

IZA DP No. 7002

House Lock and Structural Unemployment

Robert G. Valletta

November 2012

House Lock and Structural Unemployment

Robert G. Valletta

*Federal Reserve Bank of San Francisco
and IZA*

Discussion Paper No. 7002
November 2012

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

House Lock and Structural Unemployment^{*}

A recent decline in geographic mobility in the United States may have been caused in part by falling house prices, through the “lock in” effects of financial constraints faced by households whose housing debt exceeds the market value of their home. I analyze the relationship between such “house lock” and the elevated levels and persistence of unemployment during the recent recession and its aftermath, using data that covers the period through the end of 2011. Because house lock will extend job search in the local labor market for homeowners whose home value has declined, I focus on differences in unemployment duration between homeowners and renters across geographic areas differentiated by the severity of the decline in home prices. The empirical analyses rely on microdata from the monthly Current Population Survey (CPS) files and an econometric method that enables the estimation of individual and aggregate covariate effects on completed unemployment durations in “synthetic cohort” (pseudo-panel) data. The estimates indicate the absence of a meaningful house lock effect on unemployment duration.

JEL Classification: J6, R31

Keywords: unemployment, house prices, mobility

Corresponding author:

Robert G. Valletta
Federal Reserve Bank of San Francisco
101 Market Street
San Francisco, CA 94105-1579
USA
E-mail: rob.valletta@sf.frb.org

^{*} I thank Leila Bengali and especially Katherine Kuang for outstanding research assistance. For their helpful comments, I thank Andy Haughwout, Betsey Stevenson, Justin Wolfers, and other seminar attendees at the Federal Reserve System Regional Research Conference (Nov. 2010), Princeton University (April 2012), and the European Association of Labour Economists Annual Meetings (Sept. 2012). The views expressed in this paper are those of the author and should not be attributed to anyone else at the Federal Reserve Bank of San Francisco or the Federal Reserve System.

House Lock and Structural Unemployment

1. Introduction

During the recent recession and its aftermath, the rate of geographic mobility in the United States reached the lowest levels recorded in U.S. Census Bureau statistics, which began in 1948. The falling home prices that preceded and intensified during the recession may be an important cause of the reduction in geographic mobility, through the “lock in” effects of financial constraints faced by households that are “underwater”—i.e., whose housing debt exceeds the market value of their homes.¹ This in turn has led to speculation among economists and other observers that the stubbornly high unemployment rates observed in 2009-2011 were caused in part by the inability of unemployed homeowners to move to geographic areas where jobs are available (see Fletcher 2010 for media discussion and quotes from economists regarding these points). To the extent that this phenomenon exists, it represents a form of structural unemployment that is likely to raise the equilibrium unemployment rate or non-accelerating inflation rate of unemployment (NAIRU) over an extended period.

In this paper, I investigate whether systematic statistical evidence can be found to support the hypothesis that house lock has contributed to higher U.S. unemployment. The analysis relates to two existing literatures. First, previous work on home prices and mobility has found that declining home prices and the consequent increase in the incidence of underwater homes are associated with reduced geographic mobility by homeowners (e.g., Chan 2001; Englehardt 2003; Ferreira, Gyourko, and Tracy 2010). A

¹ I refer to this phenomenon as “house lock,” because it is similar to the “job lock” effect of reduced job mobility arising from the costs associated with the potential loss of employer provided health insurance (e.g., Buchmueller and Valletta 1996).

separate literature has investigated the relationship between home ownership and unemployment rates, at the individual or regional level (e.g., Oswald 1996; Munch, Rosholm, and Svarer 2006; Coulson and Fisher 2009). This literature was largely propelled by Oswald's argument that reduced mobility associated with home ownership creates labor market inefficiency and higher unemployment rates. The results from this literature are mixed, reflecting a diversity of empirical approaches and specific hypotheses.

I link these two literatures together by examining the relationship between falling home prices and individual unemployment experiences, with declining geographic mobility operating as the unobserved link between them. Past work on the relationship between home ownership and unemployment focused on differences in unemployment outcomes based on home ownership versus renting measured at the individual level or across geographic areas (countries, states, or metropolitan areas). In these approaches, it is difficult to control for systematic differences between homeowners and renters, or across geographic areas that differ in their rate of home ownership, that cause different employment outcomes. As such, the papers in this literature often adopt instrumental variable (IV) strategies or selectivity corrections to eliminate the estimation bias introduced by unobserved heterogeneity or endogeneity in the determination of home ownership and labor market status.

I extend the existing literature to assess the labor market impact of house lock in the recent housing bust, relying on direct measures of the likely extent of house lock that enable me to avoid reliance on IV strategies. Although past work has produced mixed evidence regarding homeownership effects on unemployment, the severity of the housing

bust and ongoing unemployment problem suggests the possibility of large effects in the recent episode. I focus directly on job search intensity and opportunities, as reflected in unemployment duration, rather than broader unemployment outcomes. To test for house lock effects, I take a “difference-in-difference” approach by comparing outcomes for individuals living in owner-occupied housing versus those living in rental housing across geographic areas differentiated by the intensity of the home price decline. Both groups face similar local labor market conditions, but owners face additional financial constraints that may lengthen their unemployment durations and hence overall unemployment; if this effect exists, it should be most pronounced in areas that have seen the largest declines in home prices (conditional on local labor market conditions).

My analysis of unemployment duration relies on a recently developed econometric approach (Güell and Hu 2006) applied to monthly microdata on unemployed individuals from the U.S. Current Population Survey (CPS). Unlike other approaches to the analysis of unemployment spells in repeated cross-sections, this method enables direct estimation of the influence of detailed individual characteristics and duration dependence, in addition to measures of general economic conditions and other time-varying factors, in the determination of unemployment duration (see also Valletta 2011).

The next section discusses prior research in more detail and establishes some central facts regarding recent movements in home prices and geographic mobility. Section 3 describes the CPS unemployment data and provides descriptive evidence regarding unemployment durations. The formal econometric approach for the analysis of expected duration in a conditional setting is described in Section 4. Section 5 discusses the specific empirical hypotheses, including caveats, and presents results. As

summarized in Section 6, I find no evidence to support the view that reduced geographic mobility by homeowners has made a substantive contribution to elevated unemployment during the recent recession and subsequent recovery (through 2011).

2. Home ownership, mobility, and unemployment

2.1 Existing research on homeowner mobility and homeowner unemployment

The hypothesis of higher structural unemployment arising from house lock has two primary components: (i) homeowners are relatively immobile and therefore tied to their local labor markets; this results from financial constraints or incentives associated with home ownership, which are particularly binding when home values have declined; (ii) homeowners' lower mobility precludes optimal job search in other geographic areas and increases their time spent unemployed in their current labor market.² The first link has found substantial support in the relevant literature, while support for the second link is mixed.

In regard to homeowners' geographic mobility, declining home prices have two potential effects. Price declines may increase default rates, causing mobility to rise as foreclosed homeowners seek alternative housing. On the other hand, it is likely that only a fraction of homeowners facing price declines or negative equity will default. Instead, when prices fall, some homeowners who might otherwise have chosen to sell their homes and move may choose instead to remain in their current residence, as a result of the financial constraints arising from low or negative housing equity combined with significant transaction costs (e.g., Chan 2001; Ferreira, Gyourko, and Tracy 2010;

² In this paper, the terms "home ownership" and "home owners" will be used synonymously with "individuals living in owner-occupied housing," as will similar terms for renters.

although see Schulhofer-Wohl 2012 for an opposing view). Mobility may also be suppressed for households not facing direct financial constraints, as a result of nominal loss aversion that causes them to place higher weight on capital losses than on equivalent gains (Genesove and Mayer 2001, Englehardt 2003). The empirical tests in this literature generally indicate that the default effect is strongly dominated by the effects arising from financial constraints and loss aversion, causing homeowner mobility to fall substantially in response to price declines.³ The estimated reductions in geographic mobility from this literature are large: the Chan, Englehardt, and Ferreira et al. papers cited above find that mobility is reduced by about 25-45 percent (relative to their base rates) for homeowners who face negative equity or a modest decline in nominal home prices.

The literature regarding the relationship between home ownership and unemployment is more mixed with respect to core hypotheses and findings. Oswald (1996) made the straightforward argument that financial constraints arising from transaction costs in housing markets reduce homeowners' flexibility in the labor market, resulting in less favorable labor market outcomes for them. He offered this as an explanation for his estimates of a positive correlation between unemployment rates and the proportion of homeowners across countries and regions.

Subsequent work has focused on refining the "Oswald hypothesis" and the associated empirical tests. Formal models of home ownership and job search produce

³ The distinction between equity constraints and nominal loss aversion as explanations of lower homeowner mobility is important from behavioral and policy perspectives but is inconsequential for the tests in this paper. As Englehardt (2003) notes, the role of equity constraints suggests a degree of market failure that reduces mobility below socially optimal levels, which could be usefully addressed by government policy. By contrast, nominal loss aversion is a characteristic of individual preferences and as such does not have implications for social welfare or efficiency-enhancing interventions by government agencies. The distinction does not affect the framework or findings of my paper because both mechanisms operate through declines in house prices and imply that some homeowners who might move to find a new job will not do so.

ambiguous predictions regarding the relationship between ownership and unemployment outcomes for individuals or geographic areas. Munch et al (2006) specify a model of job search that allows for transitions into employment in the local labor market or outside the local labor market. In this model, unemployed individuals living in owner-occupied housing face higher moving costs than renters, which lowers their transition rates into employment outside the local market but raises their transition rates within the local market; the overall effect on homeowner unemployment durations is ambiguous. They estimate competing risk models of the separate transition rates and find that home ownership reduces employment transitions through geographic mobility, as argued by Oswald. However, this effect is more than offset by homeowners' increased transition rates for jobs in the local labor market, implying that home ownership reduces unemployment on net.

Coulson and Fisher (2010) compare the implications of alternative search models for homeowner versus renter unemployment at the individual and aggregate (market) level. In these models, whether housing market frictions increase homeowner unemployment (relative to renters) depends on the nature of the wage-setting process and whether firm entry based on a zero-profit condition is incorporated. Moreover, the models' predictions generally vary in regard to the relationship between home ownership and unemployment at the individual and aggregate levels. The authors test these predictions using IV estimation methods applied to two sets of cross-section data: aggregate data for U.S. MSAs and individual data from the 1990 U.S. Census. None of the theoretical models does a very good job of explaining the empirically estimated relationships between ownership and unemployment. At the individual level, the

findings indicate that unemployment is lower for homeowners, although the reliance on untested exclusion restrictions raises questions about whether the endogeneity of home ownership with respect to labor market outcomes is fully purged.

To summarize, these existing literatures have found substantial support for the first component of the house lock hypothesis—reduced homeowner mobility in response to price declines—but mixed to weak evidence regarding the second component—prolonged job search and elevated unemployment experiences for homeowners. However, given the unprecedented extent of recent declines in U.S. home prices, historically low mobility rates, and persistently high unemployment rates, both of these elements may be operating with unusual force in the recent housing bust and recession. As such, even in the absence of compelling and consistent evidence from past work regarding the link between ownership and unemployment, house lock may be a significant factor in the current environment.

A number of recent papers have attempted to assess the extent of house lock and its impact on the U.S. labor market in recent years. These papers can be grouped into three broad categories: (i) macroeconomic models of the impact of the housing downturn on labor mobility and unemployment (Herkenhoff and Ohanian 2011; Karahan and Rhee 2011; Sterk 2010); (ii) empirical studies of recent changes in geographic mobility and their implications for labor market outcomes (Aaronson and Davis 2011; Donovan and Schnure 2011; Modestino and Dennett 2011; Molloy, Smith, and Wozniak 2011); (iii) direct assessments of the relationship between housing market conditions, geographic mobility, and individual labor market outcomes (Farber 2010, Schmitt and Warner 2011). While the macro models in (i) suggest small to moderate effects of recent housing market

conditions on the aggregate unemployment rate, the direct empirical assessments from (ii) and (iii) find very small or essentially no effects of house lock on the labor market. My own approach, described in more detail below, is closest in spirit and execution to the papers in (iii). However, relative to those papers, in which assessment of house lock constituted a limited portion of the paper, I use more extensive data on unemployment combined with an empirical methodology that is particularly well-suited to assess the impacts of local housing market conditions on unemployment outcomes.

2.2 House price and mobility trends

The house lock argument is predicated on a decline in home prices that reduces geographic mobility. Figure 1 illustrates the net decline and varied movement in prices during the recent housing bust (from 2005 forward), for the nation as a whole and selected metropolitan areas (Metropolitan Statistical Areas, or “MSAs”), using the repeat-sales index for single-family homes from the Federal Housing Financing Agency (FHFA, formerly called OFHEO). From the peak in early 2007 through the end of 2011, nominal home prices across the U.S. fell about 16 percent on average. However, the pattern varied substantially across MSAs. The figure shows the price series for the two MSAs at the bottom and top ends of the distribution of house price movements, among the 235 MSAs for which price series are available. The difference in the experiences of these two MSAs is quite large, with Merced (in California’s Central Valley) seeing a decline of about 63 percent from its peak in mid-2006 and Midland, Texas seeing an increase of about 20 percent. The underlying variation in price movements across the full sample of MSAs will be exploited for the empirical tests in Section 5.

As expected, the decline in home prices has been accompanied by a decline in homeowner mobility (Figures 2 and 3). The mobility calculations are from the U.S. Census Bureau, based on reported geographic moves from the Annual Demographic Supplement (conducted each March) to the monthly household survey (the Current Population Survey, or CPS). Figure 2 displays overall and group-specific mobility rates across states (Panel A) and counties (Panel B), for the period over which separate data on owners and renters is available (back to 1988). In 2009 the overall rate of cross-state mobility fell to 1.6 percent and has stayed largely stable at that low level since; as noted earlier, this is the lowest level recorded in U.S. Census Bureau statistics, which began in 1948.⁴ Overall mobility across counties also declined to historical lows in 2009. It is important to note, however, that mobility has been on a long-term downward trend, with rates in prior decades well above those from the 1990s and 2000s. Figure 2 also shows that renter mobility substantially exceeds owner mobility.⁵ Mobility for both groups fell noticeably after 2005.⁶ To account for the different population shares of owners and renters, Figure 3 illustrates the two groups' shares of total mobility (measured relative to the overall population). The two groups made a similar contribution to the decline in overall mobility after 2005, and their patterns of mobility over time are not noticeably different. The similarity in mobility patterns for owners and renters are not consistent with the house lock hypothesis—if house lock mattered, the decline in mobility probably would be more pronounced for owners—but a more definitive conclusion requires examination of data on unemployment experiences of the two groups.

⁴ See the historical data tables available here:

<http://www.census.gov/population/www/socdemo/migrate.html>

⁵ Table 1, discussed in the next section, compares the characteristics of owners and renters.

⁶ Some of the recent decline is an artifact of changes in measurement and reporting procedures used for the CPS Annual Demographic Supplement (see Kaplan and Schulhofer-Wohl 2012).

3. CPS unemployment data and descriptive analyses

3.1 CPS data and unemployment rates

The data used for the analysis of unemployment duration are constructed from the microdata files of the U.S. Current Population Survey (CPS), a monthly survey of about 60,000 households that is used for official monthly labor force tabulations and other government statistics. These data are available back to 1976, but this paper focuses on the period from January 2008 through December 2011 (with additional descriptive analyses of unemployment duration back to 2005). This period encompasses both the severe recession that began in December 2007 and the subsequent modest recovery in employment that began in 2010 and continued in 2011. Inclusion of the last few years of data, when the U.S. labor market was improving but housing market conditions remained quite weak, is important for the house lock test: the combination of growing employment and flat or declining housing prices observed during 2010-2011 is most likely to create conditions under which homeowners would like to move to accept available jobs but are unable to do so because they are underwater on their homes.

Observations were pulled for all individuals identified as unemployed in the survey, age 16 and older. The resulting data set for the primary analyses has about 3500-6700 observations on unemployed individuals per month, depending on the prevailing unemployment rate at the time; the average for the sample frame is about 5400 per month. For the difference-in-difference regression analyses discussed in Section 5, the sample is further restricted to individuals living in one of 235 MSAs, to enable matching of the FHFA housing price series available at the MSA level; this reduces the sample size

to about 3650 per month.⁷ All of the analyses below incorporate the CPS labor force sampling weights, which are designed to yield monthly samples that are representative of the U.S. labor force. In addition to the large size of the sample, the timeliness of these data is highly advantageous, enabling the near-contemporaneous analysis of ongoing labor market developments. By contrast, panel data sets that could be used for richer econometric modeling of unemployment duration are only available with significant lags.

A key variable for these analyses identifies whether the individual lives in a housing unit that is owned by a household member or is rented.⁸ Table 1 lists means for a standard set of individual control variables, with the sample divided into owner and renter groups, for the analysis period of January 2008 through June 2010. Homeowners account for about 56 percent of the unemployed population in metropolitan areas, and they generally have characteristics associated with more advantageous labor market outcomes: they are older, have higher educational attainment, are less likely to be members of racial and ethnic minority groups, and are more likely to be married. Also, although not shown in the table, the occupational distribution for prior jobs held shows that homeowners are more likely to have held positions in higher-skilled occupations, such as managerial and professional positions, than are renters.⁹

Figure 4, Panel A, confirms the expectation of lower unemployment rates for homeowners than for renters, with rates for renters generally running about twice those of

⁷ Results of the descriptive analyses are essentially the same when the restricted regression sample is used; results available on request.

⁸ The ownership category includes units for which the purchase process has been initiated but not completed. A very small third group, in which no housing payments are being made, are included with renters in the analysis.

⁹ Similar tabulations for prior industry affiliation indicate little difference between the two groups. Tabulations of occupation and industry affiliation for unemployed individuals exclude new entrants to the labor force, for whom no prior employment history exists.

owners.¹⁰ Panel B shows the unemployment rates for owners and renters measured as a percentage of the overall labor force (rather than the group-specific rates in Panel A). The owner share of overall unemployment drifted up a bit in the early 2000s as their population share rose. Most notably, Panel B shows that owners accounted for a disproportionately large share of the increase in unemployment during the recent recession. In particular, the owner contribution to the increase in the overall unemployment rate between late 2007 and late 2009 was about 0.7 percentage points larger than the renter contribution. However, the unemployment share for owners has declined somewhat more since then, again suggesting that house lock has not been a significant problem.

3.2 Descriptive analyses of unemployment duration

While the differences in unemployment rate movements between owners and renters are suggestive, understanding the behavioral contributions to these differences requires an examination of unemployment duration data. In the CPS microdata, unemployment duration is measured as the duration of ongoing (interrupted) spells at the time of the survey, rather than completed duration for individuals who have exited unemployment. This “interrupted spell” measure is used for the calculation of the BLS’s oft-cited “average duration” and “median duration” series, plus the related series that represent the proportion of individuals whose duration falls within specific intervals (e.g., less than 5 weeks, greater than 26 weeks, etc.).¹¹ These series based on interrupted spell durations are subject to well-known biases with respect to measurement of expected

¹⁰ These were calculated using the complete set of CPS labor force observations for individuals 16 and older.

¹¹ A major CPS survey redesign in 1994 substantially altered reported unemployment durations, but this is only relevant for pre- and post-1994 comparisons (Valletta 2011).

duration for an individual entering unemployment, including underestimation of its cyclical elasticity and responsiveness to labor market shocks (Carlson and Horrigan 1983; Sider 1985; Horrigan 1987).

Given the biases in measured duration based on interrupted spells, I focus the descriptive analyses on a measure of expected completed duration for an individual entering unemployment in a particular month (e.g., Sider 1985; Baker 1992a; Valletta 2011). This measure of expected duration is formed based on counts of individuals within duration intervals that correspond to the monthly sampling window for the CPS survey. These counts are used to define and estimate continuation probabilities between adjacent duration categories for “synthetic cohorts” of individuals. The continuation probabilities are then aggregated using standardized formulas to calculate the expected completed duration of unemployment for an individual entering unemployment in a particular month, under the assumption that the continuation probabilities remain the same. This method is described in detail in Appendix A.

Figure 5 compares expected unemployment duration, measured in weeks, for owners and renters, over the period just prior to the housing downturn (2005) through the end of 2011. Panels A and B divide the sample up into MSAs for which the decline in house prices was larger (greater negative) or smaller (less negative) than the national decline. If the house- lock hypothesis is correct, the increase in unemployment durations as the housing bust and economic downturn intensified should be larger for owners than for renters, particularly in MSAs that experienced the greatest decline in home prices. No support for the hypothesis is evident in these charts: although the patterns over time are noisy, both groups saw a substantial increase in unemployment duration during the

recession, and the level of duration is similar for owners and renters over the entire sample frame for both sets of MSAs.

Figure 5 therefore provides no evidence in favor of house lock during the recent downturn. However, as already noted with regard to the descriptive statistics in Table 1, homeowners possess other individual characteristics that typically lead to relatively favorable labor market outcomes. Such characteristics may exert a strong influence on the patterns of unemployment duration in the recent downturn. It is therefore important to apply a method that accounts for differences in observable covariates between owners and renters.¹²

4. Econometric Approach

The econometric approach used for the formal analyses of unemployment duration is adapted from Güell and Hu (2006; henceforth “GH”). GH developed a generalized method of moments (GMM) approach that enables estimation of detailed covariate effects on unemployment duration at the individual level, along with estimation of the conditional effects of duration dependence and the impact of time-varying factors such as labor market conditions. They also outlined a maximum likelihood (ML) alternative that is more straightforward to estimate and does not pose any notable downsides when the available duration measure is reported in precise, uniform duration intervals (e.g., weekly, as in the CPS) rather than unevenly grouped intervals. I therefore rely on the ML approach in this paper.

¹² Expected completed duration can be computed for sub-groups, but this approach quickly runs up against constraints imposed by the “curse of dimensionality” (i.e., small sample sizes when the sub-groups are defined by more than a few characteristics).

The estimator relates closely to the framework for calculating “expected completed duration” discussed in the previous section and in Appendix A. Intuitively, it is implemented by arranging separate base and continuation samples across the full range of unemployment duration intervals, referred to as “duration classes.” For example, the different duration classes could consist of individuals unemployed for 0 to 4 weeks in month $t-1$ paired with those unemployed for 5-8 weeks in month t , 5 to 8 weeks in month $t-1$ paired with 9-12 weeks in month t , 13-26 weeks in month $t-3$ paired with 27-39 weeks in month t , etc. For estimation, the characteristics of the continuation samples are compared with those of the base samples: differences in the distribution of characteristics between the base and continuation samples are used to infer the effects of the measured variables; and the declines in sample sizes between the base and continuation samples across different duration intervals reflects baseline duration dependence, which is recovered in the estimates.

Consider an example of covariate effects in this model. If individuals in the continuation samples have lower educational attainment on average than do individuals in the base samples, the regression estimates will indicate that unemployment exit rates increase with education, or equivalently that unemployment duration declines with education. These covariate effects can be constrained to be equivalent across all duration intervals, or they can be allowed to vary across duration intervals (by interacting the covariates with duration indicators). I take the former approach, for simplicity and because the estimated covariate effects in my setting are largely uniform across duration (base/continuation) pairings. The estimated covariate effects are interpreted as the

average effect of the observed covariate on continuation rates across the complete set of duration pairings.¹³

More formally, let y represent an indicator for whether an individual defined by characteristics X remains unemployed between consecutive months $t=0$ and $t=1$, which also represent the base and continuation samples in this derivation (the procedure generalizes identically to alternative duration intervals and spacings). We are interested in the conditional distribution of y , or $P(y=1|X)$. We do not observe y but instead observe \tilde{y} , which identifies whether an observation belongs to the $t=0$ or $t=1$ sample. If m_0 and m_1 represent the respective sample sizes in $t=0$ and $t=1$ (weighted using the survey weights), then the joint distribution of the observed variables X and \tilde{y} is:

$$\begin{aligned} P(X = x, \tilde{y} = 1) &= \frac{m_1}{(m_0 + m_1)} P(X = x|y = 1) \\ &= \frac{m_1}{(m_0 + m_1)} \frac{P(y = 1|X = x)P(X = x)}{P(y = 1)} \end{aligned}$$

Manipulation based on Bayes' rule and the dichotomous definition of \tilde{y} yields:

$$\begin{aligned} P(\tilde{y} = 1|X = x) &= \frac{P(X = x, \tilde{y} = 1)}{P(X = x)} = \frac{P(X = x, \tilde{y} = 1)}{P(X = x, \tilde{y} = 0) + P(X = x, \tilde{y} = 1)} \\ &= \frac{1}{1 + \frac{m_0}{m_1} \frac{P(y = 1)}{P(y = 1|X = x)}} \end{aligned}$$

¹³ In their GMM approach, Güell and Hu (2006) also impose equality of the covariate effects (except duration dependence) across duration classes, but they propose and implement an optimal GMM weighting for the estimated coefficients (except baseline duration effects) across classes.

$$= \frac{1}{1 + \alpha \frac{1}{P(y=1|X=x)}}$$

where $\alpha=(m_0/m_1)P(y=1)$. Assuming a logit specification for $P(y=1|X=x)$ yields an equation that can be estimated by maximum likelihood:

$$P(\tilde{y} = 1|X = x) = \frac{1}{1 + \alpha \frac{1 + \exp(x\beta)}{\exp(x\beta)}} = \frac{\exp(x\beta)}{\alpha + (1 + \alpha) \exp(x\beta)} \quad (1)$$

Equation (1) is essentially a logit equation for observing whether a particular observation is in the base or continuation sample, with the incorporation of a rescaling factor α that is estimated along with the β 's.

As noted by GH, the estimator is valid under the assumption that the members of the base and continuation groups are sampled from the same population, which is a feature of the stratified cross-sectional sampling scheme used for the monthly CPS.¹⁴ For my implementation, the base and continuation categories are defined to match the duration intervals used for the earlier calculation of expected completed duration, which in turn are designed to produce reliable estimates by generating cohort sizes that are sufficiently large within each duration interval (see Appendix A).¹⁵

¹⁴ This assumption holds only for observed features of the population, such as age, education, etc. The GH framework abstracts from unobserved heterogeneity, which cannot be accounted for using synthetic cohorts (unlike a true panel with repeat observations on unemployment spells).

¹⁵The primary practical difficulty in implementing this estimator is the need for identification of observations across the dual dimensions of synthetic cohorts and calendar time, for proper matching of time-varying factors such as local labor market conditions.

5. Regression Analysis

5.1 *Model specification and caveats*

For the estimation results discussed in the next sub-section, I include the full set of individual covariates listed earlier (in Table 1), plus an interaction between gender and marital status. The key variables for the house-lock test include an indicator for whether the individual lives in an owner-occupied or renter household and measures of MSA home prices. Because movements in home prices will reflect local economic conditions more generally, I also include alternative indicators of local labor market conditions as control variables, measured at a monthly frequency for each MSA.¹⁶

The essence of my test for house-lock effects on the labor market is to examine whether homeowners in MSAs that saw large house price declines experienced longer durations of unemployment than did renters and homeowners in MSAs that saw more limited declines, hence the focus will be on interaction effects between home ownership and price declines. It should be noted that the monthly CPS surveys do not track individuals who move, and as such the analyses do not directly account for mobility. This does not undermine the test for house lock in regard to unemployment durations, because the test relies specifically on homeowners' immobility: if they are affected by house lock, they will remain in the same labor market and face extended unemployment durations. Moreover, if the house lock hypothesis is correct, renters are more likely than owners to move and become employed, reducing their sample share at long durations in

¹⁶ To account for seasonal effects, a complete set of calendar month dummies (excluding one month) were also included; the results are nearly identical when these dummies are excluded.

the areas from which they moved (and not affecting measured duration in their destination MSAs).¹⁷

Despite the ability of the test to identify house lock if it exists, the lack of mobility data precludes careful analysis of the underlying adjustment mechanisms. These analyses also ignore direct job-to-job transitions and thus may be missing a reduction in such transitions for house-locked individuals, which could contribute to elevated unemployment rates. Given that direct job-to-job transitions tend to decline during periods of labor market weakness (Fallick and Fleischman 2004), it is likely that this effect, if it exists, is second-order relative to extended search by house-locked individuals.

5.2 Results

Table 2 lists the estimation results. The six columns are distinguished by the control for local labor market conditions (none, local unemployment, local employment growth) and local housing prices (dummy for MSAs with a price decline that exceeded the national decline, direct measure of the percentage change in housing prices from peak to trough). Those key variables are at the top of the table and will be discussed momentarily.

Turning first to the control variables, their coefficients are very consistent across the columns and will be discussed as a group. The equations appear to be well-specified and produce expected results. The coefficients represent each variable's association with unemployment continuation rates, hence a positive coefficient indicates that larger values

¹⁷ The inability to track movers in these data could be a problem for the house lock test if low housing and rental prices in the areas hardest hit by the housing crisis attract renter in-migrants who are willing to incur longer durations of unemployment than do renters who already live in those locations. This is a plausible scenario but is unlikely to be occurring to a substantial degree in the data, given the very low mobility rates shown in Figures 2-3.

for that explanatory variable are associated with longer unemployment durations. Individuals living in MSAs with high unemployment rates (sixth row, columns 3 and 4) or slow employment growth (seventh row, columns 5 and 6) experience significantly longer unemployment durations. Focusing on selected other coefficients on individual covariates that are statistically significant at conventional levels, younger individuals experience shorter unemployment durations than do prime-age individuals (the omitted age group is 45-54), and members of selected racial and ethnic minorities experience longer durations. Married individuals of both genders experience shorter spells.

Comparison of the coefficients across the duration categories listed at the bottom of the table indicates duration dependence. The first category, representing the baseline continuation rate between the first and second month of unemployment, is normalized to equal zero. Duration dependence is uneven across the different categories, at first increasing and then declining for longer durations. However, it shows a substantial increase between the penultimate and final categories, indicating strong positive duration dependence for very long spells of unemployment.¹⁸

Shifting back to the top of the table, the key coefficients provide no statistical evidence in favor of the house-lock hypothesis. The effects of local housing prices on unemployment durations are as expected. Individuals living in MSAs that saw especially large price drops (the MSA decline dummy in the second row) experience significantly longer unemployment durations. Similarly, the direct measure of house price changes (fourth row) indicates that individuals in MSAs in which house prices grew relatively rapidly (or declined less rapidly) experienced shorter unemployment durations. These

¹⁸ Underlying the duration dependence estimates are transition rates out of the labor force in addition to transition rates into employment; as such, these findings are not directly comparable to past work that focuses on employment transitions only.

effects of local housing market conditions hold up even when direct measures of local labor market conditions are incorporated in columns 3-6.¹⁹ By contrast, the indicators for homeownership and its interaction with the price variables indicate very small, statistically insignificant effects. The interaction terms (Owner*HPI group and Owner*%Δ HPI) reflect the direct test for house lock, and their coefficients, which are statistically indistinguishable from zero, indicate no evidence in favor of house-lock effects on unemployment duration.

These findings were subjected to two robustness tests (not displayed in a table, but available on request). First, as noted earlier, the impact of house lock is most likely to be observed when new jobs are being created and employment is growing, thereby creating incentives for individuals to move to areas where the new jobs exist. I therefore ran the same regressions focusing only on data for 2010-11, when U.S. employment was growing (compared with sharp losses in 2008-09) but housing prices remained flat to down. The results were similar to those reported in Table 2 for the period 2008-2011, providing no evidence to support the hypothesis of house-lock effects on the labor market. Second, it is important to note that the recent housing bust in the United States was accompanied by a sharp increase in home foreclosures (i.e., repossessions by the mortgage lender), which was most pronounced in areas that saw the largest decline in home prices. Because former owners who lose their homes to foreclosure are likely to be quite mobile, which will offset the house lock effect for owners who are underwater but remain in their homes, I adjusted the specification to capture the potential foreclosure effect. In particular, I grouped observations using a three-way breakdown of changes in

¹⁹ The MSA unemployment rate in columns 3-4 is endogenous with respect to unemployment durations, but the results are nearly identical when this variable is replaced with local employment growth (columns 5-6) and when both are excluded (columns 1-2).

home prices (rather than the two-way breakdown used in Table 2), based on whether homes in the MSA saw a change in prices that was much smaller (more negative), approximately equal to, or much larger than the United States as a whole. Like the two-category variable used in Table 2, the interaction of this three-category variable with the homeownership indicator yielded coefficients that were highly insignificant based on conventional statistical criteria, again providing no evidence in favor of house-lock effects on unemployment duration.

6. Conclusions

I examined the evidence in favor of house-lock effects on unemployment duration: i.e., the extent to which declining house prices during the recent U.S. housing bust reduced homeowners' geographic mobility and raised their time spent searching for a job in their local labor market. Examination of descriptive evidence regarding relative geographic mobility rates and unemployment for homeowners and renters suggested little or no impact of house lock through 2011. Similarly, more formal econometric analyses of unemployment durations that compared homeowners and renters across MSAs distinguished by the extent of home price declines also provided no evidence in favor of house-lock effects on unemployment.

The absence of a house-lock effect on recent U.S. unemployment is perhaps not surprising, given that elevated unemployment rates and modest employment growth were evident in virtually all areas of the United States during my sample frame of 2008-2011. Put simply, job seekers, whether they are homeowner or renters, faced uniformly weak employment opportunities throughout the country (with the exception of a few very

small, resource dependent states or localities). Moreover, past research on home prices and geographic mobility focused on earlier periods when foreclosure rates were well below their current levels (e.g., Chan 2001; Englehardt 2003; Ferreira, Gyourko, and Tracy 2010). By contrast, over the past few years, the mobility-inducing impact of foreclosures may have significantly offset the mobility-reducing impact of being underwater, implying that the net effect of housing market conditions on homeowner mobility has been essentially zero. Although my robustness tests included a modest, limited attempt at distinguishing between the foreclosure and underwater effects, additional research that provides more precise tests could be quite valuable.

Appendix A: Data Adjustments and Calculation of Expected Completed Duration

This appendix describes adjustments for digit preference in reported unemployment durations (a form of reporting error) and the construction of the expected completed unemployment duration series using the CPS interrupted duration measure.

Digit preference

To account for “digit preference” in the CPS unemployment duration data—the tendency for respondents to report durations as multiples of one month or half-years (i.e., multiples of 4 or 26)—I follow previous analysts by allocating a fixed share of bunched (heaped) observations to the next monthly interval. Due to greater heaping observed following the 1994 CPS survey redesign, I expanded the set of recoded durations relative to those chosen by analysts who used earlier data. In particular, I allocated 50 percent of respondents reporting the following durations of unemployment to the next weekly value: 4, 8, 12, 16, 20, 26, 30, 39, 43, 52, 56, and 78 weeks. I also reset 50 percent of the responses of 99 weeks to 100 weeks (after imposition of the top code adjustment described in the next paragraph). Sider (1985) and Baker (1992b) report that the estimated level of expected completed duration is sensitive to the allocation rule but cyclical variation and other patterns over time are not.

Calculation of Expected Completed Duration

The CPS survey collects information on the length of existing unemployment spells up to the date of the survey. The average duration measure formed from these data (and published by the BLS) will not in general correspond to the expected duration of a completed spell for a new entrant to unemployment, particularly under changing labor

market conditions such as rising unemployment (i.e., “nonsteady state” conditions). The general nonsteady-state approach to estimating expected completed duration using grouped duration data is a “synthetic cohort” approach (see e.g. Sider 1985, Baker 1992a).¹ This approach relies on the pseudo-estimation of monthly continuation rates—i.e., the probabilities that an unemployment spell will continue from one month to the next.

My application of the synthetic cohort approach to obtain nonparametric estimates of expected completed duration from grouped duration data follows M. Baker (1992a); see G. Baker and Trivedi (1985) for a more general overview. We begin with continuation probabilities, defined as the conditional probability that individuals whose unemployment spell has lasted $(j-1)$ months at time $(t-1)$ will remain unemployed into the next period:

$$f_j(t) = \frac{n(j,t)}{n(j-1,t-1)} \quad (\text{A1})$$

where $n(\cdot)$ represents the sampled number of individuals unemployed for a given number of months at the time of a particular monthly survey. In a rotating sample survey such as the CPS, the sample used to calculate the numerator and denominator differs, but under the assumption that each monthly sample represents the target U.S. population (as the CPS is constructed), this expression provides an estimate of the continuation probability for a fixed representative cohort.

¹ This is a “synthetic cohort” approach in that with a rotating monthly sample such as the CPS, the estimate of unemployment continuation probabilities is formed by comparing different groups over time, rather than by following the same individuals through time.

The product of the continuation probabilities represents the empirical survivor function, or the proportion of individuals entering unemployment at time $(t-j)$ who remain unemployed at time t :

$$G_j(t) = f_0(t)f_1(t)f_2(t)f_3(t)\dots f_j(t) \quad (\text{A2})$$

In this expression, $f_0(t)$ is the continuation probability for the entering cohort, which is defined identically as one. Assuming that the duration intervals are not all identical (e.g., not all one month), the expected completed duration in a particular month t , $D(t)$, is estimated as:

$$D(t) = 1 + \sum_{j=1}^m G_j(T_j) * (T_j - T_{j-1}) \quad (\text{A3})$$

where the T 's represent duration intervals (measured in units of the monthly sampling window) and T_m is the maximum duration measured or used.

Empirical implementation requires setting the width and number of duration intervals used for estimation. I follow Baker (1992a) in using 6 unequally spaced duration intervals and corresponding continuation probabilities; the intervals are designed to produce reliable estimates by generating cohort sizes that are sufficiently large within each interval:

- $f_1(t)$: 5-8 weeks in month t to <5 weeks in $(t-1)$
- $f_2(t)$: 9-12 weeks in month t to 5-8 weeks in $(t-1)$
- $f_3(t)$: 13-16 weeks in month t to 9-12 weeks in $(t-1)$
- $f_4(t)$: 27-39 weeks in month t to 13-26 weeks in $(t-3)$
- $f_5(t)$: 53-78 weeks in month t to 27-52 weeks in $(t-6)$
- $f_6(t)$: 100+ weeks in month t to 53-99 weeks in $(t-12)$

Note the variation in duration intervals for $f_4(t)$ - $f_6(t)$, which must be incorporated into the duration estimate based on equation (3). Then the expected completed duration is formed as:

$$D(t) = 1 + f_1 + f_2 f_1 + f_3 f_2 f_1 + 3 f_4 f_3 f_2 f_1 + 6 f_5 f_4 f_3 f_2 f_1 + 12 f_6 f_5 f_4 f_3 f_2 f_1 \quad (\text{A4})$$

where the time identifier (t) has been suppressed on the right-hand side of (4) for simplicity. $D(t)$ is defined as the expected duration of unemployment (in months) for a cohort that enters unemployment at t and faces current economic conditions throughout the unemployment spells of cohort members. For the charts displayed in this paper, I estimated expected completed duration for samples of homeowners and renters separately; estimation by group proceeds by first restricting the unemployment sample to the specified group, than estimating expected completed duration as described above.²

² Following past practice (e.g., Sider 1985), I multiplied estimates of expected duration in months by 4.3 to obtain expected duration in weeks for the charts.

References


- Aaronson, Daniel, and Jonathan Davis. 2011. "How much has house lock affected labor mobility and the unemployment rate?" *Chicago Fed Letter* Number 290 (September).
- Baker, Michael. 1992a. "Unemployment Duration: Compositional Effects and Cyclical Variability." *American Economic Review* 82(1): 313-321.
- Baker, Michael. 1992b. "Digit Preference in CPS Unemployment Data." *Economics Letters* 39 (1, May): 117-121.
- Baker, G.M., and P.K. Trivedi. 1985. "Estimation of Unemployment Duration from Grouped Data: A Comparative Study." *Journal of Labor Economics* 3(2): 153-174.
- Buchmueller, Thomas C. and Robert G. Valletta. 1996. "The Effects of Employer-Provided Health Insurance on Worker Mobility." *Industrial and Labor Relations Review* 49(3, April): 439-55.
- Carlson, John, and Michael Horrigan. 1983. "Measures of Unemployment as Guides to Research and Policy: Comment." *American Economic Review* 73(5): 1143-52.
- Chan, Sewin. 2001. "Spatial Lock-in: Do Falling House Prices Constrain Residential Mobility?" *Journal of Urban Economics* 49 (3): 567-86.
- Coulson, N. Edward, and Lynn M. Fisher. 2009. "Housing tenure and labor market impacts: The search goes on." *Journal of Urban Economics* 65: 252-64.
- Donovan, Colleen, and Calvin Schnure. 2011. "Locked in the House: Do Underwater Mortgages Reduce Labor Market Mobility?" Working Paper, Freddie Mac, May.
- Engelhardt, Gary V. 2003. "Nominal Loss Aversion, Housing Equity Constraints, and Household Mobility: Evidence from the United States." *Journal of Urban Economics* 53 (1): 171-95.
- Farber, Henry S. 2010. "Job Loss and Unemployment in the 21st Century: The Great Recession in Labor Market Perspective." Working Paper #560, Industrial Relations Section, Princeton University. December.
- Fallick, Bruce, and Charles A. Fleischman. 2004. "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows." Finance and Economics Discussion Series 2004-34, Board of Governors of the Federal Reserve System.
- Ferreira, Fernando, Joseph Gyourko, and Joseph Tracy. 2010. "Housing Busts and Household Mobility." *Journal of Urban Economics* 68: pp. 34-45.

- Fletcher, Michael A. 2010. "In slow economy, Americans stay put." *Washington Post*, July 30.
- Genesove, David, and Christopher Mayer. 2001. "Loss-aversion and Seller Behavior: Evidence from the Housing Market." *Quarterly Journal of Economics* 116(4): 1233-60.
- Güell, Maia, and Luojia Hu. 2006. "Estimating the probability of leaving unemployment using uncompleted spells from repeated cross-section data." *Journal of Econometrics* 133: 307-341.
- Herkenhoff, Kyle F., and Lee E. Ohanian. 2011. "Labor Market Dysfunction During the Great Recession." NBER Working Paper 17313 (August). Cambridge, MA: National Bureau of Economic Research.
- Horrigan, Michael. 1987. "Time Spent Unemployed: A New Look at Data from the CPS." *Monthly Labor Review* 110(7, July): 3-15.
- Kaplan, Greg, and Sam Schulhofer-Wohl. 2012. "Interstate Migration Has Fallen Less Than You Think: Consequences of Hot Deck Imputation in the Current Population Survey." *Demography* 49(3, Aug.): 1061-74.
- Karahan, Fatih, and Serena Rhee. 2011. "Geographical Reallocation and Unemployment during the Great Recession: The Role of the Housing Bust." Manuscript, Department of Economics, University of Pennsylvania.
- Modestino, Alicia Sasser, and Julia Dennett. 2011. "Are American Homeowners Locked Their Houses? The Impact of Housing Market Conditions on State-to-State Migration." Manuscript, Federal Reserve Bank of Boston, December.
- Molloy, Raven Saks, Christopher L. Smith, and Abigail Wozniak. 2011. "Internal Migration in the United States." *Journal of Economic Perspectives* 25(3): 173-196.
- Munch, Jakob R., Michael Rosholm, and Michael Svarer. 2006. "Are Homeowners Really More Unemployed?" *Economic Journal* 116: 991-1013.
- Oswald, Andrew. 1996. "A Conjecture of the Explanation for High Unemployment in the Industrialised Nations: Part I." Warwick University Economic Research Paper No. 475.
- Schmitt, John, and Kris Warner. 2011. "Deconstructing Structural Unemployment." Report, Center for Economic and Policy Research, March.
- Schulhofer-Wohl, Sam. 2012. "Negative Equity Does Not Reduce Homeowners' Mobility." FRB Minneapolis *Quarterly Review*, 35(1, February): 2-14.
- Sterk, Vincent. 2010. "Home Equity, Mobility, and Macroeconomic Fluctuations." DNB (DeNederlandsche Bank) Working Paper 265 (November).

Sider, Hal. 1985. "Unemployment Duration and Incidence: 1968-82." *American Economic Review* 75(3, June): 461-472.

Valletta, Robert G. 2011. "Rising Unemployment Duration in the United States: Composition or Behavior?" Mimeo, Federal Reserve Bank of San Francisco, May.


Figure 1: House Prices (FHFA), 2005Q1 - 2011Q4
(Normalized to 2007Q1 US peak)




Note: Gray area denotes the most recent NBER recession period.

Figure 2: Geographic Mobility Rates,
by Home Ownership (1988-2011)
(share of group)

Panel A: Across states




Panel B: Across counties




Note: Census Bureau estimates using March CPS data.

Figure 3: Geographic Mobility Rates,
by Home Ownership (1988-2011)
(share of total population)

Panel A: Across states





Panel B: Across counties



Note: Census Bureau estimates using March CPS data.


Figure 4: Unemployment Rates, by Home Ownership (Jan. 1989 - Dec. 2011)




Note: Authors' calculations from monthly CPS microdata and BLS labor force series. Gray bars denote NBER recession dates.

Figure 5: Unemployment Duration, by Home Ownership (through 12/11)

Panel A: MSAs with (HPI decline)>US



Panel B: MSAs with (HPI decline) < US



Note: Authors' calculations from monthly CPS microdata. Duration measured in expected completed form, see text for description. MSA house prices from FHFA, change measured over 2007q1-2011q2. Gray bar denotes NBER recession dates.

Table 1: Population Characteristics, by Homeowner Status
(mean values; from CPS micro data, Jan. 2008 - Dec. 2011)

<u>Variable</u>	<u>Homeowners</u>	<u>Renters</u>
Age (years)	38.7	34.0
	<u>(shares)</u>	
Education		
<High School	0.176	0.264
High School	0.349	0.362
Some College	0.283	0.258
College Grad	0.142	0.089
>College	0.051	0.027
Race		
White	0.659	0.434
Black	0.132	0.256
Hispanic	0.140	0.234
Asian	0.041	0.038
Other	0.028	0.039
Married	0.449	0.340
Female	0.407	0.447
Vet	0.071	0.049
Sample Size	147459	110243
Sample share (weighted)	0.559	0.441

Table 2: Unemployment Duration (ML estimates, data for Jan. 2008 - Dec. 2011)

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	<u>No cyclical control</u>		<u>Control for local unemp rate</u>		<u>Control for local emp growth</u>	
	<u>HPI groups</u>	<u>HPI change</u>	<u>HPI groups</u>	<u>HPI change</u>	<u>HPI groups</u>	<u>HPI change</u>
Home owner	0.0182 (0.0253)	0.0252 (0.0275)	0.0467 (0.0400)	0.0647 (0.0440)	0.0145 (0.0258)	0.0233 (0.0279)
MSA HPI decline>U.S. (dummy)	0.311** (0.0327)		0.369** (0.0698)		0.283** (0.0333)	
Owner*HPI group	-0.0116 (0.0389)		0.0126 (0.0731)		-0.0103 (0.0398)	
%ΔHPI (peak to trough)		-0.00892** (0.000922)		-0.00936** (0.00219)		-0.00795** (0.000916)
Owner*%ΔHPI		0.000394 (0.00109)		0.000792 (0.00214)		0.000521 (0.00110)
MSA unemployment rate			0.143** (0.0188)	0.137** (0.0206)		
MSA emp growth (12-month %)					-4.568** (0.413)	-4.346** (0.409)
Age 16-19	-0.475** (0.0356)	-0.480** (0.0359)	-0.664** (0.0625)	-0.663** (0.0636)	-0.481** (0.0363)	-0.487** (0.0364)
Age 20-24	-0.328** (0.0329)	-0.333** (0.0331)	-0.484** (0.0591)	-0.482** (0.0598)	-0.335** (0.0336)	-0.340** (0.0337)
Age 25-34	-0.194** (0.0301)	-0.194** (0.0304)	-0.260** (0.0520)	-0.260** (0.0516)	-0.201** (0.0307)	-0.202** (0.0308)
Age 35-44	-0.149** (0.0313)	-0.151** (0.0316)	-0.172** (0.0547)	-0.174** (0.0541)	-0.158** (0.0319)	-0.161** (0.0321)
Age 55-64	0.0480 (0.0386)	0.0510 (0.0392)	0.167* (0.0738)	0.166* (0.0734)	0.0523 (0.0397)	0.0534 (0.0399)
Age 65+	-0.0161 (0.0622)	-0.0195 (0.0632)	-0.00415 (0.119)	0.00897 (0.118)	-0.0173 (0.0639)	-0.0205 (0.0645)

(continued)

Table 2 (continued)

<u>VARIABLES</u>	(1)	(2)	(3)	(4)	(5)	(6)
	<u>No cyclical control</u>		<u>Control for local unemp rate</u>		<u>Control for local emp growth</u>	
	<u>HPI groups</u>	<u>HPI change</u>	<u>HPI groups</u>	<u>HPI change</u>	<u>HPI groups</u>	<u>HPI change</u>
Education: HS degree	0.0176 (0.0275)	0.0179 (0.0278)	0.100* (0.0501)	0.0974 (0.0498)	0.0181 (0.0280)	0.0174 (0.0281)
Some college	-0.0379 (0.0290)	-0.0342 (0.0293)	-0.0177 (0.0486)	-0.0109 (0.0481)	-0.0312 (0.0297)	-0.0284 (0.0298)
College degree	-0.0815* (0.0364)	-0.0703 (0.0368)	-0.0888 (0.0617)	-0.0724 (0.0610)	-0.0778* (0.0372)	-0.0680 (0.0374)
Graduate degree	-0.0978 (0.0537)	-0.0856 (0.0541)	-0.0992 (0.0936)	-0.0919 (0.0919)	-0.0906 (0.0551)	-0.0802 (0.0552)
Race/ethnic: Black	0.204** (0.0287)	0.218** (0.0293)	0.430** (0.0643)	0.436** (0.0694)	0.215** (0.0296)	0.226** (0.0299)
Hispanic	-0.0600* (0.0265)	-0.0758** (0.0269)	-0.151** (0.0471)	-0.144** (0.0467)	-0.0281 (0.0273)	-0.0415 (0.0276)
Asian	0.183** (0.0558)	0.193** (0.0568)	0.288** (0.106)	0.309** (0.105)	0.218** (0.0576)	0.226** (0.0582)
Other	0.0915 (0.0526)	0.0909 (0.0531)	0.242* (0.101)	0.238* (0.0998)	0.109* (0.0542)	0.108* (0.0544)
Military veteran	-0.0590 (0.0395)	-0.0591 (0.0400)	-0.0989 (0.0705)	-0.0989 (0.0699)	-0.0574 (0.0405)	-0.0576 (0.0407)
Married	-0.125** (0.0278)	-0.126** (0.0280)	-0.172** (0.0487)	-0.170** (0.0480)	-0.126** (0.0284)	-0.128** (0.0285)
Female	-0.0205 (0.0246)	-0.0205 (0.0248)	-0.0225 (0.0422)	-0.0209 (0.0416)	-0.0133 (0.0252)	-0.0141 (0.0253)
Female*married	0.0696 (0.0388)	0.0708 (0.0392)	0.120 (0.0689)	0.118 (0.0681)	0.0696 (0.0397)	0.0709 (0.0399)

(continued)

Table 2 (continued)

<u>VARIABLES</u>	(1)	(2)	(3)	(4)	(5)	(6)
	<u>No cyclical control</u>		<u>Control for local unemp rate</u>		<u>Control for local emp growth</u>	
	<u>HPI groups</u>	<u>HPI change</u>	<u>HPI groups</u>	<u>HPI change</u>	<u>HPI groups</u>	<u>HPI change</u>
Duration months 1-2 (omitted)	0	0	0	0	0	0
Duration months 2-3	0.358** (0.0377)	0.360** (0.0384)	0.848** (0.152)	0.820** (0.156)	0.369** (0.0397)	0.369** (0.0397)
Duration months 3-4	0.143** (0.0356)	0.142** (0.0360)	0.312** (0.0833)	0.304** (0.0841)	0.141** (0.0366)	0.139** (0.0368)
Duration quarters 2-3	-0.437** (0.0261)	-0.446** (0.0264)	-0.768** (0.0662)	-0.758** (0.0714)	-0.457** (0.0268)	-0.464** (0.0270)
Duration quarters 3/4-5/6	-0.241** (0.0280)	-0.250** (0.0283)	-0.545** (0.0671)	-0.533** (0.0711)	-0.239** (0.0286)	-0.247** (0.0288)
Duration years 2-3+	0.416** (0.0478)	0.413** (0.0487)	1.293** (0.444)	1.192** (0.420)	0.458** (0.0489)	0.451** (0.0488)
Alpha (scaling parameter)	0.609** (0.0199)	0.618** (0.0207)	1.013** (0.0495)	1.005** (0.0559)	0.628** (0.0208)	0.633** (0.0209)
Observations	174678	174678	174678	174678	174678	174678

** p<0.01, * p<0.05

Note: Includes month dummies (coefficients not shown). HPI refers to the FHFA house price index series, measured for the U.S. and 235 MSAs; HPI variables refer to MSA price changes from the national peak to the end of the data frame (2007Q1 to 2011Q4). Omitted categories for categorical variables are age 45-54, education<(high school degree), white. Robust standard errors in parentheses