The Great Recession and the Retirement Plans of Older Americans*

Brooke Helppie McFall†

September 23, 2011

Abstract

This paper examines the labor supply effects of the wealth losses during the stock market crash of 2008 and 2009. A life-cycle model incorporating both consumption and retirement timing implies that exogenous wealth losses should delay optimal retirement timing. Using data from the Cognitive Economics Study and the Health and Retirement Study, this paper quantifies the wealth losses suffered by older Americans in terms of the additional length of time they would have to work to maintain the pre-crash consumption plan implied by their wealth holdings and expected retirement timing. Using these measures, Tobit regressions and a novel method for reducing the impact of error-ridden observations are used to examine the relationship between this measure of wealth loss and retirement planning. Several potential sources of heterogeneity in individuals' reactions to the crash are also examined. Results imply that wealth losses of 2008 and 2009 are associated with an increase in planned retirement age on the order of a few weeks to a few months for the average older American, but up to several months for some segments of the population. These results are consistent with results of recent studies and the life-cycle model, but stand in contrast to other examinations of wealth shocks on the general population of older Americans.

1 Introduction

On October 1, 2007, the Dow Jones Industrial Average reached the 14,000 mark to close at a new all-time high. Within two weeks, closing values began a slow decline that would leave

*The support of National Institute on Aging grant 2P01AG026571 is gratefully acknowledged.
†University of Michigan. E-mail: bhelppie@umich.edu
†University of Michigan. E-mail: bhelppie@umich.edu

the Dow more than twenty percent lower by the following autumn. But the real crash was yet to come, as the weakening real estate market and the resulting failure of major banks in September and October 2008 sent stock values into a series of steep declines. It was five more months before the stock market hit bottom: on March 9, 2009, the Dow closed at 6,547.05, less than half of its October 2007 peak, and on par with closing prices from a decade earlier. For older Americans, whose stock holdings had grown to more than fifteen percent of total assets by 2006 (Gustman, Steinmeier and Tabatabai, 2010, p. 311), the stock market crash of 2008 caused large, unanticipated, and widespread losses of wealth over a period of just a few months.

In addition to its role in bank failures and the stock market crash, the weak real estate market directly impacted households, reducing housing prices by more than thirty percent from their peak in the first half of 2006 through early 2009 (S&P/Case-Shiller). Between September 2008 and May 2009, the national unemployment rate increased by more than fifty percent (Bureau of Labor Statistics, Current Population Survey), providing a further threat to older Americans' financial stability through erosion of employment security and earnings.

An intertemporal budget constraint from a simple, dynamic life-cycle model of consumption and labor supply dictates that if even a portion of the wealth losses of 2008 and 2009 are permanent, those who lost wealth must increase future earnings, decrease future consumption, or both. Since one way to increase earnings is to work longer, retirement timing is an important margin along which individuals might adjust to wealth losses.

In this paper, I use data from two nationally-representative studies— the Cognitive Economics Study and the Health and Retirement Study— to examine the impact of recent wealth losses on the retirement plans of older Americans. I begin by quantifying the wealth losses suffered by older Americans in terms of the additional length of time they would have to work to maintain the pre-crash consumption plan implied by their wealth holdings and expected retirement timing. This measure of wealth loss has the intuitive interpretation that, if individuals cared only to maintain their pre-crash consumption plans, a loss of wealth equivalent to an additional year of work would result in a one year increase in planned retirement age. I then use descriptive and regression analysis to study the relationship between this measure of wealth loss and the reported changes in retirement timing. In extensions of my basic regression analysis, I also examine several potential sources of heterogeneity in individuals' reactions to the crash.

My analyses show that older Americans plan to delay retirement in response to the crash. My preferred estimates imply that the average wealth loss between July 2008 and May 2009 is associated with an increase in expected retirement age of approximately four months, about 8.6 percent of the adjustment that would be needed to fully make up for wealth losses.

Additionally, the average wealth loss is associated with increases in the probabilities that an individual will be working full-time after reaching age 62 and after reaching age 65.

This paper is the first to use new data from pre- and post-crash surveys from the Cognitive Economics Study (CogEcon) and the Health and Retirement Study (HRS) to examine the impact of wealth shocks on the age at which older adults expect to retire. In many analyses presented in this paper, the association between changes in wealth and changes in retirement plans are statistically significant. Unlike most previous research, this paper finds statistically-significant evidence of wealth effects on retirement timing. Moreover, it is the first to examine the role of heterogeneity in reactions to wealth shocks by wealth, time to retirement, expectations about future economic conditions, cognitive ability, financial knowledge, and changes in bequest plans.

2 Background

The classic life-cycle model predicts that optimal consumption from the present until the end of life should be proportional to net worth, where net worth is defined as the net value of assets currently held plus the discounted value of future income (Modigliani and Brumberg, 1954/2005). The key components of the model that are responsible for this prediction are an individual utility function and a basic intertemporal budget constraint. The utility function drives individuals' desire to smooth consumption over time, while the budget constraint requires that the present discounted value of all future consumption must be equal to the sum of current assets and the expected present discounted value of future income flows. Assuming that individuals expect to work through the middle of their lives and retire towards the end of life, the life-cycle model implies that individuals will save while working to fund a smooth consumption path over the rest of their lives (Modigliani and Brumberg, 1954/2005).

The classic life-cycle model treats retirement as a period of life during which individuals do not work, and must live out of savings. The labor supply decision, including the decision to retire, is not explicitly modeled. Over the last three decades, however, this life-cycle model has spawned a class of dynamic, structural models in which retirement is a choice variable, the timing of which is driven by a preference for retirement leisure, a disutility of work and/or a real wage that declines as workers age. These models seek to predict how consumption, saving and labor supply decisions are affected by income, individual preferences, risk, pension and Social Security rules and other variables of interest in a public policy context [MaCurdy, 1981, Gustman and Steinmeier, 1986, Kimball and Shapiro, 2003, Blau, 2008, Low et al., 2010]. They confirm the prediction of the basic life-cycle model that permanent increases in lifetime resources result in increased future consumption, and permanent decreases in lifetime resources result in decreased future consumption.

Further, when facing an unforeseen negative shock to assets, the intertemporal budget constraint implies that individuals must increase income, reduce planned consumption or adjust both income and consumption. Similarly, an unforeseen positive shock must result in increased consumption, reduced income, or a combination of both. Thus, these dynamic models predict that unexpected changes in wage rates or other shocks to wealth should affect individuals' labor supply decisions, including their retirement timing.

Despite widespread use of these models, a large body of literature assessing the impact of changes in wages on labor supply has not clearly supported the implications of dynamic life-cycle models. Several papers using experimental data and evidence from inheritances have found that large, unforeseen monetary gains are associated with reduced labor supply, often in the form of earlier retirement [Holtz-Eakin et al., 1993, Imbens et al., 2001, Kimball and Shapiro, 2008]. However, empirical studies of the impact of broad wealth shocks due to stock market movements in the 1990s and early 2000s have generally failed to show strong associations with changes in retirement timing (Hurd and Reti, 2001, Coronado and Perozek, 2003, Kezdi and Sevak, 2004, Coile and Levine, 2005, Hurd et al., 2009; the main exception is Sevak, 2002).

While it is possible that the implications of the life-cycle hypothesis on retirement timing are not borne out in the real world, there are three other possible explanations that may have contributed to the mixed findings of broadly-representative empirical analyses in the past. First, the generally weak results in papers seeking identification based on the impact of broad economic trends may be partially attributable to the difficulty of finding sources of plausibly exogenous variation in wealth. Second, the combination of high levels of measurement error in wealth data with the relatively small changes in wealth most households have experienced in past business cycles may have caused non-findings due to attenuation bias. Third, the possibility that fixed costs and non-linearities in the underlying models may mean that reduced-form econometric models that ignore these issues have produced unreliable results. My study has advantages over existing papers in each of these three areas.

First, the economic crisis of 2008 provides a more powerful example of a plausibly exogenous wealth shock than the 1990s and early 2000s business cycle, the period that has been the focus of most similar studies. The recent downturn caused losses broadly, affecting stock values and employment across many industries, as well as the real estate sector. By contrast, the late 1990s/early 2000s business cycle was based on the protracted rise and subsequent fall of "dot-com" industries and their stocks. Because fewer older households owned significant amounts of stock at that time, the impact of stock prices was also less broad. The growth of defined contribution pension plans has greatly increased the importance of stock holdings in households' retirement savings portfolios over the past decade, resulting in non-trivial exposure of more households to the stock-market in the more recent downturn. Additionally, as with the early 2000s stock market crash, the most recent crash was largely unanticipated, resulting in a cleaner quasi-experiment than the protracted stock run-up of the late 1990s. Indeed, recent papers have suggested that wealth losses in 2008 and 2009 may affect future retirement behavior [Gustman et al., 2009, Goda et al., 2011] but the authors have not yet placed much weight on such findings. Lending support to the findings of these recent studies, summary statistics from the Cognitive Economics Study show that more than forty percent of working respondents reported that their expected age of retirement had increased "as a result of the economic crisis," and most who reported a change reported an increase of two or more years.

Second, the signal-to-noise ratio in measures of wealth changes I use in this study may be larger than in prior studies. This is due to a combination of two factors: the large magnitudes of wealth losses experienced by a large proportion of households during this economic crisis, and the fact that surveys designed in the wake of the crisis have yielded data from direct questions about wealth losses for most types of assets affected by the crisis. The former means that the true wealth changes tend to be farther from zero than has been the case over other periods. The latter leads me to believe that my measures of wealth change do not suffer from the same compounded error as true longitudinal data, and are therefore likely to be subject to less attenuation bias than purely longitudinal wealth measures.

Third, I argue in this paper that it is important to account for fixed costs and non-linearities in examining the impact of wealth shocks on retirement timing. Previous studies using reduced-form regression techniques have ignored these issues. My econometric specification takes these into account. I use a corner solution model to explicitly allow for non-zero adjustments by individuals whose fixed-costs of adjustment are not outweighed by the potential gains from adjustment, while also estimating the effect of wealth losses on the size of adjustments for individuals who do report changes in retirement timing. Additionally, I employ a quadratic term in my regression analyses to pick up potential non-linearities in the underlying optimization model. In my analysis, I also discuss the possibility that inclusion of the quadratic term may strengthen the estimated linear effect of wealth losses by reducing the impact of observations with implausibly large wealth losses.

3 Theoretical considerations

The intertemporal budget constraint in a standard life-cycle model requires that the expected present discounted value of consumption be less than or equal to the present value of assets plus the present discounted value of future income flows. In the simplest case,

$$\sum_{s=\tau}^{T} \frac{C_s}{(1+r)^s} = A_\tau + \sum_{s=\tau}^{T} \frac{Y_s}{(1+r)^s}$$

where C_s is consumption at time s, Y_s is income at time s, r is the interest rate and A_τ is assets at the time of optimization (τ). A loss of asset holdings must be accompanied by a reduction in consumption or an increase in income in order for this equality to hold. Assuming that individuals will perfectly smooth consumption, let "sustainable consumption," $SC = C_s$, $s \in [\tau, T]$, be the smoothed consumption level attainable in all periods from the reference period τ until death at time T, given A_τ, assets held at the time of optimization, and planned income path Y.

Figure 1 uses Modigliani's canonical graph to illustrate the impact of an asset loss on consumption, holding labor supply constant. Y represents labor earnings, A is accumulated wealth, and SC is the implied "sustainable consumption" that can be supported using accumulated wealth and planned future labor earnings. For a given income path, a negative asset shock necessarily translates to lower sustainable consumption. The reduction in sustainable consumption is shown by the drop in SC from the upper, dotted SC path to the lower SC path.

Now, briefly consider the implications of allowing labor supply to be a choice variable. If one assumes for simplicity, as I do in this paper, that an individual's real wage is a known constant, then labor income is only a function of the quantity of labor supplied. In theory, individuals may adjust their labor supply along the extensive margin (i.e., whether to work) or the intensive margin (i.e., how much to work). In fact, conditional on working, hours worked may be adjustable only to the extent that workers may choose between employers offering wage packages with different hours (Blundell and MaCurdy, 1999, page 1588). This means that adjustment along the intensive margin of labor supply may entail significant search costs. Empirically, according to Heckman (1993, page 118) "... the strongest empirical effects of wages and nonlabor income on labor supply are to be found at the extensive margin— at the margin of entry and exit— where the elasticities are definitely not zero." For these reasons, this paper focuses exclusively on the extensive margin of labor supply called "retirement." For tractability, I define retirement as an irreversible, complete cessation of work for pay.

While development of a dynamic structural model is beyond the scope of this paper and unnecessary to my empirical analysis, consideration of such a model is useful for deriving intuition about the expected impact of the economic crisis. In Appendix 9, I present a simplified version of a life-cycle model of consumption and labor supply developed by Kimball and Shapiro [2008, 2003]. Figure 2 presents a graphical representation of the optimal retirement choice problem, based on a life-cycle model of consumption and labor supply like that in Kimball and Shapiro[2008, 2003], that might be underlying Modigliani's static model. In this figure, the upward-sloping curve represents the marginal disutility of work, per dollar earned. The disutility of work function incorporates the costs of working associated with distaste for work, effort costs of work, and/or fixed costs of going to work. The marginal disutility of work could be increasing with age due to expectations that health will decline with age, social norms that one "should" be retired by a particular age, spousal labor force status, or other factors. The downward-sloping curve in Figure 2 represents the marginal value of wealth, assuming optimal choice of consumption path for each possible retirement age along the horizontal axis. An individual will plan to retire when the marginal utility cost of continued work is expected to exceed the marginal utility gain from the consumption funded by continued work.

After an unforeseen loss of wealth, however, the marginal value of wealth would be expected to shift upward, as in Figure 3. If retirement must take place at the originally-planned age, R_0, the wealth shock necessitates a lower level of consumption over the remainder of life. If, however, retirement is a choice variable, it can be seen that consumption could actually remain unchanged by choosing retirement age $R_{\overline{sc}}$. The value $R_{\overline{sc}}$ represents the "constant sustainable consumption retirement age," or the retirement age that would be necessary to maintain the pre-shock level of consumption. The optimal post-shock retirement age, however, is at R^*, the new intersection between the marginal disutility of work per dollar earned and the marginal utility value of wealth.

Figure 4 illustrates the result of optimal retirement choice after an asset loss within the simple Modigliani framework. Retirement at the originally-planned retirement age, R_0, requires that consumption be adjusted to absorb the entire loss of assets. By contrast, retirement at $R_{\overline{sc}}$ requires only that retirement age be adjusted, leaving consumption unchanged. The new optimal retirement age, R^*, will actually lie somewhere between R_0 and $R_{\overline{sc}}$, depending on the relative slopes in the underlying optimization problem.

4 Empirical framework

In my analyses, I regress measures of the change in expected retirement timing, $\triangle retirement\ timing$, on the change in retirement age that would be necessary if consumption were kept constant at pre-crash levels, $R_{\overline{sc}} - R_0$. This implies the base specification

$$\triangle retirement\ timing_j = \alpha + \beta_1 \left(R_{\overline{sc}} - R_0\right)_j + Z_j'\gamma + \varepsilon_j \qquad (1)$$

where R_0 is "pre-crash" retirement age, and $R_{\overline{sc}} - R_0$ is the additional number of years individual j would need to work to maintain his or her pre-crash consumption path, or the "constant sustainable consumption retirement age." The dependent variable, $\triangle retirement\ timing$, differs by dataset, and is discussed in more detail in Section 5. In some specifications, a vector of variables Z is also included to capture effects related to preferences and expectations, allowing the β coefficients to reflect adjustments due to the wealth shock. Additionally, some specifications include interactions of variables in Z with the wealth losses to explore observed heterogeneity in the relationship between wealth losses and retirement timing with respect to these variables.

If individuals adjust to wealth shocks solely along the consumption margin, I expect estimates of β_1 to be zero. If individuals adjust along the retirement age margin, I expect estimates of β_1 to be positive. In the extreme case in which individuals adjust solely along the retirement age margin, the dependent variable is the change in planned retirement age (in years), and there is no measurement error, the expected coefficient would be one.

From the earlier discussion of the optimal retirement choice problem it might seem that, holding all else constant, the change from the originally-planned retirement age, R_0, to post-shock optimal retirement age, R^*, will strictly increase as the size of the asset loss increases. There are, however, two main reasons this need not be true: discontinuities in the marginal value of wealth or marginal disutility of work curves, and fixed costs related to implementation of retirement age adjustments or re-optimization of retirement age. First, the marginal value of wealth curve need not be continuous. In particular, factors such as Medicare, Social Security and defined benefit pension rules may result in discontinuous jumps at threshold ages or levels of job tenure. Second, the marginal disutility of work curve may also contain discontinuous jumps at particular ages (for example, based on social norms that one "should" be retired by a particular age) or at ages when other events are expected to occur (for example, changes in spousal labor force status). Third, there may be fixed costs related to implementation of changes in retirement age, especially for those who were within a few months of retirement before the asset loss. These costs might include time spent revising Social Security or retirement-related paperwork, effort needed to train a different successor for one's job, or monetary costs related to maintaining one's primary residence for longer than expected (for example, losing a buyer for one's primary home or extra repair costs). Finally, there might be effort costs due to re-optimizing one's retirement age and consumption path, monetary costs due to hiring help to re-optimize, and emotional costs due to acknowledging the need to delay retirement. All of these factors would contribute to heaping at the "no change" corner solution.[1]

[1] Indeed, evidence of heaping, seen in my descriptive analysis (Section 6.1), suggests that fixed costs and

I primarily focus on results from Tobit regression specifications in this paper, since the Tobit model can provide consistent estimates of the relationship between wealth losses and observed changes in retirement age in the presence of heaping at a corner solution, at least for individuals who are not at the corner solution. The Tobit specification is:

$$\triangle retirement\ timing_j^* = \alpha + \beta_1 \left(R_{\overline{sc}} - R_0\right)_j + Z_j'\gamma + \varepsilon_j \tag{2}$$

$$\triangle retirement\ timing_j = max\left(0, \triangle retirement\ timing_j^*\right) \tag{3}$$

$$\varepsilon_j | \left(R_{\overline{sc}} - R_0\right)_j, Z_j \sim N(0, \sigma^2) \tag{4}$$

where the latent variable, $\triangle retirement\ timing^*$ can be thought of as the optimal change in retirement timing that would result in the absence of the fixed costs and discontinuities.

It is also important to consider that there may be non-linearities in the relationship between changes in retirement age and the relative magnitude of asset losses, even among individuals reporting non-zero retirement changes. The effect of a small perturbation in asset holdings on retirement timing may be expected to be well-approximated by a model that is linear in $R_{\overline{sc}} - R_0$. However, the magnitude of the losses seen in 2008 are likely, at least for some people, to have had a large enough impact on the marginal utility of wealth that the effect is not well-approximated by this model. Indeed, given that the marginal disutility of work might be increasing at an increasing rate with age, the marginal change in optimal retirement age may actually decline as the wealth shock increases. Thus, a squared term is also introduced.

$$\triangle retirement\ timing_j = \alpha + \beta_1 \left(R_{\overline{sc}} - R_0\right)_j + \beta_2 \left(R_{\overline{sc}} - R_0\right)_j^2 + Z_j'\gamma + \varepsilon_j \tag{5}$$

To implement these analyses, I need measures of planned retirement age as of mid-2008 (R_0), the "constant sustainable consumption retirement age" ($R_{\overline{sc}}$), measures of changes in retirement timing ($\triangle retirement\ timing$) and variables found in the Z vector. The next section describes data and measurement considerations related to these variables.

other discontinuities are important in predicting the adjustment of retirement plans to wealth losses in 2008 and 2009.

5 Data

The Cognitive Economics Project[2] (CogEcon) and the Health and Retirement Study[3] (HRS) are nationally-representative studies of older Americans, both of which fielded surveys in 2008 before the crash. After the stock market crash, the researchers with these studies developed "post-crash" surveys and fielded these by mid-2009. By design, the timing and content of these surveys provide excellent data to analyze the impact of a wealth shock on older Americans' retirement plans. This paper uses data from the baseline CogEcon 2008 survey, the 2009 "Post-Crash" CogEcon survey, the HRS 2006 and 2008 Core interviews, and the HRS 2009 Internet Survey. These datasets contain detailed, longitudinal data about older Americans' preferences, expectations, financial situations and expected retirement timing, both before and after the stock market crash.

The CogEcon study has a smaller sample size than the HRS, but was designed to provide a direct measure of the dependent variable suggested by theory: the change in expected retirement age. Additionally, the CogEcon data provide detailed measures of changes in wealth and other impacts of the economic crisis, and access to restricted geographic data has enabled me to use county-level unemployment rates in my analyses. By contrast, the dependent variable available for the HRS analysis is the change in the "subjective probability" of full-time work after ages 62 and 65, and the measures of changes in wealth are less complete. However, the HRS offers information about Social Security, defined benefit pension wealth, and expected bequests that is not available in the CogEcon study, and therefore provides better measures of some aspects of wealth and other margins of adjustment to wealth losses. To take advantage of the strengths of both studies, I conduct analyses using both datasets.

5.1 Cognitive Economics

The CogEcon study has been designed by a group of economists from the Survey Research Center at the University of Michigan to explore the relationship between cognitive measures and a variety of economic variables, including measures of financial knowledge, wealth holdings and how financial decisions are made. The first CogEcon survey was fielded in 2008

[2]The Cognitive Economics Survey is supported by NIA program project 2P01AG026571, "Behavior on Surveys and in the Economy Using HRS," Robert J. Willis, PI. In addition to Willis, University of Michigan faculty Gwen Fisher, Miles Kimball, Matthew Shapiro, and Tyler Shumway and graduate students Brooke Helppie McFall and Joanne Hsu had roles in designing and fielding the CogEcon study.

[3]The HRS is sponsored by the National Institute on Aging (grant number NIA U01AG009740) and is conducted by the University of Michigan.

to 1222 eligible respondents to a partner study, CogUSA.[4] The final response rate for CogEcon 2008 was 80.8 percent, with 987 respondents having submitted completed surveys. The Post-Crash survey was fielded to 939 of these respondents in May and June 2009, and attained responses from 848 responses (90.2 percent response rate).

For analyses using CogEcon data, I start with data from the 848 CogEcon participants who responded to both CogEcon 2008 and the Post-Crash survey.[5] I combine the CogEcon data with demographic and cognitive measures from the CogUSA study. The final analysis sample uses data from the 320 respondents who were working at the time of the Post-Crash survey, reported planned retirement age as of July 2008 at least as large as their age in July 2008, provided earnings information in either the 2008 or Post-Crash survey, and provided some wealth data.[6]

Table 1 presents descriptive statistics about the sample. The sample is 52% female and 23% single, with an average age of 60.6 years at the time of the Post-Crash survey. The average education level of the sample is 14.9 years. Median annual earnings in 2008 were \$52,023,[7] and the median age at which respondents reported that they had planned, as of July 2008, to retire "completely" was 66 years. In some specifications, the number of observations is further reduced by nonresponse to additional variables included in the analyses.

The CogUSA sample is slightly more educated and wealthier, and slightly less representative of minorities than the general HRS population. To correct for this, regressions presented in the main paper use weights developed to make inference with CogUSA data more representative of the general population of older Americans.

[4]The CogUSA study, formerly NGCS+HRS, was started in 2006 by a cognitive psychologist, John J. McArdle, with the goal of conducting extensive cognitive tests and gathering rich demographic and health data on a nationally-representative sample of older Americans. The CogUSA Study is sponsored by the National Institute of Aging, grant number R37 AG007137, "Assessing and Improving Cognitive Measurements in the HRS," John J. McArdle, PI.

[5]The fielding timeline of the CogEcon and HRS surveys used in this paper are illustrated in Figure 6. More than ninety percent of CogEcon respondents completed their 2008 (pre-crash) questionnaires by September 1, 2008, while ninety-five percent of completed Post-Crash surveys were submitted by July 1, 2009.

[6]Non-responses to questions about the value of particular assets are coded as zeroes in the data used for my analysis. However, item non-response rates were extremely low. For example, only 1.96 percent of respondents in my sample who indicated that they had tax-advantaged retirement accounts did not give information about the value of these accounts. Because the values of many different types of assets were added together to create the measures of total wealth upon which the main independent variable of interest is based, the I believe the underestimation of wealth from this coding is minimal.

[7]For some respondents, earnings reported in the 2008 survey, from "last year" were used. However, all earnings are converted to 2009 dollars.

5.2 Health and Retirement Study

The second dataset used in this paper is from the HRS. The HRS has fielded "Core" interviews by telephone or in person in even years since 1992. Roughly every two years since 2003, some respondents with Internet access have also been asked to complete web-based surveys. The 2009 Internet survey provides "post-crash" data for HRS respondents.

In addition to 2009 Internet survey data, I use RAND HRS data (Version J),, 2008 Tracker data, the Cross-Wave Social Security Wealth File (Version 4.0), Imputations for Employer-Sponsored Pension Wealth from Current Jobs in 2004 (Version 1.0) and table data from Gustman et al. [2010a] for pension wealth.

To be included in my sample, respondents must have submitted a 2009 Internet Survey and have completed their HRS 2008 Core interview prior to September 1, 2008.[8] This date restriction ensures that baseline wealth and retirement expectations from 2008 were measured prior to the stock market crash and the other wealth losses that occurred from fall 2008 through spring 2009. Furthermore, respondents must have been assigned a non-zero 2008 Core interview sampling weight.[9] Respondents must also have been in the labor force (working, on leave or unemployed and looking for work)[10] and under the age of 65 at the time of the 2009 Internet survey. Unfortunately, because the Internet Survey is only fielded to HRS respondents who have identified themselves as internet users in the past, the sample may be less representative of the general population than the full HRS sample.[11]

The final sample size for the "under-62" analyses is 589, while the "under-65" sample size is 594. These respondents were in the labor force, answered some questions about wealth, and responded to the questions about work after age 62 (the under-62 sample) or age 65 (the under-65 sample) that are used to create the dependent variables used in my analyses. gives some summary statistics for the HRS sample. At 55% female and 22% single, the composition of the HRS samples are quite similar to the CogEcon sample. Respondents in the HRS samples are slightly less educated than the CogEcon sample at the median and have lower mean annual earnings, but do have quite comparable median earnings. They are

[8] Ninety percent of the 2008 Core Interviews took place prior to September 2008, so a relatively small number of observations were excluded due to late 2008 Core interview timing.

[9] More than ninety percent of zero sampling weights occur due to age ineligibility. The 2008 HRS Core interview weights were developed to reweight the HRS sample to mirror the population of Americans over age fifty in 2004, so respondents who were age 50 or younger in 2004 are assigned weights of zero.

[10] Regression results are qualitatively robust to exclusion of those who were temporarily laid off or on leave, or unemployed and looking for work at the time of the Internet Survey (approximately 30 observations, depending on the analysis).

[11] For example, Hsu, Fisher and Willis (2008) find that respondents to internet surveys tend to be younger and of higher cognitive ability, even after controlling for education level, than respondents to other modes of mixed-mode survey efforts. By contrast, the CogEcon survey was fielded in both mail and internet modes, allowing individuals who were not internet-users to respond to the survey.

also younger, on average, than the CogEcon sample, because they must have been under 62 or 65 at the time of the Internet Survey to have answered questions related to my dependent variables. It might also be noted that the planned retirement ages are younger; however, most of these values are imputed, and those for whom it is not imputed may be different from the average person of comparable age.[12]

5.3 Measurement

As discussed in Section 3, the estimation equations I use are linear and Tobit regressions of the form

$$\triangle retirement\, timing_j = \alpha + \beta_1 \left(R_{\overline{sc}} - R_0\right)_j + Z_j'\gamma + \varepsilon_j \tag{6}$$

and

$$\triangle retirement\, timing_j = \alpha + \beta_1 \left(R_{\overline{sc}} - R_0\right)_j + \beta_2 \left(R_{\overline{sc}} - R_0\right)_j^2 + Z_{ij}'\gamma + \varepsilon_j \tag{7}$$

where R_0 is the "pre-crash" retirement age work variable and $R_{\overline{sc}} - R_0$ is the additional number of years an individual would need to work to attain his or her pre-crash consumption path, where $R_{\overline{sc}}$ is the "constant sustainable consumption retirement age" for individual j. In some specifications, I also interact the Z variables with the $(R_{\overline{sc}} - R_0)$ terms to explore heterogeneity in the relationship between wealth and retirement changes.

5.3.1 Dependent variables

The dependent variable used in the analyses, $\triangle$*retirement timing*, differs by dataset. Because only two years have elapsed since the stock market crash, there has not yet been time to observe changes in actual retirement behavior. In both the CogEcon and HRS analyses, I use variables measuring expected changes in retirement timing.

In the CogEcon data, the dependent variable is $R_{09} - R_0$, the difference between the "post-crash" planned retirement age and the "pre-crash" planned retirement age. This variable is derived from a series of questions in the CogEcon Post-Crash Survey about retirement timing. First, respondents were asked for their current labor force status.[13] If they were not completely retired, they were next asked at what age they planned to retire completely, yielding R_{09}. Next, respondents were asked "As a result of the economic crisis, has the age

[12] See the HRS wealth measures section for detail on this.

[13] These categories are comparable to those standard in the HRS, and include: working, unemployed, disabled, homemaker, retired, etc. Respondents who selected "retired," were then asked if they were "completely retired."

at which you plan to retire completely changed since July 2008?" If they indicated a change, they were then asked "As of July 2008, at what age were you planning to retire completely?" If no change was reported, R_0 was set equal to R_{09}. If a change was reported, the July 2008 planned retirement age was used for R_0.

The dependent variable in the CogEcon analyses has a clear interpretation in the context of the theoretical framework discussed earlier. A strength of this question series is that it directly asks respondents to report the causal impact of the economic crisis on retirement age, and so might capture fewer changes in retirement age that are unrelated to the economic crisis, compared changes that might be measured by other surveys. Furthermore, because the labor supply questions were asked toward the beginning of the survey, before questions about the impact of the crash on their wealth holdings, the question order probably helped minimize priming bias in the answers to these questions.

For the HRS data, these variables are based on responses to the "probabilistic expectations" questions related to retirement timing in the 2009 Internet Survey,

> *Thinking about work in general and not just your present job, what do you think the chances are that you will be working full-time after you reach age 62?*

and

> *Thinking about work in general and not just your present job, what do you think the chances are that you will be working full-time after you reach age 65?*

Respondents answer these questions by giving a probability between 0% and 100%. Parallel questions were asked in the HRS 2008 Core interviews, as well. The dependent variables for the HRS analyses are ${}_{08}\triangle_{09}Pr(FT62)$, the change in reported "subjective probability of full-time work after age 62" as of the 2008 HRS Core interview and the 2009 Internet survey, and ${}_{08}\triangle_{09}Pr(FT65)$, the change in reported "subjective probability of full-time work after age 65" as of the 2008 HRS Core interview and the 2009 Internet survey.

The obvious benefit of using expectations data over observed behavior is that first differences specifications yield a much larger proportion of "changes" at any particular point in time, since observed retirement transitions are binary. Expectations data are particularly useful for studying reactions to shocks, because the effects of a shock on a broad population may be observed immediately, rather than only after many years have passed.

One might be concerned about using expectations data to draw conclusions about actual future behavior, because it is conceivable that expectations are not predictive of actual behavior. However, research by Manski [2004] suggests that probabilistic expectations are actually the measures of expectations that are called for by modern economic theory. While my theoretical framework does not explicitly model uncertainty, use of dependent variables based on probabilistic expectations may provide some insight into this issue. Additionally, studies by Hurd and McGarry [1995], Hurd [2009] have validated that probabilistic expectations data about life expectancy and retirement age are predictive of actual behavior. Several studies, many using the HRS, have validated the relationships between probabilistic expectations data and actual outcomes [Hurd and McGarry, 1995, Dominitz and Manski, 1997, McGarry, 2004, Dominitz and Manski, 2005, Hurd, 2009].

Another analysis, by Hurd [2009], compared population averages of full-time work expectations with actual outcomes, and concluded that the average expected probability of full-time work after age 62 was closely related to the actual probability of full-time work after age 62. Additionally, using linear probability model estimations on individual data from the HRS, Chan and Stevens [2004] have shown that subjective retirement expectations are strongly related to later employment status, even after controlling for age, health, marital status and education. Both Chan and Stevens (2004) and Hurd (2009) have found that, as individuals approach a question's reference age (62 or 65), the predictive value of their expectations grows.

Providing support for the validity of expected retirement age measures, Benitez-Silva and Dwyer [2005] have shown that expected retirement age in earlier waves of the HRS are extremely strong predictors of expected retirement age in later waves, and could not reject that the regression coefficient on previously reported retirement age is one, after controlling for selection and reporting errors. Thus, they could not reject that retirement expectations follow the rational expectations hypothesis. They also examined the role of new information, and concluded that models of perfect foresight are not rejected with respect to most changes in economic variables.

Using correlations and linear probability models with HRS Core interview data from the early-to-mid 2000's, I have also confirmed that reported subjective probabilities of full-time work after ages 62 and the expected age of full retirement from four years before reaching age 62 are strongly predictive of actual full-time work status after age 62. The correlation coefficients were 0.38 and 0.24, respectively. A 10 percentage point increase in the subjective probability of full-time work after age 62 was associated with a 4 percentage point increase in the probability that the individual was observed to be working full-time after age 62. Each additional expected year of work was associated with a 3 percentage point increase in the probability of full-time work after age 62. All coefficients on the expectations variables were highly statistically significant. Furthermore, the correlation between the two expectations measures was 0.51. Similar analyses of the relationship between actual full-time work status after age 65 and the subjective probability of full-time work after age 65 or expected retirement age, both observed approximately six years before reaching age 65, yielded comparable results.

In sum, I argue that my use of expected retirement age and the subjective probability of full-time work as proxies for actual future retirement behavior is valid. In fact, use of expectations may actually be preferable in a natural experiment context because such measures are more directly related to the dynamic programming problem individuals are thought to solve when planning for retirement. Using expectations data captures the immediate effect of a shock on the maximization problem with current expected values of future variables. By contrast, retrospective analyses of the effect of a shock on actual retirement outcomes many years down the line may be affected by unknown future realizations of variables that are correlated with but not caused by the initial shock, some of which may be unobserved. Standard estimation procedures using observed retirement outcomes would be likely to yield biased estimates of the impact of the shock on observed retirement outcomes.

5.3.2 Wealth measures

While variables related to expected retirement timing are directly observed in the data, it is necessary to calculate and annuitize measures of total wealth in order to derive the "constant sustainable consumption retirement age" variable ($R_{\overline{sc}}$).

In this paper, I define total wealth as the discounted sum of expected future household labor earnings, household financial wealth, defined contribution pension account holdings, defined benefit plan and combination plan wealth, Social Security wealth, and net equity in homes and other real estate. The counterfactual level of total wealth, referred to throughout the paper as "pre-crash" wealth, is the level of wealth that would have been held by the household in mid-2009 if the economic crisis had not happened.[14] Total wealth after the onset of the crisis is the level of wealth held by the household in mid-2009, holding planned retirement age constant at its 2008 level. All account holdings and income streams used to calculate pre- and post-crash total wealth are in pre-tax 2009 dollars.[15]

After calculating the total wealth measures, I divide each observation of total pre-crash wealth by an individual-specific annuity price to get the pre-crash level of annual annuity income— that is, the level of sustainable consumption— that could be purchased with the

[14]Specifically, the counterfactual level of wealth is calculated as if accumulated financial wealth, pension and Social Security wealth are at their pre-crash (2008) levels, while future earnings are those expected from 2009 onward.

[15]Because each individual's Social Security income, defined benefit pension income, and distributions from non-Roth tax-advantaged retirement accounts are likely to be taxed at difficult-to-predict marginal income tax rates, I have up-weighted all other assets. These other assets are likely to be taxed at the capital gains rate (if at all). Specifically, I multiplied the value of each of these assets by $1/(1-\tau)$, where τ is set at 0.15 (the current capital gains tax rate for most assets) before summing all assets to calculate total wealth.

present discounted value of pre-crash wealth in 2009. Similarly, I use post-crash wealth to calculate the post-crash sustainable consumption level. These estimates of sustainable consumption are then used to calculate the primary independent variable of interest in this study, $R_{\overline{sc}} - R_0$, the additional number of years individuals would have to work to make up their losses completely. The details of this process are described below.

CogEcon total wealth measures I use data from both the 2008 CogEcon survey and the CogEcon Post-Crash survey to calculate pre- and post-crash household financial wealth. The post-crash data contain information about the levels of wealth held in tax-advantaged retirement accounts (for example, 401(k) plans and IRAs), and how much these had changed since July 2008. The surveys similarly solicited levels and changes of holdings in checking, savings, money market accounts, certificates of deposit, Treasury bills, cash, credit card debt, and stocks or stock mutual funds held outside of tax-advantaged retirement accounts. For financial assets for which respondents reported Post-Crash levels[16] and percent changes, I calculate the July 2008 values using the 2009 levels and changes as:

$$value_{08} = \frac{value_{09}}{(1 + (percent\ change/100))}$$

while, in cases where I have data on levels and the dollar value of the change,[17] $value_{08}$ is calculated as the sum of $value_{09}$ and the reported change since July 2008. The value of bonds holdings was only asked in 2008, so this value is assumed constant from 2008 to 2009.

The pre-and post-crash gross value and changes in the value of real estate holdings are also available in the post-crash data. Using the reported mortgage balances and the reported changes in these balances since July 2008, I also calculate pre- and post-crash net real estate holdings. Values of farm and business equity were only asked in 2008, so these values are assumed constant from 2008 to 2009.

For earnings in 2009, I use the average of inflation-adjusted 2007 and 2008 earnings if the respondent gave dollar values for both, or if the respondent gave "range card" answers for both. If the respondent reported a value for either year, but gave a "range card" answer

[16] For questions asking for the dollar amounts of earnings, assets or debts, the CogEcon surveys offered the option to give either a value or a "range letter" answer. Range letters are from a "range card," which allows respondents to choose from a set of dollar ranges, each represented by a letter. Respondent answering using a "range card" are assigned a value corresponding to the midpoint of the range. For example, respondents who indicate that they hold tax-advantaged retirement assets in the range "$100,000 to $250,000" are assigned a value of "$175,000.50." For the highest range, which is open-ended, the assigned value is 1.4 times the lower bound. Therefore, respondents indicating that they hold "More than $1,000,000" in tax-advantaged retirement assets are assigned a value of "$1,400,000."

[17] Except in the case of primary home value, questions asking about changes since 2008 gave respondents the option to answer with a percent or a dollar amount. With respect to changes in the value of their primary homes, respondents were asked by what percent the value of homes in their neighborhoods had changed.

or no answer for the other year, that year's earnings was used. For respondents who did not give a specific value in either year, I use the mid-point of 2008 earnings if the respondent gave a range answer for that year, and 2007 earnings if the respondent gave a range for 2007 earnings but gave neither a range nor a value for 2008 earnings.

Especially for individuals who are far from retirement, future earnings are an important component of total wealth. To calculate the expected present discounted value of future household earnings, it is necessary to assume a path for each respondent's future earnings over his or her life. Ideally, I would know how much paid work each respondent would be doing in each future year, and the earnings he or she would receive for that work. Furthermore, because the future is uncertain, I would also need to account for the probabilities that a person would become unemployed at a particular time, the amount of time that person would take to find a new job, the probability of re-employment after "complete retirement," and so on. Given that this is a study of older adults, and that studies of the time-path of labor earnings tend to show that earnings peak around 30 years of experience and may begin a slow decline thereafter, it is a reasonable simplification to assume constant real earnings from 2009 until retirement. That is, I assume that nominal earnings will grow at the rate of inflation, π. The expected present discounted value of earnings for individual j is therefore calculated:

$$EPDV(earnings)_j = \sum_{s=\tau}^{R_0} \left((1 - UE\ rate_j) \times \frac{earn}{(1+r)^s} \right) \tag{8}$$

where R_0 is pre-crash planned retirement age; the real interest rate, r, is 0.03;[18] and s takes on values from the individual's 2009 age to their pre-crash planned retirement age. In calculations of pre-crash wealth, $UE\ rate_j$ is the unemployment rate in May 2008 in the county of individual j's residence, taken from the U.S. Bureau of Labor Statistics Local Area Unemployment Statistics (LAUS) series and matched to the data by county-level FIPS code.[19] In calculations of post-crash wealth, Equation 8 uses county unemployment rates from May 2009.[20] The expected present discounted value of future household earnings is the

[18] Following Gustman et al. (2010), who use long-term projections from the Social Security Administration for future nominal interest and inflation rates. In their study, the nominal interest rate, i, is 5.8 percent; the inflation rate, π, is 2.8 percent.

[19] County-level unemployment statistics are not seasonally-adjusted, so I use May 2008 unemployment for pre-crash wealth calculations, and May 2009 unemployment data for post-crash calculations to net out the seasonal component of unemployment.

[20] One might worry about this simple way of including employment probabilities, since it doesn't account for the possibility that the labor market will get better, nor does it account for the fact that individual unemployment is serially correlated. However, robustness checks, in which analyses were run without using employment probabilities in calculating the expected present discounted value of earnings, show that the qualitative results are robust to inclusion or exclusion of these rates in calculating the present discounted value

sum of respondents' expected present discounted value of earnings and the expected present discounted value of their significant others' earnings, where relevant.

The CogEcon study does not contain information about Social Security wealth, a major component of older Americans' wealth. I estimate household Social Security wealth using the estimated present discounted value of Social Security benefits from the Cross-Wave Prospective Social Security Wealth Measures of Pre-Retirees (Version 4.0) [Kapinos et al., 2011]. These wealth measures are based on data provided by the Social Security Administration through 2004, and incorporate projected future earnings based on a weighted average of past earnings if the respondent had not yet reached normal retirement age by 2004. Assuming that respondents will claim benefits beginning at their normal retirement age, I estimate Social Security wealth for the CogEcon respondents using measures of individual Social Security wealth from HRS respondents of similar age in 2004 to the CogEcon respondents in 2009. Specifically, I assign CogEcon respondents the mean value of Social Security wealth from HRS respondents with the same age group by sex by occupation group. For coupled CogEcon respondents for whom occupational data are available for their spouses or partners, I estimate spouse or partner Social Security wealth similarly. I then sum the Social Security wealth estimates for both members of the household to obtain household Social Security wealth. In cases where a spouse's or partner's occupation or age are unknown, I assign the CogEcon respondent the cell mean of household Social Security wealth from similar HRS respondents.

The CogEcon data also do not include much information about defined benefit pension wealth. For those who are not yet retired, the data only contain an indicator variable that is equal to one if either the respondent or the spouse/partner has a defined benefit pension. Therefore, I estimate defined benefit pension wealth for the CogEcon respondents based on defined benefit pension wealth information in the HRS dataset, Imputations for Pension-Related Variables (Final, Version 1.0) [Health and Retirement Study, 2009]. Appendix 10 details the estimation procedure.

In sum, 2008 wealth for each individual j is measured as

$$w08_j = FW_{08,j} + NE_{08,j} + SS_{08,j} + DB_{08,j} + EPDV\ earnings_{08,j}$$

where FW_{08} is financial wealth in 2008, and includes tax-advantaged retirement accounts, checking, savings, money market accounts, certificates of deposit, Treasury bills, cash, credit card debt, stocks or stock mutual funds held outside of tax-advantaged retirement accounts and bonds. NE_{08} is net equity in real estate, businesses and farms in 2008,

of earnings. On net, I have chosen to present the results that do use the local unemployment information, since it seems important to account for the uncertainty of future income flows.

$EPDV\ earnings_{08,j}$ is the sum of the respondent's and his or her significant other's present discounted values of future earnings from 2009 until the age of retirement that was expected as of July 2008, SS_{08} is estimated Social Security wealth, and DB_{08} is estimated defined benefit pension wealth. Similarly, wealth in 2009 is measured as

$$w09_j = FW_{09,j} + NE_{09,j} + SS_{08,j} + DB_{08,j} + EPDV\ earnings_{08',j}$$

where financial wealth and net equity in real estate, businesses and farms reflect the post-crash values of these assets. Social Security and defined benefit pension wealth are assumed unchanged. The expected present discounted value of earnings is unchanged except that the county-level unemployment measure reflects May 2009 levels.

Measurement error in these wealth calculations, particularly due to imputation of Social Security and defined benefit wealth, is likely non-trivial in magnitude. However, some of this error is likely to be of second order importance because my independent variables of interest are based on *changes* in wealth, as opposed to *levels* of wealth. Specifically, the components of wealth that are most likely to be error-ridden, Social Security and defined benefit pensions, are probably quite constant, so error in these may only slightly affect the independent variable of interest, $R_{\overline{sc}} - R_0$. Additionally, by relying primarily on retrospective accounts of wealth losses from the Post-Crash survey, I believe that my change measures are subject to less measurement error than measures based on true panel data.[21] The time to planned retirement is also held constant in calculating the expected present discounted value of earnings measures for both my pre- and post-crash wealth measures. This is by design, since I later compare the reported changes in planned retirement age to the amount by which labor supply would have to increase to allow respondents to continue consuming on their pre-crash consumption path.

HRS Wealth Where possible, the HRS wealth measures are calculated in the same way as the CogEcon measures. As in the CogEcon wealth calculations, all wealth measures are in pre-tax, 2009 dollars and, following Gustman et al. (2010), income streams are converted to present discounted values using a real interest rate of 3 percent.

In the HRS data, some measures of financial wealth, including wealth held in checking, savings and money market accounts, certificates of deposit, government savings bonds or Treasury bills, other government bonds, and debts like credit card balances or other loans (subtracted), are only available in 2008 Core data. Because these components of wealth were

[21] Analyses by members of the CogEcon study team have shown that, while the distributions of wave to wave wealth changes look quite similar to wealth changes based on the retrospective accounts, the retrospective changes have lower variance and include fewer highly implausible or nonsensical changes.

not asked about in the 2009 Internet Survey, I assume that these are constant from 2008 through 2009. This assumption seems reasonable because returns to these types of assets are likely to have been quite stable relative to stock and real estate assets.

The 2009 Internet Survey did gather information about the value of IRAs and Keogh accounts, 401(k) and other employer-sponsored retirement saving plans, trusts, other mutual funds, and other stock holdings. This is important, because these types of assets are likely to include stock holdings, and were therefore subject to significant change between late 2008 and mid-2009.

As in the CogEcon data, I use the 2009 Internet Survey data to impute the levels of retirement assets, trusts, mutual funds, and other stock assets, as well as primary home equity that households held as of August 2008. In particular, the Internet Survey asks for the 2009 levels of these asset holdings and the percent change since September 2008.[22] Using this information, I calculate the September 2008 value of these assets as:

$$value_{08} = \frac{value_{09}}{(1 + (percent\ change/100))}$$

Thus, financial wealth in 2008 is calculated as the sum of wealth held in checking, savings and money market accounts, certificates of deposit, government savings bonds or Treasury bills, other government bonds, minus debts like credit card balances or other loans, plus the calculated 2008 values of IRAs and Keogh accounts, 401(k) and other employer-sponsored retirement saving plans, trusts, other mutual funds, and other stock holdings. Financial wealth in 2009 is calculated as the sum of wealth held in checking, savings and money market accounts, certificates of deposit, government savings bonds or Treasury bills, other government bonds, minus debts like credit card balances or other loans in 2008, plus the reported 2009 values of IRAs and Keogh accounts, 401(k) and other employer-sponsored retirement saving plans, trusts, other mutual funds, and other stock holdings.

The 2009 Internet Survey also contain information about the value of respondents' primary homes, as well as changes in their value. I use this information to construct the 2008 and 2009 values of primary home using the same method as was used for financial assets. Using mortgage balance information from 2009 and the 2008 Core interview, I then calculate primary home equity at each time point. A disadvantage of the 2009 Internet Survey data, relative to the CogEcon data, is that information about net real estate equity other than the primary home was not asked, and so must be imputed. For 2008 second home and other real

[22] As in the CogEcon data, respondents who didn't know or didn't want to report an exact value or percent change, but who did indicate a range, are assigned the midpoint of this range. For open-ended range responses (for example, "More than $1,000,000"), the bottom of the range is multiplied by 1.4 to get the imputed value.

estate holdings, I am able to use the net values for second homes and other real estate from the 2008 Core interview. For 2009, I use the maximum of an estimated net value in 2009 and \$0 for each,[23] where I estimate the value of real estate assets in 2009 using a Census region-specific change factor based on Case-Shiller index data and the net equity in these assets in 2008.

To get the Census region-specific change factor, I sum the housing stock for the 20 Case-Shiller statistical areas (SAs) by Census region (northeast, midwest, south, and west). Once I have the total housing stock represented by the Case-Shiller index in each Census region k, I calculate the relative weight of each statistical area l within its corresponding Census region in terms of housing stock using the equation $weight_{lk} = \frac{housing\ stock\ in\ SA_l}{housing\ stock\ in\ region_k}$. Next, I multiply the summer 2008 to summer 2009 change for each Case-Shiller SA l, $\%\triangle\, housing_l$, by the corresponding weight, where the index value for each summer is the average from June, July and August of that year. Lastly, I sum this weighted percent change in real estate prices across statistical areas within each region to get the percent change in home values within each region. That is, the region-specific change factor is:

$$\%\triangle housing_k = \sum_l (weight_{lk} \times \%\triangle housing_l)$$

I calculate the estimated net value in 2009 by multiplying a Census region-specific change factor by the total value of the home in 2008, and then subtracting the balance of any mortgages or loans using the property as collateral. Thus, the 2009 net value of first and second homes for respondent j in Census region k are calculated as:

$$net\ home09_j = max\{0,\ (gross\ home08_j \times (1 + \%\triangle housing_k)) - home\ debt08_j\}$$

For other real estate, I estimate the net value in 2009 by multiplying a Census region-specific change factor by the net value of the asset in 2008.[24]

Pension wealth estimates for defined-benefit and combination plans are the maximum of estimates from table data from Gustman et al. [2010a] and regression-based estimates from the *Imputations for Pension-Related Variables* (Final, Version 1.0) for individuals who indicated that they expected to receive defined-benefit or combination plan benefits in the future. The table data from Gustman, Steinmeier and Tabatabai (2010) are household-level

[23] This is reasonable if one assumes that respondents will strategically default on any mortgage if they want to be rid of the property and they have negative equity.

[24] This is likely to overestimate wealth from other real estate in 2009 in cases where a mortgage balance exists. However, it can be difficult to qualify for mortgages on additional real estate, and few individuals have such assets, so the impact of this issue is likely small.

estimates based on all defined benefit and combination plan pension wealth accumulated through the HRS 2006 Core interview wave. These pension data incorporate pensions from current jobs for those working at the time of the 2006 interview, last jobs for those who had changed jobs since their last interview, and all previous jobs for which pensions had been reported. The real value of defined benefit and combination plan pension wealth is assumed to have been constant since 2006.[25] Because the table data are missing for many respondents who stated in the 2008 Core interview that they expected to receive defined-benefit or combination plan benefits, I also create regression-based estimates of 2008 defined-benefit and combination plan wealth. For individuals with values from both the table data and an estimate, I use the maximum of the two estimates.

For Social Security wealth, I created regression-based estimates using the present discounted value of Social Security benefits from the Cross-Wave Prospective Social Security Wealth Measures of Pre-Retirees (Version 4.0). It was necessary to estimate Social Security wealth, rather than using the 2004 estimates directly, to account for growth in earnings and work tenure that accumulated between 2004 and 2008.

Lastly, it is important to consider future labor earnings, or human wealth, as a component of household wealth. Because I do not currently have access to the restricted geographic information about HRS respondents, I cannot use county-level unemployment rates, as I did in the CogEcon section. Instead, the expected present discounted value of earnings for individual j is calculated:

$$EPDV(earnings)_j = \sum_{s=\tau}^{R_0} \left((1 - UE\ rate_j) \times \frac{earn}{(1+r)^s} \right) \tag{9}$$

where, again, R_0 is expected age of retirement; the real interest rate, r, is 0.03;[26] and s takes on values from the individual's age in 2009 through their pre-crash planned retirement age. In the HRS data, however, $UE\ rate_j$ is the unemployment rate in May 2008 (for pre-crash $EPDV(earnings)$) or May 2009 (for post-crash $EPDV(earnings)$) in the Census division of individual j's residence, from the U.S. Bureau of Labor Statistics' Current Population Survey. It is also important to note that the expected age of retirement, R, is not asked directly of all HRS respondents. To avoid losing a majority of the size of my HRS analysis sample, I impute this age for respondents who did not answer this question by combining

[25] As soon as estimates incorporating data from the 2008 Core interview are available, I will substitute these into my analyses for the 2006 data. Using the 2006 data likely results in a downward bias of total pension wealth, since growth above the rate of inflation is likely to have occurred between 2006 and 2008.

[26] Following Gustman et al. (2010), who use long-term projections from the Social Security Administration for future nominal interest and inflation rates. In their study, the nominal interest rate, i, is 5.8 percent; the inflation rate, π, is 2.8 percent.

information from several variables. The imputation of this variable is described in detail in Appendix 11. Overall, I impute or have an actual retirement age for 99.5 percent of the 1563 working respondents in the HRS Internet Survey sample who were aged 64 or younger at the time of the 2008 Core interview, and who completed the 2009 HRS Internet Survey. The expected present discounted value of future household earnings is the sum of respondents' expected present discounted value of earnings and the expected present discounted value of their significant others' earnings, where relevant.

In sum, in the HRS sample, total wealth for both 2008 and 2009 are calculated as the sum of total financial wealth, real estate equity, defined benefit pension wealth, Social Security wealth, and the expected present discounted value of future household earnings. Both the CogEcon and the HRS wealth measures aggregate holdings in a nearly-identical set of asset types, although the way particular asset holdings are calculated does differ slightly.

Sustainable consumption Under certain conditions, introducing an annuity market is equivalent to removing uncertainty about life expectancy from the lifetime resource allocation problem (Yaari, 1965). This observation provides a convenient framework for quantifying the impact of a wealth shock in the presence of uncertain life expectancy. Once I have calculated total wealth as described above, I divide households' total pre- and post-crash wealth measures by an individual-specific annuity price to get estimates of "sustainable consumption" available to each household before and after the crash. Because it seems reasonable to assume that individuals plan for the lifetime consumption of their spouses and partners, as well as themselves, I calculate the price of an annuity that will pay:

- Households with single individuals \$1 per year, in 2009 dollars, until death.
- Households with coupled individuals \$1 per year in 2009 dollars until the death of the first member of the couple, after which \$0.67 per year will be paid until the death of the remaining member of the couple.[27]

The equation used to calculate each individual's annuity price, a_j is:

$$a_j = (1+L)\sum_{s=1}^{\infty}\left(\frac{P1P2_{j,s}}{(1+r)^s} + 0.67\frac{(P1_{j,s})\,(1-P2_{j,s})}{(1+r)^s} + 0.67\frac{(P2_{j,s})\,(1-P1_{js})}{(1+r)^s}\right)$$

where the real interest rate is set at 3 percent. The load factor L, set to 18 percent, was backed out of estimates by Mitchell et al. [1999] for average annuity payouts per dollar

[27] Research by Shapiro [2009], using the HRS, has shown that consumption drops by about a third upon the death of one spouse. At least initially, this does not appear to be due to resource constraints, but to an actual decline in costs. Hurd and Rohwedder [2010a] have also used this figure in estimating lifetime consumption paths.

premium. $P1_{j,s}$ is the probability that respondent j will be alive in s years, $P2_{j,s}$ is the probability that respondent j's spouse or partner will be alive in s years, and $P1P2_{j,s}$ is the probability that both members of the couple are still alive in s years. All survival probabilities are age- and sex-specific, and are derived from the Social Security Administration's Period Life Tables (Social Security Administration, 2006).

Change in retirement timing needed to make up wealth losses Changes in pre- and post-crash sustainable consumption can certainly help illustrate the magnitude of the effect of the crash. However, the theoretical considerations discussed earlier imply that a particularly interesting measure is how much longer individuals would have to work in order to attain the sustainable consumption levels they could have maintained if the crash had not happened. To calculate this number, I first define $R_{\overline{sc}}$ as the age until which respondents would need to work to attain the sustainable consumption paths they would have maintained given pre-crash wealth levels. $R_{\overline{sc}}$ solves the equation:

$$\sum_{s=\tau}^{R_{\overline{sc}}} \left((1 - UE\ rate_{j,09}) \times \frac{earn}{(1+r)^s} \right) = \sum_{s=\tau}^{R_0} \left((1 - UE\ rate_{j,08}) \times \frac{earn}{(1+r)^s} \right) - (w09 - w08)$$

where τ is the respondent's age in 2009, $r = 0.03$ is the real interest rate, $UE\ rate_{j,08}$ and $UE\ rate_{j,09}$ are individual j's county-specific (in CogEcon) or Census-division specific (in HRS) unemployment rates in May 2008 and May 2009, respectively, and $(w09 - w08)$ is the change in total wealth from July 2008 until May 2009, respectively. Essentially, this equation says that the present discounted value of earnings from working to $R_{\overline{sc}}$ must equal the present discounted value of earnings from working to R_0, plus the amount of wealth lost during the crash. Then, the extra number of years an individual needs to work to attain her pre-crash sustainable consumption level, denoted $R_{\overline{sc}} - R_0$, is simply the difference between $R_{\overline{sc}}$ and the originally-planned retirement age, R_0.

6 Results

6.1 Descriptive analysis

To provide a sense of the material impact of the shock on sustainable consumption levels, Table 2 displays the unweighted mean, 25th percentile, median, and 75th percentile of sustainable consumption for my CogEcon sample and the two HRS samples.[28] The medians look reasonable in magnitude, given that median household income in the United States in 2009 was around $52,000 in 2009 (U.S. Census Bureau). It can be seen that the post-crash distribution of sustainable consumption is generally lower than the pre-crash distribution, indicating a reduced sustainable standard of living, holding labor supply constant. The post-crash inter-quartile ranges have dropped by ten percent in the HRS data and eighteen percent in the CogEcon data, implying some reduction in inequality.

Table 3 illustrates that median losses in sustainable consumption are quite comparable between samples, at just under five percent for all three. The mean loss observed in the CogEcon sample is 8.7 percent, whereas the HRS samples show mean losses of about 6.7 percent.[29] These losses are not staggering, in that most people experiencing such losses are not in danger of falling into poverty as a result. However, a sustained reduction of "just" five percent in material quality of life for more than half of the individuals is not trivial, and a quarter of individuals in each sample would be facing losses of more than 11 percent of their consumption *for the rest of their lives*, all else equal.

Rather than passively accept a reduction in standard of living for the rest of one's life, some people may prefer to delay retirement. Indeed, out of the CogEcon sample, 128 respondents, or 40 percent of the sample, reported that their planned retirement age had increased by at least one year, while only five respondents (1.6 percent) reported a decrease in planned retirement age. Figure 7 displays the reported changes in retirement age since July 2008. The mean change reported by all respondents in this sample was 1.6 years, with median of 0 years and a change of 3 years at the 75th percentile. In Figure 8, I have plotted the cumulative distribution of expected retirement status over time for the CogEcon sample, with age on the horizontal axis. Of note here is that the entire distribution of expected retirement ages has shifted to the right. Whereas half of respondents expected to retire by age 65 as of 2008, the age at which half of respondents expected to retire was 66 in 2009.

The 2009 HRS Internet Survey data do not contain expected retirement age, but ask

[28] The mean estimated sustainable consumption levels from CogEcon are higher than those estimated using the HRS data. This is partly due to the fact that the retirement ages I imputed for use in calculating the present discounted value of earnings in the HRS are, on average, almost three years lower than those reported by the CogEcon respondents. The present discounted value of earnings calculated in the HRS are probably much too low for individuals planning to work much past 66, since the largest imputed HRS retirement age was 70. By contrast, the largest age reported by HRS respondents who did give retirement age was age 80, and 2.5 percent of CogEcon respondents reported expected retirement ages of 90 or older before the crash.

[29] This is partly due to changes in home value. The CogEcon respondents reported mean losses in the net value of their primary homes of 9.2 percent, around double the mean losses of just 4.4 percent reported in the HRS sample. Additionally, the CogEcon data contain respondent reports of losses in second home and other real estate wealth that were, at 17 percent of gross value, slightly higher than the HRS real estate loss estimates based on the Case-Shiller index, which averaged 13.3 percent over the nation as a whole.

for the subjective probability of full-time work after age 62 and 65. Table 5 shows that the mean subjective probability of full-time work after age 62 reported by the HRS respondents increased by 8.7 percentage points over the two years from 2006 to 2008, but just 3.5 percentage points over the one year between the 2008 Core interview and the 2009 Internet survey. The median changes in subjective probability of full-time work after age 62 ($\triangle Pr(FT62)$) over both periods were zero. At the 75th percentile, however, the changes in $\triangle Pr(FT62)$ were 20 percentage points between 2006 and 2008 (2 years), and 19.5 percentage points between the 2008 Core interview and the 2009 Internet survey (just 1 year). While the lower end and middle of the distribution of $\triangle Pr(FT62)$ appear to have followed a similar trend before and after the crash, the upper end of the distribution indicates that expectations of later work may have increased more rapidly after the crash.

Similar examination of changes in the probability of full-time work after age 65 show an even stronger trend toward delay of retirement. The mean change in the subjective probability of full-time work after age 65 was 8.1 percentage points from 2008 to 2009, compared with just 6.5 percentage points from 2006 to 2008. The median increase was 2 percentage points between 2008 and 2009, compared with a zero percentage point change from 2006 to 2008. At the 75th percentile, as well, the change between 2008 and 2009 (25 percentage points) greatly outpaced that between 2006 and 2008 (20 percentage points).

Next, I examine how the reported changes in retirement age in each sample ($\triangle retirementtiming$) compare with the extra years respondents would need to work to attain their pre-crash sustainable consumption paths ($R_{\overline{sc}} - R_0$). Table 4 presents summary statistics for $R_{\overline{sc}} - R_0$. For each sample, the first column gives the statistics for all members of the sample, while the second column is restricted to those members reporting a non-zero increase in retirement age (CogEcon) or probability of full-time work (HRS). In the CogEcon sample, the mean of $R_{\overline{sc}} - R_0$ is 3.7 years overall, and 4.1 years for those who reported an increase in their expected age of retirement. The distributions are both skewed, such that 25th percentile is 0.5 years for the full sample and 0.9 years for those reporting a change, the median is 1.6 years for the full sample and 1.7 years for those reporting a change, and the 75th percentile is 4.1 years for the full sample and 3.9 years for those reporting a change. Similarly, the means of $R_{\overline{sc}} - R_0$ in the HRS samples are 4.9 and 5 years for the full age 62 and age 65 samples, respectively. The 25th percentiles of both "full" HRS samples are 0.65 years, the medians are 1.9 years, and the 75th percentiles of the distributions of $R_{\overline{sc}}$ are both approximately 4.9 years, as well. Additionally, comparisons between the first and second columns for each of the HRS samples show that respondents who adjusted their retirement plans tend to be those who would need to work longer to make up their losses. Despite the differences in wealth measures between the CogEcon and HRS, samples, the means and medians are relatively similar across the studies. Overall, this table shows that the wealth losses from the crash, if permanent, would require quite large adjustments of retirement timing to fully make up. Furthermore, respondents who indicate an increase in expected retirement age or in the subjective probability of full-time work into their 60s tend to be those with larger wealth losses (as measured by $R_{\overline{sc}} - R_0$).

Figure 9 displays a histogram of $R_{\overline{sc}}$ from the CogEcon data, rounded to the closest integer, and R_{09}, reported post-crash planned retirement ages. In this figure, R_{09} and $R_{\overline{sc}}$ have been top-coded at age 90. Ignoring the spikes due to top-coding, the modes of both distributions are at age 65, with spikes at ages 62 and 66 and some evidence of focal answers at 60, 65, 70, 75, 80, and 85. There is a significant spike at 90. This is induced by top-coding, but is more significant for $R_{\overline{sc}}$ than R_{09}. The implication of this spike is that a non-trivial percentage of respondents would have to work beyond age 90 to fully recoup losses sustained between 2008 and 2009.

Figure 10 uses CogEcon data to compare $R_{\overline{sc}} - R_0$, the extra number of years of work needed to maintain pre-crisis standards of living, and $R_{09} - R_0$, the reported change in planned retirement ages. The distributions look relatively similar. However, the distribution of reported changes in retirement age is compressed toward zero, relative to $R_{\overline{sc}} - R_0$. The compressed distribution provides suggestive evidence that the cost of adjusting retirement age may grow with the size of adjustment. This could be true, for example, if the marginal disutility of work increases non-linearly with age, making increasing one's retirement age beyond 65 or 70 less attractive than accepting a somewhat lower material standard of living. Alternatively, the compressed distribution is also consistent with the possibility that particularly large values of $R_{\overline{sc}} - R_0$ are more likely to be the result of measurement error, and therefore do not result in large observed changes in reported retirement age. The incidence of reported changes of one year are much lower than might be expected, given the relatively large number of observations for which $R_{\overline{sc}} - R_0$ is equal to one. The gap at one year suggests that a fixed cost of adjusting retirement age may exist, as was suggested in Section 4.

In this section, I have established that the impacts of the asset losses between 2008 and 2009 are non-trivial. I have also shown evidence that retirement expectations in my sample have shifted toward later retirement. Next, I turn to regression analysis to examine the relationship between these phenomena.

6.2 Regression analysis

As discussed in Section 3, the life-cycle model featuring choice of retirement timing and consumption implies that asset shocks will affect the chosen retirement age, level of sustainable consumption, or both. Moreover, if retirement leisure is a normal good, the model implies that individuals will adjust to asset shocks, at least somewhat, along the retirement age margin. Using the empirical framework presented in Section 4, I test this implication by regressing the observed change in retirement timing ($\triangle$*retirement timing*) on the change in retirement age that would be necessary to restore the pre-crash sustainable consumption level ($R_{\overline{sc}} - R_0$). Based on my discussion in Section 4 about the possibility that both measurement error and non-linearities in the underlying optimization problem may affect the regression estimates, I also include the square of ($R_{\overline{sc}} - R_0$) in some regressions to relax the restriction that large values of $R_{\overline{sc}} - R_0$ have the same estimated marginal effect as more moderate values.

In the discussion of the empirical framework (Section 4), I have pointed out that there may be fixed costs associated with changing retirement plans, and that non-linearities in the underlying optimization problem may result in heaping at the "zero adjustment" margin. Consistent with this observation, I have shown in Section 6.1 that there are large numbers of respondents in both samples for whom no change in the retirement timing variable is observed. In the case of a mass at zero adjustment, estimates from corner solution models are more likely to be consistent than ordinary least squares estimates. The Tobit model is more restrictive than many other econometric models for corner solutions, but provides efficiency gains over multi-equation models. Because specification tests discussed in Section 6.2.4 do not provide significant cause for concern about the Tobit specification, most analyses presented in this paper use the Tobit specification.

6.2.1 CogEcon base regressions

Table 6 presents the main regression results from the CogEcon sample, in which the dependent variable is the reported change in expected retirement age, $R_{09} - R_0$. In these regressions the independent variable of interest is the change in retirement age that would be needed to make up wealth losses, $R_{\overline{sc}} - R_0$. The first column of Table 6 presents the results from an ordinary least squares regression using the CogEcon data.[30] As predicted by theory, the coefficient on $R_{\overline{sc}} - R_0$ is positive. The coefficient of 0.058 (s.e. 0.042) implies that, for each year individuals would have to work to make up wealth losses, on average they only increase expected retirement age by 0.058 years, or about three weeks. For individuals with an average value of $R_{\overline{sc}} - R_0$ (3.7 years), this translates to a predicted change in retirement age of about two and a half months. However, as is the case with many studies of the impact of wealth changes on retirement timing [McGarry, 2004, Chan and Stevens, 2004, Hurd

[30] As in all regressions using the CogEcon sample, I use CogUSA sampling weights and report robust standard errors.

et al., 2009], the effect is not statistically significantly different from zero. In Column 2, I include the variable $(R_{\overline{sc}} - R_0)^2$, to allow for possible non-linear effects of the independent variable. Here, the coefficient on the linear term is virtually unchanged (0.057, s.e. 0.038), and the coefficient on the squared term is virtually zero (0.0003, s.e. 0.0009) and imprecisely estimated, but the F-test does suggest that this model improves the fit. Taking into account both the linear and squared terms, the marginal effect of $R_{\overline{sc}} - R_0$ on the predicted change in retirement age is 0.059 years (s.e. 0.044), virtually unchanged from the linear model.

Use of the Tobit model, rather than ordinary least squares, may allow for consistent estimation in the presence of a spike at zero. Columns 3 and 4 present the results from Tobit regressions with censoring at zero.[31] In Column 3, where $R_{\overline{sc}} - R_0$ enters the regression only linearly, the coefficient on $R_{\overline{sc}} - R_0$ is 0.109 and statistically insignificant. The average marginal effect of $R_{\overline{sc}} - R_0$ on retirement timing for those who reported a change,

$$\frac{\partial E(\triangle retirement\, timing | R_{\overline{sc}} - R_0, \triangle retirement\, timing > 0)}{\partial (R_{\overline{sc}} - R_0)}$$

is 0.042 (s.e. 0.036), slightly smaller than that implied by the OLS regression. It implies a retirement age effect of just under two months for a respondent with the average value of $R_{\overline{sc}} - R_0$.

In Column 4, results are shown from a Tobit regression including $(R_{\overline{sc}} - R_0)^2$. While we might expect the coefficient on this squared term to be positive as a result of picking up a threshold effect, the Tobit regression specification explicitly models the threshold. It seems that the addition of this squared term serves, instead, to minimize the effect of very large—and possibly error-ridden—values of $R_{\overline{sc}} - R_0$ on the main estimated effect of $R_{\overline{sc}} - R_0$. Indeed, the coefficient on $R_{\overline{sc}} - R_0$ is 0.311 years (0.163), much larger than in first three specifications, and statistically significant at the 5 percent level. The coefficient on the squared term is -0.010 years (0.008), which is not statistically significantly different from zero, but appears to have improved the fit, nonetheless. Accounting for both the linear and squared terms, the average marginal effect of $R_{\overline{sc}} - R_0$ is 0.086 years (s.e. 0.043), indicating an adjustment of just over one month for each additional year or work needed to make up wealth losses. Given that the average of $R_{\overline{sc}} - R_0$ is 3.7 years, this works out to just under 3.9 months of additional work for an individual with the average increase in work years needed to attain pre-crash sustainable consumption levels.

Results from the specification used in Column 4 are also presented in graphical form in Figures 11 and 12. In Figure 11, the predicted the probability of an increase in retirement age (based on the Tobit regression) and the proportion of respondents actually reporting

[31] Setting the censoring point at one results in qualitatively similar estimates, but reduces the uncensored sample size. Thus, I present all results with the censoring point at zero.

an increase in expected retirement age are plotted over bins corresponding to ranges of the continuous variable, $R_{\overline{sc}} - R_0$. The predicted and actual probabilities of adjustment are of comparable magnitudes, and the patterns are reasonably similar. In Figure 12, the predicted increase in retirement age (based on the Tobit regression) and the average reported increase in expected retirement age are plotted over the bins. Here, it can be seen that the Tobit regression under-predicts the size of the reported changes. This is consistent with the results we would expect if attenuation bias due to measurement error in $R_{\overline{sc}} - R_0$ is a significant problem.

6.2.2 HRS base regressions

Tables 7 and 8 present the results from Tobit regressions like those in Columns 3 and 4 of Table 6, but using the HRS sample. To reduce the number of regressions presented, I restrict results presented in the rest of this paper to Tobit specifications. These are more likely to provide consistent estimates, given the spike of observations at zero, compared to linear regression specifications. In general, the implied effects from the OLS regressions on the HRS sample are very imprecisely measured and have smaller or comparable magnitudes to those estimated using the Tobit specifications.

In Table 7, the dependent variable is ${}_{08}\triangle_{09}Pr(FT62)$, while in Table 8 the dependent variable is ${}_{08}\triangle_{09}Pr(FT65)$. The coefficient sizes and marginal effects are not directly comparable to the CogEcon results. To provide a crude basis for comparison of the magnitudes of the CogEcon and HRS results, I have used 2006 and 2008 Core data to estimate the average effect of a percentage point increase in the probability of full-time work on the change in age at which HRS respondents planned to stop work completely.[32] A one percentage point increase in the probability of full-time work after reaching age 62 is associated with about a one week increase in the planned age of retirement. Similarly, a one percentage point increase in the probability of full-time work after reaching age 65 is associated with an increase of about six days in the planned age of retirement.

The first column of Table 7 presents the results from a Tobit regression of ${}_{08}\triangle_{09}Pr(FT62)$ on $R_{\overline{sc}} - R_0$. The coefficient on $R_{\overline{sc}} - R_0$ is 0.193 (s.e. 0.267). This translates to an average marginal effect of 0.079 (s.e. 0.109), meaning that a one year increase in the number of years an individual would need to work to make up losses is associated with less than a 0.1 percentage point increase in the probability of full-time work after age 62. At the mean of $R_{\overline{sc}} - R_0$ (4.9 years) the implied effect of wealth losses on retirement age is about three days. In addition to being very small, this estimate is very imprecisely estimated. Column 2 shows that including the squared term of $R_{\overline{sc}} - R_0$ in the regression increases the magnitude

[32] See Appendix 12 for the results from these regressions.

of the coefficient on the linear wealth loss measure ($R_{\overline{sc}} - R_0$). The average marginal effect of $R_{\overline{sc}} - R_0$ is now 0.245 (s.e. 0.224). This is still very small and statistically insignificant, however; it implies that a wealth loss that would take an extra year of work to make up is only associated with a quarter of a percentage point increase in the probability of full-time work after age 62. At the mean of $R_{\overline{sc}} - R_0$, the implied effect is a retirement delay of just ten days in response to a wealth loss that would take 4.9 years to make up.

Results from the specification used in Column 2 are also presented in graphical form in Figures 13 and 14. Figure 13 illustrates the predicted and observed adjustments along the extensive margin. The predicted and actual probabilities of adjustment are of comparable magnitudes, and the patterns are reasonably similar. Figure 14 illustrates the predicted and observed adjustments along the intensive margin. Here, it can be seen that the shape of the line representing the reported increases is different from the line representing predicted increases, implying that the model may not fit the data particularly well in this case. Additionally, the underprediction of adjustments by the Tobit may be indicative that measurement error is causing significant attenuation bias.

Due to the fact that most respondents' retirement ages are imputed, measurement error in $R_{\overline{sc}} - R_0$ may be an even larger concern in the HRS data than in the CogEcon data. Specifically, if retirement age is imprecisely measured, then both the earnings component of wealth and R_0 contain a lower signal-to-error ratio in the HRS for individuals with imputed values of R_0, resulting in less-precise calculated values of $R_{\overline{sc}} - R_0$. It is not clear whether these values are biased, or only subject to random error. If, however, the error is classical, regression coefficients may be attenuated.

In an attempt to reduce measurement error due to imputation of R_0, Column 3 presents the results from conducting the same regression on the subset of the HRS respondents who did report an expected retirement age in the 2008 Core interview. Using this restricted sample, the coefficient on $R_{\overline{sc}} - R_0$ is 0.49 (s.e. 0.765), and the average marginal effect is 0.191 (s.e. 0.295). The implied effect at the average wealth loss is about a week. Column 4 displays coefficients from the regression on the resticted sample including the squared term of $R_{\overline{sc}} - R_0$. The coefficients in this regression have larger magnitudes than those estimated using the full sample, and are they are similar in sign and relative magnitude to the CogEcon results, but they are still very imprecisely estimated. The marginal effect of $R_{\overline{sc}} - R_0$ in Column 4 is 0.426 percentage points (s.e. 1.094), implying that at the average wealth loss (in terms of $R_{\overline{sc}} - R_0$) of 4.9 years, the average retirement in retirement age is only about eighteen days.

In Columns 3 and 4, the estimated marginal effects are larger than the results from the full sample. This is suggestive that measurement error may be causing attenuation bias in the full sample regressions. However, it could also be the case, for example, that respondents who have better-defined retirement plans (and therefore provided a retirement age) are more reactive to wealth losses. Whether or not classical measurement error is reduced in this subsample, these results continue to imply much smaller effects than the CogEcon estimates, and the estimated effects are not statistically significantly different from zero.

In Table 8, in which ${}_{08}\triangle_{09}Pr(FT65)$ is the dependent variable, the estimates tell a slightly different story. Column 1 reports results from the regression of ${}_{08}\triangle_{09}Pr(FT65)$ on the linear measure of the extra number of years an individual would need to work to make up losses ($R_{\overline{sc}} - R_0$), for the full sample. The size of marginal effect, at 0.137 (s.e. 0.121), is almost double that from Column 1 of Table 7, but still implies that the average $R_{\overline{sc}} - R_0$ of 5 years is associated with a very small average increase in retirement age (about five days). Column 2, in which the square of $R_{\overline{sc}} - R_0$ is an additional regressor, shows that the coefficients on $R_{\overline{sc}} - R_0$ and $(R_{\overline{sc}} - R_0)^2$ have larger magnitudes than in the other HRS regressions and are statistically significant at the one percent level. The average marginal effect of $R_{\overline{sc}} - R_0$ is 0.736 (s.e. 0.235), or almost three-quarters of a percentage point increase in the probability of full-time work after age 65 for each extra year of work needed to make up wealth losses. Using the crude comparison of 6.7 days delay in retirement per percentage point increase in ${}_{08}\triangle_{09}Pr(FT65)$, the implied average effect at the mean value of $R_{\overline{sc}} - R_0$ (5 years) on retirement age is about three and a half weeks.

Results from the specification used in Column 2 are also presented in graphical form in Figures 15 and 16. Figure 15 illustrates the predicted and observed adjustments along the extensive margin. The predicted and actual probabilities of adjustment are of comparable magnitudes, and the patterns are quite similar. Figure 16 illustrates the predicted and observed adjustments along the intensive margin. Again, consistent with the attenuation bias discussion with respect to Figures 12 and 14, it can be seen that the Tobit regression under-predicts the size of the reported changes.

Columns 3 and 4 of Table 8 repeat the analyses from Columns 1 and 2, but restrict the sample to respondents who reported an expected retirement age in the 2008 Core interview. In Column 3, the sign on the coefficient on $R_{\overline{sc}} - R_0$ is negative. This is contrary to theoretical predictions, but relatively small and not statistically significant. In Column 4, the results are very similar to those in Column 2, but are less statistically significant. The average marginal effect of $R_{\overline{sc}} - R_0$ is 0.661 (s.e. 0.456), or about two-thirds of a percentage point increase in the probability of full-time work after age 65 for each extra year of work that would be needed to make up wealth losses. At the mean of $R_{\overline{sc}} - R_0$, this implies an increase of retirement age of just over three weeks.

These analyses, like those using the CogEcon data, provide some evidence that wealth losses are associated with delay of retirement. A one or two percentage point increase in the subjective probability of full-time work at age 62 or 65 may not seem particularly significant in the economic sense, but given that Hurd [2009] has found that average subjective probabilities reported by HRS respondents are close to the population average outcomes, the effect of the wealth losses may have meaningful effects on labor supply in the aggregate. Additionally, to the extent that measurement error in the HRS data is causing attenuation bias in the analyses, the aggregate labor supply effects of wealth losses may be much larger.

6.2.3 Comparison of CogEcon and HRS findings

Table 9 provides a summary of the regression results from the CogEcon sample and the two HRS samples. My preferred specifications, Tobit regressions including the number of additional years it would be necessary to work to maintain pre-crash sustainable consumption levels and the square of that number show that changes in planned retirement age are, indeed, positively associated with the impact of wealth losses. In particular, the fourth row of Table 9 summarizes the results from my preferred specification for the CogEcon sample, and shows that the average marginal effect translates to an average increase of about four months for individuals who suffered the mean wealth loss, in terms of years of work needed to make up losses suffered in 2008 and 2009. In the HRS data, the marginal effects of average wealth losses appear to explain 1 to 2 percentage points of the increase in the probabilities of full-time work after age 62, and 3 to 4 percentage points of the increase in the probabilities of full-time work after age 65, at least for those respondents who gave non-zero changes in retirement timing. In contrast to the results from the CogEcon regression results, HRS regression results summarized in rows 6, 8, 10 and 12 seem to imply smaller planned delays of retirement, possibly on the order of 1.5 to 3.5 weeks. However, these estimates may be attenuated due to measurement error.

Figures 17 through 22 allow for a visual comparison of the results from the different datasets. Figures 17 and 18 present the results from the preferred CogEcon specification presented in Column 4 of Table 6, a Tobit regression including $R_{\overline{sc}} - R_0$ and $(R_{\overline{sc}} - R_0)^2$. The first of these presents a plot of the average marginal effect of $R_{\overline{sc}} - R_0$ within each sub-group of $R_{\overline{sc}} - R_0$.

These show that the marginal effect of needing to work an extra year to makeup losses is relatively flat for low and moderate levels of wealth losses and appears to decline for the largest losses, relative to the years of work needed to make up losses. While inclusion of $(R_{\overline{sc}} - R_0)^2$ in regressions might typically be expected to create an inverse U-shaped plot of the marginal effects of wealth losses on retirement age, only very large values of wealth change exhibit the expected pattern. An interesting interpretation of this pattern is that the turned-down shape exhibited by this graph may be an illustration of the attenuating impact of measurement error in particularly large values of this calculated variable. When the quadratic term is included in the regressions, the largest values of $R_{\overline{sc}} - R_0$ receive less weight in the estimation of the coefficient on the linear measure of wealth loss, thereby reducing attenuation bias in the estimate of the coefficient on the linear measure of $R_{\overline{sc}} - R_0$.

Figure 18 plots predicted changes in retirement age over different levels of $R_{\overline{sc}} - R_0$. These changes display roughly the expected pattern: smaller increases for individuals with no or low losses in wealth, somewhat larger increases in retirement age for those who are more affected by the crash, and then a slight dip in the predicted effect on those with the largest values of $R_{\overline{sc}} - R_0$. Figures 19 and 20 are parallel graphs for the preferred full-sample HRS regressions using ${}_{08}\triangle_{09}Pr(FT62)$ as the dependent variable, and Figures 21 and 22 present results from the preferred full-sample regression using ${}_{08}\triangle_{09}Pr(FT65)$ as the dependent variable. The marginal effects graphs are all very similar in shape, as are the predicted outcome graphs. Overall, the CogEcon and HRS samples appear to tell very similar stories. However, the predicted outcome graphs also show that wealth losses (in the form of $R_{\overline{sc}} - R_0$) may not be whole story.

6.2.4 Comparison of Tobit with Cragg's two-tiered model

The Tobit model is quite restrictive. A single underlying mechanism determines both the marginal effects of variables at the observed outcome and whether the observed outcome is at a corner solution. Two-tiered models relax this restriction by allowing different equations for the intensive and extensive margins. Cragg (1971) suggests a two-tiered model consisting of a probit and a truncated normal regression. While the Tobit model offers greater efficiency than a two-tiered model—an important consideration, given the small sample sizes used in this study—it is not consistent if misspecified. Below, I present comparisons between the Tobit, probit and Cragg's alternative, as well as results from two separate specification tests, to affirm my use of the Tobit model in this paper.

Table 10 presents the estimates from Tobit, probit and truncated normal regressions on the CogEcon sample. These estimates allow comparisons between the Tobit and probit models, and between the Tobit and Cragg models. In both the Tobit and truncated normal regressions, the dependent variable is censored at zero. For the probit, the dependent variable is an indicator variable equal to one if planned retirement age increased, and zero otherwise.

As with the Tobit, the coefficients on $R_{\overline{sc}} - R_0$ are positive in both the probit and truncated normal regressions, and the coefficients on $(R_{\overline{sc}} - R_0)^2$ are negative. A simple test for whether the Tobit may be misspecified is to compare the Tobit estimates, normalized by the estimated standard error of the regression, to the probit estimates. If they are of different signs or of very different magnitudes, this may suggest that the Tobit may be inappropriate (Wooldridge, 2002, pp. 533-534). Comparing the estimated Tobit and probit coefficients, it can be seen that the estimate $\frac{\beta_{Tobit}}{\sigma_{Tobit}}$ for $R_{\overline{sc}} - R_0$, 0.0498 is very similar to the estimate of β_{probit}, 0.0473. For $(R_{\overline{sc}} - R_0)^2$, the comparable estimates are -0.0016 for the Tobit to -0.0021 for the probit.

Table 11 presents the results from similar regressions on the HRS samples. In the ${}_{08}\triangle_{09}Pr(FT62)$ estimates (Columns 1-3), although none of the coefficients are distinguishable from zero at standard levels of significance, the normalized Tobit estimates again appear to be somewhat similar to the probit estimates. The estimate $\frac{\beta_{Tobit}}{\sigma_{Tobit}}$ for $R_{\overline{sc}} - R_0$, 0.0222 is similar to the estimate of β_{probit}, 0.0374. For $(R_{\overline{sc}} - R_0)^2$, the comparable estimates are -0.0006 for the Tobit to -0.0014 for the probit. In the ${}_{08}\triangle_{09}Pr(FT65)$ regressions (Columns 4-6), the coefficient estimates for $R_{\overline{sc}} - R_0$ and $(R_{\overline{sc}} - R_0)^2$ are statistically different from zero at the 1 percent significance level, and of the same signs in both the Tobit and probit specifications. In the truncated normal regression, the estimates are also reasonably similar to the Tobit estimates, but not statistically significant. As with the CogEcon results, the normalized Tobit estimates again appear to be very similar to the probit estimates. The estimate $\frac{\beta_{Tobit}}{\sigma_{Tobit}}$ for $R_{\overline{sc}} - R_0$, 0.064, is very similar to the estimate of β_{probit}, 0.065. For $(R_{\overline{sc}} - R_0)^2$, the comparable estimates are -0.0021 for the Tobit to -0.0023 for the probit.

The Cragg model nests the Tobit in the special case that $\frac{\beta_{truncated}}{\sigma_{truncated}} = \gamma_{probit}$. Using the log-likelihoods from maximium likelihood estimation of the Tobit and Cragg models, a likelihood-ratio test can be used to test the null hypothesis that the Tobit is nested in the Cragg model against the alternative that it is not. Rejection of the null hypothesis would suggest that the Tobit model is misspecified. For my preferred Tobit specifications, in which I regress retirement timing on $R_{\overline{sc}} - R_0$ and $(R_{\overline{sc}} - R_0)^2$, this likelihood-ratio test is calculated $-2(ln\mathcal{L}_{Tobit} - (ln\mathcal{L}_{probit} + ln\mathcal{L}_{truncated}))$, and has a $\chi^2(4)$ distribution. Because LR statistics based on weighted samples are generally not valid,[33] I conduct LR tests for my preferred specifications and their Cragg model alternatives using results from regressions conducted without weights. The full results from these regressions can be seen in Appendix 13.

For the CogEcon sample, the $\chi^2(4)$ test statistic of 1.78 implies a p-value of 0.78, failing to reject the null hypothesis. Similarly, for the HRS analysis with ${}_{08}\triangle_{09}Pr(FT65)$ as the dependent variable, the $\chi^2(4)$ test statistic of 5.77 (p-value 0.22) also fails to reject the null. In the case of the specification with the weakest results, the HRS analysis with ${}_{08}\triangle_{09}Pr(FT62)$ as the dependent variable, the $\chi^2(4)$ test statistic is 9.89 (p-value 0.04), rejecting the null hypothesis at the 5 percent significance level. Despite the rejection of the null in the last of these tests, I continue to present Tobit results for the ${}_{08}\triangle_{09}Pr(FT62)$ analyses to maintain

[33] See Wooldridge (2002) page 539.

comparability with other results in this paper.

Together, the proportionality results and likelihood-ratio tests do not raise significant concern that the Tobit model is misspecified. Furthermore, the imprecisely-estimated truncated normal regression results underscore the importance of the efficiency gain from the Tobit in yielding precise estimates for the small samples used in this study.

6.2.5 Robustness of estimates to alternate measures of total wealth

The approach used in wealth calculation to this point implicitly assumes that individuals optimize retirement and consumption plans subject to the constraint that they decumulate household assets down to zero by the time of death. However, Hurd and Smith [2002] estimated that the median HRS respondent of a decade ago would leave between $50,000 and $100,000 in the form of bequests. Not surprisingly, their estimates of mean expected bequests were even higher, ranging from $165,000 for individuals born before 1924 to more than $250,000 for those born between 1942 and 1947. Furthermore, research by Lusardi and Mitchell (2007) has shown that a vast majority of HRS homeowners do not think it likely that they will sell their homes to finance retirement, implying that respondents expect to retain a significant amount of primary home equity.

In Tables 12 and 13, I present the results from robustness checks, in which my "preferred" baseline regressions[34] are run using $R_{\overline{sc}} - R_0$ that have been calculated with alternate measures of total wealth. In Table 12, I present results based on exclusion of primary home equity.

In Table 13, I have excluded estimated expected bequests.[35] While neither the CogEcon study nor the HRS gather expected bequests directly, the HRS Core interviews ask probabilistic expectations questions about the probability of leaving at least $10,000 ($Pr(B \geq \$10k)$) and at least $100,000 ($Pr(B \geq \$100k)$). The 2009 Internet Survey also asked about the probability of leaving at least $500,000 ($Pr(B \geq \$500k)$). I generated point estimates of expected bequests in 2008, and subtracted this amount from both the 2008 and 2009 wealth figures before calculating $R_{\overline{sc}} - R_0$.

The general story told by my baseline results is unchanged under these alternate specifications.

[34] Tobit regressions of change in retirement timing on $R_{\overline{sc}} - R_0$ and $(R_{\overline{sc}} - R_0)^2$. For the HRS analyses, I use the "full" samples.

[35] The generation of point estimates for expected bequests is described in detail in Appendix 14.

6.2.6 Heterogeneity

I next explore heterogeneity in individuals' responses to wealth losses, and to the crash in general. In this section, I explore several possible ways in which individuals' reactions to a similar wealth loss may differ. First, rates of time preference and risk aversion may have affected the magnitudes of wealth levels, but are also likely to be associated with the reactions to wealth losses. Thus, it is interesting to explore the relationship between wealth *levels* and reactions to wealth *losses*. Second, different retirement horizons carry different implications for the costs of changing (or not changing) retirement plans. Specifically, those closest to retirement have less time over which to smooth consumption, and may be more likely to delay retirement due to the crash. Third, optimal reactions to comparable losses of wealth may differ by individual according to expectations about the economic recovery. Those who think that the economy will be slow to recover may be more reactive to wealth losses. Fourth, the effort needed to re-optimize one's retirement and consumption path may affect both the decision to change retirement age and the precision with which one calculates a new optimal retirement age. I use measures of financial knowledge and cognitive ability from the CogEcon and CogUSA studies to examine whether these factors are related to changes in retirement plans. Fifth, if individuals' pre-crash plans did not involve fully decumulating their assets (that is, if they were planning to leave a bequest), they may have had an additional margin over which to adjust to their wealth losses. Using information about expected bequest plans in the HRS, I examine the relationship between expected bequests, wealth losses and retirement plans. The findings in this section are suggestive that individuals' preferences, expectations and abilities are important factors to consider when examining the relationship between wealth losses and retirement plans.

Wealth levels and changes in planned retirement In an examination of the role of uncertainty in wealth accumulation in the HRS, Lusardi (1998) has found empirical support for some of the predictions of a life-cycle model with uncertainty. In particular, she has found that households that are more risk-averse or have longer planning horizons (implying lower discount rates) tend to accumulate more wealth. I expect that levels of wealth and, therefore, the incidence of wealth losses in 2008 and 2009, are correlated with a tendency to make up more of a wealth loss with longer work, as opposed to lower consumption. At the same time, households with the highest wealth may be less reactive to losses than those farther down the distribution, because the marginal value of consumption is likely to be lower for these individuals. Thus, I expect the marginal effect of my measure of wealth losses on retirement age to be most pronounced for individuals near the middle of the wealth distribution.

Table 14 presents the results from Tobit regressions including pre-crash wealth terciles.

In column 1 are the results from including different intercepts for each wealth tercile in the regression of $R_{09} - R_0$ on $R_{\overline{sc}} - R_0$ and $(R_{\overline{sc}} - R_0)^2$. The coefficients on the wealth tercile measures are statistically significant and of the expected signs. The large, negative coefficient on the third (highest) wealth tercile implies that wealthier individuals are less likely to change their retirement plans, compared to households with less wealth. The marginal effect of $R_{\overline{sc}} - R_0$, an extra year of work needed to attain the pre-crash consumption level, also reflects this pattern: at 0.169 (s.e. 0.066) and 0.181 (s.e. 0.062), the average marginal effects for households in the lowest two wealth terciles are quite comparable to one another, and much larger than the average marginal effect among the wealthiest households (0.058, s.e. 0.023). The marginal effects are equivalent to between 2 and 9 months of adjustment in retirement age for each year one would have to work to attain one's pre-crash sustainable consumption level. Additionally, these are all statistically significantly different from zero at the 5 percent level or the 1 percent level, as well as different from across terciles ($\chi^2(2) = 6.29$, p-value=0.04).

In column 2, I also interact the pre-crisis wealth terciles with the $R_{\overline{sc}} - R_0$ terms. Now the coefficients are much more imprecisely estimated, and the coefficients on the $R_{\overline{sc}} - R_0$ variables are virtually zero. However, the average marginal effects are similar in magnitude to those in column 1 and are, again, statistically significantly different from one another across terciles ($\chi^2(2) = 5.31$, p-value=0.07). For the lowest two wealth terciles, these effects are equivalent to about between 2 months of adjustment in retirement age for each year one would have to work to attain one's pre-crash sustainable consumption level; for the top wealth tercile, the average marginal effect implies a change in retirement age of about three weeks for each year of $R_{\overline{sc}} - R_0$.

The results presented in Table 14 provide some support for the hypothesis that those at the top of the wealth distribution are less reactive to wealth losses, possibly because of a lower marginal value of wealth.

Expectations and changes in planned retirement A recent structural life-cycle model by Low et al. [2010] illustrates the importance of incorporating risk into life-cycle models. Their model predicts that increased job destruction and wage variation have strong impacts on welfare, and that individuals are willing to pay significant amounts to avoid these risks. In the option value framework of Stock and Wise [1990], low expectations or uncertainty about the future increase the option value of continued work, resulting in later planned retirement. In their analysis of the determinants of retirement expectations, Chan and Stevens [2004] include controls for future expectations about job losses to try to control for changes in the probability of full-time work due to factors outside of individuals' control. They find that the ease of finding a new job is positively related to the subjective probability of full-time work after reaching age 62 or 65.

The option value of continuing work beyond one's originally planned retirement age may be highest for individuals who were already close to retirement in 2008. Those who are closest to retirement are likely to be the most reactive to their wealth losses, since uncertainty about when and to what extent the stock, labor and housing markets would rebound may lead these individuals to continue working until the uncertainty surrounding the recession has been resolved. I do, however, expect that uncertainty about future labor market, stock market and real estate returns will still be related to changes in retirement age, even for those not close to planned retirement, because continued work provides insurance against negative asset shocks regardless of time to retirement. Tables 15, 16 and 17 present results from Tobit regressions of $R_{09} - R_0$ on $R_{\overline{sc}} - R_0$, $(R_{\overline{sc}} - R_0)^2$, and several variables related to the option value of continued work.

In column 1 of Table 15, indicators of time from 2009 to individuals' pre-crash retirement ages (less than two years, two to five years, five to ten years and more than ten years) are included in the base regression. At 0.316 (s.e. 0.17) and -0.011 (s.e. 0.01), the coefficients on $R_{\overline{sc}} - R_0$ and $(R_{\overline{sc}} - R_0)^2$ are similar to the preferred estimates in column 4 of Table 6. However, the coefficients on the indicators of time to retirement show that, the farther away from one's 2008 planned retirement age, the smaller the change in planned retirement age. Although the average marginal effects of $R_{\overline{sc}} - R_0$ are not statistically significantly different from one another across groups, these do decline monotonically as time to retirement increases, dropping from 0.119 (s.e. 0.06), or around 44 days per year of $R_{\overline{sc}} - R_0$ for those within two years of retirement, to 0.056 (s.e.0.03), or around 20 days per year of $R_{\overline{sc}} - R_0$ for those more than ten years from retirement.

In column 2, these indicator variables are interacted with $R_{\overline{sc}} - R_0$ and $(R_{\overline{sc}} - R_0)^2$. The coefficient on $R_{\overline{sc}} - R_0$ is now larger, at 0.584 years (s.e. 0.24), while the coefficient on $(R_{\overline{sc}} - R_0)^2$ is similar to the other analyses, at -0.02 (s.e. 0.01). However, the interaction terms with the indicators of years to retirement negate this effect for all but those closest to retirement. The average marginal effect of $R_{\overline{sc}} - R_0$ is 0.216 (s.e. 0.07) for those closest to retirement, equivalent to about 2.6 months, but much smaller and very imprecisely estimated for the other groups. Thus, those closest to retirement are reacting the most (2.5 months) to each year of work needed to attain pre-crash consumption, while those farther from retirement may be reacting to the asset losses by delaying retirement by just a few days (for those 2 to 5 years out) to a month (for those 5 to 10 years out) per year needed to attain pre-crash consumption.

In Table 16, column 1 displays the results of the Tobit regression of $R_{09} - R_0$ on $R_{\overline{sc}} -$

R_0, $(R_{\overline{sc}} - R_0)^2$ and variables indicating stock market, labor market and housing market optimism.[36] The coefficients on $R_{\overline{sc}} - R_0$ and $(R_{\overline{sc}} - R_0)^2$ are similar but smaller in magnitude than in the base specification (Table 6, column 4), but labor market and stock market optimism are associated with much smaller changes in retirement age. In this regression, the average marginal effect of $R_{\overline{sc}} - R_0$ is 0.05 (0.035), or about 2 weeks' increase in retirement age, compared to optimism about the labor and stock markets being associated with 0.427 (s.e. 0.41) and 0.73 (s.e. 0.38) year decreases in retirement age, respectively. The coefficient on housing market optimism is close to zero.

Interacting the stock market optimism variable with $R_{\overline{sc}} - R_0$ and $(R_{\overline{sc}} - R_0)^2$ yields the results seen in column 2. The coefficient on $R_{\overline{sc}} - R_0$ is larger than in previous specifications, at 0.39 (s.e. 0.14), but the average marginal effect is 0.05 (s.e. 0.04), or 2 weeks, slightly smaller than in other specifications. However, the average marginal effect of $R_{\overline{sc}} - R_0$ among those who are not optimistic about the stock market is 0.13 (s.e. 0.04), or an increase in retirement age of 1.5 months per year of $R_{\overline{sc}} - R_0$. This is statistically significantly different from the average marginal effect for those who are optimistic (-0.11, s.e. 0.076). That is, individuals who were more certain that the stock market would be higher in one year were much less reactive to wealth losses.

Including labor market optimism instead of stock market optimism yields substantively similar results, though these are not statistically significantly different by group. See column 3 of Table 16 for details. Housing market optimism, by contrast, is virtually unrelated to the reported changes in retirement age (see column 4).

In addition to labor market optimism, an additional measure that might be related to the option value of keeping one's job is the local unemployment rate: if the local unemployment situation worsens, especially contemporaneously with financial and real estate asset losses, the option value model predicts that the value of continued work will increase. Table 17 presents results from regressions including a categorical variable for the change in county unemployment rate between May 2008 and May 2009.[37] The labor market performed ex-

[36]Stock market optimism is coded as one if a respondent answers that there is more than a 50% chance to the question "By next year at this time, what are the chances that mutual fund shares invested in blue chip stocks like those in the Dow Jones Industrial Average will be worth more than they are today?" and zero otherwise. Labor market optimism is coded as one if a respondent answers that there is more than a 50% chance to the question "Two years from now, what is the percent chance that jobs will be easier to find than they are right now?" Similarly, the housing market optimism variable is from the question "We are interested in how the value of your home will change in the future. What is the percent chance that one year from now your home will be worth more than today?" The results of the regressions are very similar when using the 0% to 100% scale instead of the indicator variables, but the "optimism" indicator variables are slightly more powerful. Given the rounding common in subjective probability questions, and the frequency of focal answers at 50%, I think the indicator variables are also easier to interpret and less subject to measurement error.

[37]Because county-level unemployment data are not seasonally-adjusted, I have used unemployment rates

tremely poorly over the year ending in May 2009: just 39 percent of the CogEcon sample resided in counties that experienced an increase in the unemployment rate of less than 3 percentage points, while 21 percent resided in counties that experienced increases in unemployment of 3 to 4 percentage points, and 40 percent resided in counties that experienced increases in unemployment of more than 4 percentage points. I created a categorical variable for the change in unemployment rate to reflect each of these three groups. Column 1 presents results from a regression in which the categorical change in unemployment variable is added to the base specification. Indicators for this variable are not statistically significant, and do not greatly change the results from the base specification (Table 6, Column 4). Indeed, the average marginal effect of one year of $R_{\overline{sc}} - R_0$ is 0.08 (s.e. 0.04), implying that a wealth loss that would take one year to make up is associated with an increase in retirement age of about one month; the average marginal effects are extremely similar across categories of the unemployment variable. The effects of the categorical variable are small and statistically indistinguishable from zero. In Column 2, the interaction of the change in unemployment rate indicators with $R_{\overline{sc}} - R_0$ and $(R_{\overline{sc}} - R_0)^2$ does change the coefficients somewhat from the base specification. However, the average marginal effect of wealth loss is still virtually unchanged from the base specification, though it is no longer statistically significant, and the average marginal effects are neither substantively nor statistically significantly different from one another across categories of the unemployment variable.[38]

In this section, I have shown that those closer to retirement are likely the most reactive to wealth losses from the crash. Stock and labor market expectations are also related to reported changes in retirement age, with greater pessimism being associated with a stronger relationship between wealth losses and changes in retirement age. However, and perhaps surprisingly, changes in the local unemployment rate do not appear to change individuals' reactivity to wealth losses. In the next set of regressions, I turn to the roles of ability and knowledge in determining individuals' reactions to wealth shocks.

Cognitive ability, knowledge and changes in planned retirement In a 2008 book chapter, Clark and D'Ambrosio assert that developing a retirement plan requires understanding of certain financial relationships. Two relationships that they claim are easy to understand are that for a given desired consumption level, retiring earlier requires greater saving, and that for given retirement timing, individuals must save more to attain higher income in retirement. However, they note that some decisions, such as deciding on re-

---

from exactly one year apart, with the end date coinciding with the CogEcon 2009 survey fielding.

[38]Results from parallel analyses using a continuous measure of change in unemployment rate yielded similar (non-)results. Given the small sample size and the clear interpretation of a categorical variable, I have opted to not to present these.

quired saving levels and portfolio allocation, require difficult calculations. They assert that most workers do not have adequate financial knowledge to choose the retirement age and consumption and savings paths that maximizes lifetime utility. The CogEcon data contain measures that allow me to test whether responses to the wealth shock are related to financial knowledge or cognitive ability.

First, the CogEcon data contain a measure of financial knowledge from a battery of 25 questions. Columns 1 and 2 in Table 18 show the Tobit results from a regression of $R_{09} - R_0$ on $R_{\overline{sc}} - R_0$, $(R_{\overline{sc}} - R_0)^2$ and financial knowledge indicators for whether the respondents' scores on this battery were in the bottom, middle or top tercile.

In the first column, it can be seen that the coefficients on $R_{\overline{sc}} - R_0$ and $(R_{\overline{sc}} - R_0)^2$ are largely unchanged, relative to the baseline specification in (Table 6, column 4). Additionally, being in the highest financial knowledge tercile is associated with a 2.5 year smaller change in retirement age, compared to those in the lowest financial knowledge tercile.

In column 2, financial knowledge indicators are interacted with the other variables. Those who are least financially knowledgeable appear to be most reactive to each additional year of work needed to attain pre-crash sustainable consumption. Indeed, the average marginal effect of $R_{\overline{sc}} - R_0$ is 0.344 years (s.e. 0.039) for the lowest tercile of financial knowledge, 0.093 years (s.e. 0.052) for the middle tercile, and -0.019 years (s.e. 0.042) for the highest tercile. It should be noted, however, that the level of wealth and the level of financial knowledge are positively correlated with one another (correlation coefficient is 0.13), so it is not clear whether financial knowledge or wealth is behind this association.

Second, the CogUSA data that are linked to the CogEcon data contain a measure of fluid intelligence called the number series score. In columns 3 and 4, I present the Tobit results including indicators for whether respondents' scores on the number series test are in the bottom, middle or top tercile. Fluid intelligence does not appear to be related to reactions to the shock. The marginal effects of $R_{\overline{sc}} - R_0$ do not differ substantively or statistically by number series tercile.

While financial literacy and ability may both affect basic financial planning decisions, I do not find evidence that ability affects reactions to the economic crisis in a systematic way. Better financial literacy is associated with less drastic reactions to wealth losses, but the interpretation of this finding is unclear because financial knowledge is also correlated with both pre-crash wealth and the incidence of the crash.

Expected bequest behavior If individuals were planning to leave a bequest before the crash, they may have had an additional margin over which to adjust to their wealth losses. If bequests are a normal good that enter directly into the utility function, one might expect that individuals who planned to leave a bequest might reduce the bequest they expected to leave in reaction to the crash. If however, individuals do not view bequests as fungible, we might expect a larger change in retirement age for individuals who do not revise their bequest downward.

In Tables 19 and 20, I use categorical variables representing the change in the probability that a respondent will leave a bequest of at least $100,000 to examine the relationship between expected bequests, wealth losses and retirement plans. The CogEcon dataset does not contain data on expected bequests, so I conduct these analyses using the HRS samples.

In Table 19, I present results from regressions including indicators for whether the subjective probability of leaving a bequest of $100,000 or more decreased, remained unchanged, or increased between the 2008 Core interview and the 2009 Internet Survey. In column 1, I simply add these indicators of change in bequest plans to the base specification (Table 7, column 2). The estimated coefficients on $R_{\overline{sc}} - R_0$ and $(R_{\overline{sc}} - R_0)^2$ are virtually identical to those in the base HRS age-62 estimates, as are the average marginal effects and implied change in the probability of full-time work after age 62. However, unchanged or increased probabilities of leaving at least $100,000 as a bequest are actually negatively related to the change in probability of full-time work after age 62, indicating that respondents who adjust their work expectations are also likely to adjust their bequest intentions. Rather than being substitutable margins of adjustment, individuals who react to the crash appear to adjust along both margins. In column 2, the indicator variables are also interacted with $R_{\overline{sc}} - R_0$ and $(R_{\overline{sc}} - R_0)^2$, and yield similar results. In columns 3 and 4, the dependent variable is the change in the probability of full-time work after age 65. Results from these regressions tell a similar story.

Table 20 repeats the analyses in Table 19 using a binary indicator for whether the probability of leaving a bequest of $100,000 or more decreased by at least 15 percentage points. These results tell a similar story to those in Table 19. These results show that individuals who adjust their work expectations also tend to alter their bequest intentions. However, individuals who alter their bequest intentions do not tend to be more reactive to wealth losses, in terms of the way they adjust their work plans (that is, the average marginal effects do not differ across groups).

These analyses seem to suggest that individuals who react to the crash in terms of their labor supply plans also tend to adjust their bequest intentions. That is, they are re-optimizing along both bequest and retirement age margins.

7 Conclusion

Economists have theorized that a negative income or wealth shock will cause individuals to re-optimize their consumption and retirement plans. In particular, a negative wealth or income shock is expected to produce a delay in expected retirement timing. In contrast to the clear predictions of most life-cycle models, many researchers have summarized empirical estimates of the impact of cyclical wealth effects on retirement timing as providing mixed or weak evidence with respect to these predictions [Coile and Levine, 2005, Hurd et al., 2009, Goda et al., 2011]. It is certainly the case that analyses of "boom" years in Hurd and Reti (2001), Coile and Levine (2006) and Hurd, Reti and Rohwedder (2009) show little to no impact of wealth changes on retirement timing. However, results presented in this paper and the results of other recent studies, plus work by Sevak (2002), Coronado and Perozek (2003), and analyses of "bust" years both in Coile and Levine (2006) and Hurd, Reti and Rohwedder (2009) all provide some support for the life-cycle model.

Based on existing empirical evidence and the analyses presented in this paper, I conclude that there is a positive relationship between wealth losses and retirement age consistent with the implications of life-cycle models. It is likely that the weak-to-zero estimated effects seen in many studies stem from measurement problems, failure to take into account the fixed costs of adjusting retirement plans and, possibly, asymmetries in the effects of wealth losses versus gains due to non-linearities in the underlying choice problem. Through use of novel data and improved econometric specification, this paper improves on each of these factors. Additionally, results in this paper show that it is interesting to consider the role of heterogeneity in preferences, expectations and other individual characteristics in examining the role of exogenous wealth shocks on retirement timing, as the estimated wealth effects often differ between individuals in expected ways.

This paper uses quasi-experimental pre- and post-crash data from the Cognitive Economics and Health and Retirement Studies to examine the impact of wealth losses between summer 2008 and summer 2009 on the retirement plans of older Americans. Calculations based on new survey data estimate that the stock and housing crises, together with rapidly rising unemployment, reduced the sustainable material standard of living of the typical (median) pre-retirement older American by about 5 percent between summer 2008 and summer 2009; average losses were almost twice as large. The additional number of years the median respondent would need to work to make up these losses is 1.6 years in the CogEcon data and 1.9 years in the HRS dataset, while the average increases needed to make up losses are 3.7 years in the CogEcon data and 4.9 to 5 years in the HRS data. Descriptive analyses show that the economic crisis did result in increases in planned retirement age: just over

40 percent of respondents in the CogEcon sample reported changing the age at which they planned to retire completely by at least a year "as a result of the economic crisis," while HRS respondents' probabilities of full-time work in their sixties also increased appreciably between 2008 and 2009. This finding is consistent with the finding of a 2009 Center for Retirement Research at Boston College survey, which found that 40 percent of workers age 45 to 59 reported that they were planning to retire later than they had planned prior to the downturn (Sass, Monk and Haverstick, 2010). If one believes that CogEcon respondents were able to correctly answer how their retirement plans had changed "as a result of the economic crisis," the CogEcon data imply that the economic crisis caused large increases in planned retirement age.

Consistent with the implications of the life-cycle hypothesis, Tobit regressions yield statistically-significant estimates reflecting a positive association between wealth losses and increases in expected retirement age. These estimates of the impact of wealth losses on retirement age, while not clearly causal, compare favorably with other recent studies. Estimates from my baseline reduced-form regression specification using the CogEcon data show that a loss in wealth that would take one additional year of work to regain is associated with an average of about a month's change in retirement age. For an older American with an average wealth loss, in terms of the number of years of work it would take to make up the loss from the crash, my estimates imply about that about four months' increase in retirement age may be attributable to the wealth loss. These estimates are roughly in line with simulation results presented in Gustman et al. [2009], which predicts an average increase in retirement age on the order of one and a half months as a result of wealth losses during the economic crisis.

Estimation results using Health and Retirement Study data also show an association between wealth losses and retirement expectations, with the average wealth loss implying a 1 and 1/4 percentage point increase in the probability of full-time work after age 62, and an increase of about 4 percentage points in the probability of full-time work after age 65, the latter statistically-significant at the 1 percent level. These estimates are similar to results from Goda, Shoven and Slavov's (2010) analysis based on HRS data from 2006 and 2008, which imply that a 40 percent decline in the S&P 500 (the average decline between the HRS 2008 Core and the 2009 Internet Survey) would be associated with a 5 percentage point increase in the probability of full-time work after age 62 or just over 1/2 percentage point increase in the probability if full-time work after age 65, the former statistically-significant at the 5 percent level. While emphasizing that wealth effects are likely outweighed in aggregate by the increased retirement rates of older unemployed workers, recent studies by Coile and Levine (2009) and Bosworth and Burtless (2010) have also found a negative relationship between recent wealth losses and retirement rates.

While my estimates of the impact of wealth losses on retirement timing are almost certainly subject to significant attenuation bias due to measurement error, the gap between the average reported change in retirement age in the CogEcon data (1.6 years) and the much smaller amount that can be explained by wealth losses also implies that heterogeneity in preferences, expectations about the future and other individual characteristics may also be important in determining the impact of wealth losses on changes in planned retirement age. My analyses suggest that wealth effects may be larger for individuals with moderate levels of wealth, and smaller for those with the highest levels of wealth. For individuals close to their pre-crash planned retirement ages, a year needed to regain lost wealth is associated with a larger increase in planned retirement age than for individuals who are farther from retirement. Additionally, individuals who are more optimistic about the rebound of the stock and labor markets over the next 1-2 years are less reactive to wealth losses and the crash, and individuals with more financial knowledge are less reactive to wealth losses than those with less financial knowledge. Interestingly, I did not find evidence that individuals in the worst labor markets were more likely to plan to hold on to their jobs for longer than individuals in better labor markets. It also appears to be the case that individuals are adjusting along more than one margin: respondents who adjusted their labor supply expectations were also likely to report decreased probabilities of leaving large bequests.

8 Figures and tables

Figure 1: Life-cycle saving and consumption

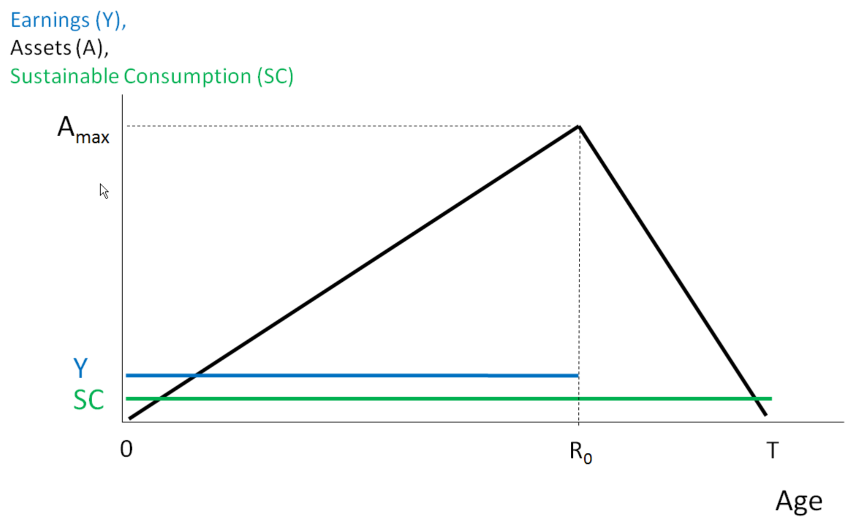

Figure 2: Optimal retirement choice

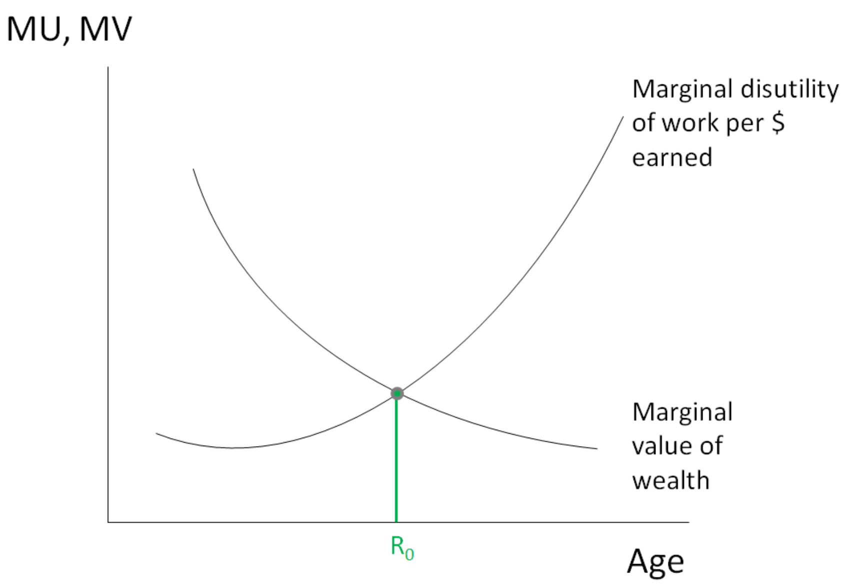

Figure 3: Optimal retirement choice after a wealth shock

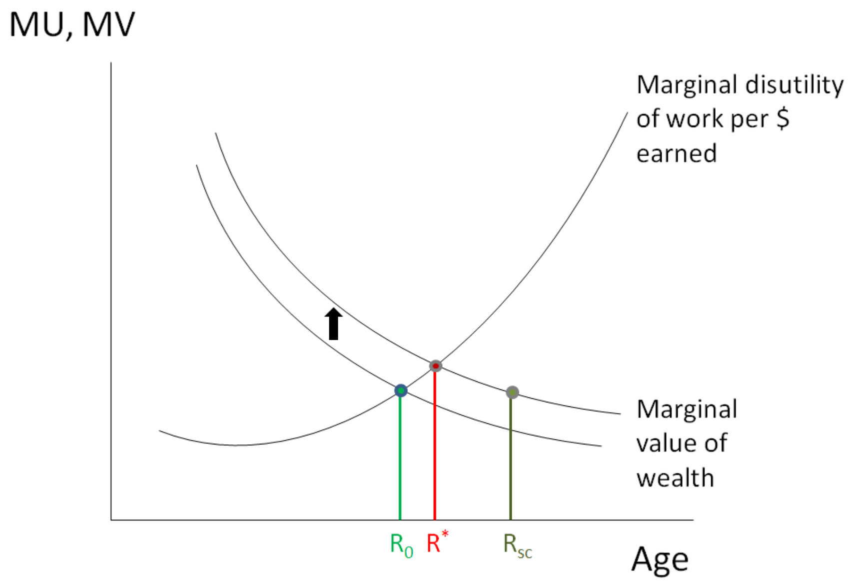

Figure 4: Life-cycle saving and consumption with variable retirement timing

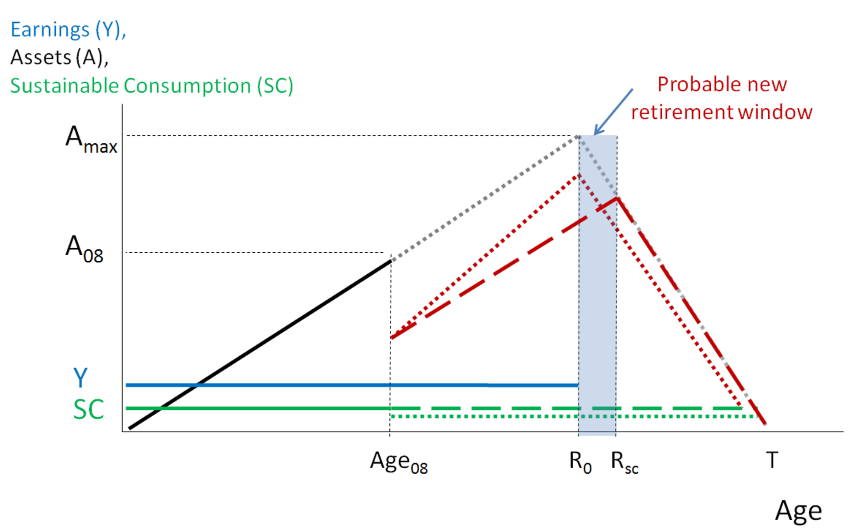

Figure 5: Dow Jones Industrial Average closing values

Source: Yahoo! Finance

Figure 6: Timeline of surveys and the Dow Jones Industrial Average

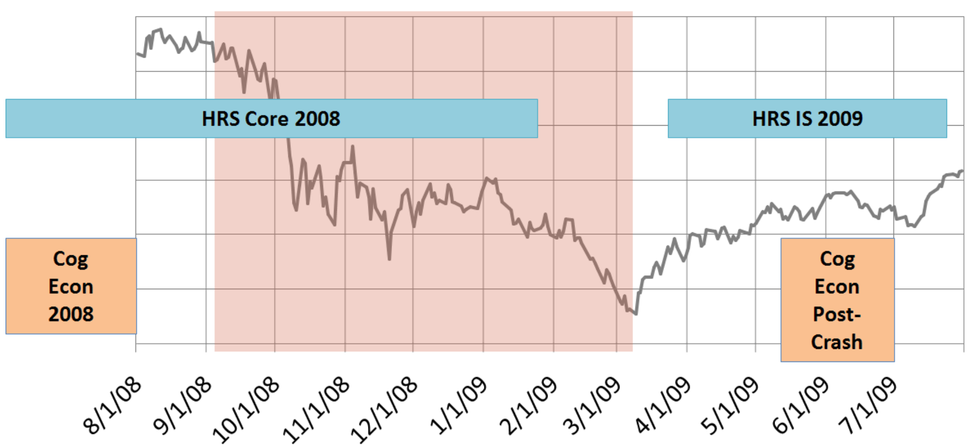

Figure 7: Changes in retirement age owing to crash (CogEcon sample)

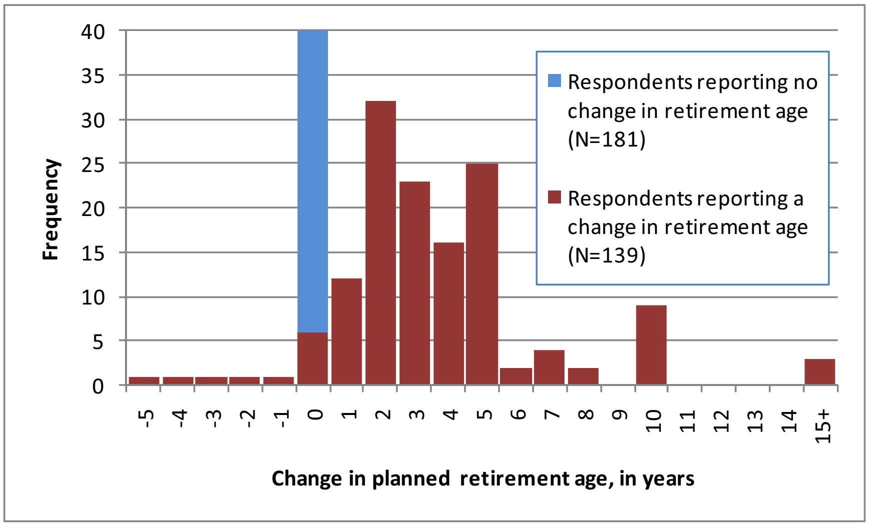

Figure 8: Cumulative distribution of expected retirement ages (CogEcon sample)

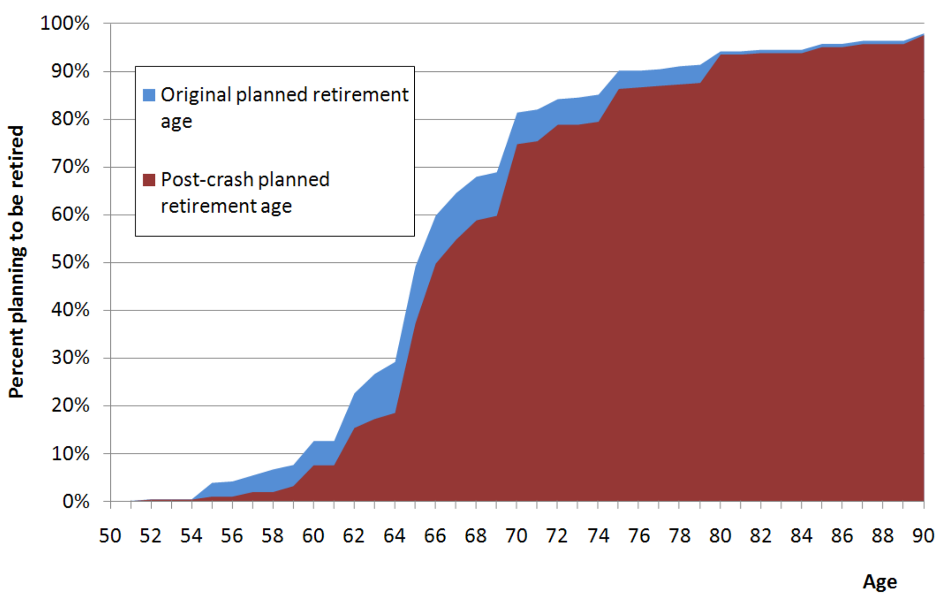

Figure 9: Comparison of retirement age needed to attain pre-crash consumption path and planned post-crash retirement age (CogEcon sample)

Figure 10: Comparison of changes in retirement age needed to attain pre-crash consumption path and reported changes (CogEcon sample)

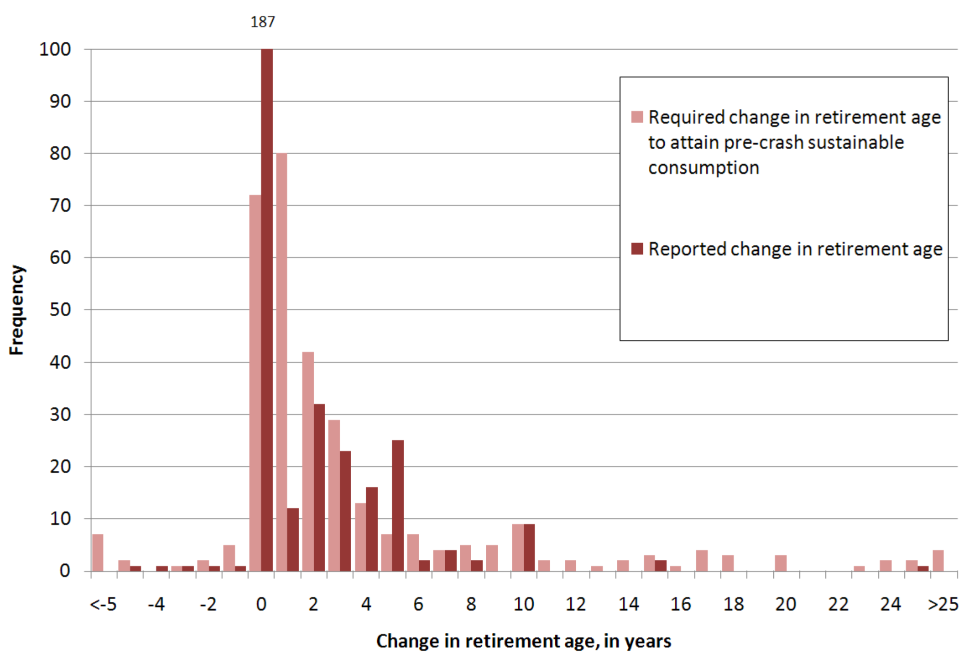

Figure 11: Extensive margin: Tobit prediction versus observed probability of increase in planned retirement age (CogEcon sample)

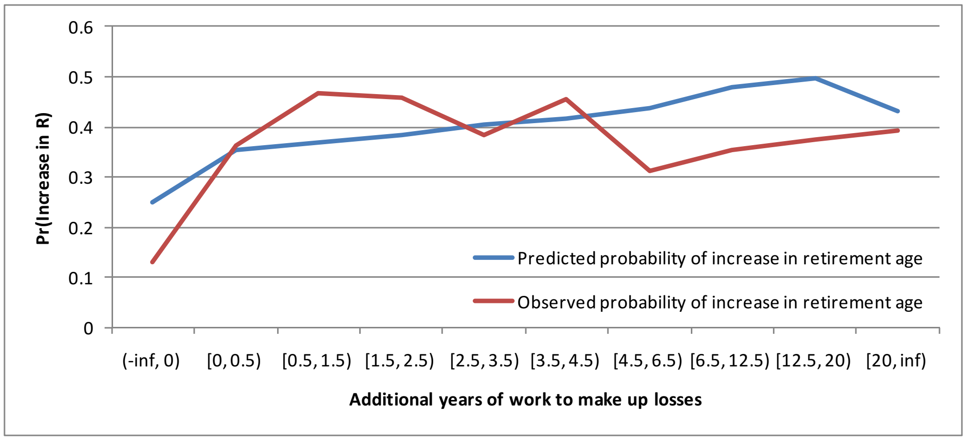

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is probability of an increase in the planned age of retirement. Lines are plotted by connecting the average for each bin.

Figure 12: Intensive margin: Tobit prediction versus observed increase in planned retirement age (CogEcon sample)

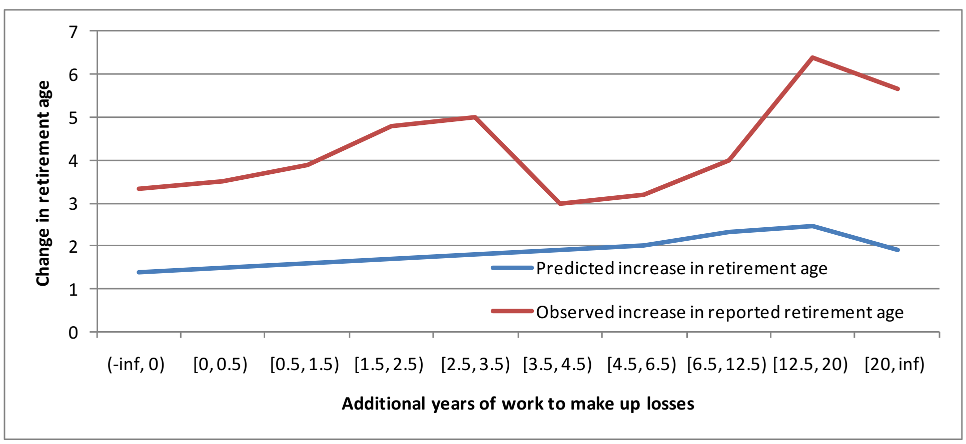

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis represents years of increase in the planned age of retirement. Lines are plotted by connecting the average for each bin.

Figure 13: Extensive margin: Tobit prediction versus observed probability of increase in Pr(FT62) (HRS <62 sample)

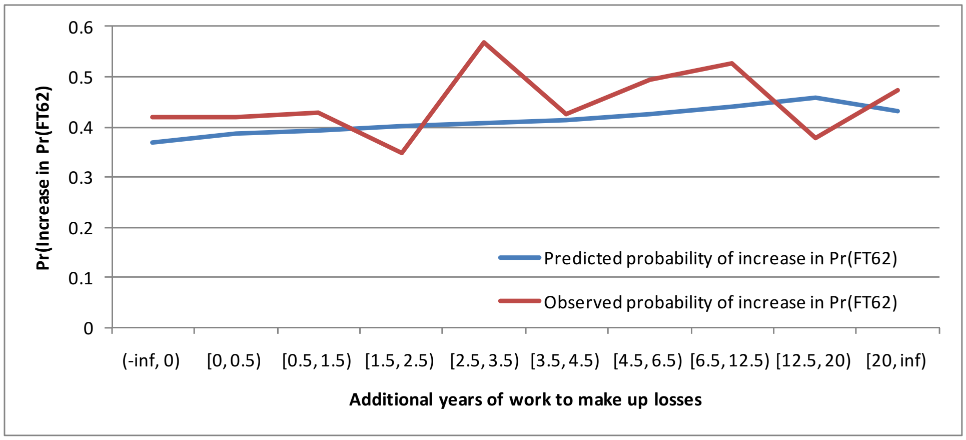

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is probability of an increase in the subjective probability of full-time work after age 62 per year of $R_{\overline{sc}} - R_0$. Lines are plotted by connecting the average for each bin.

Figure 14: Intensive margin: Tobit prediction versus observed increase in Pr(FT62) (HRS <62 sample)

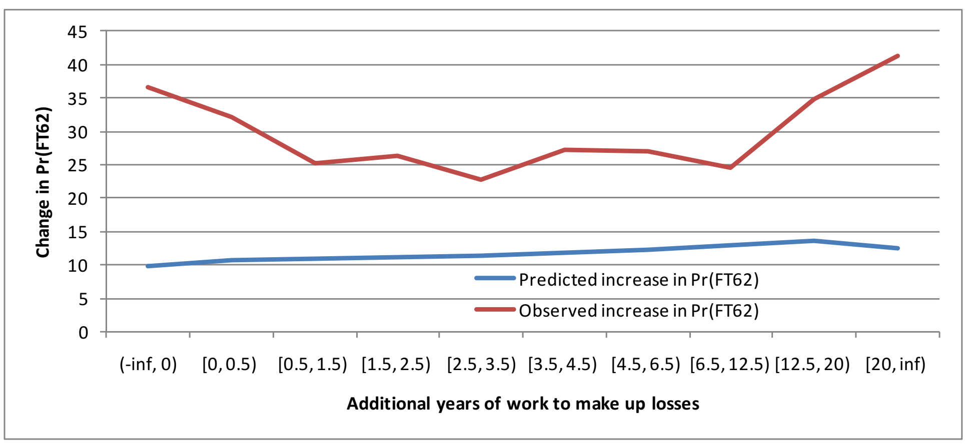

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is the increase in the subjective probability of full-time work after age 62 per year of $R_{\overline{sc}} - R_0$. Lines are plotted by connecting the average for each bin.

Figure 15: Extensive margin: Tobit prediction versus observed probability of increase in Pr(FT65) (HRS <65 sample)

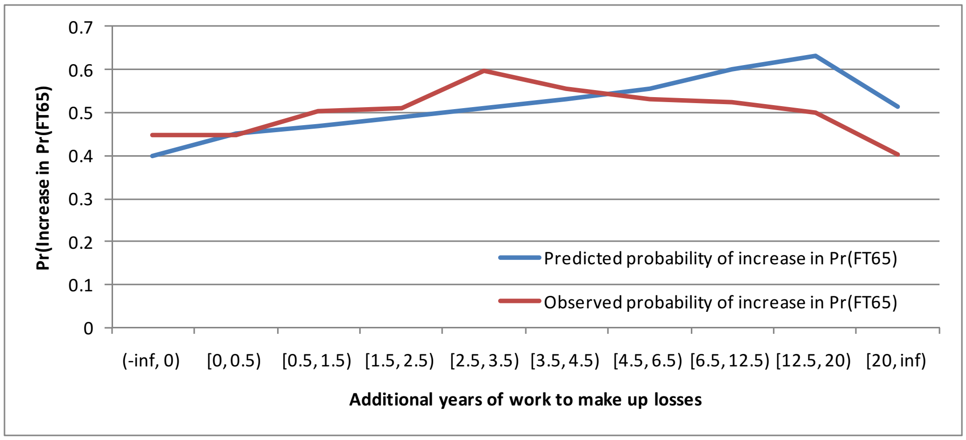

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is probability of an increase in the subjective probability of full-time work after age 65 per year of $R_{\overline{sc}} - R_0$. Lines are plotted by connecting the average for each bin.

Figure 16: Intensive margin: Tobit prediction versus observed increase in Pr(FT65) (HRS <65 sample)

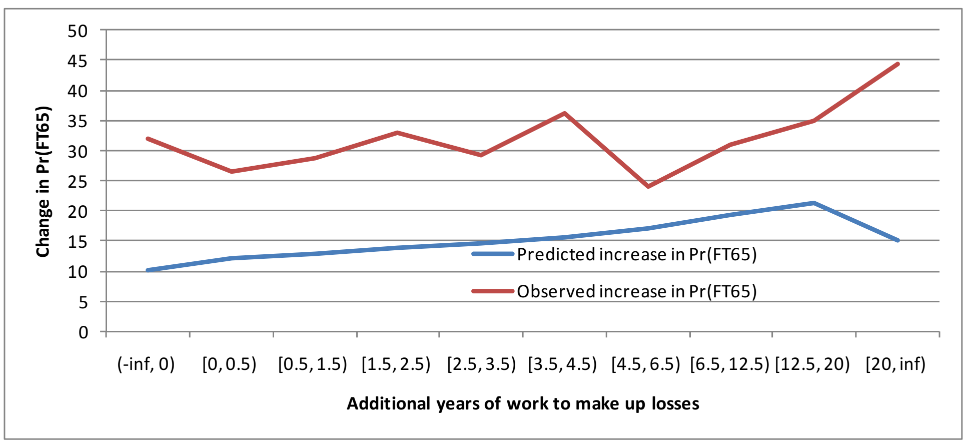

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is the increase in the subjective probability of full-time work after age 65 per year of $R_{\overline{sc}} - R_0$. Lines are plotted by connecting the average for each bin.

Figure 17: Average marginal effects by $R_{\overline{sc}} - R_0$ group (CogEcon sample)

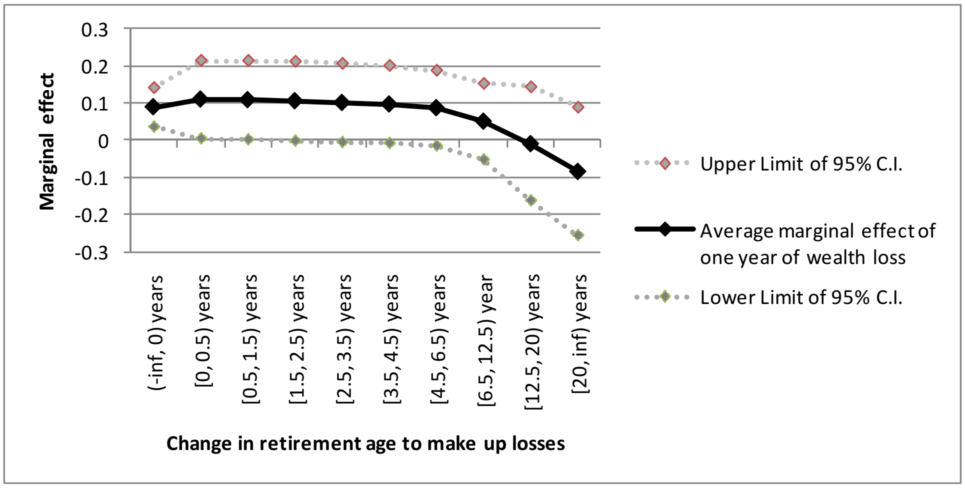

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Estimates based on results from regression shown in Column 4 of Table 6.

Figure 18: Average predicted change in retirement age by $R_{\overline{sc}} - R_0$ group (CogEcon sample)

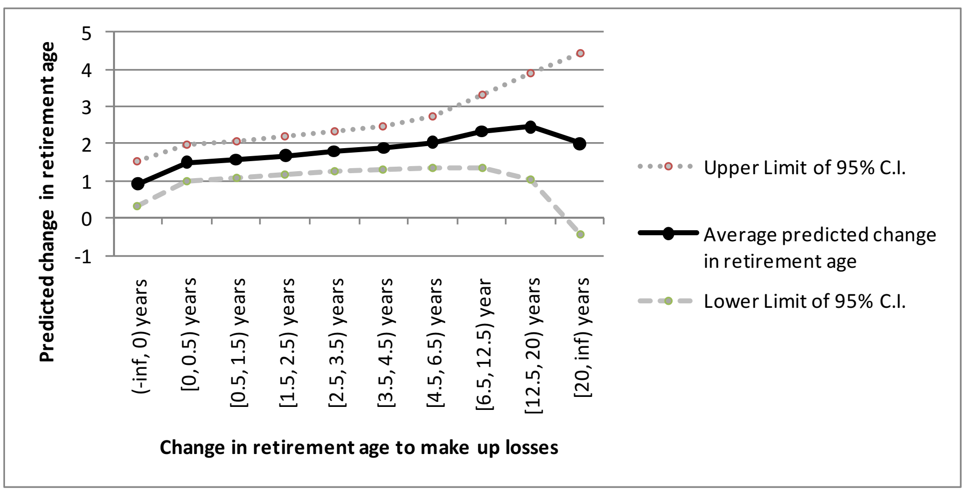

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Estimates based on results from regression shown in Column 4 of Table 6.

Figure 19: Average marginal effects by $R_{\overline{sc}} - R_0$ group (HRS <62 sample)

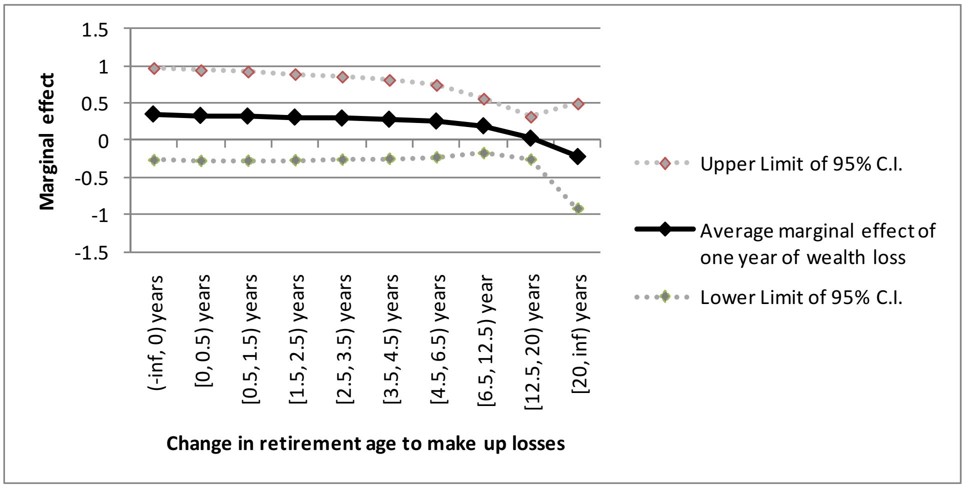

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is percentage point change in the probability of full-time work after age 62 per year of $R_{\overline{sc}} - R_0$. Estimates based on results from regression shown in Column 2 of Table 7.

Figure 20: Average predicted change in probability of full-time work by $R_{\overline{sc}} - R_0$ group (HRS <62 sample)

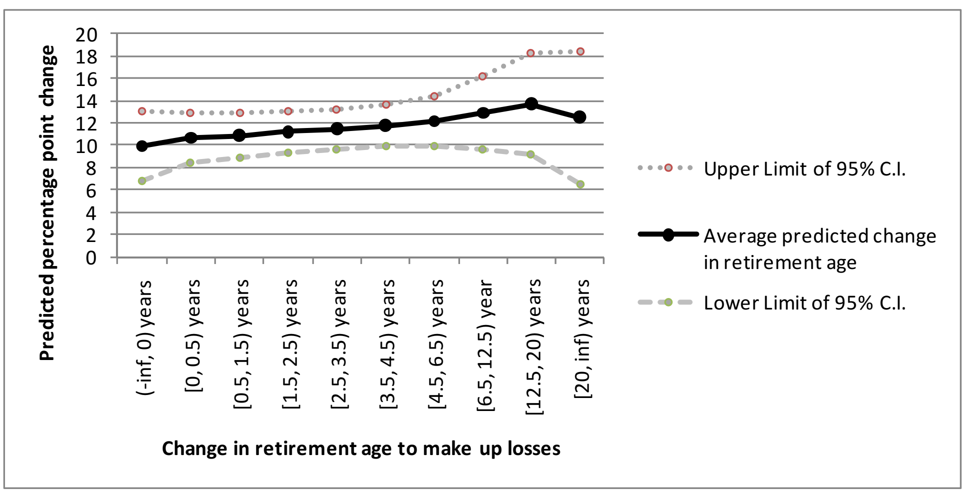

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is percentage point change in the probability of full-time work after age 62 per year of $R_{\overline{sc}} - R_0$. Estimates based on results from regression shown in Column 2 of Table 7.

Figure 21: Average marginal effects by $R_{\overline{sc}} - R_0$ group (HRS <65 sample)

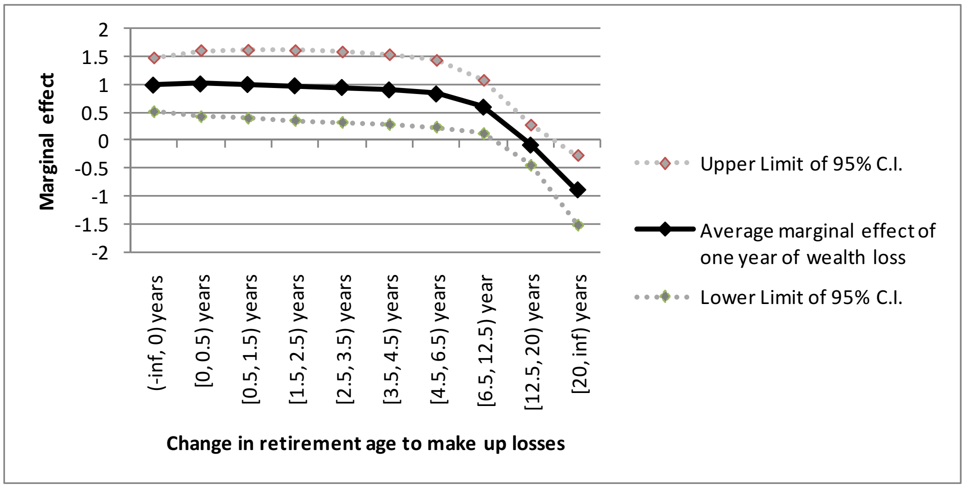

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is percentage point change in the probability of full-time work after age 62 per year of $R_{\overline{sc}} - R_0$. Estimates based on results from regression shown in Column 2 of Table 8.

Figure 22: Average predicted change in probability of full-time work by $R_{\overline{sc}} - R_0$ group (HRS <65 sample)

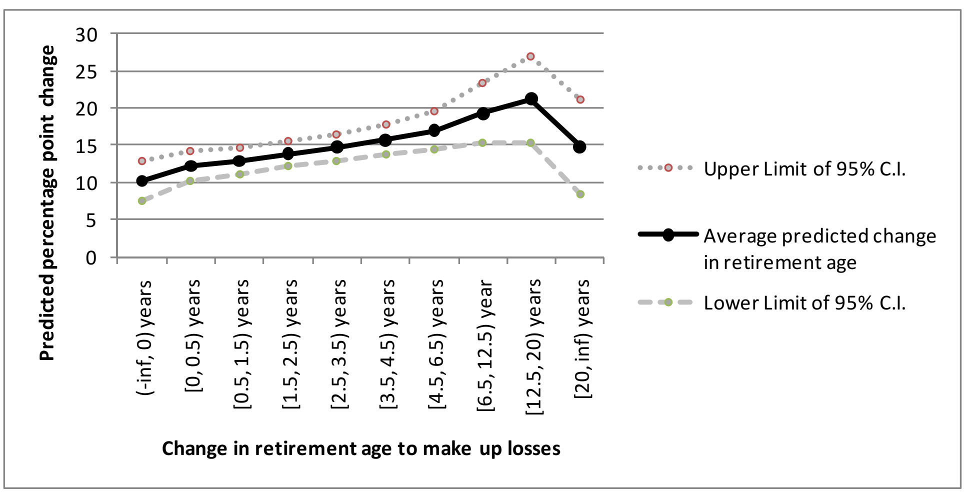

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is percentage point change in the probability of full-time work after age 62 per year of $R_{\overline{sc}} - R_0$. Estimates based on results from regression shown in Column 2 of Table 8.

Table 1: Descriptive statistics

CogEcon sample (N=320)	Mean	Median	St. Dev.
Proportion Female	0.52	–	–
Proportion Single	0.23	–	–
Education (years)	14.93	16	2.01
Annual Earnings	$79,880	$52,023	$238,967
Age at Post-Crash Survey	60.61	59.88	6.30
Planned Retirement Age as of 2008	67.79	66	9.30
HRS <62 sample (N=589)	**Mean**	**Median**	**St. Dev.**
Proportion Female	0.55	–	–
Proportion Single	0.22	–	–
Education (years)	14.62	15	1.99
Annual Earnings	$59,943	$46,000	$69,216
Age at Post-Crash Survey	58.44	58.41	1.80
Planned Retirement Age as of 2008 (imputed)	64.34	65	2.66
Planned Retirement Age as of 2008 (not imputed, N=136)	63.57	64	3.15
Sample: HRS <65 (N=594)	**Mean**	**Median**	**St. Dev.**
Proportion Female	0.55	–	–
Proportion Single	0.22	–	–
Education (years)	14.60	15	1.99
Annual Earnings	$59,886	$46,000	$69,966
Age at Post-Crash Survey	58.50	58.50	1.86
Planned Retirement Age as of 2008 (imputed)	64.38	65	2.66
Planned Retirement Age as of 2008 (not imputed, N=136)	63.63	64	3.12

Table 2: Sustainable consumption levels, pre- and post-crash

Pre-Crash Sustainable Consumption			
Sample:	CogEcon	HRS <62	HRS <65
Mean	$99,071	$78,015	$77,660
25th %	$40,083	$41,954	$41,826
Median	$63,112	$63,639	$63,853
75th %	$99,101	$94,092	$94,557
Post-Crash Sustainable Consumption			
Sample:	CogEcon	HRS <62	HRS <65
Mean	$90,523	$71,288	$70,939
25th %	$37,351	$40,268	$40,099
Median	$58,440	$58,702	$58,806
75th %	$91,994	$87,726	$87,726
Observations	320	589	594

Table 3: Changes in sustainable consumption levels, 2008 to 2009

Sample:	CogEcon	HRS <62	HRS <65
Mean	-8.65%	-6.67%	-6.70%
25th %	-13.65%	-11.02%	-11.02%
Median	-4.62%	-4.96%	-4.97%
75th %	-1.86%	-1.99%	-1.99%
Observations	320	589	594

Table 4: Extra work years needed to make up lost wealth ($R_{\overline{sc}} - R_0$)

Sample:	CogEcon		HRS <62		HRS <65	
	All	$\triangle R$>0	All	$\triangle Pr(FT62)$>0	All	$\triangle Pr(FT65)$>0
Mean	3.72	4.10	4.92	5.02	4.99	4.84
25th %	0.52	0.89	0.65	0.74	0.65	0.74
Median	1.64	1.66	1.88	1.98	1.88	2.06
75th %	4.11	3.90	4.91	5.44	4.95	5.15
St. Dev.	7.49	6.17	7.99	7.80	8.09	7.52

Table 5: Changes in subjective probabilities of full-time work in HRS, 2006-2008 and 2008-2009

	$\triangle Pr(FT62)$		$\triangle Pr(FT65)$	
	2006 to 2008	2008 to 2009	2006 to 2008	2008 to 2009
Mean	8.7 p.p.	3.5 p.p.	6.5 p.p.	8.1 p.p.
Median	0 p.p.	0 p.p.	0 p.p.	2 p.p.
75th %	20 p.p.	19.5 p.p.	20 p.p.	25 p.p.
Observations	580	580	585	585

Table 6: Impact of wealth losses on retirement age (CogEcon sample)

	(1)	(2)	(3)	(4)
Specification:	OLS	OLS	Tobit	Tobit
$R_{\overline{sc}} - R_0$	0.058	0.057	0.110	0.311*
	(0.043)	(0.038)	(0.091)	(0.163)
$(R_{\overline{sc}} - R_0)^2$	–	0.000	–	-0.010
	–	(0.001)	–	(0.008)
Constant	1.434***	1.422***	-2.184***	-2.384***
	(0.224)	(0.241)	(0.829)	(0.861)
Sigma	–	–	6.258***	6.245***
	–	–	(0.947)	(0.950)
Observations	320	320	320	320
Number uncensored obs.	–	–	128	128
R^2	0.017	0.017	–	–
Pseudo-R^2	–	–	0.003	0.006
Log-Likelihood	–	–	-459.90	-458.70
F-test (H_0 : Coefs. jointly 0)	1.89	3.30	1.45	1.89
Prob >F	0.17	0.04	0.23	0.15
Marginal effect at	0.058	0.059	0.043	0.086**
mean of $R_{\overline{sc}} - R_0$ (3.721)	(0.042)	(0.044)	(0.036)	(0.043)

Notes: Dependent variable is reported change in retirement age. All analyses include CogUSA sampling weights. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The censoring point for Tobit regressions is 0.

Table 7: Impact of wealth losses on probability of full-time work after age 62 (HRS <62 sample)

Specification:	(1) Tobit	(2) Tobit	(3) Tobit	(4) Tobit
$R_{\overline{sc}} - R_0$	0.193	0.871	0.494	1.474
	(0.267)	(0.854)	(0.765)	(2.596)
$(R_{\overline{sc}} - R_0)^2$		-0.025		-0.032
		(0.030)		(0.081)
Constant	-10.28***	-11.46***	-18.96**	-21.20**
	(2.733)	(3.076)	(8.733)	(10.670)
Sigma	39.13***	39.15***	56.68***	56.75***
	(2.357)	(2.367)	(5.725)	(5.720)
Observations	589	589	139	139
Number uncensored obs.	247	247	56	56
Pseudo-R^2	0.000	0.001	0.001	0.001
Log-Likelihood	-1.24×10^7	-1.24×10^7	-2.99×10^7	-2.99×10^7
F-test (H_0: Coefs. jointly 0)	0.524	0.618	0.417	0.283
Prob >F	0.470	0.539	0.519	0.754
Marginal effect at	0.079	0.246	0.191	0.426
mean of $R_{\overline{sc}} - R_0$ (4.919)	(0.109)	(0.224)	(0.295)	(0.651)

Notes: Dependent variable is the change in the probability of full-time work after age 62, ${}_{08}\triangle_{09}Pr(FT62)$. Censoring point is zero in all regressions. All analyses include 2008 Core sampling weights. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 8: Impact of wealth losses on probability of full-time work after age 65 (HRS <65 sample)

Specification:	(1) Tobit	(2) Tobit	(3) Tobit	(4) Tobit
$R_{\overline{sc}} - R_0$	0.276	2.291***	-0.168	2.210
	(0.243)	(0.757)	(0.415)	(1.451)
$(R_{\overline{sc}} - R_0)^2$		-0.0743***		-0.0797*
		(0.026)		(0.045)
Constant	-1.673	-5.044*	-1.068	-6.360
	(2.368)	(2.735)	(4.639)	(5.850)
Sigma	36.08***	35.78***	35.62***	35.52***
	(1.797)	(1.823)	(3.928)	(4.031)
Observations	594	594	140	140
Number uncensored obs.	298	298	73	73
Pseudo-R^2	0.001	0.003	0.000	0.003
Log-Likelihood	-1.44×10^7	-1.44×10^7	-3.39×10^6	-3.38×10^{-6}
F-test (H_0: Coefs. jointly 0)	1.286	4.578	0.164	1.803
Prob >F	0.257	0.011	0.686	0.169
Marginal effect at	0.137	0.736***	-0.08	0.661
mean of $R_{\overline{sc}} - R_0$ (4.989)	(0.121)	(0.235)	(0.198)	(0.456)

Notes: Dependent variable is the change in the probability of full-time work after age 65, ${}_{08}\triangle_{09}Pr(FT62)$. Censoring point is zero in all regressions. All analyses include Core 2008 sampling weights. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 9: Comparison of results from different specifications and samples

Dataset	Sample	Specification	Dependent variable	Includes $(R_{\overline{sc}} - R_0)^2$	Marginal effect[a]	Mean of $(R_{\overline{sc}} - R_0)$	Implied[b] effect on $(R^* - R_0)$
CogEcon	–	OLS	$(R_{09} - R_0)$	No	0.058	3.72	2.6 months
CogEcon	–	OLS	$(R_{09} - R_0)$	Yes	0.059	3.72	2.7 months
CogEcon	–	Tobit	$(R_{09} - R_0)$	No	0.043	3.72	2 months
CogEcon	–	Tobit	$(R_{09} - R_0)$	Yes	0.086**	3.72	3.9 months
HRS	Full	Tobit	${}_{08}\triangle_{09}Pr(FT62)$	No	0.079	4.92	3 days
HRS	Full	Tobit	${}_{08}\triangle_{09}Pr(FT62)$	Yes	0.246	4.92	10 days
HRS	Rest.	Tobit	${}_{08}\triangle_{09}Pr(FT62)$	No	0.191	4.92	8 days
HRS	Rest.	Tobit	${}_{08}\triangle_{09}Pr(FT62)$	Yes	0.426	4.92	18 days
HRS	Full	Tobit	${}_{08}\triangle_{09}Pr(FT65)$	No	0.137	4.99	5 days
HRS	Full	Tobit	${}_{08}\triangle_{09}Pr(FT65)$	Yes	0.736***	4.99	25 days
HRS	Rest.	Tobit	${}_{08}\triangle_{09}Pr(FT65)$	No	-0.08	4.99	-3 days
HRS	Rest.	Tobit	${}_{08}\triangle_{09}Pr(FT65)$	Yes	0.661	4.99	22 days

[a] Average estimated effect of a wealth loss that would take one additional year of work to make up, in terms of sustainable consumption.
[b] HRS marginal effects are translated to time metric by multiplying by 8.5 and 6.7 days per percentage point for the age-62 and age-65 regressions, respectively.

Table 10: Comparison of Tobit, probit and Cragg models (CogEcon sample)

	(1)	(2)	(3)
Specification:	Tobit	Probit	Truncated
Dependent variable:	$\triangle R$	$I_{\triangle R>0}$	$\triangle R$
$R_{\overline{sc}} - R_0$	0.311*	0.047*	0.575
	(0.163)	(0.025)	(0.680)
$(R_{\overline{sc}} - R_0)^2$	-0.010	-0.002*	-0.005
	(0.008)	(0.001)	(0.023)
Constant	-2.384***	-0.339***	-6.463
	(0.861)	(0.108)	(15.633)
Sigma	6.245***		7.003
	(0.950)		(4.894)
Observations	320	320	128
Log-Likelihood	-458.7	-187.5	-268.3

Notes: Dependent variable in Tobit and truncated normal specifications is reported change in retirement age. In probit specification, dependent variable is an indicator that is equal to one if retirement age increased, and zero otherwise. Censoring point for Tobit and truncated regressions is 0. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 11: Comparison of Tobit, probit and Cragg models (HRS samples)

	Probability of full-time work after age 62			Probability of full-time work after age 65		
	(1)	(2)	(3)	(4)	(5)	(6)
Specification:	Tobit	Probit	Truncated	Tobit	Probit	Truncated
Dependent variable:	$\triangle Pr(FT62)$	$I_{>0}$	$\triangle Pr(FT62)$	$\triangle Pr(FT65)$	$I_{>0}$	$\triangle Pr(FT65)$
$R_{\overline{sc}} - R_0$	0.871	0.037	-2.766	2.291***	0.065***	1.648
	(0.854)	(0.023)	(3.568)	(0.757)	(0.024)	(1.257)
$(R_{\overline{sc}} - R_0)^2$	-0.025	-0.001*	0.151	-0.074***	-0.002***	-0.030
	(0.030)	(0.001)	(0.123)	(0.026)	(0.001)	(0.045)
Constant	-11.464***	-0.288***	-46.400	-5.044*	-0.134*	3.624
	(3.076)	(0.077)	(30.778)	(2.735)	(0.076)	(7.792)
Sigma	39.146***		50.211***	35.783***		31.651***
	(2.367)		(8.801)	(1.823)		(3.369)
Observations	589	589	247	594	594	298
Log-Likelihood	-1.24×10^7	-3.43×10^6	-8.94×10^6	-1.44×10^7	-3.51×10^6	-1.08×10^7

Notes: In probit specification, dependent variable is an indicator that is equal to one if the probability of full-time work past the reference age (62 in Column 2, and 65 in Column 5) increased between 2008 and 2009, and zero otherwise. Censoring point for Tobit and truncated regressions is 0. Standard errors in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 12: Robustness check excluding primary home from total wealth

	(1)	(2)	(3)
Sample	CogEcon sample	HRS <62 sample	HRS <65 sample
Dependent variable	$R_{09} - R_0$	$\triangle Pr(FT62)$	$\triangle Pr(FT65)$
Specification:	Tobit	Tobit	Tobit
$R_{\overline{sc}} - R_0$	0.203	0.560	1.911**
	(0.143)	(1.114)	(0.902)
$(R_{\overline{sc}} - R_0)^2$	-0.006	(0.009)	-0.069**
	(0.006)	(0.038)	(0.031)
Constant	-2.343***	-10.92***	-3.041
	(0.874)	(3.120)	(2.631)
Sigma	6.270***	39.16***	35.97***
	(0.956)	(2.344)	(1.868)
Observations	320	591	595
Number uncensored obs.	128	248	299
Pseudo-R^2	0.003	0.001	0.002
Log-Likelihood	-459.7	-1.25×10^7	-1.44×10^7
F-test (H_0 : Coefs. jointly 0)	1.09	0.61	2.57
Prob >F	0.34	0.54	0.08
Mean of $R_{\overline{sc}} - R_0$	5.14	3.96	4.01
Avg. marginal effect at mean	0.054	0.197	0.674
	(0.035)	(0.321)	(0.324)
Implied p.p. change at mean	–	0.78	2.70
Implied change in retirement age (in days) at mean	100.97	6.63	18.11

Notes: CogEcon analyses include CogUSA sampling weights; HRS analyses include 2008 Core sampling weights. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The censoring point for all regressions is 0.

Table 13: Robustness check excluding expected bequests from total wealth

	(1)	(2)
Sample	HRS <62 sample	HRS <65 sample
Dependent variable	${}_{08}\triangle_{09}Pr(FT62)$	${}_{08}\triangle_{09}Pr(FT65)$
Specification:	Tobit	Tobit
$R_{\overline{sc}} - R_0$	0.832	2.260***
	(0.854)	(0.757)
$(R_{\overline{sc}} - R_0)^2$	-0.023	-0.0722***
	(0.030)	(0.026)
Constant	-11.40***	-4.987*
	(3.050)	(2.716)
Sigma	39.12***	35.77***
	(2.365)	(1.817)
Observations	588	593
Number uncensored obs.	247	298
Pseudo-R^2	0.001	0.003
Log-Likelihood	-1.24×10^7	-1.44×10^7
F-test (H_0 : Coefs. jointly zero)	0.62	4.48
Prob >F	0.54	0.01
Mean of $R_{\overline{sc}} - R_0$	4.76	4.82
Avg. marginal effect at mean	0.242	0.739
	(0.226)	(0.237)
Implied p.p. change at mean	1.15	3.56
Implied change in retirement age (in days) at mean	9.79	23.87

Notes: HRS analyses include 2008 Core sampling weights. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The censoring point for all regressions is 0.

Table 14: Regressions with wealth terciles (CogEcon sample)

	(1)	(2)
Wealth Tercile Indicator		
2nd	1.52	0.795
	(1.08)	(1.19)
3rd (highest wealth)	-3.532**	-2.408
	(1.52)	(1.66)
$(R_{\overline{sc}} - R_0)$	0.489***	0.0161
	(0.19)	(0.49)
Wealth Tercile Indicator $\times (R_{\overline{sc}} - R_0)$		
2nd		0.0161
		(0.49)
3rd (highest wealth)		-0.197
		(0.51)
$(R_{\overline{sc}} - R_0)^2$	-0.0148*	-0.0341
	(0.01)	(0.03)
Wealth Tercile Indicator $\times (R_{\overline{sc}} - R_0)^2$		
2nd		0.023
		(0.03)
3rd (highest wealth)		0.0165
		(0.03)
Constant	-2.228**	-2.077**
	(0.95)	(1.01)
Sigma	5.933***	5.795***
	(0.93)	(0.93)
Observations	320	320
Number uncensored obs	128	128
Pseudo-R^2	0.0272	0.034
Log-Likelihood	-448.8	-445.7
F-test: All jointly=0	4.071	2.552
Prob > F	0.003	0.011
Mean of $(R_{\overline{sc}} - R_0)$, by Wealth Tercile		
1st	1.37	1.37
2nd	3.17	3.17
3rd (highest wealth)	6.81	6.81
Marginal effect of 1 yr of $(R_{\overline{sc}} - R_0)$, by Wealth Tercile		

Table 2.14: Regressions with wealth terciles (CogEcon sample) *(continued)*

	(1)	(2)
1st	0.169**	0.18
	(0.066)	(0.14)
2nd	0.181***	0.238***
	(0.062)	(0.09)
3rd (highest wealth)	0.058**	0.03
	(0.023)	(0.04)

Notes: Results from Tobit regressions, with dependent variable reported change in retirement age, censored from below at zero. All analyses include CogUSA sampling weights. Robust standard errors in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$. In column 1, marginal effects at each tercile are statistically significantly different from one another at the 5 percent level ($\chi^2(2) = 6.29$, p-value=0.04). In column 2, marginal effects at each tercile are statistically significantly different from one another at the 10 percent level ($\chi^2(2) = 5.31$, p-value=0.07).

Table 15: Regressions with pre-crash time to retirement (CogEcon sample)

	(1)		(2)	
	Estimate	Std. Error	Estimate	Std. Error
2-5 years until R_0	0.116	(1.24)	1.087	(1.63)
5-10 years until R_0	-2.862**	(1.27)	-2.366	(1.72)
More than 10 years until R_0	-3.524**	(1.36)	-2.666	(1.66)
$(R_{\overline{sc}} - R_0)$	0.316*	(0.17)	0.584**	(0.24)
(2-5 years until R_0)$\times (R_{\overline{sc}} - R_0)$			-0.649	(0.63)
(5-10 years until R_0)$\times (R_{\overline{sc}} - R_0)$			0.016	(0.52)
(More than 10 years until R_0)$\times (R_{\overline{sc}} - R_0)$			-0.241	(0.30)
$(R_{\overline{sc}} - R_0)^2$	-0.011	(0.01)	-0.018	(0.01)
(2-5 years until R_0)$\times (R_{\overline{sc}} - R_0)^2$			0.026	(0.02)
(5-10 years until R_0)$\times (R_{\overline{sc}} - R_0)^2$			-0.014	(0.02)
(More than 10 years until R_0)$\times (R_{\overline{sc}} - R_0)^2$			0.0004	(0.01)
Constant	-0.474	(0.86)	-1.508	(1.13)
Sigma	6.185***	(1.00)	6.112***	(0.97)
Observations	320		320	
Number uncensored obs	128		128	
Log-Likelihood	-451.7		-449.1	
Pseudo-R^2	0.021		0.027	
F-test: All coefficients jointly=0	3.394		3.242	
Prob >F	0.005		0.000	

Table 2.15: Regressions with pre-crash time to retirement (CogEcon sample) *(continued)*

	(1)		(2)	
	Estimate	Std. Error	Estimate	Std. Error
Mean of $(R_{\overline{sc}} - R_0)$: less than 2 yrs until R_0	3.21			
Mean of $(R_{\overline{sc}} - R_0)$: 2-5 yrs until R_0	5.11		5.11	
Mean of $(R_{\overline{sc}} - R_0)$: 5-10 yrs until R_0	4.17		4.17	
Mean of $(R_{\overline{sc}} - R_0)$: >10 yrs until R_0	2.51		2.51	
Avg. marginal effect if <2 yrs until R_0	0.119**	(0.06)	0.216***	(0.07)
Avg. marginal effect if 2-5 yrs until R_0	0.1*	(0.05)	0.01	(0.18)
Avg. marginal effect if 5-10 yrs until R_0	0.067**	(0.03)	0.12	(0.10)
Avg. marginal effect if >10 yrs until R_0	0.056*	(0.03)	0.05	(0.04)

Notes: Results from Tobit regressions, with dependent variable reported change in retirement age, censored from below at zero. Excluded category is "less than 2 years until retirement." Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Both regressions use 320 observations, of which 128 are uncensored. All analyses include CogUSA sampling weights. In column 1, marginal effects for all groups are not statistically significantly different from one another at standard levels ($\chi^2(3) = 3.38$, p-value=0.34). In column 2, marginal effects for all groups are not statistically significantly different from one another at standard levels ($\chi^2(3) = 4.4$, p-value=0.22). However, in columns 1 and 2, the marginal effects for the group closest to retirement are statistically significantly different from the marginal effects for the group farthest from retirement at the 10 and 5 percent levels, respectively.

Table 16: Regressions with expectations (CogEcon sample)

Optimism Indicators:	(1) All	(2) Stock Market	(3) Labor Market	(4) Housing Market
$(R_{\overline{sc}} - R_0)$	0.2	0.394***	0.441**	0.28
	(0.13)	(0.14)	(0.21)	(0.22)
$(R_{\overline{sc}} - R_0)^2$	-0.00708	-0.0124*	-0.015	-0.00826
	(0.01)	(0.01)	(0.01)	(0.01)
Stock Market Optimism	-1.966*	-0.0824		
	(1.08)	(1.20)		
Labor Market Optimism	-1.127		-1.127	
	(1.14)		(1.14)	
Housing Market Optimism	0.0606			-0.0121
	(1.11)			(1.48)
(Optimism Indicator)$\times$ $(R_{\overline{sc}} - R_0)$		-0.837**	0.928	-0.277
		(0.35)	(1.14)	(0.40)
(Optimism Indicator)$\times$ $(R_{\overline{sc}} - R_0)^2$		0.0221	-0.248	-0.00103
		(0.02)	(0.16)	(0.02)
Constant	-0.782	-1.273*	-1.728**	-2.056**
	(0.74)	(0.76)	(0.87)	(0.88)
Sigma	5.171***	5.067***	6.072***	6.289***
	(0.52)	(0.49)	(0.95)	(0.96)
Observations	291	293	307	305
Number uncensored obs	116	117	124	123
Pseudo-R^2	-380.1	-379	-433	-438.9
Log-Likelihood	0.0134	0.0219	0.0243	0.00574

Table 2.16: Regressions with expectations (CogEcon sample) *(continued)*

Optimism Indicators:	(1) All	(2) Stock Market	(3) Labor Market	(4) Housing Market
F-test: All coefficients jointly=0	1.859	3.181	1.926	0.616
Prob > F	0.102	0.00821	0.0899	0.688
Mean of $(R_{\overline{sc}} - R_0)$ at Optimism Indicator=1		4.36	3.46	3.99
Mean of $(R_{\overline{sc}} - R_0)$ at Optimism Indicator=0		3.46	3.91	3.69
Avg. marginal effect at Optimism Indicator=1		-0.111	0.092	0.060
		(0.076)	(0.144)	(0.052)
Avg. marginal effect at Optimism Indicator=0		0.129***	0.132**	-0.397
		(0.042)	(0.058)	(0.455)

Notes: All regressions are Tobits with censoring point set at zero. Dependent variable is the reported change in retirement age. All analyses include CogUSA sampling weights. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. In column 2, the average marginal effect for the group that is optimistic about the stock market is statistically significantly different from that of the pessimistic group at the 1 percent level ($\chi^2(1) = 7.75$, p-value=0.005). In column 3, the average marginal effect for the group that is optimistic about the labor market is not statistically significantly different from that of the pessimistic group ($\chi^2(1) = 0.06$, p-value=0.80). In column 4, the average marginal effect for the group that is optimistic about the housing market is not statistically significantly different from that of the pessimistic group ($\chi^2(1) = 0.75$, p-value=0.39).

Table 17: Regressions with unemployment rate change groups (CogEcon sample)

	(1)	(2)
	(0.159)	(0.137)
$(R_{\overline{sc}} - R_0)^2$	-0.0102	0.00006
	(0.00776)	(0.0101)
$\triangle UE \in [3, 4]$ p.p.	0.587	2.054
	(1.271)	(1.594)
$\triangle UE > 4$ p.p.	0.438	0.710
	(1.299)	(1.535)
$(\triangle UE \in [3, 4]$ p.p.$)\times(R_{\overline{sc}} - R_0)$		-0.334
		(0.570)
$(\triangle UE > 4$ p.p.$)\times(R_{\overline{sc}} - R_0)$		0.159
		(0.382)
$(\triangle UE \in [3, 4]$ p.p.$)\times(R_{\overline{sc}} - R_0)^2$		-0.0037
		(0.0242)
$(\triangle UE > 4$ p.p.$)\times(R_{\overline{sc}} - R_0)\,2$		-0.0179
		(0.0162)
Constant	-2.637**	-2.916**
	(1.126)	(1.192)
Sigma	6.240***	6.107***
	(0.941)	(0.919)
Observations	320	320
Number uncensored obs	128	128

Table 2.17: Regressions with unemployment rate change groups (CogEcon sample) *(continued)*

	(1)	(2)
$(R_{\overline{sc}} - R_0)$	0.307*	0.269*
Log-Likelihood	-459	-456
Pseudo-R^2	0.006	0.013
F-test	1.040	1.019
Prob > F	0.387	0.421
Mean for $\triangle UE < 3$ p.p. group	2.23	2.23
Mean for $\triangle UE \in [3,4]$ p.p. group	5.33	5.33
Mean at for $\triangle UE > 4$ p.p. group	4.27	4.27
Avg. marginal effect for$\triangle UE < 3$ p.p. group	0.083**	0.099
	(0.042)	(0.050)
Avg. marginal effect for $\triangle UE \in [3,4]$ p.p. group	0.084*	-0.037
	(0.045)	(0.167)
Avg. marginal effect for$\triangle UE > 4$ p.p. group	0.086*	0.110
	(0.045)	(0.096)

Notes: All regressions are Tobits with censoring point set at zero. Dependent variable is the reported change in retirement age. Excluded category is ($\triangle UE < 50\%$) All analyses include CogUSA sampling weights. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The marginal effects are not statistically significantly different by group at any standard level of significance ($\chi^2(2) = 0.59$, p-value=0.75, and $\chi^2(2) = 0.20$, p-value=0.91, respectively).

Table 18: Regressions with financial knowledge and fluid intelligence (CogEcon sample)

Tercile Indicators Used:	(1) Financial Knowledge (FK)	(2) Financial Knowledge (FK)	(3) Fluid Intelligence (FI)	(4) Fluid Intelligence (FI)
$(R_{\overline{sc}} - R_0)$	0.373**	1.168**	0.324**	0.708
	(0.17)	(0.49)	(0.16)	(0.48)
$(R_{\overline{sc}} - R_0)^2$	-0.0129	-0.0524**	-0.0107	-0.0285
	(0.01)	(0.02)	(0.01)	(0.02)
2nd Tercile Indicator	-0.261	0.625	2.098	2.282
	(1.25)	(1.46)	(1.67)	(1.95)
3rd Tercile Indicator	-2.495*	-0.112	0.634	1.091
	(1.41)	(1.57)	(1.42)	(1.80)
(2nd Tercile Indicator)$\times (R_{\overline{sc}} - R_0)$		-0.861*		-0.483
		(0.51)		(0.51)
(3rd Tercile Indicator)$\times (R_{\overline{sc}} - R_0)$		-1.197**		-0.429
		(0.55)		(0.58)
(2nd Tercile Indicator)$\times (R_{\overline{sc}} - R_0)^2$		0.0447*		0.0262
		(0.02)		(0.02)
(3rd Tercile Indicator)$\times (R_{\overline{sc}} - R_0)^2$		0.0484*		0.0186
		(0.03)		(0.02)
Constant	-1.629	-2.538**	-3.377**	-3.638**
	(1.04)	(1.24)	(1.58)	(1.80)

Table 2.18: Regressions with financial knowledge and fluid intelligence (CogEcon sample) *(continued)*

	(1)	(2)	(3)	(4)
Tercile Indicators Used:	Financial Knowledge (FK)	Financial Knowledge (FK)	Fluid Intelligence (FI)	Fluid Intelligence (FI)
Sigma	6.143***	6.002***	6.172***	6.124***
	(0.95)	(0.94)	(0.90)	(0.90)
Observations	295	295	320	320
Number uncensored obs.	120	120	128	128
Pseudo-R^2	-434.4	-430.8	-456.9	-455.9
Log-Likelihood	0.0112	0.0193	0.00977	0.0119
F-test	1.914	1.448	1.482	0.802
Prob > F	0.108	0.176	0.207	0.601

Table 2.18: Regressions with financial knowledge and fluid intelligence (CogEcon sample) *(continued)*

	(1)	(2)	(3)	(4)
Tercile Indicators Used:	Financial Knowledge (FK)	Financial Knowledge (FK)	Fluid Intelligence (FI)	Fluid Intelligence (FI)
Mean at 1st tercile	3.29	3.29	3.85	3.85
Mean at 2nd tercile	4.41	4.41	3.35	3.35
Mean at 3rd tercile	4.48	4.48	3.93	3.93
Avg. marginal effect at 1st tercile	0.124**	0.344***	0.076*	0.160
	(0.056)	(0.039)	(0.039)	(0.102)
Avg. marginal effect at 2nd tercile	0.101**	0.093*	0.109**	0.092
	(0.043)	(0.052)	(0.052)	(0.060)
Avg. marginal effect at 3rd tercile	0.072**	-0.019	0.083*	0.069
	(0.034)	(0.042)	(0.042)	(0.085)

Notes: All regressions are Tobits with censoring point set at zero. Dependent variable is the reported change in retirement age. Excluded category is ($\triangle UE < 50\%$) All analyses include CogUSA sampling weights. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The marginal effects are not statistically significantly different by group at any standard level of significance ($\chi^2(2) = 0.59$, p-value=0.75, and $\chi^2(2) = 0.20$, p-value=0.91, respectively).

Table 19: Regressions with three categories of bequest change (HRS samples)

	(1)	(2)	(3)	(4)
Dependent variable:	$\triangle Pr(FT62)$	$\triangle Pr(FT62)$	$\triangle Pr(FT65)$	$\triangle Pr(FT65)$
$(R_{\overline{sc}} - R_0)$	0.930	0.768	2.326***	0.932
	(0.86)	(1.33)	(0.77)	(1.15)
$(R_{\overline{sc}} - R_0)^2$	-0.028	-0.008	-0.0758***	-0.011
	(0.03)	(0.05)	(0.03)	(0.04)
${}_{08}\triangle_{09}Pr(Beq\,\$100k) = 0$	-12.39**	-9.694	-5.310	-3.629
	(5.43)	(7.40)	(4.92)	(6.32)
${}_{08}\triangle_{09}Pr(Beq\,\$100k) > 0$	-3.335	-1.610	-5.216	-9.230
	(4.53)	(6.83)	(3.89)	(5.97)
$({}_{08}\triangle_{09}Pr(Beq\,\$100k) = 0) \times (R_{\overline{sc}} - R_0)$		0.088		0.592
		(2.19)		(1.97)
$({}_{08}\triangle_{09}Pr(Beq\,\$100k) > 0) \times (R_{\overline{sc}} - R_0)$		0.187		3.206*
		(2.02)		(1.79)
$({}_{08}\triangle_{09}Pr(Beq\,\$100k) = 0) \times (R_{\overline{sc}} - R_0)^2$		-0.038		-0.051
		(0.08)		(0.08)
$({}_{08}\triangle_{09}Pr(Beq\,\$100k) > 0) \times (R_{\overline{sc}} - R_0)^2$		-0.030		-0.139**
		(0.07)		(0.06)

Table 2.19: Regressions with three categories of bequest change (HRS samples) *(continued)*

	(1)	(2)	(3)	(4)
Dependent variable:	$\triangle Pr(FT62)$	$\triangle Pr(FT62)$	$\triangle Pr(FT65)$	$\triangle Pr(FT65)$
Constant	-7.611**	-8.364**	-2.079	-0.707
	(3.51)	(4.09)	(3.10)	(3.54)
Sigma	38.69***	38.49***	35.40***	35.05***
	(2.38)	(2.34)	(1.80)	(1.75)
Observations	567	567	572	572
Number uncensored obs	240	240	288	288
Log-Likelihood	-12100000	-12100000	-14000000	-13900000
Pseudo R^2	0.002	0.003	0.004	0.007
F-test: All coefficients jointly=0	1.555	0.961	2.768	2.025
Prob > F	0.185	0.465	0.027	0.042

Table 2.19: Regressions with three categories of bequest change (HRS samples) *(continued)*

	(1)	(2)	(3)	(4)
Dependent variable:	$\triangle Pr(FT62)$	$\triangle Pr(FT62)$	$\triangle Pr(FT65)$	$\triangle Pr(FT65)$
Mean $R_{\overline{sc}} - R_0$ at ${}_{08}\triangle_{09}Pr(Beq\ \$100k) < 0$	4.29	4.29	4.50	4.50
Mean $R_{\overline{sc}} - R_0$ at ${}_{08}\triangle_{09}Pr(Beq\ \$100k) = 0$	5.09	5.09	5.05	5.05
Mean $R_{\overline{sc}} - R_0$ at ${}_{08}\triangle_{09}Pr(Beq\ \$100k) > 0$	5.74	5.74	5.67	5.67
Avg. marginal effect at ${}_{08}\triangle_{09}Pr(Beq\ \$100k) < 0$	0.294	0.306	0.827***	0.434
	(0.26)	(0.38)	(0.27)	(0.39)
Avg. marginal effect at ${}_{08}\triangle_{09}Pr(Beq\ \$100k) = 0$	0.200	0.140	0.693***	0.430
	(0.17)	(0.36)	(0.24)	(0.48)
Avg. marginal effect at ${}_{08}\triangle_{09}Pr(Beq\ \$100k) > 0$	0.245	0.224	0.681***	1.193***
	(0.21)	(0.40)	(0.22)	(0.39)

Notes: All regressions are Tobits with censoring point set at zero. Dependent variable is the reported change in retirement age. Excluded category is ($\triangle Pr(Beq\ \$100k)$ <0) All analyses include CogUSA sampling weights. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The marginal effects are not statistically significantly different by group at any standard level of significance.

Table 20: Regressions with two categories of bequest change (HRS samples)

	(1)	(2)	(3)	(4)
Dependent variable:	$\triangle Pr(FT62)$	$\triangle Pr(FT62)$	$\triangle Pr(FT65)$	$\triangle Pr(FT65)$
$(R_{\overline{sc}} - R_0)$	0.901	0.689	2.327***	2.775***
	(0.85)	(1.00)	(0.76)	(0.88)
$(R_{\overline{sc}} - R_0)^2$	-0.027	-0.018	-0.0764***	-0.0992***
	(0.03)	(0.03)	(0.03)	(0.03)
${}_{08}\triangle_{09}Pr\,(Beq\,\$100k) \geq 15\,p.p.$	7.022*	6.115	8.869**	10.34**
	(4.22)	(6.22)	(3.69)	(5.17)
$({}_{08}\triangle_{09}Pr\,(Beq\,\$100k) \geq 15\,p.p.) \times (R_{\overline{sc}} - R_0)$		0.880		-1.832
		(1.99)		(1.69)
$({}_{08}\triangle_{09}Pr\,(Beq\,\$100k) \geq 15\,p.p.) \times (R_{\overline{sc}} - R_0)^2$		-0.040		0.086
		(0.07)		(0.06)

Table 2.20: Regressions with two categories of bequest change (HRS samples) *(continued)*

	(1)	(2)	(3)	(4)
Dependent variable:	$\triangle Pr(FT62)$	$\triangle Pr(FT62)$	$\triangle Pr(FT65)$	$\triangle Pr(FT65)$
Constant	-13.07***	-12.85***	-7.627**	-7.849**
	(3.43)	(3.72)	(3.15)	(3.33)
Sigma	38.76***	38.75***	35.21***	35.05***
	(2.39)	(2.39)	(1.76)	(1.71)
Observations	567	567	572	572
Number uncensored obs	240	240	288	288
Log-Likelihood	-12100000	-12100000	-13900000	-13900000
Pseudo R^2	0.002	0.002	0.005	0.006
F-test: All coefficients jointly=0	1.291	0.915	0.000	0.000
Prob > F	0.277	0.471	0.002	0.005

Table 2.20: Regressions with two categories of bequest change (HRS samples) *(continued)*

	(1)	(2)	(3)	(4)
Dependent variable:	$\triangle Pr(FT62)$	$\triangle Pr(FT62)$	$\triangle Pr(FT65)$	$\triangle Pr(FT65)$
Mean $R_{\overline{sc}} - R_0$ at $\triangle Pr\,(Beq\,\$100k) < 15\,p.p.$	5.11	5.11	5.06	5.06
Mean $R_{\overline{sc}} - R_0$ at $\triangle Pr\,(Beq\,\$100k) \geq 15\,p.p.$	4.52	4.52	4.82	4.82
Avg. marg. effect at $\triangle Pr\,(Beq\,\$100k) < 15\,p.p.$	0.237	0.193	0.705***	0.825***
	(0.21)	(0.25)	(0.22)	(0.27)
Avg. marg. effect at $\triangle Pr\,(Beq\,\$100k) < 15\,p.p.$	0.285	0.459	0.845***	0.447
	(0.26)	(0.50)	(0.27)	(0.48)

Notes: All regressions are Tobits with censoring point set at zero. Dependent variable is the reported change in retirement age. Excluded category is ($\triangle Pr(Beq\,\$100k)$ <0) All analyses include CogUSA sampling weights. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The marginal effects are not statistically significantly different by group at any standard level of significance.

References

Hugo Benitez-Silva and Debra S. Dwyer. The rationality of retirement expectations and the role of new information. *Review of Economics and Statistics*, 87(3):587–592, 2005. ISSN 0034-6535.

David M. Blau. Retirement and consumption in a life cycle model. *Journal of Labor Economics*, 26(1):35–71, January 2008.

Richard Blundell, Thomas Macurdy, Orly Ashenfelter, and David Card. Chapter 27: Labor supply: A review of alternative approaches. In *Handbook of Labor Economics*, volume Volume 3, Part 1, pages 1559–1695. Elsevier, 1999. ISBN 1573-4463.

Barry Bosworth and Gary Burtless. Recessions, wealth destruction, and the timing of retirement. *Retirement Research Center at Boston College Working Paper Series*, WP 2010-22, November 2010.

Sewin Chan and Ann Huff Stevens. Do changes in pension incentives affect retirement? A longitudinal study of subjective retirement expectations. *Journal of Public Economics*, 88 (7-8):1307–1333, July 2004. ISSN 0047-2727.

Cortney C Coile and Phillip B Levine. Bulls, bears, and retirement behavior. *Industrial and Labor Relations Review*, 59, 2005.

Courtney C. Coile and Phillip B. Levine. Recessions, reeling markets, and retiree well-being. *National Bureau of Economic Research Working Paper Series*, No. 16066, June 2010.

Julia L Coronado and Maria G Perozek. Wealth effects and the consumption of leisure: Retirement decisions during the stock market boom of the 1990s. *FEDS Working Paper Series*, No. 2003-20, 2003.

John G. Cragg. Some statistical models for limited dependent variables with application to the demand for durable goods. *Econometrica*, 39(5):829–844, 1971. ISSN 0012-9682.

J Dominitz and CF Manski. Perceptions of economic insecurity: Evidence from the Survey of Economic Expectations. *Public Opinion Quarterly*, 61:261–287, 1997.

J Dominitz and CF Manski. Measuring and interpreting expectations of equity returns. *National Bureau of Economic Research Working Paper Series*, No. 11313, 2005.

Gopi Shah Goda, John B. Shoven, and Sita Nataraj Slavov. Does stock market performance influence retirement expectations? *National Bureau of Economic Research Working Paper Series*, No. 16211, July 2010.

Gopi Shah Goda, John B Shoven, and Sita Nataraj Slavov. What explains changes in retirement plans during the Great Recession? *American Economic Review*, 101(3):29–34, May 2011.

A Gustman and T Steinmeier. A structural retirement model. *Econometrica*, 54, 1986.

Alan L. Gustman, Thomas L. Steinmeier, and Nahid Tabatabai. What the stock market decline means for the financial security and retirement choices of the near-retirement population. *National Bureau of Economic Research Working Paper Series*, No. 15435, October 2009.

Alan L. Gustman, Thomas L. Steinmeier, and Nahid Tabatabai. *Pensions in the Health and Retirement Study*. Harvard University Press, May 2010a. ISBN 0674048660.

Alan L. Gustman, Nahid Tabatabai, and Thomas L. Steinmeier. Table data: Pensions in the Health and Retirement Study (Final version 1.0), May 2010b.

Health and Retirement Study. Imputations for Employer-Sponsored Pension Wealth from Current Jobs in 2004. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740). Ann Arbor, MI., July 2009.

Health and Retirement Study. 2006 Core Data (Final Release, Version 2.0). Produced and distributed by the University of Michigan with funding from the National Institute on Aging (Grant number NIA U01AG009740). Ann Arbor, MI., October 2010a.

Health and Retirement Study. 2008 Core Data (Final Release, Version 1.0). Produced and distributed by the University of Michigan with funding from the National Institute on Aging (Grant number NIA U01AG009740). Ann Arbor, MI., October 2010b.

Health and Retirement Study. Cross-Wave Tracker 2008 (Early Version 1.0). Produced and distributed by the University of Michigan with funding from the National Institute on Aging (Grant number NIA U01AG009740). Ann Arbor, MI., October 2010c.

Health and Retirement Study. 2009 Internet Survey (v.1.0) public use dataset. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (Grant number NIA U01AG009740). Ann Arbor, MI., October 2010d.

James J. Heckman. What has been learned about labor supply in the past twenty years? *The American Economic Review*, 83(2):116–121, May 1993.

Douglas Holtz-Eakin, David Joulfaian, and Harvey S. Rosen. The Carnegie conjecture: Some empirical evidence. *The Quarterly Journal of Economics*, 108(2):413–435, May 1993. ISSN 00335533.

MD Hurd, M Reti, and S Rohwedder. The effect of large capital gains or losses on retirement. In David Wise, editor, *Developments in the Economics of Aging*, National Bureau of Economic Research Conference Report, pages 127–172. University of Chicago Press, Chicago, IL, 2009.

Michael Hurd and Kathleen McGarry. Evaluation of the subjective probabilities of survival in the health and retirement study. *Journal of Human Resources*, 30:S268–S292, 1995.

Michael Hurd and Monica Reti. The effects of large capital gains on work and consumption: Evidence from four waves of the HRS. Final report prepared for the US Department of Labor, October 2001.

Michael Hurd and James P. Smith. Expected bequests and their distribution. *National Bureau of Economic Research Working Paper Series*, No. 9142, September 2002.

Michael D. Hurd. Subjective probabilities in household surveys. *Annual Review of Economics*, 1(1):543–564, 2009. ISSN 1941-1383.

Michael D. Hurd and Susann Rohwedder. The effect of the risk of out-of-pocket spending for health care on economic preparation for retirement. *Michigan Retirement Research Center Research Brief*, September 2010, September 2010a.

Michael D. Hurd and Susann Rohwedder. The effects of the economic crisis on the older population. *Michigan Retirement Research Center Working Paper Series*, WP 2010-231, November 2010b.

Guido W. Imbens, Donald B. Rubin, and Bruce I. Sacerdote. Estimating the effect of unearned income on labor earnings, savings, and consumption: Evidence from a survey of lottery players. *The American Economic Review*, 91(4):778–794, September 2001. ISSN 00028282.

Kandice Kapinos. Cross-Wave Prospective Social Security Wealth Measures of Pre-Retirees. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (Grant number NIA U01AG009740). Ann Arbor, MI., March 2010.

Kandice Kapinos, Charles Brown, Michael Nolte, Helena Stolyarova, and David Weir. Prospective Social Security Wealth Measures of Pre-retirees, Public release, Version 4.0, May 2011.

Gabor Kezdi and Purvi Sevak. Economic adjustment of recent retirees to adverse wealth shocks. *Michigan Retirement Research Center Working Paper Series*, WP 2004-075, 2004.

Miles S. Kimball and Matthew D. Shapiro. Social security, retirement and wealth: Theory and implications. *Michigan Retirement Research Center Working Paper Series*, WP 2003-054, June 2003.

Miles S. Kimball and Matthew D. Shapiro. Labor supply: Are the income and substitution effects both large or both small? *National Bureau of Economic Research Working Paper Series*, No. 14208, July 2008.

Hamish Low, Costas Meghir, and Luigi Pistaferri. Wage risk and employment risk over the life cycle. *American Economic Review*, 100(4):1432–67, 2010.

Thomas E. MaCurdy. An empirical model of labor supply in a Life-Cycle setting. *The Journal of Political Economy*, 89(6):1059–1085, December 1981.

Charles F. Manski. Measuring expectations. *Econometrica*, 72(5):1329–1376, 2004. ISSN 0012-9682.

Kathleen McGarry. Health and retirement: Do changes in health affect retirement expectations? *Journal of Human Resources*, XXXIX(3):624–648, July 2004.

Olivia S. Mitchell, James M. Poterba, Mark J. Warshawsky, and Jeffrey R. Brown. New evidence on the money's worth of individual annuities. *The American Economic Review*, 89(5):1299–1318, December 1999. ISSN 00028282.

Franco Modigliani and Richard Brumberg. Utility analysis and the consumption function: An interpretation of Cross-Section data. In Francesco Franco, editor, *The Collected Papers of Franco Modigliani*, volume 6, pages 3–45. MIT Press, Cambridge, Massachusetts, 2005. (Originally published in 1954).

RAND Corporation. RAND HRS Data (Version J). Public use dataset. Produced by the RAND Center for the Study of Aging, with funding from the National Institute on Aging and the Social Security Administration. Santa Monica, CA., 2010.

Steven A. Sass, Courtney Monk, and Kelly Haverstick. Workers' response to the market crash: Save more, work more? *Center for Retirement Research at Boston College*, Issue Brief 10-3, February 2010.

Purvi Sevak. Wealth shocks and retirement timing: Evidence from the nineties. *Michigan Retirement Research Center*, WP-027, 2002.

Matthew D. Shapiro. Buffering shocks to Well-Being late in life. *Michigan Retirement Research Center Working Papers*, No. 2009-211, October 2009.

Social Security Administration. Period Life Table, Actuarial Life Tables. http://www.ssa.gov/OACT/STATS/table4c6.html, 2006.

S&P/Case-Shiller. Home Price Index Data, seasonally-adjusted, June 2010. URL `http://www.standardandpoors.com/indices/sp-case-shiller-home-price-indices/`.

James H. Stock and David A. Wise. Pensions, the option value of work, and retirement. *Econometrica*, 58(5):1151–1180, September 1990. ISSN 00129682.

U.S. Bureau of Labor Statistics. Labor Force Statistics from the Current Population Survey, seasonally-adjusted unemployment data., n.d.a.

U.S. Bureau of Labor Statistics. Local Area Unemployment Statistics by Census Division, seasonally-unadjusted unemployment rates for 2008-2009, n.d.b.

U.S. Bureau of Labor Statistics. Local Area Unemployment Statistics, County Laborforce Data, seasonally-unadjusted unemployment rates for 2008-2009, n.d.c.

U.S. Census Bureau. American community survey 5-Year estimates, 2005-2009., n.d. URL `http://factfinder.census.gov/`.

Jeffrey M. Wooldridge. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press, 2002.

Menahem E. Yaari. Uncertain lifetime, life insurance, and the theory of the consumer. *The Review of Economic Studies*, 32(2):137–150, April 1965. ISSN 0034-6527.

Yahoo Finance. Dow Jones Industrial Average Historical Price Data, n.d. URL `http://finance.yahoo.com/`.

9 Model of optimal retirement choice

Underlying the simple Modigliani model is a more complicated model of retirement and consumption choice. A simplified version of a model of optimal consumption and retirement timing from Miles Kimball and Matthew Shapiro (2003) posits that, at any point in time τ each individual chooses future consumption path, C_t, and labor market participation path, χ_t, from time τ until known time of death, T, according to

$$\max_{C_t,\chi_t} \int_{\tau}^{T} \left\{ e^{-\rho(T-t)} \left(\frac{C_t^{1-\frac{1}{\theta}}}{1-\frac{1}{\theta}} - (e^{\alpha-\zeta t})\chi_t \right) \right\} dt \tag{10}$$

subject to

$$\dot{A} = rA_t + \omega_t \chi_t - C_t \tag{11}$$

where

$$\chi_t = \begin{cases} 0\, if\ working\ at\ time\ t \\ 1\, if\ not\ working\ at\ time\ t \end{cases} \tag{12}$$

and ρ is the rate of time preference, θ is the coefficient (or inverse?) of relative risk aversion, and α and ς are "disutility of work" parameters, all individual-specific. Additionally, A_t denotes assets at time t and ω_t is wage at time t. Defining λ_t as the shadow value of wealth, the current-value Hamiltonian is

$$\mathcal{H} = \frac{C_t^{1-\frac{1}{\theta}}}{1-\frac{1}{\theta}} - e^{\alpha-\zeta t}\chi_t + \lambda_t \left[rA_t + \omega_t \chi_t - C_t \right] \tag{13}$$

- check margins– seems in wrong rows (see m's comments)which implies the following first-order conditions:

$$h_c = 0 \Leftrightarrow C_t^{-1/\theta} = \lambda_t \tag{14}$$

$$h_A = \rho\lambda_t - \dot{\lambda}_t \Leftrightarrow r\lambda_t = \rho\lambda_t - \dot{\lambda}_t \tag{15}$$

$$\dot{A} = rA_t + \omega_t \chi_t - C_t \tag{16}$$

Letting χ, the decision to work, be characterized by

$$\chi_t = \begin{cases} 0\,if \quad \lambda_t\omega_t \geq e^{\alpha+\zeta t} \\ 1\,if \quad \lambda_t\omega_t \leq e^{\alpha+\zeta t} \end{cases} \tag{17}$$

it must be that the optimal time of retirement, R, solves

$$\omega_R\lambda_R = e^{\alpha+\zeta t} \tag{18}$$

Now, given the first-order condition for assets, h_A, it can be shown that

$$\lambda_R = \lambda_t e^{(\rho-r)(R-t)} \tag{19}$$

Plugging this into the equation for $\omega_R\lambda_R$ from above,

$$\omega_R\lambda_t e^{(\rho-r)(R-t)} = e^{\alpha+\zeta t} \tag{20}$$

gives the result that an individual is indifferent between working and not working when the marginal disutility of continuing to work is equal to the marginal utility gained from continuing to work.

Taking logs of both sides and solving for R yields the equation for the optimal retirement time,

$$R = \frac{ln(\lambda_t) + ln(\omega_R) - (\rho - r)t - \alpha}{\varsigma - \rho + r} \tag{21}$$

Note that $\partial R/\partial ln(\lambda_t) > 0$. This implies that the higher the marginal increase in current utility from relaxing the budget constraint, the later a person will retire. In the context of this paper, I expect that a negative shock to accumulated assets, such as losses from a stock or housing market bust, or losses in future income flows, will cause an increase in an individual's optimal retirement age.

10 Imputation of defined benefit pension wealth for CogEcon

I impute defined benefit pension wealth estimates for the CogEcon respondents based on defined benefit pension wealth information in the HRS dataset *Imputations for Pension-Related Variables* (Final, Version 1.0), according to the following:

1. For CogEcon respondents who indicated that they (and their spouse/partner, if in a

relationship) do not have a defined benefit pension, I assign a defined benefit pension value of $0.

2. For single CogEcon respondents who indicated that they do have a defined benefit pension, I assign the inflation-adjusted cell mean (age group by sex by occupation group) of defined benefit plan wealth, using the defined benefit plan value calculated using the HRS respondents' expected retirement age. I match the cell means to CogEcon respondents who were in the age range in 2009 that the HRS respondents were in in 2004. So, for example, a female CogEcon respondent in an "Education, Training and Library" occupation who was aged between 45 and 49 in 2009 would be assigned the inflation-adjusted cell mean defined benefit pension wealth of female HRS respondents with defined benefit pensions in an "Education, Training and Library" occupation who were aged between 45 and 49 in 2004.

3. For coupled CogEcon respondents who indicated that they or their partner have a defined benefit pension, but for whom the CogEcon data don't contain the information about the spouse or partner's occupation or age, I assume only the respondent has a defined benefit pension, and assign an estimated defined benefit pension value using the same method as that used for single CogEcon respondents.

4. For coupled CogEcon respondents who indicated that they or their partner have a defined benefit pension, and for whom I have occupation and age data for both members of the couple, I calculate the age group by sex by occupation probabilities that each person has a defined benefit pension (the number with non-zero defined benefit wealth values over the total number of respondents in that sex by age by occupation group in the 2004 core HRS data). Then, I use the same method as described in items 2 and 3 to match CogEcon respondents to the cell means of defined benefit pensions from comparable HRS respondents. Next, I multiply each partner's cell mean by his or her probability of having a defined benefit pension and sum across both individuals in the household.

11 Derivation of expected retirement age in HRS sample

Unfortunately, the expected age of retirement is not asked directly of all HRS respondents. Instead, I derive this age by combining information from several variables, as follows:

1. If a respondent's retirement plans include stopping work altogether, I use the planned age for stopping work as the expected age of retirement.

2. If there is no planned age of retirement, I predict retirement age from a linear regression of the expected age of retirement on the probabilities of full-time work after age 62 and age 65 given by the respondent in 2006 and 2008, plus the respondent's age and labor force status (full-time, part-time or partly retired) at the 2008 interview. The adjusted r-squared from this regression is 0.424.

3. If there is still no expected age of retirement, I predicted retirement age from a regression of expected age of retirement on the probabilities of full-time work after age 65 given by the respondent in 2006 and 2008, and on the respondent's age and labor force status at the 2008 interview. The adjusted r-squared from this regression is 0.361.

4. If there is still no expected age of retirement, I predicted retirement age from a similar regression to (b), using 2008 data only. The adjusted r-squared from this regression is 0.385.

5. If there is still no expected age of retirement, I predicted retirement age from a regression of expected age of retirement on the probabilities of full-time work after age 62 and 65 given by the respondent in 2006, and on the respondent's age and labor force status at the 2008 interview. The adjusted r-squared from this regression is 0.262. (10 observations)

6. If there is still no expected age of retirement, I use age 65 as the expected retirement age for these individuals. Age 65 is the mean, median and mode of the expected retirement age for individuals under age 65 in 2008 who expected to completely stop working, and thus seems like a reasonable estimate for those who do not give enough information to allow for an estimated retirement age.

12 Regression estimates used in comparisons of CogEcon and HRS results

Using the final HRS dataset, I regressed the change in reported retirement age between Core 2006 and Core 2008, $R_{08} - R_{06}$, on the change in the probabilities of full-time work reported in 2006 and 2008, ${}_{08}\triangle_{09}Pr(FT62)$ and ${}_{08}\triangle_{09}Pr(FT65)$. These regressions only include those respondents who actually reported planned or expected age of retirement in both the 2006 and 2008 surveys, so the sample size is quite small. The results from these regressions are shown below. To calculate the expected change in retirement age for a one percentage point change in the probability of full-time work, I multiplied each estimated coefficient by 365.25, the number of days in a year. For the subset of individuals in my final regression sample, these regressions yield estimates of an 8.5 day increase in retirement age for a one percentage point increase in the probability of full-time work after age 62, and a 6.7 day increase in retirement age for a one percentage point increase in the probability of full-time work after age 65.

Table 21: Regression estimates used in comparisons between CogEcon and HRS results

	(1)	(2)
${}_{08}\triangle_{09}Pr(FT62)$	0.0232	
	(0.02)	
${}_{08}\triangle_{09}Pr(FT65)$		0.0183***
		(0.01)
Constant	0.367	0.3
	(0.23)	(0.21)
Observations	71	83
R-squared	0.069	0.094
Implied change per 1 p.p. increase:	8.5 days	6.7 days

Note: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

13 Specification comparison results without sampling weights

The following tables present the regression results on which the likelihood ratio tests discussed in Section 6.2.4 are based. In each table, the samples have been restricted to include only observations that are also included in my preferred regression specifications that I present in the main portion of this study.

Table 22: Comparison of Tobit, probit and Cragg models (CogEcon sample)

	(1)	(2)	(3)
Specification:	Tobit	Probit	Truncated
Dependent variable:	$\triangle R$	$I_{\triangle R>0}$	$\triangle R$
$R_{\overline{sc}} - R_0$	0.231*	0.034	0.522
	(0.133)	(0.024)	(0.412)
$(R_{\overline{sc}} - R_0)^2$	-0.009	-0.002	-0.015
	(0.006)	(0.001)	(0.016)
Constant	-1.863***	-0.299***	-2.288
	(0.550)	(0.088)	(3.203)
Sigma	5.809***		5.606***
	(0.412)		(1.067)
Observations	320	320	128
Log-Likelihood	-519.5	-213.6	-305.0

Notes: Dependent variable in Tobit and truncated normal specifications is reported change in retirement age. In probit specification, dependent variable is an indicator that is equal to one if retirement age increased, and zero otherwise. Censoring point for Tobit and truncated regressions is 0. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. LR Test statistic ($\sim \chi^2(4)$), $-2(LL_{Tobit} - (LL_{Probit} + LL_{Truncated}))$, is 1.78 (p-value 0.78).

14 Estimating expected bequests in the HRS

To generate point estimates of expected bequests, I first averaged responses from 2004, 2006 and 2008 for each individual to reduce measurement error (this calculation yielded $Pr(B \geq \$10k)_{avg}$ and $Pr(B \geq \$100k)_{avg}$). Next, I calculated each individual's total bequeathable wealth ($beq\ w$) as the sum of financial wealth, real estate and business assets (future earnings, Social Security wealth, and defined benefit pension wealth were excluded from the bequeathable wealth calculation).

I then took the average of $(1 - Pr(B \geq \$10k)_{avg})$ across all individuals to get the population average probability of leaving less than \$10,000 in wealth, $1 - Pr(B \geq \$10k)_{pop}$. Next, I took the average of ($Pr(B \geq \$10k)_{avg}$-$Pr(B \geq \$100k)_{avg}$) across all individuals with

Table 23: Comparison of Tobit, probit and Cragg models (HRS samples)

	Probability of full-time work after age 62			Probability of full-time work after age 65		
	(1)	(2)	(3)	(4)	(5)	(6)
Specification:	Tobit	Probit	Truncated	Tobit	Probit	Truncated
Dependent variable:	${}_{08}\triangle_{09}Pr(FT62)$	$I_{>0}$	${}_{08}\triangle_{09}Pr(FT62)$	${}_{08}\triangle_{09}Pr(FT65)$	$I_{>0}$	${}_{08}\triangle_{09}Pr(FT65)$
$R_{\overline{sc}} - R_0$	0.698	0.033	-3.089	1.747**	0.047**	1.361
	(0.811)	(0.022)	(2.984)	(0.744)	(0.022)	(1.460)
$(R_{\overline{sc}} - R_0)^2$	-0.017	-0.001	0.161	-0.060**	-0.002**	-0.015
	(0.029)	(0.001)	(0.107)	(0.027)	(0.001)	(0.051)
Constant	-10.804***	-0.265***	-40.826	-3.078	-0.064	-1.215
	(2.934)	(0.071)	(33.652)	(2.549)	(0.070)	(10.132)
Sigma	40.530***		50.520***	38.185***		36.454***
	(2.056)		(9.647)	(1.743)		(3.956)
Observations	589	589	247	594	594	298
Log-Likelihood	-1472	-399.4	-1068	-1713	-408.8	-1301

Notes: In probit specification, dependent variable is an indicator that is equal to one if the probability of full-time work past the reference age (62 in Column 2, and 65 in Column 5) increased between 2008 and 2009, and zero otherwise. Censoring point for Tobit and truncated regressions is 0. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. For the ${}_{08}\triangle_{09}Pr(FT62)$ analyses, the LR Test statistic ($\sim \chi^2(4)$), $-2(LL_{Tobit} - (LL_{Probit} + LL_{Truncated}))$ is 9.89 (p-value 0.04). For the ${}_{08}\triangle_{09}Pr(FT65)$ analyses, the LR Test statistic ($\sim \chi^2(4)$), $-2(LL_{Tobit} - (LL_{Probit} + LL_{Truncated}))$ is 5.77 (p-value 0.22).

at least \$10,000 in wealth to get the population average probability of leaving between \$10,000 and \$100,000 in wealth, $(Pr(B \geq \$100k))_{pop100}$ for individuals with more than \$100,000 but less than \$500,000 in bequeathable wealth. Next, I estimated a linear regression of $(Pr(B \geq \$500k))$ on the 2009 values of $Pr(B \geq \$10k)$ and $Pr(B \geq \$100k)$, bequeathable wealth in 2009, plus the square of each of these, for individuals with at least \$500,000 in bequeathable wealth in 2009, and applied the estimated equation to $Pr(B \geq \$10k)_{avg}$, $Pr(B \geq \$100k)_{avg}$, and 2008 wealth to predict $(Pr(B \geq \$500k))_{08}$. Then, I applied these predictions to calculate $(Pr(B \geq \$100k) - Pr(B \geq \$500k))_{pop}$ for individuals with more than \$500,000 in bequeathable wealth.

Finally, I used the following equation to create point estimates that were plausible, given bequeathable wealth, and also increasing with the subjective probability measures of leaving a bequest:

$$E(bequest) = \begin{cases} \left(\frac{1-Pr(B\geq\$10k)_{avg}}{1-Pr(B\geq\$10k)_{pop}}\right) \times beq\,w & if\ beq\,w < \$10k \\ \left(\frac{(Pr(B\geq\$10k)-Pr(B\geq\$100k))_{avg}}{(Pr(B\geq\$10k)-Pr(B\geq\$100k))_{pop}}\right) \times beq\,w & if\ beq\,w \in [\$10k, \$100k) \\ \left(\frac{(Pr(B\geq\$100k))_{avg}}{(Pr(B\geq\$100k))_{pop}}\right) \times beq\,w & if\ beq\,w \in [\$100k, \$500k) \\ \left(\frac{(Pr(B\geq\$500k))_{avg}}{(Pr(B\geq\$500k))_{pop}}\right) \times beq\,w & if\ beq\,w \in [\$500k, inf) \end{cases}$$

The estimated values of $E(bequest)$ have a mean of \$368,000 and a median of \$140,000. The 25th percentile observation is \$36,000, and the 75th percentile observation is \$322,000. These estimates seem reasonably in line with Hurd and Smith [2002] and Hurd and Rohwedder [2010b], but each individual's expected bequest is feasible given his or her own wealth. These other studies were interested in population statistics, so feasibility of the individual estimates was not important to their estimation strategy.

The standard deviation is \$1,601,000. (All rounded to the nearest \$1,000.) These range from 20 percent of total wealth at the 25th percentile to 67 percent of total wealth at the 75th percentile. The mean is 46 percent of bequeathable wealth, and the median is 43 percent. In terms of bequeathable wealth, the inter-quartile range is from 20 percent to 100%, with mean 59 percent and median 67 percent.