Labour Economics 25 (2013) 86–97

Contents lists available at ScienceDirect

Labour Economics

journal homepage: www.elsevier.com/locate/labeco

House lock and structural unemployment☆

Robert G. Valletta *

Federal Reserve Bank of San Francisco, 101 Market Street, San Francisco, CA 94105-1530, USA

HIGHLIGHTS

- The recent sharp decline in U.S. house prices may have affected unemployment.
- Underwater homeowners may extend job search in their local labor market.
- I test for such “house lock” effects using duration data from CPS cross-sections.
- The results suggest that house lock has not constrained the aggregate labor market.

ARTICLE INFO

Article history:
Received 2 November 2012
Received in revised form 2 April 2013
Accepted 3 April 2013
Available online 12 April 2013

JEL classification:
J6
R31

Keywords:
Unemployment
House prices
Mobility

ABSTRACT

A recent decline in internal migration 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 for the years 2008–11. Because house lock is likely to 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 on an econometric method that enables the estimation of individual and aggregate covariate effects on unemployment durations using repeated cross-section data. I do not uncover systematic evidence to support the house-lock hypothesis.

Published by Elsevier B.V.

1. Introduction

During the 2007–09 U.S. recession and its aftermath, the rate of internal migration across geographic areas 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 this 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. Like other recent authors, I will refer to this phenomenon as “house lock.” The possibility of widespread house lock in recent years 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 (Fletcher, 2010). To the extent that this phenomenon exists, it represents a form of structural unemployment that may persist after the U.S. economy has fully recovered from the recent recession, implying a higher equilibrium or “natural” rate of unemployment.

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 share of underwater homes are associated with reduced geographic mobility by homeowners (e.g., Chan, 2001; Engelhardt, 2003; Ferreira et al., 2010). A separate literature has investigated the relationship between home ownership and unemployment rates, at the individual or regional level (e.g., Oswald, 1996; Munch et al., 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.

☆ I thank Leila Bengali and Katherine Kuang for their 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). I also thank the editor and the two anonymous referees for their comments. 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.

* Tel.: +1 415 974 3345; fax: +1 415 977 4084.
E-mail address: rob.valletta@sf.frb.org.

0927-5371/$ – see front matter. Published by Elsevier B.V.
http://dx.doi.org/10.1016/j.labeco.2013.04.002

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 variation in unemployment outcomes and ownership status 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 minimize 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. 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. Other recent papers have attempted to empirically assess the extent of house lock by examining geographic mobility, making only indirect inferences about labor market outcomes (e.g., Donovan and Schnure, 2011; Modestino and Dennett, 2013). By contrast, I assess house lock through direct examination of the relationship between housing market conditions and labor market outcomes, focusing on unemployment duration. 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, with a further comparison made across geographic areas differentiated by the intensity of the home price decline. Compared with renters, homeowners face additional financial constraints in weak housing markets 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.

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 local economic conditions, conditional on detailed individual characteristics and duration dependence (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 presents my empirical framework, focusing on identification. Section 4 describes the CPS unemployment data and provides descriptive evidence regarding unemployment rates and durations for owners and renters. The formal econometric approach for the analysis of unemployment duration is described in Section 5, and Section 6 presents the estimation results. To summarize briefly, 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. The final section summarizes the findings and discusses their broader implications.

2. Home ownership, mobility, and unemployment

2.1. Prior 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.[1] Substantial support for the first link has been found 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 choose to accept the reputational and credit costs arising from default rather than the immediate financial cost of selling their home at a loss.[2] 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 et al., 2010; although see Schulhofer-Wohl, 2012 for an opposing view).[3] Mobility may also be suppressed for households not facing direct financial constraints, as a result of nominal loss aversion that causes them to place greater weight on capital losses than on equivalent gains (Genesove and Mayer, 2001; Engelhardt, 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.[4] 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 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. Oswald offered this as an explanation for his estimates of a positive correlation between unemployment rates and the proportion of homeowners across countries and sub-national regions.

Subsequent work has focused on refining the "Oswald hypothesis" and the associated empirical tests. Formal models of home ownership and job search produce ambiguous predictions regarding the relationship between ownership and unemployment outcomes for individuals or geographic areas. Munch et al. (2006) specified a model of job search that allows for transitions into employment within or outside the local labor market. In their model, unemployed individuals living in owner-occupied housing face higher moving costs than renters, which lowers owners' 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.

[1] In this paper, the terms "homeownership" and "homeowners" will be used synonymously with "individuals living in owner-occupied housing," as will similar terms for renters.

[2] This is commonly referred to as "strategic default;" see for example Guiso et al. (2009).

[3] Chan (2001) and Ferreira et al. (2010) discuss a number of features of housing markets that create financial constraints for low or negative equity homeowners who wish to move. These features include the possibility that they are unable to finance the transaction costs of selling their home, or that negative equity and rising interest rates may require owners to put up additional cash beyond standard closing costs for purchasing a new home.

[4] 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 Engelhardt (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.

Coulson and Fisher (2009) compare the implications of alternative search models for homeowner versus renter unemployment at the individual and aggregate 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. It is important to note that analyses of the link between homeownership and unemployment have focused on equilibrium or steady-state relationships between the ownership choice and job search. By contrast, I am examining a period of unusual dislocation in the U.S. housing and labor markets. Given the large declines in U.S. home prices since 2007, historically low geographic mobility rates, and persistently high unemployment rates, the link between homeownership and unemployment may be unusually strong in recent years.

A number of 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) calibrated simulation models of the impact of the housing downturn on geographic mobility and unemployment (Herkenhoff and Ohanian, 2011; Karahan and Rhee, 2011; Sterk, 2010); (ii) empirical studies of recent changes in geographic mobility and their potential implications for labor market outcomes (Aaronson and Davis, 2011; Donovan and Schnure, 2011; Modestino and Dennett, 2013; Molloy et al., 2011); and (iii) empirical assessments of the relationship between housing market conditions, geographic mobility, and labor market conditions (Farber, 2012; 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 in Section 3 below, is closest in spirit and execution to the papers in (iii). However, relative to those papers, which focused on mobility and did not directly examine the links between housing market conditions and individual labor market outcomes, I rely on an empirical framework that enables direct assessment of the influence of local housing market conditions on unemployment duration.

2.2. Housing market and geographic mobility trends

The house-lock hypothesis is predicated on a decline in home prices that reduced geographic mobility. Fig. 1 illustrates the net decline and varied movement in prices during the recent housing bust (year 2005 forward), for the nation as a whole and selected metropolitan areas (Metropolitan Statistical Areas, or "MSAs"). The series displayed are from the repeat-sales index for single-family homes compiled by the Federal Housing Financing Agency (FHFA).[5] From

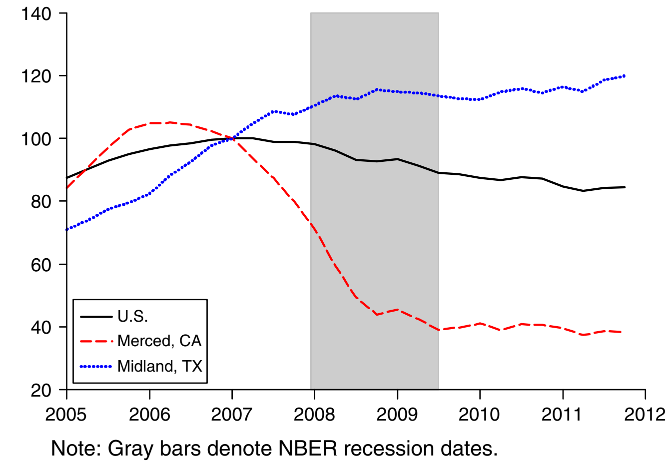

Fig. 1. House prices (FHFA), 2005Q1–2011Q4. (Normalized to 100 in 2007Q1).

the peak in early 2007 through the end of 2011, nominal home prices across the U.S. fell about 16 percent on average.[6] 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. Price changes for other MSAs lie along the full spectrum between these extremes. The underlying variation in price movements across the full sample of MSAs will be exploited for the empirical tests in Section 6.

The recent decline in home prices affected large numbers of households in the United States. Home ownership rates peaked at about 69 percent in 2005. They fell subsequently, to 66 percent by the end of 2011, but homeowners still account for the majority of American households.[7] Price declines caused large numbers of homeowners to slip underwater (to owe more on their home than its market price). Available estimates indicate that at the end of 2011 nearly one in four U.S. residential properties with a mortgage were underwater.[8] On a base of about 75 million owner-occupied housing units, this implies 17 million underwater households, which is large relative to the 13–14 million unemployed individuals at the time. Given wide dispersion in price declines across states, the share of underwater mortgages also varied widely, reaching as high as about 60 percent in Nevada.

These adverse housing market conditions were offset somewhat by government policy responses to mitigate the crisis. The U.S. central bank (the Federal Reserve) dropped its target interest rates essentially to zero, causing mortgage rates to drop to historical lows and supporting mortgage refinance activities. This had only limited benefits for underwater homeowners, however; most have limited financial

[5] Among the various reliable sources of U.S. home prices, the FHFA index provides series for the widest available range of metropolitan areas. Additional information is available on their website: http://www.fhfa.gov/Default.aspx?Page=14.

[6] The last episode of a sustained decline in U.S. housing prices began in the 1920s and continued through the Great Depression. Limited data on home prices is available for this period, and White (2009) argues that available series underestimate the amplitude of price movements during the 1920s and 1930s. Adjusting for likely biases, he concludes that boom and bust cycle during this period was comparable to that observed in the run-up to the most recent U.S. recession. By contrast, unemployment rates in the Great Depression reached approximately 25 percent, much higher than the maximum in the recent downturn (10.0 percent in late 2009). Home ownership rates were much lower in the Great Depression than currently (around 45 percent, compared with 65–70 percent recently), and internal migration rates also were quite low (Molloy et al., 2011).

[7] Data on homeownership rates and counts are from the U.S. Census Bureau: http://www.census.gov/housing/hvs/.

[8] The most commonly cited source of data on underwater homes in the U.S. is the real estate data provider CoreLogic, which estimated the share of underwater homes at 22.8 percent in the fourth quarter of 2011 (press release at http://www.corelogic.com/about-us/news/corelogic-reports-negative-equity-increase-in-q4-2011.aspx).

assets other than the value of their home and have been unable to take advantage of low refinancing rates, particularly given more conservative mortgage lending standards in the aftermath of the housing crash. More generally, the federal government instituted a variety of mortgage modification programs intended to help underwater homeowners maintain ownership by reducing their current payments, typically through changes in the mortgage interest rate or time profile of payments. By providing a financial cushion to struggling homeowners and making it easier for them to stay in their homes even if they are unemployed, such measures will tend to intensify rather than mitigate house-lock effects (as implied by the model and simulations of Herkenhoff and Ohanian, 2011). Despite these measures, U.S. foreclosure rates in the past few years greatly exceeded any previously recorded highs in data extending back at least to the early 1970s.

Fig. 2 shows that geographic mobility has declined over the period approximately corresponding to the U.S. housing market downturn. The mobility calculations displayed in this figure are from the U.S. Census Bureau, based on reported geographic moves from the Annual Demographic Supplement to the monthly CPS survey (conducted each March). Fig. 2 (Panel A) displays overall and group-specific mobility rates across states for the period over which separate data on owners and renters is available (back to 1988).[9] In 2009 the overall rate of interstate mobility fell to 1.6 percent and subsequently has stayed largely stable. As noted earlier, this is the lowest level recorded in U.S. Census Bureau statistics, which began in 1948.[10]

Panel A of Fig. 2 shows that renter mobility substantially exceeds owner mobility on average.[11] Mobility for both groups fell noticeably after 2005.[12] However, Panel B, which displays the difference in mobility rates between renters and owners from Panel A, shows that mobility fell more for renters than for owners between 2005 and 2007, near the beginning of the housing downturn. Since then, the relative mobility rates of owners and renters have remained largely constant. While the decline in owner mobility is broadly consistent with house lock, the larger decline in mobility for renters than owners is not. To draw any firm conclusions, however, it is necessary to examine direct evidence on labor market outcomes, which I describe in subsequent sections.

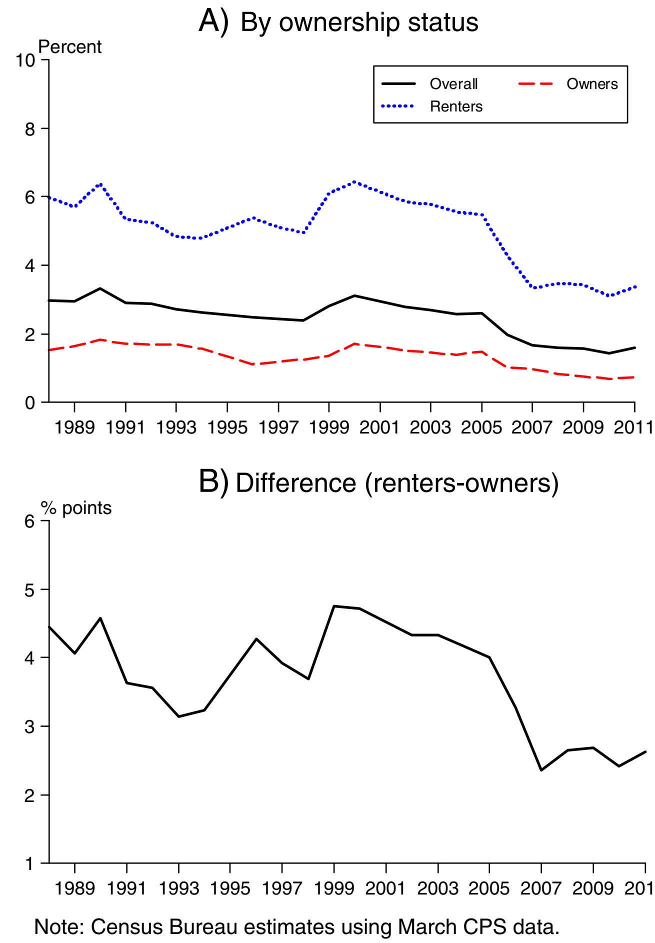

Fig. 2. Geographic mobility rates, annual (1988–2011) (Share of group).

3. Empirical framework and identification

As discussed in Section 2.1, existing theoretical models and empirical analyses have not reached definitive conclusions regarding the relationship between homeownership and unemployment. Moreover, past models focused primarily on equilibrium or steady-state relationships between homeownership, geographic mobility, and unemployment. By contrast, I am focusing on a period of severe dislocation in housing and labor markets. Given these considerations, rather than specifying a precise theoretical model, I will describe in broad terms the economic environment and individual choices that are relevant for my analysis.

My analyses will focus on unemployed individuals, distinguishing between homeowners and renters. Consider an initial situation in which individuals have chosen to live in one of two different geographic areas, based on their expected income streams and local amenities. I assume that the areas are sufficiently far apart that individuals must live and work in the same area. Now consider an adverse economic shock that has uneven effects across these two areas, increasing unemployment and reducing expected income and housing prices more in area 1 than area 2.[13] This creates incentives for individuals to move from area 1 to area 2. However, the large decline in home prices in area 1 may cause a relatively large number of homeowners to slip underwater and face higher moving costs, as suggested by the empirical research cited in the previous section. Compared with renters, who are largely unaffected by these changes in home prices, locked-in homeowners are likely to extend their job search and experience longer unemployment durations in the local labor market.

Given these considerations, the empirical analysis will focus on job search in the local labor market for unemployed individuals. I assume that homeowners and renters can exit unemployment through job-finding in the local labor market or searching in an alternative geographic area with better labor market conditions. In the monthly CPS data that I will use for my empirical analysis, either choice will correspond to the end of a measured spell of unemployment and hence a reduction in expected unemployment duration (as discussed in more detail below).

A key estimation obstacle is the likelihood of systematic unobserved as well as observed differences between owners and renters that

[9] Mobility across counties, which includes moves both within and across states, exhibits similar patterns over time. I focus here on interstate migration because it is more likely to occur for job-related reasons than is intrastate migration (Molloy et al., 2011).

[10] U.S. geographic mobility has been on a long-term downward trend, with rates in prior decades well above those from the 1990s and 2000s. See the historical data tables available here: http://www.census.gov/population/www/socdemo/migrate.html.

[11] This partly reflects other characteristics of the owner and renter groups, as discussed in conjunction with Table 1 in Section 4.1 below.

[12] 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). However, other data sources also show a decline in mobility over this period (see e.g. Modestino and Dennett, 2013).

[13] Like the variation in housing price movements described in Section 2.2, wide variation in employment shocks and unemployment rates was evident across states and metropolitan areas during the recent recession and its aftermath. Peak state unemployment rates ranged as high as about 14 percent in Michigan and Nevada and 12½ percent in California, compared with peaks slightly above 8 percent in a few large states such as Texas and Minnesota and only 4 to 5 percent in smaller states such as North Dakota and Nebraska.

affect their labor market outcomes. I therefore will not rely on simple differences in unemployment outcomes between owners and renters in my analyses. Instead, I will use the effects of housing market conditions on renter unemployment duration as a comparative baseline for owners, which corresponds to a "difference-in-difference" approach.

The following equation represents the difference-in-difference estimation framework that will be implemented in subsequent sections.

$$D_{ijt} = \beta X_{ijt} + \gamma Z_{jt} + \delta_1 O_{ijt} + \delta_2 H_{jt} + \delta_3 \left(O_{ijt} \cdot H_{jt}\right) + \varepsilon_{ijt} \tag{1}$$

The subscripts (*i, j, t*) refer to person, location (metropolitan area), and the date (the CPS data are available at a monthly frequency). D represents the duration of unemployment or a related measurement, which will be specified in more detail in Section 5. X and Z refer to person-specific and location-specific characteristics, β and γ are coefficient vectors to be estimated, and ε is a random error term reflecting unobserved factors that are assumed to be uncorrelated with the other explanatory variables in the model.

The key variables are indicators for the individual's home ownership status O and the severity of the house price decline in the individual's metropolitan area H, along with the interaction between these two variables. The coefficient on the interaction variable $O{\cdot}H$ represents the difference in conditional unemployment duration between homeowners in areas with large price declines and those in areas with small price declines, relative to the same difference for renters.[14] If house lock exists, it will be most readily observable for homeowners in geographic areas that experienced the most adverse conditions in housing markets, through the channel of increased moving costs for underwater homeowners that preclude moving to areas where labor market conditions are more favorable. However, it is important to norm this comparison relative to a control or placebo group that is not directly affected by the decline in home prices (renters).

The primary advantage of this difference-in-difference approach is that it addresses the most prominent threats to identification of a house-lock effect, including potential endogeneity and selection effects in the relationship between homeownership and labor market outcomes (see e.g., Lovenheim and Mumford, forthcoming). Rather than relying on instrumental variables to purge any unobserved correlation between ownership status and labor market outcomes (e.g., Coulson and Fisher, 2009), this approach eliminates the influence of unobservables by comparing differences across sub-groups of homeowners to the same difference across sub-groups of renters, with the sub-groups defined by the degree of their exposure to the housing downturn.

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.[15] My empirical framework accommodates this drawback by focusing on job search and unemployment durations in the local labor market; individuals who move are treated as exiting local unemployment. Put differently, my empirical 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, whereas individuals (renters) who are not affected by house lock will find jobs relatively quickly or move to another area, reducing measured duration in the area from which they moved. 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 unlikely to be a widespread phenomenon for two reasons: (i) geographic mobility for purposes of job search is relatively infrequent (Molloy et al., 2011); (ii) the decline in housing prices in recent years generally was not accompanied by a similar-sized decline in rents, implying limited incentives for relocation by renters.

Despite the ability of the empirical test to identify house lock through immobility and prolonged unemployment, the lack of mobility data precludes 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.

4. CPS unemployment data and descriptive analyses

4.1. CPS unemployment data

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. I use data for 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, thereby incorporating a period of employment expansion during which house lock may be a significant constraint on homeowners' desired location decisions and unemployment outcomes.

Observations were included for all individuals age 16 and older identified as unemployed in the survey and who live in one of the 235 MSAs for which the FHFA housing price series described in Section 2.2 are available. 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.

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.[16] Table 1 lists means for a standard set of individual control variables, with the sample divided into owner and renter groups. To facilitate comparisons that are relevant for the difference-in-difference analyses, the owner and renter samples are further stratified based on whether the MSA in which each individual lives experienced price declines that exceeded or fell short of the U.S. average (16 percent, as noted in Section 2.2). The figures in the table indicate that homeowners 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.[17] On the other hand, average unemployment duration for owners and renters tends to be quite similar (as discussed in more detail in the next section).

The additional breakdown in Table 1 into sub-samples defined by the extent of housing price decline reveals that homeowners and renters in areas with large declines have characteristics that are similar to those in areas with small declines. The primary exception

[14] Madrian (1994) and Buchmueller and Valletta (1996) demonstrate this proposition in the context of job mobility analyses.

[15] Panel data sets that could be used for direct analysis of geographic mobility and labor market outcomes, such as the Panel Study of Income Dynamics (PSID) or Survey of Income and Program Participation (SIPP) provide relatively small sample sizes and limited geographic detail compared with the CPS repeated cross-section data that I use.

[16] 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.

[17] 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. 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.

Table 1
Characteristics of unemployed, by Homeowner Status and House Price Change (MSA price change compared with U.S. average change) (mean values; from CPS micro data, Jan. 2008–Dec. 2011).

	Homeowners		Renters	
	(HPI decline) > US	(HPI decline) < US	(HPI decline) > US	(HPI decline) < US
Individual characteristics				
	(Means)			
Unemployment duration (weeks)	30.7	26.2	29.3	25.4
Age (years)	39.1	37.8	34.2	32.8
	(Shares)			
Education				
<High school	0.181	0.193	0.262	0.273
High school	0.330	0.358	0.348	0.372
Some college	0.299	0.269	0.280	0.263
College grad	0.144	0.132	0.083	0.071
>College	0.045	0.048	0.027	0.022
Race/ethnicity				
White	0.611	0.693	0.420	0.468
Black	0.110	0.155	0.203	0.316
Hispanic	0.195	0.107	0.287	0.159
Asian	0.049	0.021	0.043	0.021
Other	0.035	0.025	0.047	0.037
Married	0.462	0.435	0.362	0.304
Female	0.409	0.412	0.439	0.459
Veteran (military)	0.075	0.075	0.052	0.056
MSA characteristics				
	(Means)			
Change in HPI (%) (2007Q1–2011Q4)	−35.9	−4.8	−37.4	−4.7
Unemployment rate (%)	10.2	8.1	10.5	8.0
Emp growth (12-month %)	−1.6	−0.8	−1.4	−0.7
Sample size	29945	38369	23600	29545
Sample Share	0.247	0.316	0.194	0.243

Note: Unemployed individuals living in MSAs, from monthly CPS files. All tabulations weighted by CPS individual labor force weights.

is the share of individuals with Hispanic ethnicity, who are more likely to live in areas that experienced large declines in home prices (e.g., states in the southwestern region of the country). The tabulations listed at the bottom of the table indicate that MSA-level labor market conditions, as reflected in the local unemployment rate and pace of employment growth, are less favorable in areas that experienced a larger decline in housing prices. Overall, the descriptive statistics in Table 1 suggest that the owner and renter sub-samples grouped by the extent of house price declines are quite similar to one another, although the local housing and labor market conditions differ substantially between these groups. This is reassuring with respect to the difference-in-difference design to be used below, as it suggests that the unobservable within-group differences for owners and renters are likely to be limited.

4.2. Descriptive analyses of unemployment rates and duration

Fig. 3, Panel A, confirms the expectation of lower unemployment rates for homeowners than for renters, based on the individual characteristics listed in Table 1. Unemployment rates for renters typically are about twice those of owners.[18] Panel B displays the difference in unemployment rates between the renter and owner groups. This gap increased substantially during the recent recession and only began to decline noticeably in 2011. Like the mobility series discussed earlier (Fig. 2, Section 2.2), this evidence is not supportive of the house-lock hypothesis, which would lead to an increase rather than decrease in the relative unemployment rate for homeowners. Given the very sharp differences in the characteristics of owners and renters, however, this evidence is far from definitive.

The focus of my analysis is unemployment duration measured in the CPS, which reflects time spent in job search. In the CPS microdata, unemployment duration is measured as the duration of ongoing spells at the time of the survey ("stock-based sampling"), rather than duration for individuals who are tracked as their unemployment spell proceeds ("flow-based sampling"). The stock-based duration 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). These series are subject to well-known biases with respect to measurement of expected duration for an individual entering unemployment, particularly over-sampling of long spells, along with 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 the stock-based duration data from the monthly CPS, 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," consisting of groups with the same unemployment duration who are followed over time (rather than following individuals over time, as in a true cohort analysis). 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.

Fig. 4 compares expected unemployment duration in weeks, for owners and renters, over the period beginning just prior to the housing

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

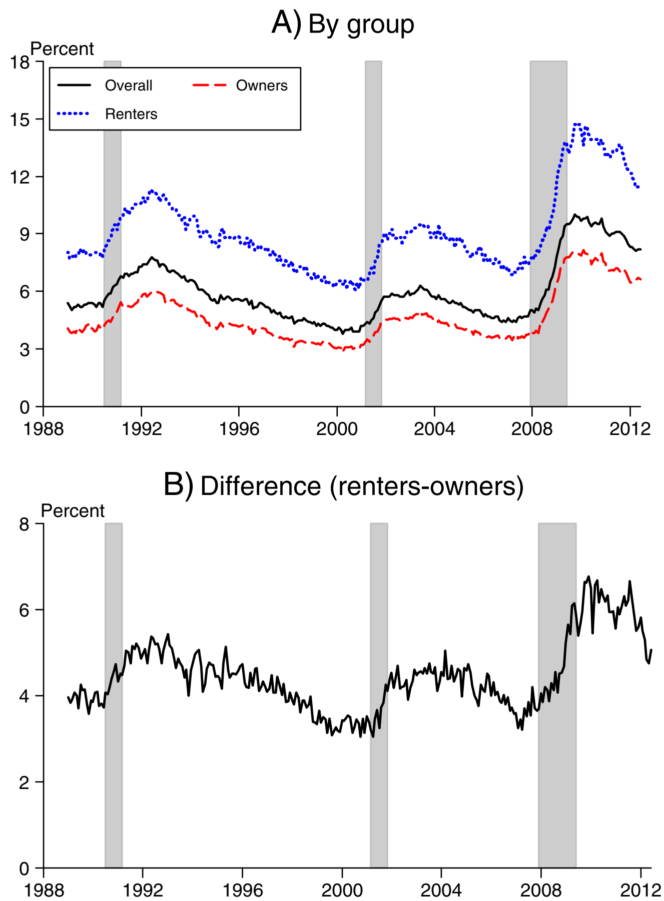

Note: Author's calculations from monthly CPS microdata and BLS labor force series. Gray bars denote NBER recession dates.

Fig. 3. Unemployment rates, monthly (Jan. 1989–Dec. 2011).

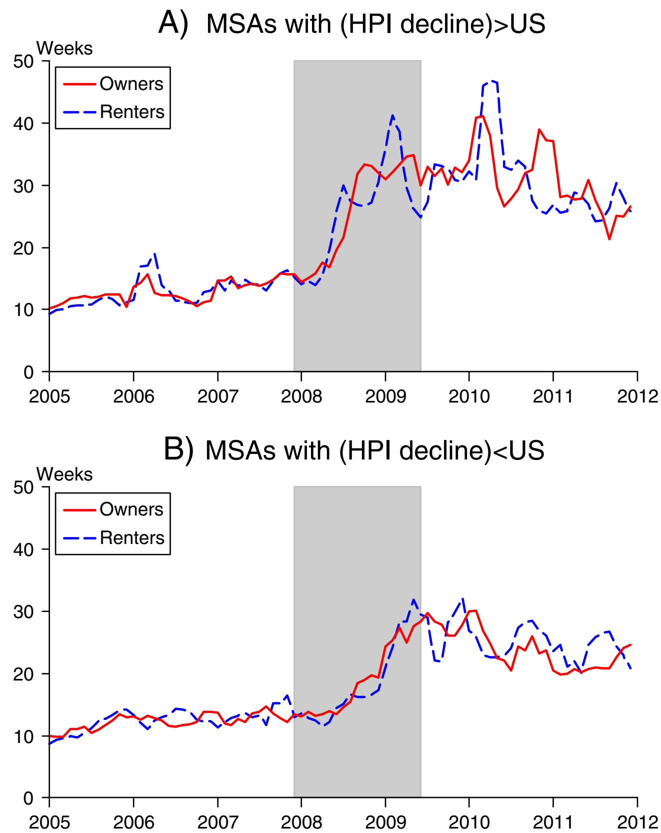

Note: Author's calculations from monthly CPS microdata (seasonally adjusted 3-month moving averages). Duration measured in expected completed form, see text for description. MSA house prices from FHFA, change measured over 2007q1-2011q2. Gray bar denotes NBER recession dates.

Fig. 4. Unemployment duration, by home ownership (through Dec. 2011).

downturn (2005) through the end of 2011. Similar to the breakdown in Table 1, Panels A and B divide the sample up into MSAs for which the decline in house prices exceeded or fell short of the overall U.S. decline. The comparison provided in this figure is analogous to the difference-in-difference test identified by Eq. (1) in Section 3, but without conditioning the estimation on the effects of any other observable covariates. If the house-lock hypothesis is correct, the increase in unemployment durations as the housing bust and economic cycle continued should be larger for owners than for renters, particularly in MSAs that experienced the greatest decline in home prices. No support for the house-lock hypothesis is evident in these charts: the owner and renter groups each experienced 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.

Fig. 4 therefore provides no evidence that would enable rejection of the null hypothesis of no house lock. 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.[19]

[19] 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).

5. Econometric approach

The econometric approach used for the formal analyses of unemployment duration in this paper is adapted from Güell and Hu (2006; henceforth "GH"). GH developed an approach to duration analysis in repeated cross-sections such as the CPS survey that enables estimation of the effects of individual covariates as well as duration dependence and time-varying factors such as local labor market conditions. While they focused primarily on a generalized method of moments (GMM) estimator, 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). I therefore implement the ML approach in this paper.

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.

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.

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:

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

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:

$$P(\tilde{y} = 1|X = x) = \frac{1}{1 + \alpha \frac{1 + \exp(x\beta)}{\exp(x\beta)}} = \frac{\exp(x\beta)}{\alpha + (1 + \alpha)\exp(x\beta)} \tag{2}$$

Eq. (2) 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.[20] 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).[21]

6. Regression analysis and results

6.1. Model specification

For the estimation results discussed in the next sub-section, I include the full set of individual covariates listed in Table 1, plus an interaction between gender and marital status and a complete set of calendar month dummies (less one) to account for seasonal effects. 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. As discussed in Section 3, my specific test for house lock is based on the estimated coefficient for the interaction between the homeowner indicator and measures of movements in local home prices. This interaction coefficient represents the difference in conditional unemployment duration between homeowners in areas with large price declines and those in areas with small price declines, minus the same difference for renters. The incremental effect of home price declines on homeowners versus renters is interpreted as representing factors that uniquely constrain homeowner mobility and increase their unemployment durations in the local labor market.

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. The local labor market indicators are the unemployment rate and the pace of growth in payroll (wage and salary) employment growth. Although local unemployment rates and payroll employment growth are closely related, they are measured from separate surveys of households and business establishments and do not always exhibit the same patterns of co-movement across different MSAs.

6.2. Primary 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 housing price movements (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). The estimated coefficients for the homeownership indicator and the other key variables are listed 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 estimated relationship with unemployment continuation rates; a positive coefficient indicates that larger values are associated with higher continuation rates and longer unemployment durations. Individuals living in MSAs with high unemployment rates (columns 3 and 4) or slow employment growth (columns 5 and 6) experience significantly longer unemployment durations. Focusing on selected other coefficients 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 of unemployment.

Comparison of the coefficients across the duration categories listed at the bottom of the table indicates duration dependence. The

[20] This assumption holds only for observed features of the population, such as age, education, etc. The GH framework abstracts from unobserved individual heterogeneity, which cannot be incorporated into estimation using repeated cross-sections (unlike a true panel or cohort analysis with repeat observations on unemployment spells).

[21] The 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.

Table 2
Unemployment duration (ML estimates, data for Jan. 2008–Dec. 2011).

	(1)	(2)	(3)	(4)	(5)	(6)
	No cyclical control		Control for local unemp rate		Control for local emp growth	
Variables	HPI groups	HPI change	HPI groups	HPI change	HPI groups	HPI change
Home owner	0.0182	0.0252	0.0467	0.0647	0.0145	0.0233
	(0.0253)	(0.0275)	(0.0400)	(0.0440)	(0.0258)	(0.0279)
MSA HPI decline > U.S.	0.311**		0.369**		0.283**	
(dummy variable)	(0.0327)		(0.0698)		(0.0333)	
Owner*HPI group	−0.0116		0.0126		−0.0103	
	(0.0389)		(0.0731)		(0.0398)	
%ΔHPI (peak to trough)		−0.00892**		−0.00936**		−0.00795**
		(0.000922)		(0.00219)		(0.000916)
Owner*%ΔHPI		0.000394		0.000792		0.000521
		(0.00109)		(0.00214)		(0.00110)
MSA unemployment rate			0.143**	0.137**		
			(0.0188)	(0.0206)		
MSA emp growth (12-month %)					−0.046**	−0.043**
					(0.004)	(0.004)
Age 16–19	−0.475**	−0.480**	−0.664**	−0.663**	−0.481**	−0.487**
	(0.0356)	(0.0359)	(0.0625)	(0.0636)	(0.0363)	(0.0364)
Age 20–24	−0.328**	−0.333**	−0.484**	−0.482**	−0.335**	−0.340**
	(0.0329)	(0.0331)	(0.0591)	(0.0598)	(0.0336)	(0.0337)
Age 25–34	−0.194**	−0.194**	−0.260**	−0.260**	−0.201**	−0.202**
	(0.0301)	(0.0304)	(0.0520)	(0.0516)	(0.0307)	(0.0308)
Age 35–44	−0.149**	−0.151**	−0.172**	−0.174**	−0.158**	−0.161**
	(0.0313)	(0.0316)	(0.0547)	(0.0541)	(0.0319)	(0.0321)
Age 55–64	0.0480	0.0510	0.167*	0.166*	0.0523	0.0534
	(0.0386)	(0.0392)	(0.0738)	(0.0734)	(0.0397)	(0.0399)
Age 65+	−0.0161	−0.0195	−0.00415	0.00897	−0.0173	−0.0205
	(0.0622)	(0.0632)	(0.119)	(0.118)	(0.0639)	(0.0645)
Education: HS degree	0.0176	0.0179	0.100*	0.0974	0.0181	0.0174
	(0.0275)	(0.0278)	(0.0501)	(0.0498)	(0.0280)	(0.0281)
Some college	−0.0379	−0.0342	−0.0177	−0.0109	−0.0312	−0.0284
	(0.0290)	(0.0293)	(0.0486)	(0.0481)	(0.0297)	(0.0298)
College degree	−0.0815*	−0.0703	−0.0888	−0.0724	−0.0778*	−0.0680
	(0.0364)	(0.0368)	(0.0617)	(0.0610)	(0.0372)	(0.0374)
Graduate degree	−0.0978	−0.0856	−0.0992	−0.0919	−0.0906	−0.0802
	(0.0537)	(0.0541)	(0.0936)	(0.0919)	(0.0551)	(0.0552)
Race/ethnic: Black	0.204**	0.218**	0.430**	0.436**	0.215**	0.226**
	(0.0287)	(0.0293)	(0.0643)	(0.0694)	(0.0296)	(0.0299)
Hispanic	−0.0600*	−0.0758**	−0.151**	−0.144**	−0.0281	−0.0415
	(0.0265)	(0.0269)	(0.0471)	(0.0467)	(0.0273)	(0.0276)
Asian	0.183**	0.193**	0.288**	0.309**	0.218**	0.226**
	(0.0558)	(0.0568)	(0.106)	(0.105)	(0.0576)	(0.0582)
Other	0.0915	0.0909	0.242*	0.238*	0.109*	0.108*
	(0.0526)	(0.0531)	(0.101)	(0.0998)	(0.0542)	(0.0544)
Military veteran	−0.0590	−0.0591	−0.0989	−0.0989	−0.0574	−0.0576
	(0.0395)	(0.0400)	(0.0705)	(0.0699)	(0.0405)	(0.0407)
Married	−0.125**	−0.126**	−0.172**	−0.170**	−0.126**	−0.128**
	(0.0278)	(0.0280)	(0.0487)	(0.0480)	(0.0284)	(0.0285)
Female	−0.0205	−0.0205	−0.0225	−0.0209	−0.0133	−0.0141
	(0.0246)	(0.0248)	(0.0422)	(0.0416)	(0.0252)	(0.0253)
Female*married	0.0696	0.0708	0.120	0.118	0.0696	0.0709
	(0.0388)	(0.0392)	(0.0689)	(0.0681)	(0.0397)	(0.0399)
Duration months 1–2	0	0	0	0	0	0
(omitted)						
Duration months 2–3	0.358**	0.360**	0.848**	0.820**	0.369**	0.369**
	(0.0377)	(0.0384)	(0.152)	(0.156)	(0.0397)	(0.0397)
Duration months 3–4	0.143**	0.142**	0.312**	0.304**	0.141**	0.139**
	(0.0356)	(0.0360)	(0.0833)	(0.0841)	(0.0366)	(0.0368)
Duration quarters 2–3	−0.437**	−0.446**	−0.768**	−0.758**	−0.457**	−0.464**
	(0.0261)	(0.0264)	(0.0662)	(0.0714)	(0.0268)	(0.0270)
Duration quarters 3/4–5/6	−0.241**	−0.250**	−0.545**	−0.533**	−0.239**	−0.247**
	(0.0280)	(0.0283)	(0.0671)	(0.0711)	(0.0286)	(0.0288)
Duration years 2–3+	0.416**	0.413**	1.293**	1.192**	0.458**	0.451**
	(0.0478)	(0.0487)	(0.444)	(0.420)	(0.0489)	(0.0488)
Alpha (scaling parameter)	0.609**	0.618**	1.013**	1.005**	0.628**	0.633**
	(0.0199)	(0.0207)	(0.0495)	(0.0559)	(0.0208)	(0.0209)
Observations	174,678	174,678	174,678	174,678	174,678	174,678

** $p < 0.01$, * $p < 0.05$.
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.

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.[22]

Shifting back to the top of the table, we see that conditional on observable characteristics owners and renters experience no statistically significant differences in unemployment duration. However, all individuals living in MSAs that saw especially large declines in home prices experience significantly longer unemployment durations (odd-numbered columns, second row of coefficients and standard errors). Similarly, the direct measure of house price changes indicates that individuals in MSAs in which house prices grew relatively rapidly (or declined less rapidly) experienced shorter unemployment durations (even-numbered columns, fourth row). These effects of local housing market conditions are maintained when direct measures of local labor market conditions are incorporated in columns 3–6, and they apply equally to owners and renters.

By contrast with the estimate for local housing market conditions, the key coefficients on the interaction between the homeownership indicator and the price variables provide no evidence of house-lock effects on unemployment duration. These coefficients are uniformly small and statistically indistinguishable from zero. The zero coefficients on the interaction variables between ownership status and the extent of decline in local home prices indicates that local housing market conditions affect unemployment duration for renters to the same degree as owners, suggesting that there are no unique financial constraints associated with homeownership that impede unemployment exits.

6.3. Robustness checks

The results from the primary specifications discussed above 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 longer period of 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. Former owners who lost 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 therefore adjusted the specification in an attempt to capture the potential foreclosure effect. In particular, I grouped observations using a three-way breakdown of changes in 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. The areas with the largest price declines are likely to see the highest foreclosure rates, whereas the middle category corresponds more closely to areas with large numbers of underwater homeowners but fewer foreclosures. 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 support of house-lock effects on unemployment duration.

7. Conclusions

I examined whether evidence can be found to support the hypothesis of house-lock effects on unemployment duration: i.e., whether declining house prices during the recent U.S. housing bust reduced homeowners' geographic mobility and raised their time spent searching for jobs in their local labor markets. Descriptive evidence regarding relative geographic mobility rates and unemployment for homeowners and renters provided no systematic support for the house-lock hypothesis. In particular, geographic mobility has declined similarly for homeowners and renters in recent years, suggesting that the two groups face similar incentives and opportunities for mobility. In addition, unemployment rates for homeowners have not been noticeably elevated relative to renters in recent years, again suggesting that the labor market environment faced by homeowners does not reflect any unusual features.

More formal econometric analyses of unemployment durations, based on a difference-in-difference framework that compared homeowners and renters across MSAs distinguished by the extent of home price declines, also provided no evidence in favor of the house-lock hypothesis. Although homeowners in areas that saw large price declines experienced longer unemployment spells than homeowners in areas with more limited price declines, the same is true for renters, indicating that financial constraints associated with homeownership that may reduce geographic mobility are not a unique barrier to unemployment exits.

The absence of a house-lock effect on recent U.S. unemployment should not be surprising, given that unemployment rates were quite elevated 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 in almost every area of the country in recent years. Moreover, past research that found links between home prices and geographic mobility focused on earlier periods when foreclosure rates were well below their current levels (e.g., Chan, 2001; Engelhardt, 2003; Ferreira et al., 2010). By contrast, in recent years U.S. foreclosure rates have reached new historical highs. An increase in geographic mobility prompted by foreclosures may have significantly offset reductions in mobility associated with home price declines more generally, implying that the net effect of housing market conditions on homeowner mobility has been essentially zero.

The distinction between mobility induced by foreclosures and immobility associated with house lock may help reconcile the widespread anecdotal evidence regarding house lock with the absence of an overall effect on labor markets. Indeed, house lock may be quantitatively important in recent years, and an alternative empirical test that enables separate identification of the foreclosure effect and the underwater effect may provide a more accurate rendering of the recent relationship between conditions in housing markets and the labor market. 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.

For now, my results fit well with other recent work that also finds little or no impact of house lock on the U.S. labor market during the "Great Recession" of 2007–09 and its aftermath (e.g., Donovan and Schnure, 2011; Modestino and Dennett, 2013; Farber, 2012; Schmitt and Warner, 2011). This line of research plays an important role in the continuing debate over the sources of the slow labor market recovery in the United States, which largely revolves around cyclical and structural explanations. Limited direct support has been found

[22] 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.

in favor of structural explanations, of which house lock is one example (for alternatives, see Daly et al., 2012; Lazear and Spletzer, 2012). This limited affirmative evidence for factors such as house lock suggests that the U.S. labor market has been suffering from a persistent shortfall in aggregate demand rather than widespread structural impediments.

Appendix A. Duration data adjustments and 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 stock-sampled 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).[23] 23This approach relies on the estimation of monthly continuation rates—i.e., the probabilities that an unemployment spell will continue from one month to the next —using grouped duration data.

My application of the synthetic cohort approach to obtain nonparametric estimates of expected completed duration from grouped duration data follows M. Baker (1992a); see 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:

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

where $n(.)$ 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.

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*:

$$G_j(t) = f_0(t) f_1(t) f_2(t) f_3(t) \ldots f_j(t) \tag{A2}$$

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:

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

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 Eq. (A3). Then the expected completed duration is formed as:

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

where the time identifier (t) has been suppressed on the right-hand side of (A4) 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. Following past practice (e.g., Sider, 1985), I multiplied estimates of expected duration in months by 4.3 to obtain expected duration in weeks.

[23] 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.

References

Aaronson, Daniel, Davis, Jonathan, 2011. How much has house lock affected labour mobility and the unemployment rate? Chicago Fed Letter 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., Trivedi, P.K., 1985. Estimation of unemployment duration from grouped data: a comparative study. Journal of Labour Economics 3 (2), 153–174.
Buchmueller, Thomas C., Valletta, Robert G., 1996. The effects of employer-provided health insurance on worker mobility. Industrial & Labor Relations Review 49 (3, April), 439–455.

Carlson, John, Horrigan, Michael, 1983. Measures of unemployment as guides to research and policy: comment. American Economic Review 73 (5), 1143–1152.
Chan, Sewin, 2001. Spatial lock-in: do falling house prices constrain residential mobility? Journal of Urban Economics 49 (3), 567–586.
Coulson, N. Edward, Fisher, Lynn M., 2009. Housing tenure and labour market impacts: the search goes on. Journal of Urban Economics 65, 252–264.
Daly, Mary, Hobijn, Bart, Şahin, Ayşegül, Valletta, Robert G., 2012. A Search and Matching Approach to Labor Markets: Did the Natural Rate of Unemployment Rise? Journal of Economic Perspectives 26 (3, Summer), 3–26.
Donovan, Colleen, Schnure, Calvin, 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), 171–195.
Fallick, Bruce, Fleischman, Charles A., 2004. Employer-to-Employer Flows in the U.S. Labour Market: The Complete Picture of Gross Worker Flows. Finance and Economics Discussion Series, 2004-34. Board of Governors of the Federal Reserve System.
Farber, Henry S., 2012. Unemployment in the great recession: did the housing market crisis prevent the unemployed from moving to take jobs? American Economic Review 102 (3, May), 520–525.
Ferreira, Fernando, Gyourko, Joseph, Tracy, Joseph, 2010. Housing busts and household mobility. Journal of Urban Economics 68, 34–45.
Fletcher, Michael A., 2010. In slow economy, Americans stay put. Washington Post (July, 30).
Genesove, David, Mayer, Christopher, 2001. Loss-aversion and seller behavior: evidence from the housing market. Quarterly Journal of Economics 116 (4), 1233–1260.
Güell, Maia, Hu, Luojia, 2006. Estimating the probability of leaving unemployment using uncompleted spells from repeated cross-section data. Journal of Econometrics 133, 307–341.
Guiso, Luigi, Sapienza, Paola, Zingales, Luigi, 2009. Moral and Social Constraints to Strategic Default on Mortgages. NBER Working Paper 15145 (July). National Bureau of Economic Research, Cambridge, MA.
Herkenhoff, Kyle F., Ohanian, Lee E., 2011. Labour Market Dysfunction During the Great Recession. NBER Working Paper 17313 (August). National Bureau of Economic Research, Cambridge, MA.
Horrigan, Michael, 1987. Time spent unemployed: a new look at data from the CPS. Monthly Labour Review 110 (7, July), 3–15.
Kaplan, Greg, Schulhofer-Wohl, Sam, 2012. Interstate migration has fallen less than you think: consequences of hot deck imputation in the current population survey. Demography 49 (3, Aug), 1061–1074.
Karahan, Fatih, Rhee, Serena, 2011. Geographical Reallocation and Unemployment during the Great Recession: The Role of the Housing Bust. Manuscript Department of Economics, University of Pennsylvania.
Lazear, Edward P., Spletzer, James R., 2012. The United States Labor Market: Status Quo or a New Normal? Presented at "The Changing Policy Landscape," 2012 Federal Reserve Bank of Kansas City Economic Policy Symposium, Jackson Hole, WY.
Lovenheim, Michael F., Mumford, Kevin J. 2013. Do family wealth shocks affect fertility choices? Evidence from the housing market. Manuscript, Cornell University, November The Review of Economics and Statistics (Forthcoming).
Madrian, Brigitte G., 1994. Employment-based health insurance and job mobility: is there evidence of job-lock? Quarterly Journal of Economics 109 (1), 27–54.
Modestino, Alicia Sasser, Dennett, Julia, 2013. Are American Homeowners Locked into Their Houses? The Impact of Housing Market Conditions on State-to-State Migration. Regional Science and Urban Economics 43 (2, March), 322–337.
Molloy, Raven Saks, Smith, Christopher L., Wozniak, Abigail, 2011. Internal migration in the United States. Journal of Economic Perspectives 25 (3), 173–196.
Munch, Jakob R., Rosholm, Michael, Svarer, Michael, 2006. Are homeowners really more unemployed? The 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, Warner, Kris, 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.
Sider, Hal, 1985. Unemployment duration and incidence: 1968–82. American Economic Review 75 (3, June), 461–472.
Sterk, Vincent, 2010. Home Equity, Mobility, and Macroeconomic Fluctuations. DNB (DeNederlandsche Bank) Working Paper 265 (November).
Valletta, Robert G., 2011. Rising Unemployment Duration in the United States: Composition or Behavior? Mimeo, Federal Reserve Bank of San Francisco (May).
White, Eugene N., 2009. Lessons from the Great American Real Estate Boom and Bust of the 1920s. NBER Working Paper 15573 (December). National Bureau of Economic Research, Cambridge, MA.