Department of Economics
Seminar Series

Courtney Coile
Wellesley College

"The Market Crash and Mass Layoffs: How the Current Economic Crisis May Affect Retirement"
and
"Recessions, Reeling Markets, and Retiree Well-Being"

Friday,
October 22, 2010
3:30 p.m.
210 Middlebush

The Market Crash and Mass Layoffs: How the Current Economic Crisis May Affect Retirement

Courtney C. Coile
Department of Economics
Wellesley College
and
NBER
ccoile@wellesley.edu

Phillip B. Levine
Department of Economics
Wellesley College
and
NBER
plevine@wellesley.edu

August 2010

Acknowledgements: The authors thank Wellesley College for financial support and seminar participants at Harvard, MIT, and the NBER Summer Institute for helpful comments.

"Those Golden Years Have Lost Their Glow; With Home Values Down, Costs Up and Their 401(k)s Declining, Some Seniors Have Had To Rethink Retirement." (*Los Angeles Times*, September 21, 2008)

"Will You Retire?; New Economic Realities Keep More Americans In the Workforce Longer." (*Washington Post*, October 15, 2008)

"Economic Crisis Scrambles Retirement Math: The 401(k) Model of Saving is Under Duress as Stocks Slide. Home Equity Losses Don't Help." (*Christian Science Monitor*, March 4, 2009)

I. INTRODUCTION

One casualty of the financial and economic crisis that began in the fall of 2008 may be workers' carefully laid retirement plans. The popular press recognized this from the start of the crisis, as the headlines listed above make clear. Front page stories of lost retirement savings and plunging home values are commonplace. With diminished retirement savings and less home equity to draw on, the story goes, expected retirement income has shrunk, forcing older individuals to stay in the labor force longer. Workers interviewed for these stories wondered when or if they would ever be able to retire.

Amidst these concerns, another news story appeared briefly in spring 2009 indicating that Social Security benefit claims have risen sharply since the crisis began, suggesting an increase in retirements rather than a decrease (Dorning, 2009). A subsequent report (Johnson and Mommaerts, 2010) indicated that new Social Security retirement awards continued to surge through 2009. Although the number of Americans turning age 62, and thereby becoming eligible for Social Security retirement benefits, rose 9 percent between 2008 and 2009, the number of new retirement benefit awards rose 20 percent for men.

But why are more workers retiring now if their expected retirement income is going down? The answer may lie in another aspect of the crisis, the weak labor market. The

unemployment rate has more than doubled and the economy has shed millions of jobs since the crisis began. Some of those workers struggling to stay employed or find new jobs are surely nearing retirement age. For the unfortunate ones who are not able to maintain or find employment, retirement may be the only solution, despite its involuntary nature.

The net effect of the current financial and economic crisis on retirement is thus far from clear, as plunging stock and home values would be expected to lead to a decrease in retirements while a weak labor market would be expected to lead to an increase. The purpose of this paper is to examine this issue. We use 30 years of data from the March Current Population Survey (CPS) to estimate models relating retirement decisions to changes in stock, housing, and labor markets over time and (where possible) across geographic locations. We also use the Survey of Consumer Finances (SCF) to provide a descriptive analysis of the impact of falling stock prices on older household's expected retirement income. We then use our regression estimates to predict the net effect of the current crisis on retirement.

Our analysis indicates that the retirement decisions of workers between ages 62 and 69 with more education are affected by long-run fluctuations in stock market returns. We also find that labor market conditions are an important determinant of retirement decisions. When the unemployment rate rises, more workers between ages 62 and 69 retire, particularly those with less education. Workers between ages 55 and 61 are not found to be responsive to either type of market fluctuation. Individuals do not seem to respond to fluctuations in the housing market regardless of their age. On net, we predict that the increase in retirement brought about the recent rise in unemployment will be almost 50 percent larger than the decrease in retirement brought about by the stock market crash.

Overall, our findings suggest that the plight of those who are forced to retire early as a result of weak labor market conditions merits greater attention. These results have potentially important distributional implications as well. It is often those on the bottom of the economic ladder who are being hurt by retiring prematurely due to labor market factors and those at the top who may not be able to retire as planned due to stock losses. Our results also have implications beyond the current economic crisis, as they suggest that the past literature on retirement has paid too little attention to the important role of labor market conditions in the retirement decision.

The remainder of our analysis proceeds as follows. In the following section, we document trends in the environment surrounding retirement decisions, including stock returns, housing prices, and the labor market. Next, we review the relevant literature and discuss the data and methods we use in the remainder of the analysis. We then present our results regarding the impact of changes in stock, housing, and labor markets, respectively, on retirement decisions. Finally, we simulate the net effect of recent market events on retirement and discuss the policy implications of our findings.

II. BACKGROUND

In this section, we present trends in stock, housing, and labor markets to review recent activity and summarize earlier events that may be less well remembered. We also discuss the conditions under which fluctuations in these markets may affect retirement behavior.

A. Trends in the Stock Market

Annual changes in the value of the stock market, as captured by the S&P 500 Index, are shown in Figure 1A. This figure reports real annual changes (adjusted for inflation) based on December monthly average values. The figure illustrates the tremendous year-to-year volatility

in aggregate stock prices. The pattern in the 1980s and early 1990s is one of two good years with 10 to 20 percent annual returns followed by a bad year with zero or negative returns. Since then, the market has experienced more prolonged booms and busts, including two five-year rallies in the late 1990s and mid 2000s, as well as a multi-year bear market early in this decade. The market fell by 40 percent in real terms in 2008, the sharpest decline in recent history.

One can see how these dramatic turnarounds in stock markets have captured the public's attention. The question at hand, though, is whether they alter retirement decisions. Given that there has always been substantial year-to-year variability in stock prices, is it sensible to expect a single year's market performance to drive behavior?

The market return over a longer period of time could potentially play a more important role in retirement decisions. In Figure 1B, we display five-year and ten-year market returns (again calculated using December monthly average values). This figure shows that there is substantial variability in longer-term returns over time. In the 1980s and early 1990s, the five-year real return was consistently about 50 percent. After that real returns rose, hitting almost 200 percent in the year 2000 before collapsing to small or negative values. Ten-year returns are higher, but the patterns are similar.

These statistics suggest that market returns could have a significant impact on retirement behavior. One worker approaching retirement age could have tripled the value of his portfolio over a five year period, while another worker's portfolio remained constant or even shrank. If workers have considerable resources invested in the stock market, a boom or a bust in the period leading up to traditional retirement ages could play a key role in the decision of when to retire. We later explore the level of stock ownership among the population and various subgroups.

B. Trends in the Housing Market

Although the volatility in the housing market is less dramatic, home values also exhibit substantial fluctuations over time. Figure 2 displays annual changes in real house prices from 1987 to 2008 based on the Case-Shiller (CS) Index for 10 large cities across the country and from 1976 to 2008 based on data from the Office of Federal Housing Enterprise Oversight Home Price Index (OFHEO) for the entire country.[1] The figure shows that housing market returns are considerably more serially correlated than stock returns. In the late 1980s and early-to-mid 1990s, home values did not keep pace with inflation. In the decade that followed, however, prices rose continuously, with annual growth rates in the Case-Shiller Index of over 10 percent in some years. House prices have fallen sharply since 2006, dropping almost 20 percent in 2008.

These statistics suggest that home prices could also affect retirement decisions. Depending on their year of birth, individuals may have doubled their home equity or had it cut in half as they approach traditional retirement ages. If workers had substantial home equity to begin with and are willing to draw down this equity during retirement, a substantial increase in home equity could accelerate retirement while a substantial drop could delay it.

C. Trends in the Labor Market

Figure 3 presents the cyclical variation in the labor market, as measured by the monthly unemployment rate for workers age 16 and over. As we describe subsequently, older workers have a lower unemployment rate, but the pattern over time is very similar to that for all workers. The highest unemployment rate in recent times was 10.8 percent in 1982. Subsequent recessions in the early 1990s and the early 2000s were less severe, with the unemployment rate reaching highs of 7.8 percent and 6.3 percent, respectively. In the current crisis, the unemployment rate is climbing rapidly; as of August of 2009, it had reached 9.7 percent. Aside from these recessions,

[1] We discuss these two indices in more detail below. Annual returns in the CS Index are calculated as the change in the December values. Annual returns in the OFHEO Index are calculated as the change in the fourth quarter values.

the unemployment rate has been at a low level, around 4.5 percent, for much of the period since the mid-1990s.

As with our earlier discussions of stock and housing markets, labor market conditions around traditional retirement ages may matter. Workers are twice as likely to be unemployed now as they were a few years ago. In times when obtaining a new job is difficult, older individuals who are laid off or unemployed for other reasons may be more likely to retire. This may be especially true for workers age 62 and up, who generally have access to Social Security.

As this discussion has made clear, there are reasons to believe that variations in stock prices, house prices, and the labor market have the potential to alter retirement behavior. It is also clear that there are important conditions for these behavioral responses. Lower stock and housing prices may lead to fewer retirements if individuals nearing retirement have sufficient stock holdings and home equity and plan to consume it during retirement. Higher unemployment rates may lead to more retirements if older individuals are unable to find work and withdraw from the labor force instead. Furthermore, for market fluctuations to affect aggregate retirement rates, the relevant elasticities must be large enough to generate behavioral responses by more than just a handful of older individuals. In the end, the retirement responses to fluctuations in stock, housing, and labor markets are empirical questions. In the remainder of this paper, we attempt to answer these questions.

III. PREVIOUS LITERATURE

Much of the existing retirement literature has focused on Social Security, private pensions, and health. While these factors may be important in explaining long-run trends, such as the steep decline in older men's labor force participation since World War II and the recent

reversal of that trend, they are unlikely to explain dramatic changes in retirement behavior in any given year, such as those that might result from the current crisis. In this section, we focus on those parts of the retirement literature that are most directly relevant to our analysis.

A. Financial Shocks

Economic theory suggests that individuals should respond to negative stock market shocks by reducing their consumption of normal goods (including leisure) and delaying retirement. Articles in the popular press have similarly asserted that this will be the effect of the current crisis. Nevertheless, there is little empirical research to support this hypothesis.

In an earlier paper (Coile and Levine, 2006), we use methods similar to those described below to address this issue. We treat the stock market boom and bust of the late 1990s and early 2000s as a quasi-experiment and explore whether groups with more stock assets were more likely to retire during the boom and less likely to retire during the bust. We find no evidence of this pattern. We also argue that individuals would have to have been implausibly sensitive to market fluctuations for the observed rise in retirement in the year 2000 to have been the result of that year's market crash. Our findings are consistent with those obtained by Hurd, et al. (2009). They are unable to find support for the notion that "households which had large (financial) gains retired earlier than they had anticipated or that they revised their retirement expectations compared with workers in households that had no large gains."[2]

There are two possible explanations for the lack of an effect. The first is that the number of people who experienced large unexpected wealth gains from market fluctuations is relatively small, as Coile and Levine (2006) argue. The second is that the effect of unexpected wealth on

[2]Sevak (2001) reached a different conclusion, finding that men in defined contribution (DC) pension plans increased their retirement rates by more than men in defined benefit (DB) pensions during the stock market boom of the late 1990s. However, this study is limited by an inability to control for differences in retirement trends between the two groups, a deficiency that is overcome in Coile and Levine (2006) by the use of the boom and bust as a double experiment.

labor supply is fairly small. This view is supported by Coronado and Perozek (2003), who find that being a stockholder during the boom of the late 1990s is associated with retiring 6 months earlier than expected, but that each additional $100,000 of unexpected gains is associated with retiring only two weeks earlier than expected. Hurd, et al. (2009) are also sympathetic to this argument, citing evidence from lotteries.

B. The Role of Housing

As with stock market shocks, economic theory suggests that unanticipated losses in home equity should lead households to retire later. However, shocks to home equity will only affect retirement behavior if households routinely consume their housing wealth in retirement. In fact, studies suggest that this is not the case. For instance, Venti and Wise (2004) find that most households do not sell their homes until they experience an event such as the death or entry into a nursing home of a spouse. This finding has led some authors to argue that many households treat their home equity as a "buffer stock" of wealth against the risk of shocks late in life. If so, then it seems unlikely that home price fluctuations will affect retirement behavior, although many recent stories in the popular press have asserted that this is the case. The effect of housing wealth on retirement has not been directly addressed in the previous literature. We provide an empirical analysis of this question below.

C. Labor Market Shocks

A small body of literature has established that job loss is relatively common for older workers (Farber, 2008; Munnell, et. al., 2006). For instance, Farber (2008) reports that 10 to 12 percent of private-sector workers between the ages of 50 and 64 experienced permanent and involuntary job losses when labor markets were weak during the 1991 to 1993 and 2001 to 2003 periods, while displacement rates of around 8 percent (over a three year period) were observed

during the expansions of the mid-to-late 1990s and the middle 2000s. Previous studies have found that job loss among older workers has long-lasting negative consequences for employment and wages (Chan and Stevens 1999, 2001, and 2004; von Wachter, 2007). Chan and Stevens (1999) estimate that the employment rate of displaced older workers two years after a job loss is 25 percentage points lower than that of similar non-displaced workers and that the median reemployed worker earns 20 percent less than at his old job.

More directly related to the question we seek to address here is our earlier work (Coile and Levine, 2007). Using similar methods and data to that described subsequently, we find that retirement transitions are cyclically sensitive, a result supported by Von Wachter (2007), Hallberg (2008), Friedberg et. al. (2008) and Munnell et. al. (2008). We estimate that changes in rates of retirement between the peak and trough of a business cycle are comparable to those brought about by moderate change in financial incentives to retire or to the threat of a health shock, factors that have traditionally received far more attention in the literature. We also find that Social Security interacts with labor market conditions in affecting retirement transitions, as the effect of the unemployment rate on retirement appears only as workers become eligible for benefits. We expand upon this discussion later in our analysis.

D. Contribution of this Research

The current analysis builds on the previous literature, including our own past work, in several ways. First, we update and extend our analyses of the effect of stock market and labor market fluctuations on retirement. Second, we provide a new analysis of the effect of housing market fluctuations on retirement, a question not addressed in the previous literature. Third, we use these various estimates to predict the net effect of the current crisis on retirement. Finally, we discuss the distributional consequences and policy implications of our findings.

IV. DATA SOURCES

The main data requirement for our analysis is a way to measure retirements for large numbers of workers over time. Beyond data on their labor market activity, we also need information on workers' asset holdings, including both financial assets and home equity. This section of the paper will describe the sources of data we use.

A. Measuring Retirement

Our main source of data for measuring retirements is the Current Population Survey (CPS). The CPS is the leading survey of labor market activity in the United States. The monthly CPS survey asks a sequence of questions about the respondent's involvement in the labor market around the time of survey and also collects demographic data. In March of each year, the "Annual Social and Economic Supplement" (previously called the "Annual Demographic Survey") is administered as a supplement to the regular monthly CPS. Each March CPS provides sample sizes of between 130,000 and 215,000. Although we only are interested in the data for workers around the age of retirement, the large size of each sample coupled with the annual nature of the survey provides us with a tremendous amount of information. For instance, when we pool data from the 1980 through 2008 March surveys for individuals between the ages of 55 and 69, we obtained a sample of nearly 600,000 individuals.

For our purposes, one key attribute of the March CPS is that it enables us to identify retirement transitions.[3] To do so, we make use of information on the labor market activity of

[3] As we describe subsequently, we define retirement as complete labor force withdrawal. However, we recognize that retirement could be defined in other ways, for example, as the initial claim of retirement benefits or as departure from a "career" job. In fact, several studies have found that it is quite common for workers to leave a career job and work for a period of time at a less demanding "bridge" job before completely withdrawing from the labor force; see

respondents in the preceding calendar year, including weeks worked, usual hours worked per week, and weeks spent looking for work. Combining this retrospective information along with that obtained in the regular monthly survey, we can define a retirement to occur when an older worker reports being in the labor force for 13 or more weeks during the preceding year, but is out of the labor force on the March survey date.[4] When we restrict our sample to those in the labor force last year in this way, we are left with a final sample size of over 300,000. Of these workers, we observe that about 9 percent retire in the following year according to our definition.[5] State of residence is available in the March CPS, which we can use to merge in state-level data on unemployment rates and house prices.

B. Measuring Home Prices

We use two sources of home price data. The first is the S&P/Case-Shiller Home Price Index, which is available monthly for 20 metropolitan areas (MSAs) beginning in 1987. The

---

Cahill, et. al. (2006) for a recent contribution. The data available to us leads us to focus on a definition of complete labor force withdrawal. However, an analysis of these other types of retirement transitions would be a fruitful area for future research.

[4] A second way that we could use CPS data is by taking advantage of the longitudinal structure of the CPS to create a short panel of information for each respondent. This panel can be created by matching CPS information for some respondents in one March CPS with that from the CPS in the following March. The procedure for doing so is reported in Madrian and Lefgren (1999). These data offer about one-third the sample size as the regular CPS. An advantage of these data, though, is that we can create a definition of retirement for workers who have been more committed to the labor market and out of the labor force for a longer period of time. We have used these data as well and obtained findings qualitatively similar to those reported subsequently. We have chosen not to report them for expediency.

We can also use these matched March CPS data to examine the likelihood of labor market reentry following retirement, as we have defined using the regular March CPS. With the matched data, we use contemporaneous and retrospective labor market activity in the first survey year to define a retirement and contemporaneous labor market activity in the second survey year. Although we do find some reentry, it tends to be lower after a recession. We also find that the more highly educated are the ones who are most likely to reenter and we cannot distinguish differences by educational attainment in terms of the cyclical sensitivity of reentry. We conclude from this that reentry is not uncommon, but that our results are unlikely to be driven by temporary labor force withdrawals.

[5] Using matched March CPS data, described in the preceding footnote, we can also estimate the likelihood that a worker who retires according to our definition regains employment in the following year. Our estimates suggest that 16 percent of those 55 to 69 and 13 percent of those 62 to 69 who retired in the preceding year found employment again in the following year.

index uses a "repeat sales pricing" methodology, where data on sale prices of individual single-family homes is collected from county records and matched to each home's previous sales price, then a weighted aggregate index is created based on the change in sales prices of these homes. We convert the index to real values using the Consumer Price Index (CPI) to calculate real changes in house prices. We calculate the percent change in the index from one March to the next, as our definition of retirement in the CPS is essentially based on changes in labor market activity between one March and the next, and relate retirement decisions in a given year to housing returns over the previous 12 months.

The second data source is the Office of Federal Housing Enterprise Oversight (OFHEO) Home Price Index. This index is available quarterly at the MSA level starting in 1975. The OFHEO index is also based on changes in the value of individual homes over time, but is calculated using Fannie Mae and Freddie Mac data on mortgages originated by these entities during home purchase and refinancing transactions. We use first-quarter data and again relate retirement decisions to home price appreciation in the previous 12 months.

In comparing the two indices, the OFHEO index has the advantage that we are able to merge home price information to the CPS data for half of our sample (essentially all observations with valid MSA data), while the comparable figure for the Case-Shiller data is only 15 percent. However, the Case-Shiller index displays more variation over time, as shown in Figure 2, which can be attributed to several differences in the construction of the two indices, including the fact that the OFHEO index does not include foreclosures. As we report below, results using the two indices are very similar.

C. Measuring Asset Values

The primary source of wealth data in the United States is the Survey of Consumer Finances (SCF). The survey has been conducted every three years since 1983, most recently in 2007, with a sample of roughly 4,500 households per survey. The survey oversamples high net worth households to obtain a more accurate estimate of aggregate wealth holdings. The survey collects detailed data on assets and income, including data on asset allocation within retirement accounts. We use the SCF to generate information on the stock holdings of older households, using sample weights to obtain statistics that are representative of the population.

V. METHODOLOGY

Although the specific methods we use depend on whether we are addressing stock market wealth, housing wealth, or unemployment, the general approach is similar. To avoid repeating ourselves, we first present the basic methodological framework and then provide details regarding the ways in which we modify it for each specific application.

A. Framework

Our goal is to determine whether different types of market conditions alter retirement decisions. Underlying our analysis is a regression model where the dependent variable is an indicator for whether an older worker retired in a particular year as a function of the market conditions he faces along with other explanatory variables, mainly demographic factors like race/ethnicity, gender, level of education, etc.[6] We also include a full set of exact age dummies, which essentially converts our retirement regression into a hazard model with a nonparametric

[6] We have also experimented with models that allow the impact of improving market conditions to differ from the negative of the impact of declining market conditions, but we found no evidence of an asymmetric effect. The models that we estimate have binary dependent variables for retirement. We report the results of linear regression models because they are easier to interpret, but we have also estimated probit models, which yielded derivatives that were similar.

baseline hazard. We use the same CPS data to provide information on retirement behavior as well as the explanatory variables (other than the market conditions).

For each analysis we exploit quasi-experimental variation in the data, which we believe is able to plausibly generate causal conclusions regarding the impact of conditions in each market on retirement behavior. Quasi-experimental variation relies on changes over time in the explanatory variables occurring in some locations or for some groups but not in other places or for other groups. Those individuals who experienced no change act as a quasi-control group for those in a quasi-treatment group who experienced a change. Comparing differences in outcomes over time between the two groups provides a means to identify the effect of the change. Statistically, this approach is referred to as a difference-in-difference method as the change, or difference, within one group is differenced from the change in the other group to estimate the effect.

In practice, this approach is generally implemented using panel data, estimating regression models that include specific market conditions (stock market, housing market, and labor market) along with relevant fixed effects when possible. One set of fixed effects would represent a vector of state of residence dummy variables that can hold constant any longstanding differences in behavior between workers who live in different areas of the country.[7] A second set of fixed effects would represent dummy variables for the time periods included in the analysis. These time fixed effects would hold constant broader social and economic conditions that may be changing over time and that might alter outcomes for all individuals. What remains to be estimated once these fixed effects are included is the difference in outcomes that take place over time between the groups. The coefficient on the market conditions variable, our key

[7] If the quasi-treatment and quasi-control groups were identified by a characteristic other than location, for example education, then the dummy variables for each education group would serve to hold constant any longstanding differences in behavior between workers in different education groups, as the state dummies do in this discussion.

explanatory variable, is this estimate. We will apply this general approach in all of our subsequent analyses.

Before providing a discussion of the application of our approach to each specific market, it is appropriate to discuss how we intend to measure market conditions and why we have made those choices. For the stock market, we use the one-, five-, and ten-year percentage change in the S&P 500 Index. For the housing market, we use the one- and five-year percentage changes in the relevant housing price index. For the labor market we use the unemployment rate. Coupling these measures with the retirement rate means that we are mixing flows (retirement) with changes (stock and house prices) and levels (the unemployment rate). We believe that the measures we have chosen do the best job of capturing each type of economic activity for the purpose at hand. First, we use transitions into retirement rather than the number of retirees at a given point in time because the former captures behavior that is occurring now, while the latter includes those who retired some time ago and thus is unlikely to be responsive to current market conditions. Second, we use the change in stock and house prices because it seems likely that retirement will be more responsive to the price changes than levels. If prices are high but stagnant, the earlier run-up in the market should have already been captured in retirement expectations; changes in behavior are more likely to be generated by changes in prices. Finally, we use the unemployment rate rather than the change in the rate because the former seems more likely to be relevant for retirement decisions. If the unemployment rate rises from 6 to 8 percent and then stays there, jobs are not secure, and older workers may continue to get laid off, even if the unemployment rate is unchanged.

B. Application to Changes in Financial Wealth

Not everyone holds financial wealth. As we document later, some segments of society have little financial wealth. Because changes in stock market conditions should have little or no direct bearing on retirement decisions for those who do not own stocks, these individuals can be thought of as a quasi-control group. We can compare the effect of stock market fluctuations on retirement for those without financial wealth to the effect for those with significant financial wealth to estimate the impact of the market on retirement.

In practice, the CPS data we use to measure retirement do not include data on financial wealth. Instead, we first divide individuals by educational attainment. As we report later, individuals with no more than a high school degree typically have very limited stock holdings and can act as a quasi-control group for college graduates, whose holdings are more extensive. If the more educated are estimated to retire at a differentially higher rate in response to higher stock market prices, this would support the hypothesis that market conditions matter.

In these specifications, we are unable to include a complete vector of year fixed effects because the stock market variables available to us vary only over time and not across locations. Instead, we capture broader movements in retirement behavior over time by including quadratic time trends in our regression model, allowing the trends to differ by group. This model enables us to identify the impact of stock market changes by estimating whether retirement behavior deviates from a quadratic trend in years in which market returns are higher. To support a causal effect, estimates would need to be greater for the more highly educated.

B. Application to Changes in Housing Wealth

Our use of quasi-experimental variation and difference-in-difference methods is somewhat different when we analyze changes in housing wealth. We first consider the variation available to us as a result of differences in house price changes by location. In the extreme, we

could think about individuals who live in locations where housing prices have remained flat (in real terms).[8] They would represent a control group to compare to those in locations where prices rose or fell. Dividing individuals in this way is a bit unrealistic, however, since housing prices tend to fluctuate everywhere at least some of the time.

Nevertheless, we can use the same methods and somewhat modify our interpretation. In reality, what we have are groups who were more affected than others in the sense that housing prices change by more in some locations at some points in time than others. Implementing the difference-in-difference method with location and time fixed effects enables us to estimate whether there are greater changes in retirement behavior in areas with greater changes in home prices. This method still holds constant longstanding differences in retirement behavior across locations and trends in retirement behavior over time that affect the population as a whole. The experimental analogy does not work quite as well here, but the general approach is the same and yields results that plausibly can be interpreted as causal.

We can further expand upon this approach by incorporated a "third difference" as well. As with financial wealth, home equity varies across individuals. While we are not able to identify the exact amount of home equity held by each individual, we can identify home ownership status in the CPS, allowing us to use those with no equity as a true quasi-control group. If we find that homeowners increase retirement by more than renters in response to an equivalent increase in housing prices, this would be consistent with the hypothesis that home equity affects retirement and provide further support for a causal interpretation of our findings.

C. Application to Changes in Labor Market Conditions

[8] In reality, since we are interested in unanticipated housing gains or losses, what should matter for retirement is not so much the total amount of the gain but the amount that was unexpected, so the ideal control group would be one where housing prices rose no more or less than expected. While it is plausible that expectations about house price appreciation may vary by location, we have no data to guide us on this point, so we must treat all gains or losses in all locations as (equally) unanticipated.

The methods available to evaluate the impact of changes in labor market conditions, as measured by the unemployment rate, are similar to those for housing wealth. The unemployment rate changes in some places at some points in time more than others and we rely on that variation just like we described with changes in housing prices.[9] We can also estimate difference-in-difference models separately for different demographic groups, including by educational attainment. Less-skilled workers tend to be more sensitive to labor market conditions (Hoynes, 2000), so we would expect any impact of an economic downturn on retirements to be larger for this group. Following the previous literature, we use less-educated as a proxy for less-skilled workers. Therefore, we can use the differential responsive in retirement to labor market conditions across educational attainment categories as a further test of a causal effect.

D. Why Three Separate Analyses?

A final important conceptual issue relates to our use of three separate models for the three markets rather than one regression model that would include all three measures of market conditions. While in principle we could use the latter approach, in reality there are important differences across the three analyses that make running separate analyses preferable, in our view. First, as just discussed, we are unable to use year fixed effects in the stock market analysis; running one joint model would prohibit us from using them in the analyses of the other markets as well. Second, data on housing prices is only available for about half of the CPS sample (those with non-missing MSA information), so estimating a single model would reduce the power of our estimates in the other analyses as well. Finally, testing our hypotheses involves comparing

[9] The use of state level unemployment rates introduces some measurement error because those data come from surveys that contain sampling variability. The BLS states "The average magnitude of the over-the-year change in an annual average state unemployment rate that is required in order to be statistically significant at the 90-percent confidence level is about 0.5 percentage point." In a linear probability model with classical measurement error, this should introduce some attenuation bias. To gauge the sensitivity to this problem, we also estimated models using the national unemployment rate rather than the state unemployment rate, including a trend and trend squared rather than year fixed effects. The results of this analysis were quite similar to those reported subsequently, suggesting the attenuation bias described earlier is unlikely to be a major issue.

coefficients across different groups in the different analyses (e.g., by homeowner status in the housing regressions vs. by educational attainment in the stock and labor market analyses). Thus we believe that conducting three separate analyses provides us with the best opportunity to analyze the effects of each market on retirement. We do, however, conduct some specification checks, discussed further below, to verify that our key results are robust to the inclusion of the other market variables.

VI. IMPACT OF LOST STOCK MARKET WEALTH

A. Descriptive analysis

Before proceeding with our econometric analysis, we begin with a descriptive analysis of individual stock holdings using data from the 2007 SCF, the most recent data available. Our goals in this analysis are to determine the level of stock holdings and the differences in holdings across population subgroups and to get a sense of whether the level of stock holdings may be sufficient to influence individuals' retirement behavior if the market rises or falls.

Table 1 presents information on stock holdings for households headed by individuals between the ages of 55 and 64, who are likely to be contemplating retirement in the near future. The results in Table 1 indicate that the typical household's stock holdings are very small. In fact, the median values of directly held stocks, stock-based mutual funds, and retirement accounts (including DC pension plans and Individual Retirement Accounts, or IRAs) that include stocks are zero or very close to it. For all stock-based investments combined, the median value of holdings is just $8,000. The 75th percentile of this distribution is just under $100,000. One needs to look very high in the distribution in order to find households with very large levels of stock holdings.

As previewed earlier, stock ownership is strongly correlated with education. The share of households with any stock-based investments is 46 percent for high school graduates vs. 78 percent for college graduates. Furthermore, those with high levels of wealth are heavily concentrated among more highly educated individuals. For example, the 75th percentile of the distribution of all stock-based investments is just $28,500 for high school graduates vs. $271,300 for college graduates.

Despite the relatively low levels of stock holdings for most households in 2007, stock holdings are even lower at the beginning of our sample period. Similar calculations from the 1989 SCF (not reported on Table 1) indicate that the share of households with any stock-based investments rose from 36% in 1989 to 58% in 2007, while the median value conditional on holding any stock-based investments rose from $30,000 to $78,000. Increases for the college-educated group were similar in absolute terms (though smaller relative to the original values), with the share of stock owners rising from 60% to 78% and the median value conditional on holding stock-based assets rising from $70,000 to $125,000.

Table 2 presents a descriptive analysis of the impact that the recent stock market crash will have on future retirement income based on the 2007 stock holdings reported in Table 1. We begin by listing different levels of stock holdings ranging from none to $500,000 in Column 1. In Column 2, we identify the fraction of households headed by an individual between ages 55 and 64 that have stock holdings at that level or lower. About 42 percent have no holdings at all and 75 percent have $100,000 or less; 8 percent have $500,000 or more. In Column 3, we approximate the loss experienced by households at each the stock threshold, assuming that their portfolios fell by 50 percent. We then make the simplifying assumption that households consume 5 percent of their wealth per year to approximate the lost retirement income resulting

from the market crash. This is reported in Column 4; Column 5 divides this figure by 12 to get monthly statistics.

The results of this analysis suggest that if households divide this lost wealth over their remaining retirement years, the change in income would be modest for most of them. Those with $100,000 of stock holdings would lose $2,500 per year or $208 per month as a result of the stock market crash. These are not insignificant values, but the losses are likely to represent a small percentage of retirement income. The losses are, obviously, even smaller for those with less invested in stocks, a group that includes 75 percent of older households in 2007.

Our conclusion from this analysis is that, based on our assumptions, there are relatively few older households that lost enough money in the recent stock market crash that their retirement income will be substantively diminished.[10] Alternative assumptions, however, may lead one to predict a larger retirement response. Individuals could plan to consume a larger share of their savings just after they retire, for example to generate retirement income until Social Security benefits are available. If so, the relatively small amounts of stock holdings that most households have could lead to a substantial shock to retirement income, if just in the short-run. This could generate a larger retirement response. In the end, this is an empirical question that we will address using the regression techniques described earlier.

B. Econometric analysis

In our econometric analysis, we estimate regression models using data from the March CPS, where the dependent variable is an indicator variable for retirement and the key explanatory variable is the change in the S&P 500 Index. As we discussed earlier, we consider

[10] Gustman, Steinmeier, and Tabatabai (2010) come to a similar conclusion using even more detailed wealth data (including Social Security and DB pension wealth) available in the Health and Retirement Survey. In their analysis, they conclude that the share of wealth associated with stocks tends to be so small that even a dramatic decline in the stock market is unlikely to have retirement implications for many workers.

the one-year change, the five-year change, and the ten-year change because the time frame over which individuals respond to market fluctuations is not clear. We implement the quasi-experimental approach described earlier where we estimate the response to market changes across groups that differ by their likelihood of holding substantial amounts of stock. For instance, more educated respondents would be predicted to respond more strongly to market fluctuations. We also estimate models separately for those 55 to 61 and those 62 to 69, since 62 is the age at which individuals are first eligible for Social Security benefits and that eligibility may alter responses.[11]

The results of this analysis are presented in Table 3. Each cell in this table represents the results of a separate regression for the demographic groups previously identified. Based on the results reported here, there is some evidence supporting the notion that stock market fluctuations alter retirement behavior. This finding is strongest for those with more education who are between 62 and 69 and in response to long-term market fluctuations.[12] For workers in this age group, the coefficients on short-run fluctuations are positively signed, though there is no systematic pattern across educational attainment groups and coefficients are small in magnitude relative to the mean retirement rate. For example, a one-standard deviation (or 16 percentage point) increase in the one-year return increases the retirement rate of college graduates by 0.4 points, or 3.2 percent relative to the mean retirement rate of 11.7 percent. For the ten-year return, however, the pattern across educational groups is consistent with what we would predict

[11] We have also estimated regression models in which the effect of stock market fluctuations is allowed to vary over time, to allow for the possibility that the response has strengthened as the number of workers with stock market assets and the value of those assets has risen. We fail to find consistent evidence in support of this hypothesis.

[12] Even among the college-educated, heterogeneity exists in the level of stock holdings, which means that the results reported here reflect the impact for the average college graduate. Clearly, some college graduates have very high levels of stock holdings and the impact may be even larger for them.

and the magnitude of the coefficients is greater.[13] A one-standard deviation (or 77 percentage point) increase in the ten-year return increases the retirement rate of college graduates by 1.5 points, or 12.9 percent relative to the mean. Despite the relatively small sample sizes, we find that the effect for college graduates is statistically different (at the 10% level) from that for high school dropouts or high school graduates, though not different from the effect for those with some college, and also statistically different than the effect for all non-college graduates collectively.

For those workers age 55 to 61, few coefficients are statistically significantly different from zero and there is no systematic pattern in coefficients across education groups.[14] There is a positive and significant effect of 5- and 10-year returns for households with some college, but the fact that households headed by a college graduate have substantially greater stock holdings yet do not respond to these return measures makes us doubt that this results for the former group reflect a causal effect of stock returns. Point estimates on short-run (one-year) fluctuations are mainly wrong-signed.

VII. IMPACT OF LOST HOUSING WEALTH

Next, we turn to our econometric analysis of the effect of housing market fluctuations on retirement. As discussed above, this analysis is largely similar to the stock market analysis, except that we now have a true quasi-control group, renters. We thus compare results by home

[13] We have also estimated similar regression models distinguishing workers by whether or not they are covered by a private pension. The type of pension held (DB versus DC) or the dollar amount of their holdings is not reported, but those with pensions are likely to have greater stock market wealth than those without, forming another type of quasi-experiment. Results by pension status are not shown in the interest of space, but are consistent with the results by education group, in that they are more in line with our expectations for older workers than for younger workers and for long-term fluctuations than for short-run fluctuations. These results are available from the authors on request.

[14] The standard errors in these models, as well as those for the housing and labor markets, are clustered by state. We have experimented with clustering by year and using unclustered (robust) standard errors in the stock market regressions, and the results are quite similar to those reported in Table 3.

ownership status rather than education level. As before, we examine the effect of the market return over different time periods, one and five years. As discussed earlier, we use two price indices to measure the variation in home prices, the Case-Shiller Index and the OFHEO index, and identify the effects of home prices on retirement based on geographic differences in home price changes over time.

The results of this exercise are presented in Table 4. When we group all households together, the evidence that home price fluctuations affect retirement is weak.[15] In the models that use the Case-Shiller data, the coefficients are wrong-signed for 62 to 69 year olds (the group that was more responsive to stock market fluctuations) and are insignificant for all age groups and time horizons. In the models using the OFHEO data, the coefficients are larger for the older group and right-signed, but also insignificant. Results from the models that estimate the effect separately by home ownership status are largely similar. Once again, the specifications using the Case-Shiller data are not supportive of the hypothesis that home prices affect retirement, while those using the OFHEO data are more in line with our expectations, in that the coefficients on home price changes are uniformly positive for homeowners and larger for the older group, but statistically insignificant. Overall, we are unable to find support for the hypothesis that retirement is responsive to home price fluctuations. This finding is consistent with the previous literature suggesting that most households do not consume their home equity in retirement. Given our results, we make no attempt to include any changes in retirement resulting from home price fluctuations in the simulations of the effect of current market conditions on retirement presented below.

[15] One possible explanation for this finding is that individuals respond to the difference between actual and expected home price appreciation rather than to actual appreciation. As we have no individual-level data on expected home price appreciation, we calculate MSA-specific quadratic trends in real house prices and use these to calculate unexpected appreciation. We fail to find that retirement is responsive to this measure.

VIII. IMPACT OF LABOR MARKET CONDITIONS

A. Descriptive Analysis

Before reporting our econometric results, we begin by presenting a descriptive analysis designed to gauge the magnitude of the potential retirement response brought about by a weak labor market. Are there enough unemployed older workers and is the likelihood of their labor force withdrawal sufficiently large that we would be able to identify whether a labor market shock would generate an aggregate retirement effect?

To begin to address this issue, we first examine the level of unemployment among older workers and how this varies over the business cycle, using official statistics from the Bureau of Labor Statistics and our own calculations from the CPS. Older workers are less likely than the average worker to be unemployed. Unemployment rates for all workers cycle around a value in the vicinity of 6 percent, while the comparable figure for those 55 to 69 is more like 4 percent. The actual number of older workers who experience some unemployment over a given period (like a year), though, is greater than that. The unemployment rate is a point-in-time measure rather than a longer window available in a retrospective measure. Our calculations indicate that the number of workers 55 to 69 experiencing some unemployment over the past year is a number more like 8 percent with cyclical swings similar to those in the official unemployment rate.[16]

Earlier in this paper, we argued that we did not expect much of an aggregate retirement response to lost stock market wealth since so few individuals hold much wealth. Yet the number of people affected by labor market shocks is probably not a lot different. We would therefore

[16] See Levine (1993) for a comparison of retrospective and contemporaneous measures of unemployment.

only observe a bigger effect of labor market fluctuations on retirement if older workers who experience unemployment are quite likely to retire. In fact, this is what the evidence shows.

We first provide some descriptive evidence on this point using data from the Displaced Worker Survey (DWS), another supplement to the CPS. A displaced worker is someone who lost their job because of a plant closing, slack demand, or because their position was abolished. We calculate the rate at which workers displaced within the last three to five years withdrew from the labor force by the survey date. For those workers aged 20 to 54, roughly 10 percent withdrew. For those aged 55 to 69, roughly 30 percent withdrew. These withdrawals of older workers would be defined as a retirement, based on the operational definition of the term used in this analysis. Thus workers are very likely to retire in response to a job displacement.

Using our March CPS data directly, we can also distinguish retirement rates between unemployed older workers and others. Figures 4 and 5 present the results from such an analysis. In Figure 4, we present retirement hazard rates by age over the 1980 to 2007 sample period and in Figure 5 we present retirement hazard rates by year over the 55 to 69 age range.[17] In both figures, solid (dashed) lines represent the retirement rates for workers who experienced no (some) unemployment in the year preceding the survey. At all ages and in all years it is clear that unemployed older workers have higher retirement rates. These results along with those from the DWS are not conclusive in showing that unemployment "causes" increased retirement rates because workers who experience a job displacement or unemployment may be more likely to withdraw from the labor force for other reasons. Nonetheless, we view this evidence as supportive of a relationship between unemployment and retirement among older workers.

[17] An older worker who is in the labor force in, say, 2003, and withdraws by the March 2004 survey is said to retire in the year 2003. We define that worker's age according to the March 2004 reported age less one to approximate age in 2003.

Figure 5 provides additional evidence that unemployment may serve as a constraint that forces workers into retirement. For workers who experience no unemployment, there is a noticeable trend towards lower retirement rates over time. Annual retirement rates for these workers are about 10 percent in the beginning of the sample period, but begin to decline in the early 1990s, reaching a level of 6 percent by 2007. This pattern is consistent with the recent trend towards greater labor force participation among older workers. Interestingly, no such pattern exists among workers experiencing some unemployment. For them, retirement rates remain roughly constant (albeit a bit noisy due to smaller sample sizes) at around 16 percent. This suggests that whatever factors are driving many workers to choose to remain in the labor force longer are not influencing the behavior of unemployed older workers.[18] This would lead one to believe that other constraints may be dominating their behavior. Again, this evidence is merely suggestive that unemployment may play an important role in the retirement process for some workers. We move on to discuss the results of our econometric analysis next.

B. Econometric Analysis

The results of our econometric analysis are reported in Table 5. In the left part of the table we show results for the full sample as well as separate estimates for workers ages 55 to 61 and 62 to 69. On the whole, we find evidence that older workers' retirement behavior is responsive to changes in labor market conditions. A one percentage point increase in the unemployment rate increases the annual retirement rate by 0.18 percentage points. The average retirement rate is 9 percent per year, so this translates into a two percent increase relative to the mean. In the current crisis, the unemployment rate has risen by around five percentage points so

[18] Friedberg and Webb (2003) argue that the shift from DB to DC pensions can explain some of this increase; Gustman and Steinmeier (2008) make a similar argument with respect to changes in Social Security rules.

far. Our estimates suggest this would increase retirements by 0.9 percentage points, or ten percent relative to the mean retirement rate.

Breaking up our sample by age, we find that the entire effect is driven by those who are 62 to 69.[19] For 55 to 61 year old workers our results indicate a small and statistically insignificant effect of higher unemployment rates. For workers between the ages of 62 and 69, we find that a one percentage point increase in the unemployment rate would generate a 0.36 percentage point increase in the retirement rate. The five percentage point jump in the unemployment rate experienced recently is predicted to increase the rate of retirement by 1.8 percentage points, or 12 percent relative to the average retirement rate of 15.6 percent.

As in past analyses, we also estimate our models by education group; we report these results in the right part of Table 5. We find that high school graduates' retirement decisions are most responsive to a weak labor market.[20] For them, a five percentage point increase in the unemployment rate would generate a 1.8 percentage point increase in the retirement rate, a 19 percent increase relative to the mean. More-educated workers do not increase their retirement

[19] We have also estimated all models for both men and women separately. For the labor market regressions, we find some evidence that the impact of unemployment on retirement may be larger for women than for men. The strength of the evidence, however, is somewhat limited by the power of the analysis. For instance, in the aggregate, we find that women are more likely to retire in response to a cyclical downturn. On the other hand, the impact of a downturn on retirements among high school graduates is statistically significantly different from zero for both men and women, but not significantly different from each other. The same pattern holds true for the older group of workers as well (ages 62-69). Because we are unable to strongly determine differences in responsiveness by gender, we have chosen to group men and women together. For the stock market regressions, there are essentially no statistically significant differences between the stock return coefficients for men and women and no consistent pattern of greater responsiveness by either group. For the housing market regressions, we find some evidence of greater responsiveness by men, but only in the models using the 5-year changes in the Case-Shiller index.

[20] F-tests on the joint significance of the coefficients on the unemployment rate interacted with education level rejects the hypothesis that these coefficients are statistically identical (p-value = .019). When we test whether the unemployment rate coefficient for high school graduates is different than that for the other education groups, we find that the differences in the coefficients are significant against high school dropouts at the 5% level (p-value = .045), significant against those with some college at the 10% level (p-value = .066), and not quite significant at the 10% level (p-value = .115) against college graduates. It is our impression that these results are strong enough to conclude that there likely is a difference in the impact of labor market conditions across educational attainment categories.

significantly (in either a statistical or economic sense) in response to rising unemployment rates. Based on this evidence, we conclude that changes in labor market conditions have an important effect on retirement decisions, particularly for high school graduates.[21]

IX. OVERALL IMPACT ON RETIREMENT

The results that we have presented suggest that the stock market may cause some workers to delay retirement. In particular, in response to long-term declines in the value of stocks, highly educated workers between the ages of 62 and 69 appear to respond by reducing their likelihood of retirement. We find no support for the idea that declining housing values will have much impact in retirement. A likely explanation for this fact, as past research would suggest, is that few older workers use their housing wealth to finance retirement consumption. The impact of a sharply contracting labor market appears to be a relevant, and apparently overlooked, factor in forecasting coming retirement trends.

Taken together, our results suggest that retirements in the near term are likely to fall because of the long-term decline in stock prices, be largely unaffected by the decline in housing prices, and rise because of the increase in the unemployment rate.[22] The net effect is uncertain

[21]One interesting finding is that the retirement rates of high school dropouts do not appear to be affected by labor market conditions despite the fact that their employment is highly cyclically sensitive. The greater cyclical sensitivity in their employment, however, does not necessarily need to translate into a higher likelihood of retirement. It could be the case that the workers whose retirements are most affected are, for instance, manufacturing workers (high school graduates) who lose relatively well paying jobs during a recession, are unable to find jobs of similar quality, and retire as a result. By contrast, those at the very bottom of the distribution may have no alternative other than to keep looking for work because they have so few resources. This point is worthy of further study.

[22] As noted above, we have chosen to conduct three separate analyses of the three markets rather than one joint regression. We take several steps to confirm that our key results are not affected by this choice. First, all our models include the unemployment rate as a control variable, and the unemployment coefficient obtained in the models presented on Tables 3 and 4 (though not included on those tables) is quite similar to that reported in Table 5. Second, we have re-estimated the models on Table 3 including housing prices and the models on Table 4 including stock prices, and the pattern of results we obtain from this exercise is very similar to the original results.

because the effect of the long-term decline in stock prices and the rapidly rising unemployment rate tend to offset each other.

To assess the relative magnitudes of the two effects, we conduct a simulation exercise designed to estimate the number of individuals in a birth cohort likely to be affected by the recent changes in the stock market and the labor market. The results of this analysis are reported in Table 6. We begin by using data from the 2005 through 2007 American Community Survey (ACS) to estimate the size of the labor force by exact age. We find that there are 2.8 million individuals in the labor force at exact age 55, a figure that declines to 1.4 million at age 62, 800,000 at age 65, and 400,000 at age 69. Then we apply to these data age-specific hazard rates that we estimate using the March CPS data to arrive at the number of retirements we would predict over the course of a typical year at each exact age. These statistics represent a baseline of the "typical" number of expected retirements per year. In total, about 2 million workers between the ages of 55 to 69 would be expected to retire per year, on average.

The remainder of the table simulates the impact of the changes in retirement brought about by the weak labor market and the plunging stock market. We use the results presented in Tables 3 and 5 to implement this. In both cases, we use the regression coefficients relating changes in market conditions to changes in retirement rates that were estimated separately for workers ages 55 to 61 and 62 to 69. For the stock market, we focus on the ten-year change in the S&P 500 index and simulate the effect of a 110 point drop in the return, which is equivalent to moving from the average ten-year return during the past thirty years (62 percent) to the ten-year return experienced in the period ending in 2008 (-48 percent). For the labor market we estimate the impact of a five percentage point increase in the unemployment rate, approximating the actual rise in that rate from the low point of 4.4 percent in March 2007 to 9.4 percent in May

2009. The product of these changes in market conditions and the age-specific coefficient estimates from Tables 3 and 5 yields estimates of the change in hazard rates. We apply these estimates to the baseline hazard rate to obtain "adjusted" hazard rates. The product of the adjusted hazard rates and the actual number of workers in the labor force at each age provides an estimate of the adjusted number of individuals retiring. Taking the difference between these new estimates of the number of annual retirements and the number in the base case provides an estimate of the impact of the changes in market conditions on retirement.

The results presented in Table 6 suggest that 86,000 workers who otherwise would have retired will not do so as a result of the declining stock market that year. As that return converges back to normal rates, the annual number of delayed retirements will decline. As a simple example, suppose that it took five years for the market to revert to normal long-term rates of return at a linear rate. In this case, our simulations suggest that 258,000 workers would delay retirement over the course of the market downturn.

On the other hand, our estimates indicate that 126,000 workers will be forced into retirement this year as a result of the weak labor market. Similarly assuming a linear return to normal labor market conditions over a five-year period, we project that 378,000 workers will be forced to retire early as a result of the recession. Importantly, these results indicate that almost 50 percent more workers will be forced to retire because of the weak labor market than will be forced to work longer because they cannot afford to retire. On net, we predict that almost 120,000 additional retirements will occur as a result of the economic crisis.

We would further argue that the impact of a weak labor market on older workers' well-being may well be more significant than that of a weak stock market even if the number of older workers affected by each were similar. As we highlighted earlier, those workers forced to stay in

the labor force because the falling stock market reduced their retirement nest egg tend to be from wealthier households. The plunging stock market cannot hurt those without large stock holdings in the first place. For these workers, the alternative to retirement may be to work for another two or three years so that they have fewer years of retirement to finance and may replenish some of their lost wealth with additional savings. We do not mean to diminish this cost for those workers. Nevertheless, our results suggest that the weak labor market has its greatest impact on less educated workers who have fewer resources in the first place. Workers who are unable to replace labor earnings lost due to a job displacement by extending their working lives are likely to have lower levels of consumption for the rest of their lives. For instance, they may need to claim Social Security earlier than planned in order to make ends meet. Although the adjustment to Social Security benefits for early claiming is designed to be roughly actuarially fair, the worker's annual flow of income from this source is reduced if he retires earlier, increasing the household's risk of poverty in old age. The cost to these individuals appears to us to be greater than that experienced by workers with substantial stock holdings who are forced to work a few extra years to make up for stock losses.

X. CONCLUSIONS

Taken as a whole, our results indicate that the public discussion regarding the impact of the recent economic crisis on retirement is off target. Some relatively wealthier workers will be forced to delay retirement, but a larger number of workers with fewer economic resources will be forced into retirement because of their inability to find new jobs. These workers may need to start collecting retirement benefits now to make ends meet, resulting in lower income in retirement and an increased risk of poverty in old age. Indeed, the fact that Social Security

claims have risen sharply since the recession began suggests this response has already begun. Despite a wealth of media attention to the effect of the economic crisis on older workers, the risks they face as a result of weak labor markets have gone largely unnoticed.

More generally, our findings suggest that the role of labor market conditions in workers' retirement decisions has not received sufficient attention from economists. For example, while our earlier work suggests that the impact of unemployment on retirement is comparable in magnitude to that of poor health, the amount of research exploring the impact of health on retirement decisions dwarfs that on labor market conditions.

Our finding that labor market conditions are an important determinant of retirement decisions may also have important implications for public policy. One example of this is the debate over raising the Social Security normal and early retirement ages. With individuals living longer and drawing more Social Security benefits over their lifetimes than in the past, one possible reform to help address the financial shortfalls in the Social Security system is to raise the retirement ages. In the past, a common criticism regarding such proposals is that they will harm those individuals who are forced to retire involuntarily because of poor health. A substantial body of evidence exists supporting the notion that poor health is an important prelude to retirement for some older workers (Currie and Madrian, 1999). Our findings indicate that unemployment may be another involuntary mechanism that leads to retirement. The concerns of of older workers with weak labor market prospects may need additional consideration in the design of policies for workers nearing retirement age.

REFERENCES

Cahill, Kevin.E., Michael D. Giandrea, & Joseph F. Quinn (2006). Retirement patterns from career employment. *The Gerontologist*, 46(4), 514-523.

Chan, Sewin and Ann Huff Stevens (1999). "Employment and Retirement Following a Late-Career Job Loss." *American Economic Review*, Vol. 89, No. 2, pp. 211-216.

Chan, Sewin and Ann Huff Stevens (2001). "Job Loss and Employment Patterns of Older Workers." *Journal of Labor Economics*, Vol. 19, No. 2, pp. 484-521.

Chan, Sewin and Ann Huff Stevens (2004). "How Does Job Loss Affect the Timing of Retirement?" *Contributions to Economic Analysis & Policy*, Vol. 3, No. 1, article 5.

Coile, Courtney C. and Phillip B. Levine (2006). "Bulls, Bears, and Retirement Behavior." *Industrial and Labor Relations Review* 59(3): 408-429.

Coile, Courtney C. and Phillip B. Levine (2007). "Labor Market Shocks and Retirement: Do Government Programs Matter?" *Journal of Public Economics*. 91(10): 1902-1919.

Coronado, Julie and Maria Perozek (2003). "Wealth Effects and the Consumption of Leisure: Retirement Decisions During the Stock Market Boom of the 1990s," Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series #2003-20.

Currie, Janet and Brigitte C. Madrian (1999). "Health, Health Insurance, and the Labor Market," in Orley Ashenfelter and David Card (eds.). *Handbook of Labor Economics, Volume 3*. New York: Elsevier.

Dickerson, Marla. "Those Golden Years Have Lost Their Glow; With Home Values Down, Costs Up and Their 401(K)S Declining, Some Seniors Have Had To Rethink Retirement. *Los Angeles Times*. September 21, 2008, p. A1.

Dorning, Mike. "Early retirement claims increase dramatically." *Los Angeles Times.* May 24, 2009.

Employee Benefit Research Institute. (2009). *EBRI Databook on Employee Benefits*. Washington, DC: Employee Benefit Research Institute.

Farber, Henry S. (2008). "Job Loss and the Decline in Job Security in the United States." Industrial Relations Section Working Paper 520.

Friedberg, Leora and Anthony Webb. (2003). "Retirement and the Evolution of the Pension Structure," NBER Working Paper No. 9999.

Friedberg, Leora, Michael Owyang, and Anthony Webb (2008). "Identifying Local Differences in Retirement Patterns," Center for Retirement Research at Boston College Working Paper #2008-18.

Goss, Stephen C. 2009. *Applications for Social Security Retired Worker Benefits in Fiscal Year 2009.* Washington, DC: Social Security Administration, Office of the Chief Actuary. May 28, 2009.

Gustman, Alan and Thoman Steinmeier. (2008). "How Changes in Social Security May Affect Recent Retirement Trends," NBER Working Paper No. 14105.

Gustman, Alan L., Thomas L. Steinmeier, and Nahid Tabatabai (2010). "What the Stock Market Decline Means for the Financial Security and Retirement Choices of the Near-Retirement Population," Journal of Economic Perspectives 24(1):161-182.

Hallberg, Daniel. 2008. "Economic Fluctuations and Retirement of Older Employees," Institute for Labour Market Policy Evaluation, number 2008-2.

Hoynes, Hilary (2000). "The Employment, Earnings, and Income of Less Skilled Workers Over the Business Cycle," in Rebecca M. Blank and David E. Card (eds.), *Finding Jobs: Work and Welfare Reform*. New York: Russell Sage.

Hurd, Michael D., Monika Reti, Susann Rohwedder. 2009. "The Effect of Large Capital Gains or Losses on Retirement", in David Wise (ed.), *Developments in the Economics of Aging*. Chicago: University of Chicago Press.

Johnson, Richard W. and Corina Mommaerts. 2010. *Social Security Retirement Benefit Awards Hit All-Time High in 2009.* Washington, DC: Urban Institute.

Levine, Phillip B. (1993). "A Comparison of Contemporaneous and Retrospective Measures of Unemployment from the Current Population Survey." *Monthly Labor Review*, 116(5): 33-38.

Madrian, Brigitte, and Lars John Lefgren. "A Note on Longitudinally Matching Current Population Survey (CPS) Respondents." National Bureau of Economic Research, technical working paper t0247, November 1999.

Munnell, Alicia, Steven Sass, Mauricio Soto, and Natalia Zhivan (2006). "Has the Displacement of Older Workers Increased?" Center for Retirement Research Working Paper 2006-17.

Alicia H. Munnell, Mauricio Soto, Robert K. Triest, and Natalia A. Zhivan (2008). "How Much Do State Economics and Other Characteristics Affect Labor Force Participation of Older Workers?," Center for Retirement Research at Boston College Working Paper #2008-12.

Sevak, Purvi. 2001. "Wealth Shocks and Retirement Timing: Evidence from the Nineties." Michigan Retirement Research Center, Working Paper WP00D1.

Trejos, Nancy. "Will You Retire?; New Economic Realities Keep More Americans In the Workforce Longer." *Washington Post*, October 15, 2008, p. H1.

Trumbull, Mark. "Economic Crisis Scrambles Retirement Math: The 401(K) Model of Saving is Under Duress as Stocks Slide. Home Equity Losses Don't Help." *Christian Science Monitor*, March 4, 2009

Venti, Steven F. and David A. Wise (2004). Aging and Housing Equity: Another Look," in David A. Wise (ed.), *Perspectives on the Economics of Aging.* Chicago: University of Chicago Press.

von Wachter, Till (2007). "The Effect of Economic Conditions on the Employment of Workers Nearing Retirement Age." Center for Retirement Research Working Paper 2007-25.

Table 1: Equity Holdings of Households Age 55-64 by Education Group, 2007 SCF

Category	% with Holdings	Median Conditional on Holding	Values among All Households at Percentile: 25th	50th	75th	90th	95th
All							
Directly-Held Stocks	0.213	24,000	0	0	0	25,000	125,000
Stock Mutual Funds (Non-Ret.)	0.140	97,000	0	0	0	45,000	191,000
Stocks in Retirement Accounts	0.500	66,500	0	20	66,500	230,000	447,500
Any Stocks	0.583	78,000	0	8,000	97,500	357,620	752,000
Less than High School							
Directly-Held Stocks	0.054	270	0	0	0	0	50
Stock Mutual Funds (Non-Ret.)	0.019	3,000	0	0	0	0	0
Stocks in Retirement Accounts	0.214	10,000	0	0	0	10,000	70,000
Any Stocks	0.214	10,000	0	0	0	10,000	70,000
High School							
Directly-Held Stocks	0.127	9,000	0	0	0	500	14,000
Stock Mutual Funds (Non-Ret.)	0.069	50,000	0	0	0	0	38,000
Stocks in Retirement Accounts	0.366	33,800	0	0	15,000	88,000	188,800
Any Stocks	0.460	35,000	0	0	28,500	130,000	212,500
Some College							
Directly-Held Stocks	0.156	3,500	0	0	0	2,000	15,000
Stock Mutual Funds (Non-Ret.)	0.060	45,000	0	0	0	0	20,000
Stocks in Retirement Accounts	0.503	60,000	0	20	61,600	160,000	224,000
Any Stocks	0.558	65,000	0	4,000	73,500	197,150	319,500
College Graduate							
Directly-Held Stocks	0.342	60,000	0	0	13,000	154,000	500,000
Stock Mutual Funds (Non-Ret.)	0.260	107,000	0	0	4,700	200,000	385,000
Stocks in Retirement Accounts	0.668	85,000	0	27,000	159,600	480,000	775,800
Any Stocks	0.775	125,000	3,250	65,100	271,300	846,000	1,865,000

Note: data are weighted to be representative of the U.S. population.

Table 2: Equity Losses of SCF Households Age 55-64 in 2008 Market Crash

Stock Assets in 2007 SCF (1)	% of Sample w/ assets at/below (2)	Asset Loss (3)	Lost Annual Retirement Income (4)	Lost Monthly Retirement Income (5)
0	0.417	0	0	0
25,000	0.587	12,500	625	52
50,000	0.654	25,000	1,250	104
100,000	0.751	50,000	2,500	208
250,000	0.869	125,000	6,250	521
500,000	0.920	250,000	12,500	1,042

Notes:

1. Assets are assumed to have dropped by 50% in value since 2007 SCF.
2. Lost retirement income is calculated by assuming that household will consume 5% of wealth each year.

Table 3: Effect of Stock Market Fluctuations on Retirement by Age, March CPS

Measures of Stock Market Performance	All	High School Dropout	High School Graduate	Attended Some College	College Graduate
Age 55-61					
Mean of Dependent Variable	0.059	0.075	0.062	0.058	0.044
% change S&P 500 - 12 Mo. (* 100)	-0.0004	0.0140	-0.0026	-0.0041	-0.0035
	(0.0041)	(0.0111)	(0.0056)	(0.0087)	(0.0068)
% change S&P 500 - 5 Year (* 100)	0.0039	0.0001	0.0024	0.0105	0.0014
	(0.0017)	(0.0046)	(0.0024)	(0.0041)	(0.0026)
% change S&P 500 - 10 year (* 100)	0.0022	-0.0021	-0.0013	0.0080	0.0028
	(0.0014)	(0.0044)	(0.0020)	(0.0032)	(0.0027)
Sample Size	210,807	42,020	72,495	43,828	52,464
Age 62-69					
Mean of Dependent Variable	0.159	0.194	0.165	0.146	0.120
% change S&P 500 - 12 Mo. (* 100)	0.0153	0.0043	0.0113	0.0265	0.0206
	(0.0083)	(0.0172)	(0.0157)	(0.0167)	(0.0161)
% change S&P 500 - 5 year (* 100)	0.0013	-0.0062	0.0003	-0.0022	0.0109
	(0.0027)	(0.0071)	(0.0067)	(0.0070)	(0.0066)
% change S&P 500 - 10 year (* 100)	0.0054	-0.0057	-0.0023	0.0120	0.0184
	(0.0040)	(0.0099)	(0.0065)	(0.0065)	(0.0080)
Sample Size	97,408	24,297	33,271	18,019	21,821

Note: Each cell entry represents a separate regression that also includes age dummies, race and ethnicity, gender, marital status, children less than 18, education, unemployment rate, state fixed effects, and a quadratic year trend. Regressions are weighted by sample weights. Standard errors are clustered at the state level. Reported coefficients show the effect of a one hundred percentage point change in the S&P 500 (e.g., a doubling of real stock values).

Table 4: Impact of Real House Price Fluctuations on the Likelihood of "Retiring" in March CPS, by Age
(standard errors in parentheses)

	Case-Shiller Data				OFHEO Data			
	12-Month Change		5-Year Change		12-Month Change		5-Year Change	
	62 to 69	55 to 61	62 to 69	55 to 61	62 to 69	55 to 61	62 to 69	55 to 61
Mean of Dependent Variable	0.134	0.052	0.134	0.052	0.143	0.054	0.143	0.054
% Change in Index (* 100)	-0.0251	0.0039	-0.0107	0.0256	0.0590	0.0025	0.0165	0.0071
	(0.0728)	(0.0410)	(0.0171)	(0.0075)	(0.0567)	(0.0369)	(0.0149)	(0.0099)
% Change Index * Owner (* 100)	-0.0234	0.0150	-0.0097	0.0276	0.0835	0.0156	0.0211	0.0076
	(0.0775)	(0.0417)	(0.0165)	(0.0074)	(0.0567)	(0.0428)	(0.0142)	(0.0108)
% Change Index * Renter (* 100)	-0.0184	-0.0293	-0.0156	0.0196	-0.0434	-0.0555	-0.0038	0.0055
	(0.0636)	(0.0584)	(0.0255)	(0.0113)	(0.1028)	(0.0441)	(0.0266)	(0.0118)
Homeowner	0.0197	-0.0056	0.0226	-0.0083	0.0037	-0.0038	0.0058	-0.0030
	(0.0099)	(0.0049)	(0.0109)	(0.0053)	(0.0057)	(0.0032)	(0.0061)	(0.0034)
Sample Size	14,784	33,126	11,709	27,310	46,993	106,808	45,008	102,436

Notes: Every column and each panel represents the results from a different regression in models where the dependent variable is an indicator for retirement and the key independent variables are those listed. Additional explanatory variables include: age dummies, race and ethnicity, gender, marital status, children less than 18, education, unemployment rate, MSA fixed effects, and year fixed effects. Regressions are weighted by sample weights. Standard errors are clustered at the MSA level. Reported coefficients show the effect of a one hundred percentage point change in the house price index (e.g., a doubling of real house values).

Table 5: Impact of Labor Market Conditions on the Likelihood of "Retiring," by Age and Educational Attainment
(standard errors in parentheses, sample size in brackets)

	Age 55 to 69	Age 62 to 69	Age 55 to 61	HS Dropout	HS Graduate	Some College	College Graduate
Mean of Dependent Variable	0.090	0.156	0.059	0.118	0.094	0.084	0.067
Coefficient on	0.018	0.036	0.010	0.006	0.035	0.001	0.008
Unemployment Rate (*10)	(0.006)	(0.014)	(0.007)	(0.014)	(0.011)	(0.013)	(0.016)
Sample Size	308,215	97,408	210,807	66,317	105,766	61,847	74,285

Notes: Each cell entry represents the coefficient on the unemployment rate in a separate regression that also includes age dummies, race and ethnicity, gender, marital status, children less than 18, education, and state and year fixed effects. Regressions are weighted by sample weights. Standard errors are clustered at the state level. Reported coefficients show the effect of a ten point change in the unemployment rate.

Table 6: Simulated Impact of Economic Crisis on Retirements
(all numbers in thousands)

	Baseline Statistics			Impact of Decline in Long-Term Stock Market Return		Impact of Increased Unemployment	
Age	Number in Labor Force	Hazard Rate	Number Retiring	Adjusted Hazard Rates	Adjusted Number Retiring	Adjusted Hazard Rates	Adjusted Number Retiring
55	2,805	0.045	127	0.043	120	0.046	129
56	2,600	0.049	126	0.046	120	0.049	127
57	2,489	0.054	134	0.051	128	0.054	135
58	2,420	0.054	131	0.052	125	0.055	132
59	2,172	0.060	131	0.058	126	0.061	132
60	1,908	0.079	152	0.077	147	0.08	152
61	1,551	0.086	133	0.083	129	0.086	134
62	1,391	0.162	225	0.154	215	0.180	250
63	1,189	0.138	164	0.131	156	0.156	185
64	1,035	0.130	134	0.123	127	0.148	153
65	794	0.194	154	0.187	149	0.212	169
66	641	0.163	104	0.156	100	0.181	116
67	578	0.158	92	0.151	88	0.176	102
68	515	0.161	83	0.154	79	0.179	92
69	433	0.154	67	0.147	64	0.172	75
total	22,522		1,957		1,871		2,083
Impact on Retirement					-86		126

Notes: The baseline number of workers in the labor force comes from the 2005-2007 American Community Survey. The baseline hazard rates are estimated from the March CPS

Figure 1A: Annual Real Percentage Change in S&P 500

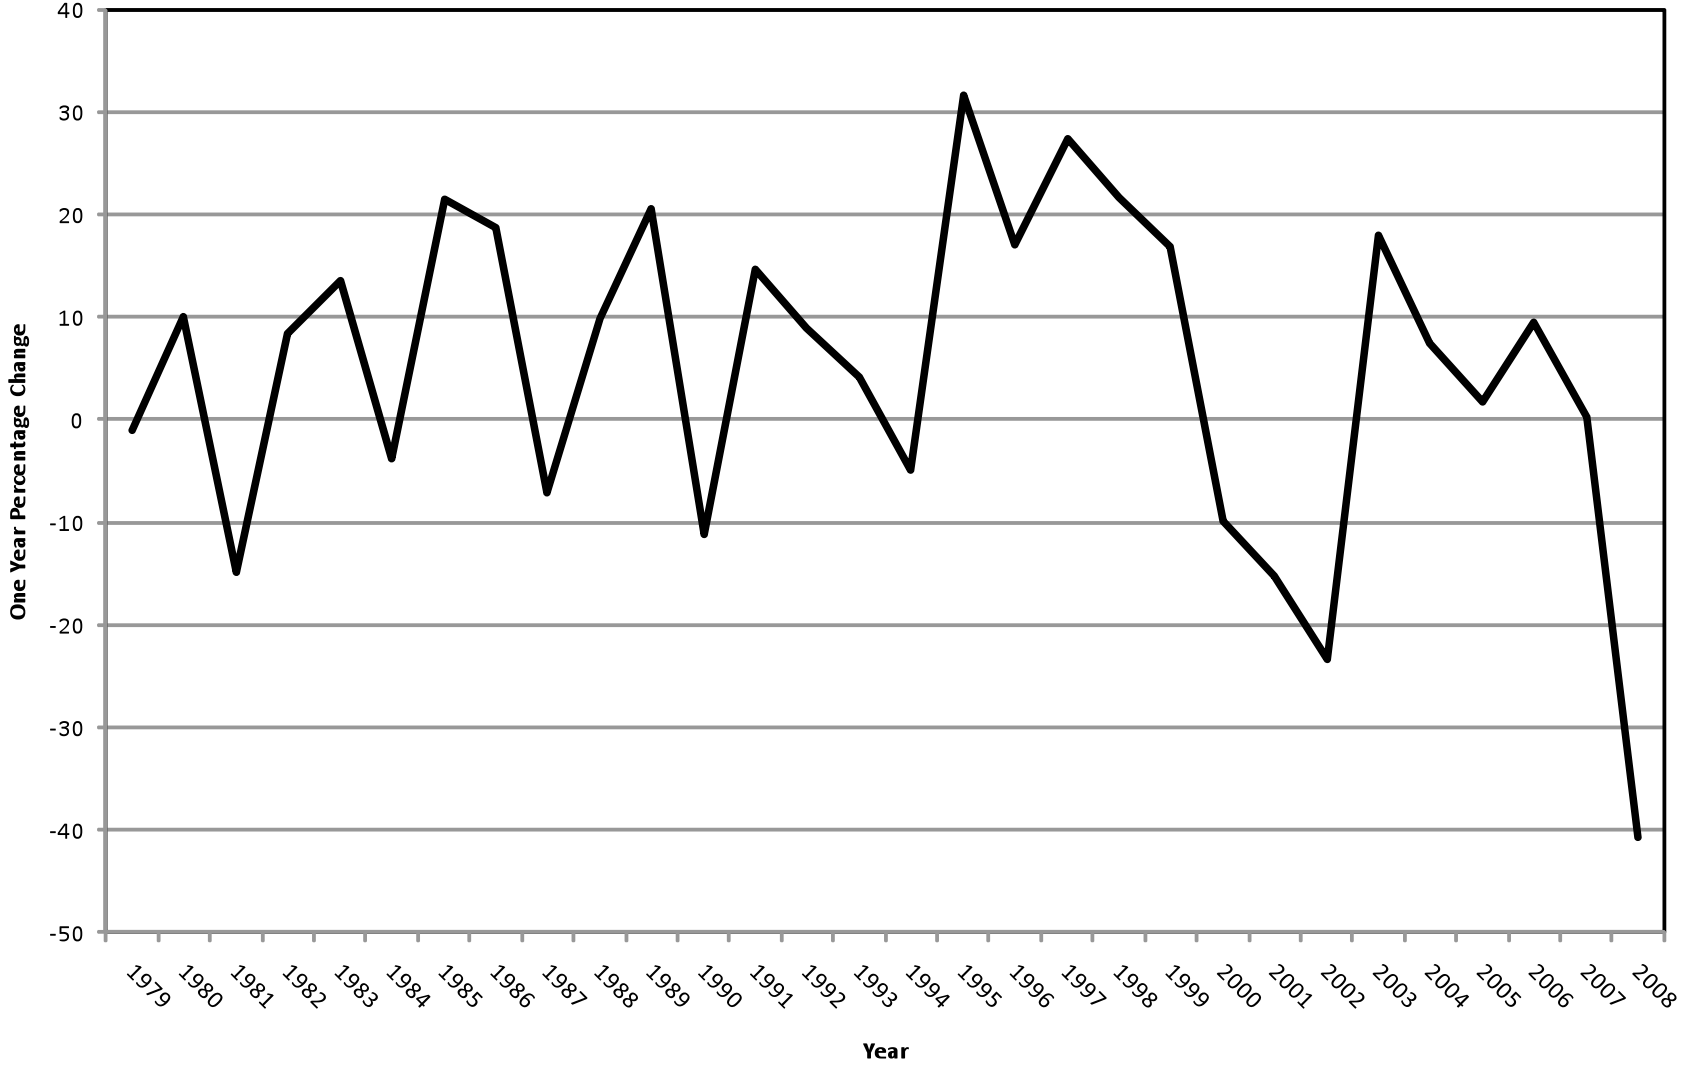

note: annual percentage change is calculated using December to December monthly averages

Figure 1B: Five and Ten Year Real Percentage Change in S&P 500

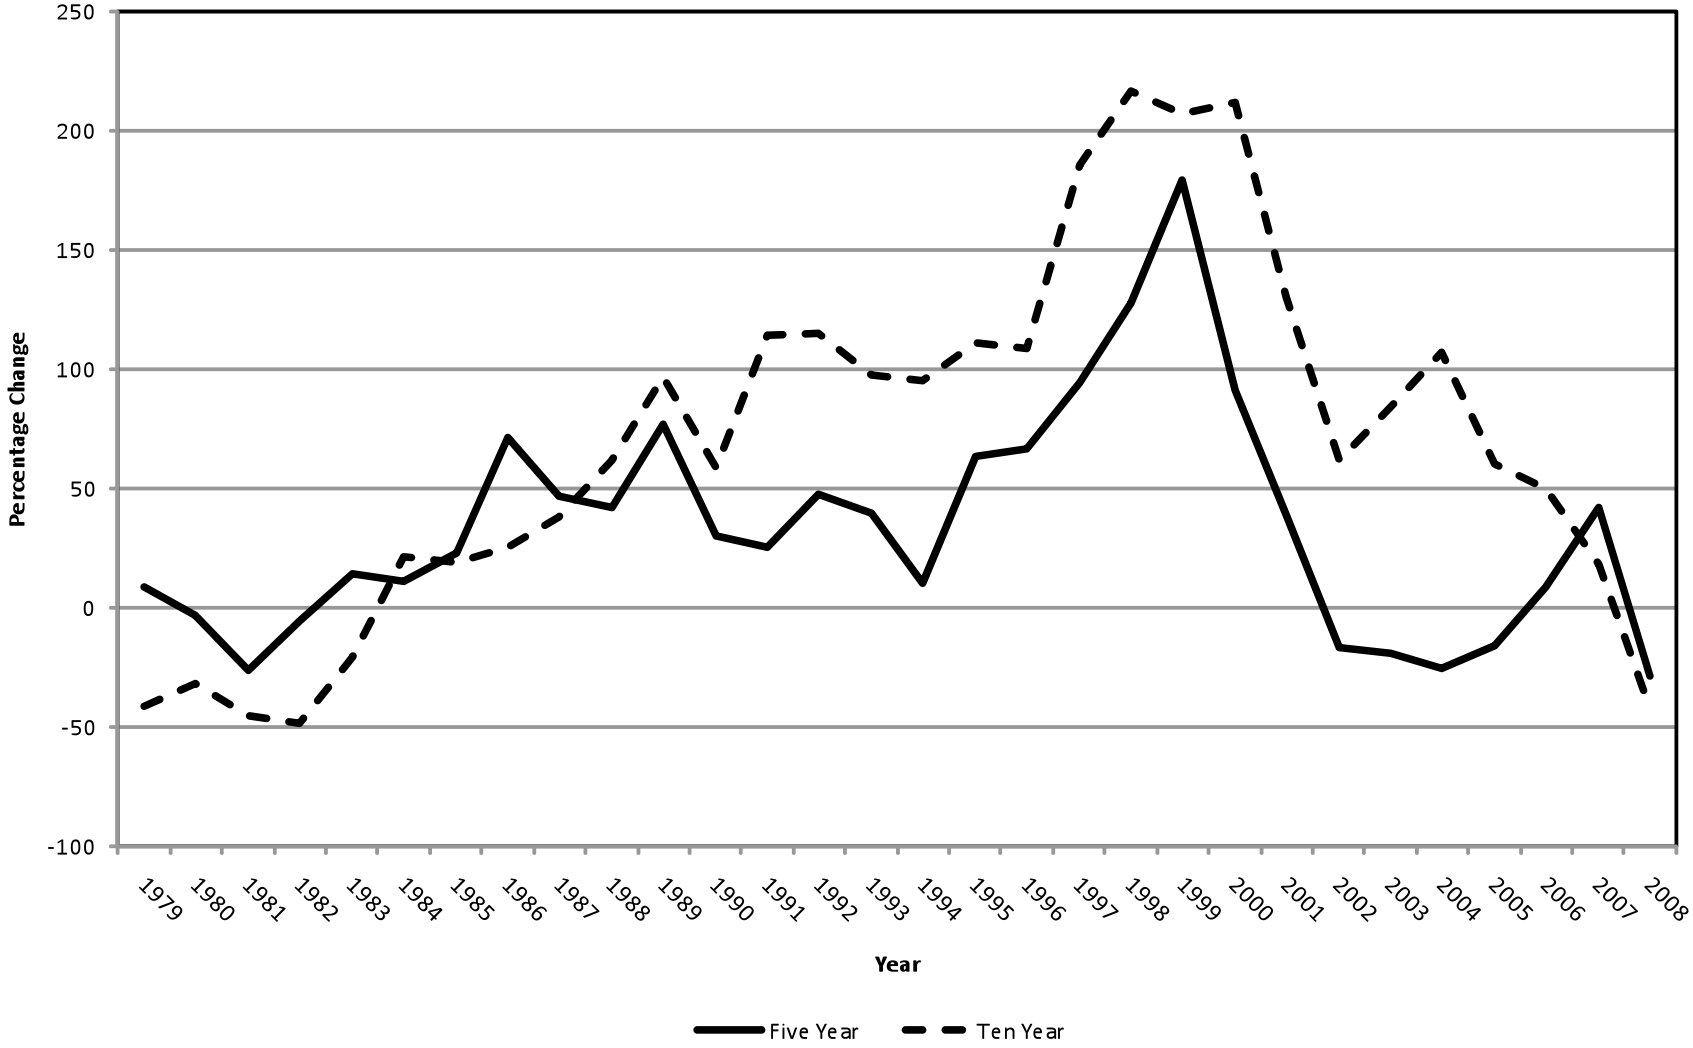

note: percentage changes are calculated using December to December monthly averages.

Figure 2: Annual Percentage Change in Real House Prices

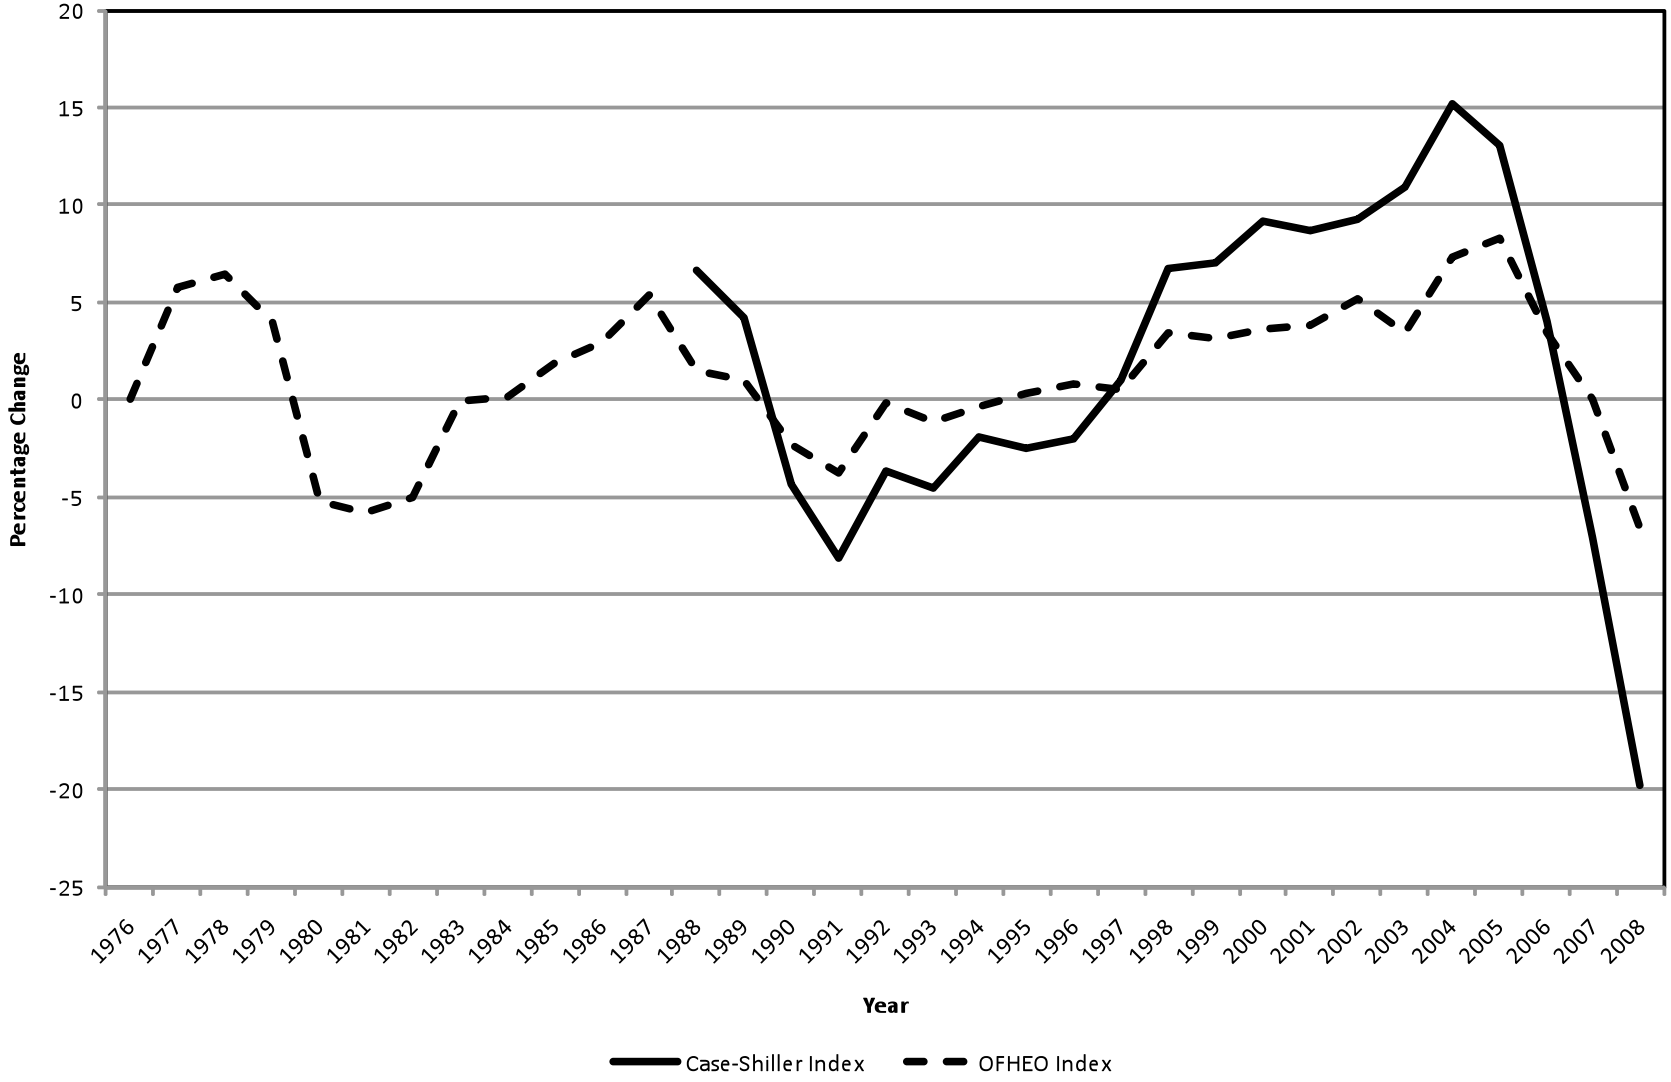

note: annual percentage change is calculated using December values for Case-Shiller Index and 4th quarter values for OFHEO Index.

Figure 3: U.S. Unemployment Rate

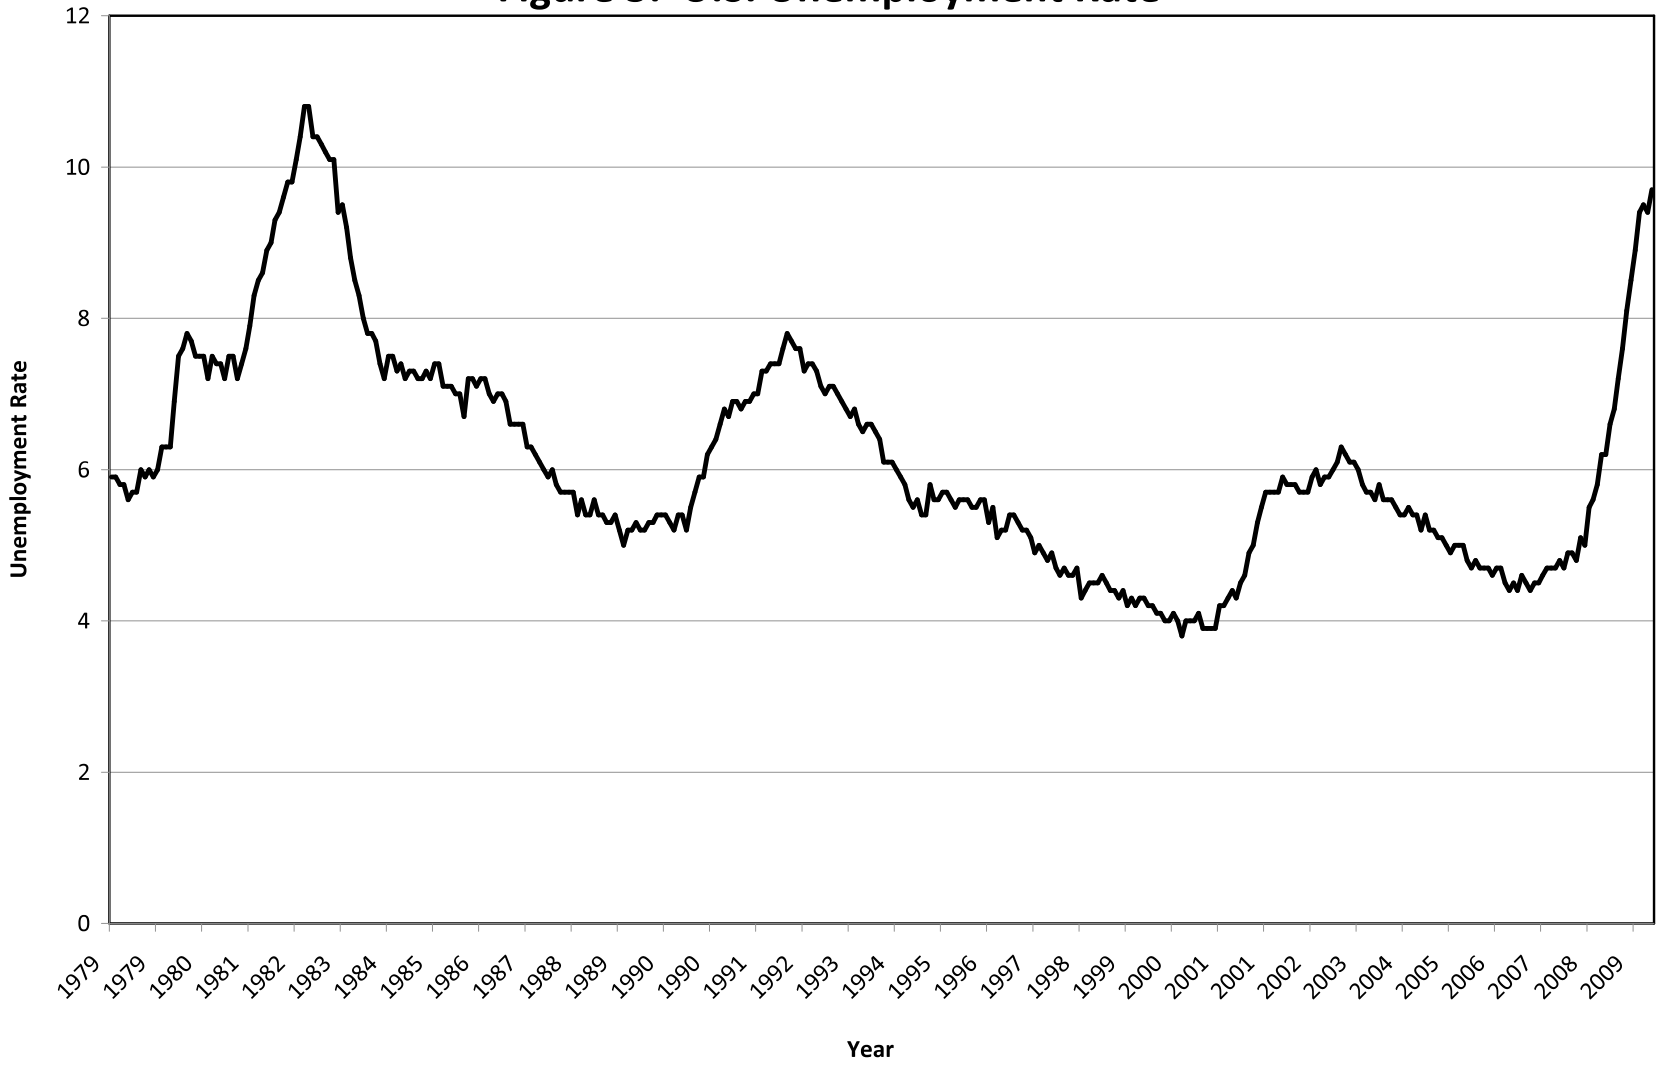

Figure 4: Empirical Retirement Hazard Rates by Age and Unemployment Status, 1980 to 2007, March CPS Data

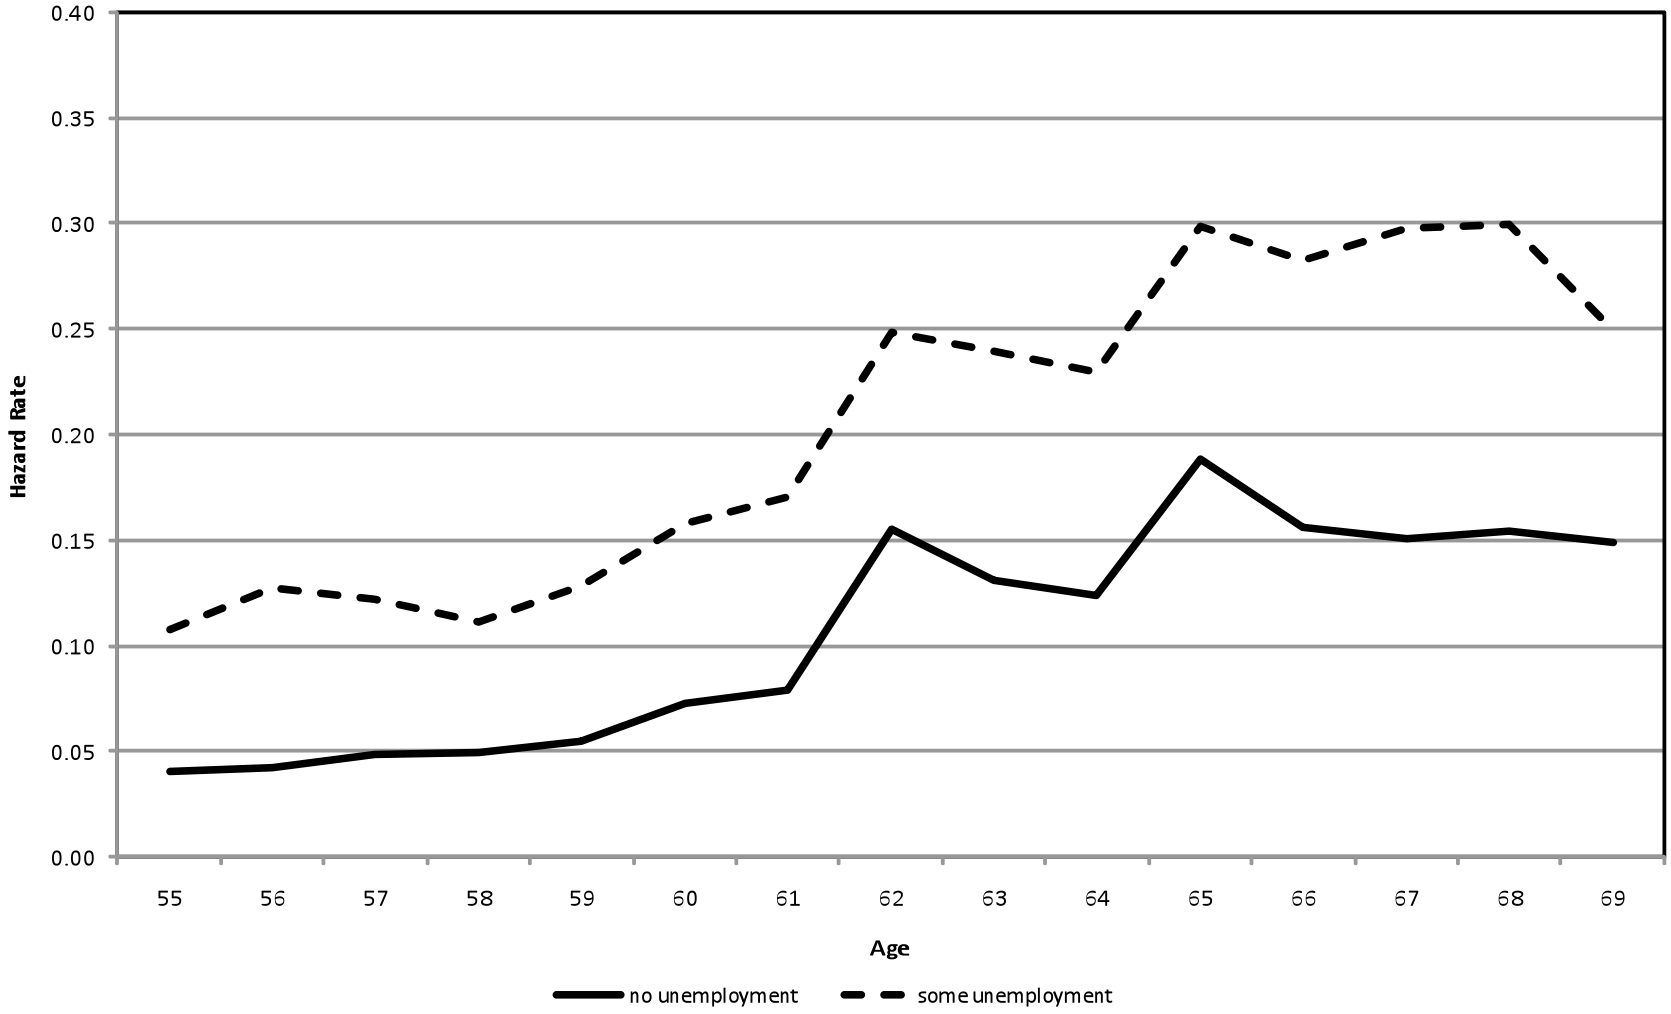

source: authors' calculations from March CPS Data.

Figure 5: Empirical Retirement Hazard Rates over Time, Workers Age 55-69, March Current Population Survey

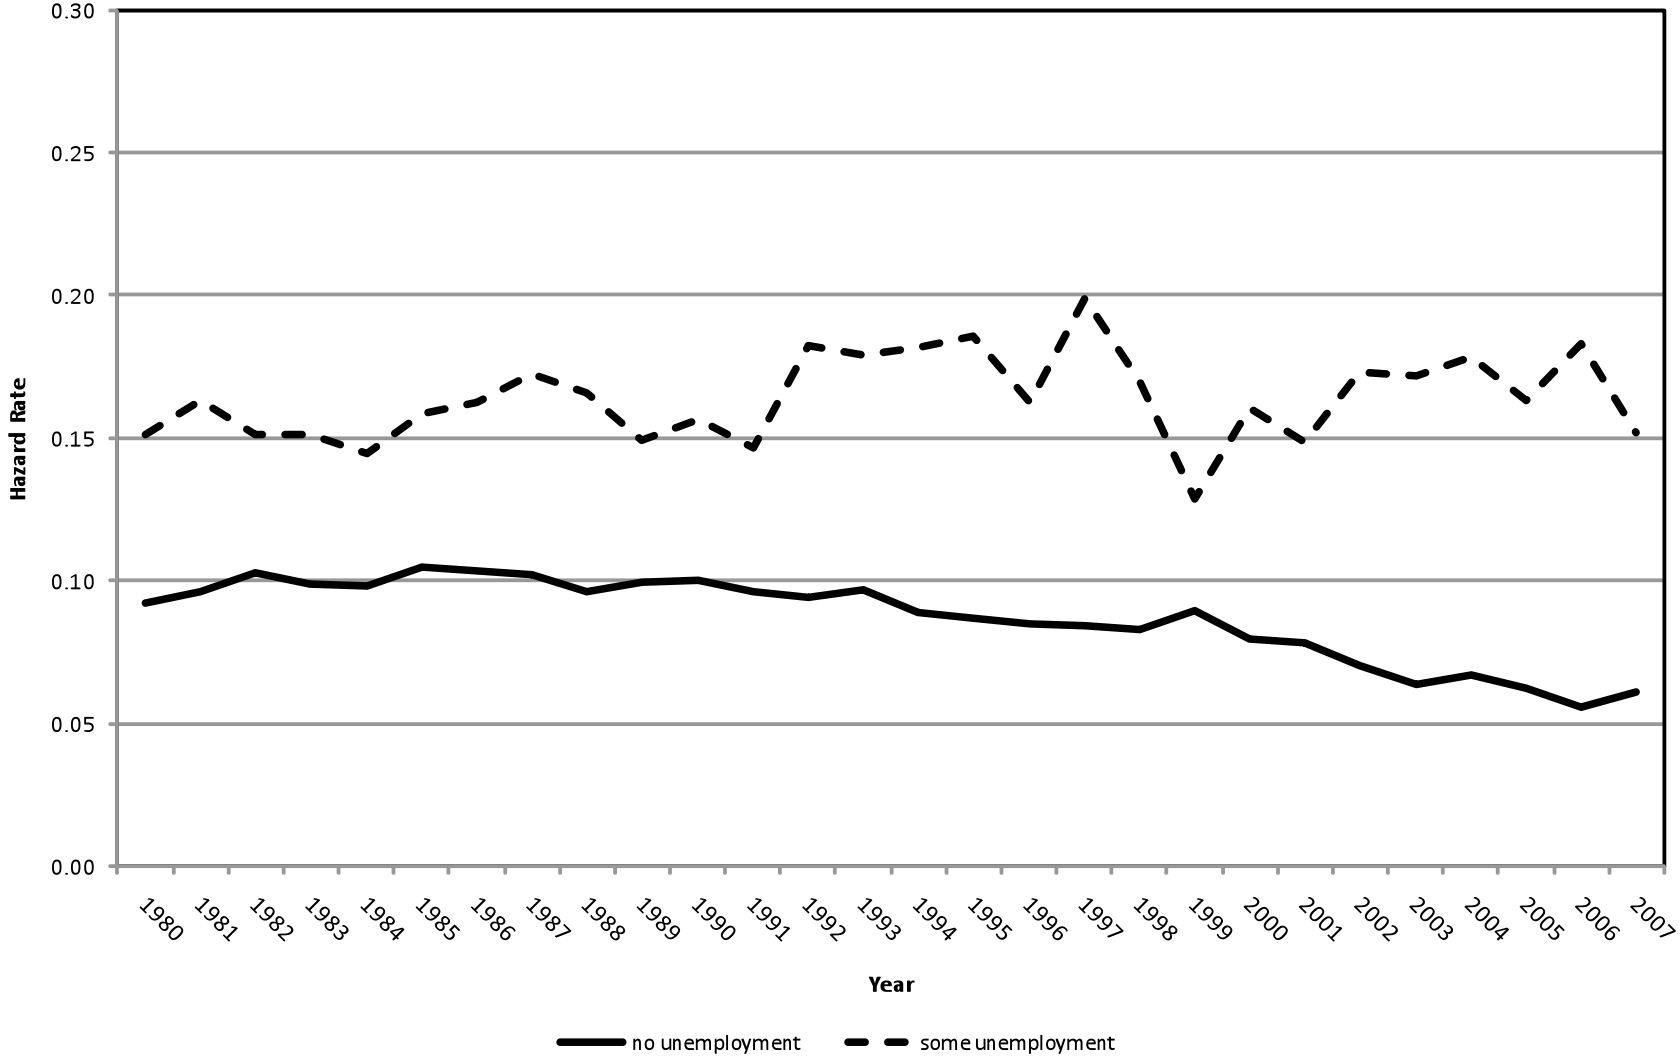

source: authors' calculations from March CPS Data.