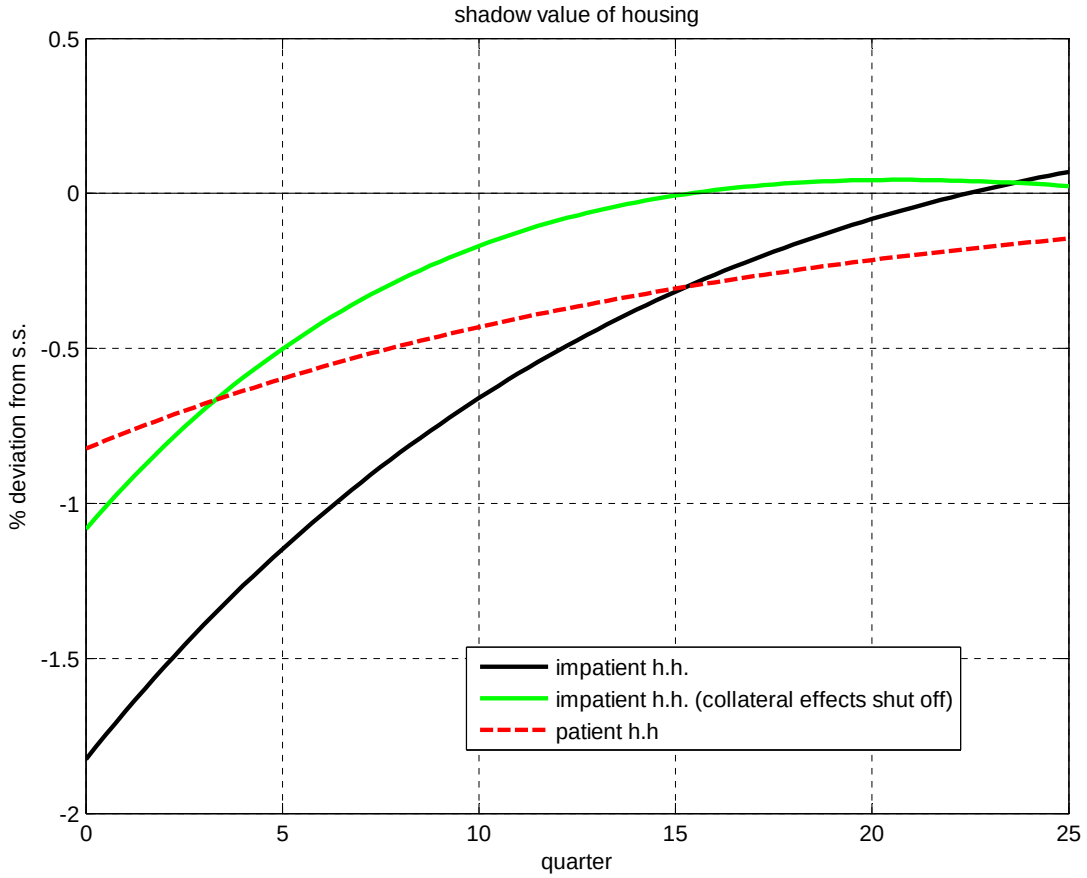
Notes to Figures 7 and 8: Log deviations from the steady state. Model simulations for the high-leverage economy ($\nu = 0.25, \chi = 0.85$) and the low-leverage economy ($\nu = 0.15, \chi = 0.75$). "Counterfactual" denotes simulations with housing preference shock innovations set to zero from 2005 onwards. For the sake of comparability, simulated data are HP-detrended as in Figure 1. The smoothing parameter value is $81{\cdot}10^5$.

Figure 9: Responses to a housing preference shock in the benchmark model : impatient households.



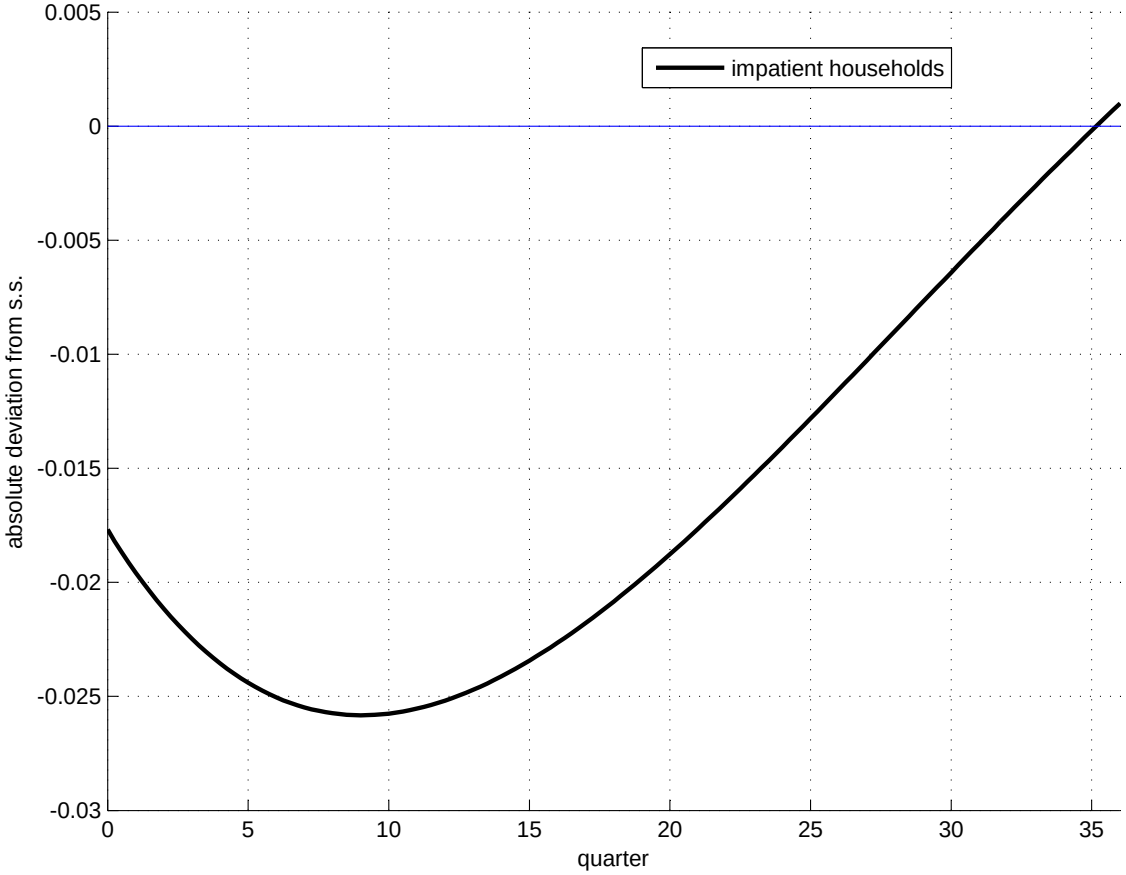
Notes: Responses of housing stock and debt of the impatient household (Panel A), and of the mobility cutoff (Panel B) to a negative housing preference shock. The reconstructed responses in Panel B are obtained using Equation (6), but with the indicated variables kept at their steady-state levels.

Figure 10: Responses to a housing preference shock in the benchmark model.



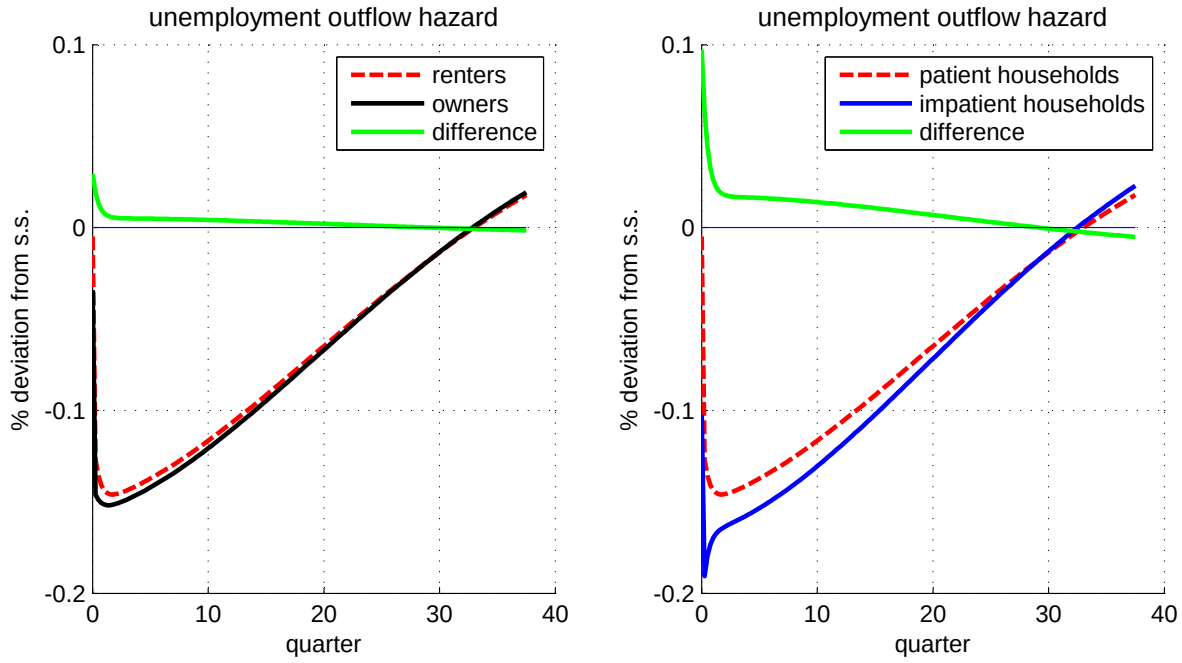
Notes: Responses of the shadow value of housing for the patient and impatient households to a negative housing preference shock. The figure also shows the IRF for the shadow value of the impatient households when collateral effects are shut off. This IRF is constructed from equation (21), but with the second present value term kept equal to its steady-state level.

Figure 11: Mortgage portability option.



Notes: Responses of the term $\lambda_{cc,t}\left(\chi p_{h,t} h_t - d_{t-1}\right) / \sum_{t=0}^{159} \beta^t$ to a negative housing preference shock that reduces house prices by twenty percent. The absolute value of the maximum decline represents the minimum cost that would prevent agents from exercising a mortgage portability option.

Figure 12: Renters and owners.



Notes: Responses to a housing preference shock that reduces house prices by one percent.