The Impact of Hypothetical Wealth Shocks on Retirement Timing

Brooke Helppie-McFall
University of Michigan

Joanne W. Hsu
Federal Reserve Board

Matthew Shapiro
University of Michigan

November 5, 2017[*]

[*] Corresponding author: Brooke Helppie-McFall, e-mail bhelppie@umich.edu.

The support of National Institute on Aging grant P01 AG026571 is gratefully acknowledged. The analysis and conclusions set forth are those of the authors and do not indicate concurrence with other members of the research staff or the Board of Governors of the Federal Reserve.

1. Introduction

The classic life-cycle model predicts that optimal consumption from the present until death will be proportional to total wealth, including the present discounted value of all income flows (Modigliani and Brumberg, 1954/2005). A dynamic, structural life-cycle model that expands on the classic life-cycle model to include retirement timing as a choice variable, like that developed in Kimball and Shapiro (2003, 2008), provides a clearer picture of the model that might underlie Modigliani's static model.

In such a model, the optimal choice of retirement problem can be illustrated by a graph over possible retirement ages, containing an upward-sloping marginal disutility of work (or utility of retirement leisure) line and a downward-sloping line representing the marginal value of wealth. The latter represents the marginal value of wealth given the optimal choice of consumption path at any retirement age along the x-axis. The intersection of these lines determines optimal retirement age.

A shock to wealth shifts the marginal value of wealth curve and, thus, optimal retirement timing—for a wealth loss, the curve shifts upward and increases the optimal retirement age; for a wealth gain, the curve shifts downward and decreases the optimal retirement age. The relative slopes of the curves dictate the extent to which a wealth shock will result in a shift of retirement age versus acceptance of a different consumption level.

Economists and policymakers have long been interested in predicting the impact of wealth shocks on retirement behavior. The economic implications of the life-cycle and related models are clear, yet a large body of existing empirical research has failed to validate these models' predictions of how individuals respond to wealth shocks. In the last several years, a spate of research has examined the impact of the Great Recession on the retirement plans and behavior of older Americans. Some have suggested that the wealth losses experienced in this time period may have had a small causal effect on future or planned retirement timing (for example, Gustman et al., 2009, 2015; Goda et al., 2011; McFall, 2011; Ondrich and Falevich, 2016). However, it is unclear whether these effects would generalize to wealth losses in the future. In particular, if survey questions about hypothetical wealth shocks could be shown to help predict changes in retirement behavior, this would be a useful tool for predicting the impact of policy proposals that would affect retiree finances, such as changes in tax treatment of retirement account savings and/or withdrawals, changes in Social Security benefit amounts, or changes in Medicare deductibles.

McFall (2011a,b) used a novel approach to examine the causal impact of wealth losses on retirement timing. Using surveys fielded before and after the financial crisis of late 2008 and early 2009, she was able to directly measure the impact of the crisis on wealth and retirement expectations. She calculated the change in retirement timing that would be needed to “make up” losses such that respondents could sustain the long-term standard of living implied by their pre-crisis wealth and used Tobit regressions to examine the relationship between this change in retirement timing and respondents’ stated changes in retirement plans. The Tobit econometric specification that is well-suited to data that included significant heaping at a corner solution, likely due to fixed costs of adjustment and non-linearities in the retirement age optimization problem.

In this paper, we follow methodology from McFall (2011a,b) to examine the impact of *hypothetical* wealth shocks on survey respondents’ expectations about future labor supply. Using data from two different surveys with distinct samples, we quantify the impact of hypothetical wealth shocks on each individual, in terms of consumption levels that could be sustained before and after these shocks, and changes in expected labor supply (specifically, the probability of working full-time past the age of 65).

We then examine the relationship between hypothetical wealth losses and planned changes in expected labor supply using descriptive and regression analyses. We find a modest but statistically-significant causal effect of hypothetical wealth losses on the subjective probability of full-time work after age 65. The effects are similar in magnitude to those found by McFall in her analysis of the impact of actual wealth losses on the same outcome variable immediately after the stock market crash in the Great Recession. While the estimated effects of hypothetical wealth gains have the expected signs—that is, they imply that a wealth gain would result in a lower chance of full-time work after age 65—they are only statistically-significant in analyses using one sample (ALP). These findings carry implications for predicting the labor market effects of policy changes affecting the finances of older adults.

2. Data

In this section, we discuss the data sources, the analysis samples, and the construction of key variables of interest. We utilize data from two different surveys in this paper.

First, we use data from the 2011 and 2013 waves of the Cognitive Economics Study (CogEcon), an innovative panel study of a national sample of persons 51 and older that began in 2008, with a post-Great Recession follow-up in 2009 and additional surveys in 2011 and 2013. This ongoing study is designed to explore the relationship between cognitive measures and a variety of economic variables, including financial knowledge, asset and debt holdings, and details of financial decision-making. These surveys were fielded by web for respondents with internet access and by mail questionnaires for those without internet access.

Second, we also employ data from a web survey fielded to the American Life Panel (ALP), a survey panel of over 5,000 American adults administered by the RAND Corporation. The ALP was developed to mirror the composition of the adult population in the United States across many measureable characteristics. The selectivity effects of the web survey mode are mitigated by giving potential participants without Internet access a device with which they may access the Internet and complete surveys. The ALP survey used data from survey MS 307: Decision Marking, which was fielded to 1,503 panel members between February 27 through August 15, 2013, mostly over the age of 50, plus basic demographic data from ALP baseline data.

The ALP survey used for this research consists of a subset of questions originally fielded in the 2011 Cognitive Economics Survey. Therefore, both sources of data contain variables measuring income, earnings, retirement plans and measures of wealth; furthermore, the question wording was largely identical. For both data sources, we restrict our analysis to individuals under the age of 65 who were still working at the time of the survey and held tax-advantaged retirement accounts whose values could rise and fall. This analysis sample was chosen due to the relevance of the outcome variables we analyze to this population.

2.1 Sample

The analysis samples are comprised of data from the respondents:

- At least of age 50 but younger than 65
- In the labor force—employed, on temporary leave, or unemployed and looking for work
- Reporting non-zero earnings for the year prior to the survey (2012 for ALP and CogEcon 2013; 2010 for CogEcon 2011)
- Reporting a non-zero amount of defined contribution retirement savings
- Reporting a subjective probability of full-time work after reaching age 65
- For whom the independent variable of interest could be calculated. That is, it must have been possible to calculate a non-infinite amount by which retirement age would need to change to provide the same level of sustainable consumption after the hypothetical shock as before it. For example, the independent variable of interest was impossible to calculate if the respondent's

earnings were so small relative to the shock changing retirement timing could not offset the loss or gain in a normal lifetime.[2]

For the hypothetical loss scenario in the ALP, the final sample size is 378. For the hypothetical gain scenario in the ALP, the final sample size is 402. The CogEcon analysis sample is 254 for both sub-samples.

2.2 Outcome Variables: Retirement Expectations

Three questions in this survey provide the unique focus of this analysis. After answering a series of questions about income, labor and earnings, respondents were asked questions on their subjective expectations about a number of topics.[3] For respondents under the age of 65 and who reported working full-time at the time of the survey, or a chance of working full-time again in the future, respondents were asked:

Question F15*: "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 who reported in the wealth section that they—or their spouse or partner—held any wealth in "tax-advantaged retirement accounts" were then asked two additional questions,

Question F16*: "Now, suppose that you find out tomorrow that the value of your retirement accounts has decreased by $[amount]. In this situation, what do you think the chances are that you would be working full-time after you reach age 65?"*

---

[2] As an obvious but extreme example, consider a case in which the respondent's earnings were $5/year and he/she gained or lost $50,000. No amount of adjustment of retirement age within the probable (as determined by non-zero values in the SSA Life Table) lifespan of any human would result in a sustainable consumption level that was equal to the pre-shock value.

[3] These questions nearly identical to subjective expectations questions asked on the Health and Retirement Study.

Question F17: "Suppose, instead, that you find out tomorrow that the value of your retirement accounts has increased by $[amount]. In this situation, what do you think the chances are that you would be working full-time after you reach age 65?"

In web mode (all of ALP and most of CogEcon respondents), [amount] was a rounded number equal to about 30 percent of the total value of tax-advantaged retirement holdings reported earlier in the survey by each respondent.[4] The 30 percent shock was chosen to approximate the median loss incurred by CogEcon respondents as a result of the Great Recession (McFall, 2011).[5] For mail mode CogEcon respondents, the survey simply asked about a $50,000 shock to retirement accounts.[6] These questions, therefore, were designed to test whether respondents would report changes to retirement plans as a result of a hypothetical loss, and to compare those responses to actual responses to real-life losses. Additionally, the questions enable us to examine whether there or not responses to losses versus gains are symmetric.

Figure 1: Response Scale from CogEcon Mail Survey

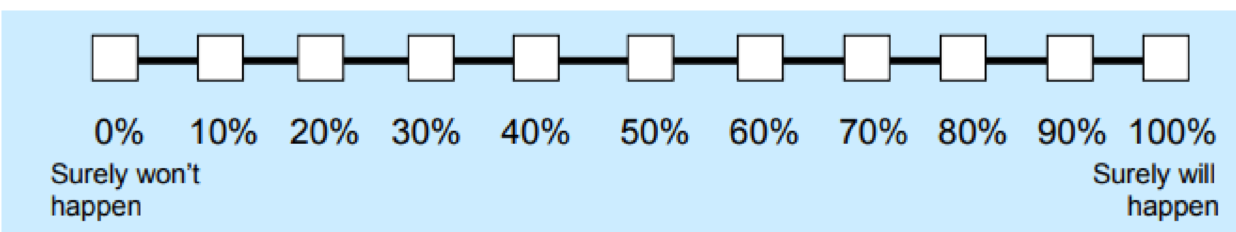

[4] For example, a respondent with $25,000 in retirement account holdings was asked about 0.3x$25,000, or $9,000. For the approximately 5 percent of respondents who gave a range card response for the retirement wealth, the amount asked about was a range calculated by multiplying both the upper and lower limits by 0.3.

[5] That is, the median respondent in the Cognitive Economics survey in 2009 reported a wealth loss (due to losses in real estate, future labor market earnings, and financial holdings) that was approximately 30 percent of his/her pre-recession tax-advantaged retirement account holdings.

[6] $50,000 was approximately 1/3 of the median value of DC retirement accounts reported by CogEcon respondents who held such accounts in 2008. For the loss scenario, respondents with less than $50,000 in retirement accounts were asked to imagine that the value of these accounts had decreased to zero.

The answer options to all three subjective probability questions are radio buttons laid out horizontally, with answer options every ten percentage points, beginning with “0%” and ranging to “100%.” The mail questionnaire presents these questions in a similar manner (see Figure 1, above).

The dependent variables in our regression analyses are the difference in probability of full-time work after age 65 from respondents’ actual financial situations to the hypothetical shocks to their retirement holdings. We calculate our dependent variables as:

$$\Delta \Pr(Work\ FT\ after\ 65)_{loss} = \Pr(Work\ FT\ after\ 65)_{loss} - \Pr(Work\ FT\ after\ 65)$$

$$\Delta \Pr(Work\ FT\ after\ 65)_{gain} = \Pr(Work\ FT\ after\ 65)_{gain} - \Pr(Work\ FT\ after\ 65)$$

2.3 Explanatory Variables: Magnitude of wealth shocks, relative to earnings

In this paper, the explanatory variables of interest are (1) the number of years by which retirement age would need to change to make sustainable consumption after a shock equal to sustainable consumption before any shock, and (2) that number squared. For the wealth shocks, these variables use the amounts asked about in the hypothetical wealth shock questions (F16 and F17). We calculate these explanatory variables for both a wealth loss, yielding variable ($R_{SC}^{loss} - R_0$), and a wealth gain, yielding the variable ($R_{SC}^{gain} - R_0$). These variables express the magnitude of wealth shocks in terms of (a) the level of consumption that could be sustained for that specific individual, and (b) each individual’s specific earning potential from working, in order to have a basis of comparison across households with very different economic and financial circumstances.

Calculating these regressors is a multi-step process. We begin by calculating time to retirement: we subtract respondents’ exact age at the time of the survey from the age at which they report planning to “retire completely.” We assume respondents will retire on their birthday in the year they turn their planned retirement age. This step is necessary for calculating the present discounted value of future earnings component of wealth.

Next, we calculate respondents' wealth holdings as described above, including the present discounted value of expected future income flows. The surveys both contain measures of financial wealth inside "tax-advantaged retirement accounts" as well as holdings of financial assets outside of retirement accounts. The latter category includes bank accounts, cash, money market funds, CDs, T-bills, stocks, mutual funds, corporate or government bonds, and anything else respondents chose to include. The data also contain respondents' reports of business wealth and "other" assets such as limited partnerships, hedge funds, commodities, timber or mineral rights, and precious metals. Additionally, the data contain information about the value of owned primary residences and other real estate, as well as the balances of mortgages or other loans against these assets.

Respondents report labor earnings from the calendar year prior to the survey for themselves and, where applicable, their spouses or partners. Using years to retirement (R); labor earnings; probability respondent is alive in the year $t + s$, where t refers to the year the survey data was collected ($P1_{t+s}$)[7]; and a real interest rate of 2.9 percent, we calculate the present discounted value of all future earnings (in 2016 dollars).

$$EPDV(earnings)_j = \sum_{s=1}^{R} \left(earnings_{j,12} \times \left(\frac{1}{1+r} \right)^s \times P1_{j,s} \right)$$

Our measure of wealth holdings is the sum of financial wealth, real estate wealth, net business wealth, and the present discounted value of future earnings. [8] For respondent j in the ALP, this is calculated:

$$Wealth_j = financial\ wealth_j + net\ real\ estate_j + net\ business\ wealth_j + EPDV(earnings)_j$$

[7] Calculated using the Social Security Period Life Table 2014.

[8] The ALP does not include information about credit card debt, non-real estate debt, Social Security or defined benefit pension wealth that would be necessary for a full picture of the resources available to respondents as they age. While it would be desirable to present summary statistics on total wealth and sustainable consumption estimates accounting for these sources of debt and wealth, these sources of debt and wealth do not impact the change in retirement age needed to equalize pre- and post-shock sustainable consumption. Our analytical approach therefore does not rely on these variables, and neither our regression results nor their interpretation are affected by the lack of these variables.

Because CogEcon collects more comprehensive measures of wealth, for that analysis we are additionally able to compute the expected present discounted value of flows from Social Security and defined benefit (DB) pensions from both the respondent and their spouse (if any). We compute these values using the value of Social Security benefits and DB pension benefits received in the past calendar year for those who have started receiving those benefits. For those who have not, we use values from questions on what respondents expect to receive from Social Security or any DB pension benefits they expect to claim in the future. Furthermore, we can also include information on credit card debt, the net value of businesses, and other assets. For respondent j in CogEcon, wealth holdings, expressed in 2016 dollars are computed as

$$Wealth_j = financial\ wealth_j + net\ real\ estate_j + net\ business\ wealth_j + EPDV(earnings)_j + EPDV(Social\ Security)_j + EPDV(DB\ pensions)_j - credit\ card\ debt_j - other\ debt_j$$

With measures of total wealth from ALP and CogEcon in hand, we next calculate sustainable consumption as the total wealth figure divided by the price of the inflation-adjusted annuity that could be purchased with a respondent's wealth holdings. The annuity price is calculated using a load factor (L) of 0.18, real interest rate (r) of 2.9 percent and gender- and year-specific probabilities that respondent and spouse/partner are living ($(P1)$ and $(P2)$, respectively). For each "dollar" benefit of the inflation-adjusted annuity, this annuity would actually pay out one dollar if both members of a couple are alive, but that figure would drop by one-third if one member of a couple passes away.[9] If the respondent was single at the time of the survey, the annuity is structured to pay out one dollar per year until the respondent passes away.

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

[9] This mirrors the spending pattern found by Shapiro (2009), which shows that consumption drops by about one-third after the death of a spouse.

Sustainable consumption, then, is calculated as:

$$SC_j = Wealth_j / a_j$$

Using the shock amounts asked about in the hypothetical wealth loss (F16) and gain (F17) questions, we then find the variables of interest, $(R_{SC}^{loss} - R_0)$ and $(R_{SC}^{gain} - R_0)$. These are the change in retirement age that would be needed to hold sustainable consumption at its pre-shock level, even after a gain or loss. For the "loss" variable, we numerically solve the following equation for the number of years by which the retirement age would need to change to hold sustainable consumption constant after such a wealth shock (R_{SC}^{loss}):

$$\sum_{s=1}^{R_{SC}^{loss}} \left(earnings_{j,12} \times \left(\frac{1}{1+r}\right)^{s} \times P1_{j,s} \right) = \left[\sum_{s=1}^{R_0} \left(earnings_{j,12} \times \left(\frac{1}{1+r}\right)^{s} \times P1_{j,s} \right) \right] + (loss)$$

where $(loss)$ is the hypothetical loss amount. For the "gain" variable, the equation is the same, replacing all "loss" variables with "gain" and changing the sign prior to $(loss)$, which is then solved for R_{SC}^{gain}.

$$\sum_{s=1}^{R_{SC}^{gain}} \left(earnings_{j,12} \times \left(\frac{1}{1+r}\right)^{s} \times P1_{j,s} \right) = \left[\sum_{s=1}^{R_0} \left(earnings_{j,12} \times \left(\frac{1}{1+r}\right)^{s} \times P1_{j,s} \right) \right] - (gain)$$

These numbers are then employed to calculate our independent variables of interest: $(R_{SC}^{loss} - R_0)$ and $(R_{SC}^{gain} - R_0)$.

3. Empirical Framework

Our empirical framework and its justification is based on McFall (2011). In our multivariate analysis, we use a Tobit specification model to regress measures of the change in the subjective probability of full-time work after reaching age 65 due to the hypothetical wealth shock on the number of years by which retirement would have to shift to hold sustainable consumption at pre-shock levels and that number squared. This model is specified:

$$\Delta\Pr(Work\ FT\ after\ 65)_j^* = \alpha + \beta_1(R_{SC} - R_0)_j + \beta_2{(R_{SC} - R_0)_j}^2 + \varepsilon_j$$

$$\Delta\Pr(Work\ FT\ after\ 65)_j = \max\left(0, \Delta\Pr(Work\ FT\ after\ 65)_j^*\right)$$

$$\varepsilon_j|(R_{SC} - R_0)_j, Z_j \sim N(0, \sigma^2)$$

The Tobit specification is appropriate because it can provide consistent estimates even in the presence of heaping at a corner solution (such as our data display), at least for individuals not at the corner solution. The latent variable, $\Delta\Pr(Work\ FT\ after\ 65)_j^*$, can be thought of as the optimal change in the probability of full-time work after age 65 that would result in the absence of fixed costs and discontinuities.

Considering the likelihood of further non-linearities in the relationship between changes in labor force expectations and the relative magnitudes of asset losses, our preferred specifications include a squared term of the original regressor. This is important if there is an increasing marginal disutility of work. For example, individuals facing a loss might be willing to adjust along the labor supply dimension less and less as they get further from their original plan. The squared term may help absorb the impact of regressor values that are so large that recouping losses is practically unattainable, either due to life expectancy or personal or strong cultural resistance to working to extremely advanced ages.

4. Descriptive Analyses

Table 1 provides summary statistics about the ALP (panels A and B) and CogEcon (panels C and D) samples, separately for the hypothetical loss and gain sub-samples, respectively.[10] Because the sub-samples of each survey overlap almost completely, it is reassuring to see that the statistics for these are only slightly different. In both ALP samples, slightly over one half of respondents are female, and just under one third are single (defined as identifying as “single” or “partnered, but not planning a financial

[10] Descriptions of the ALP subsample answering the hypothetical gain question are nearly identical to those answering the loss question, as the two subsamples overlap almost completely. For the CogEcon, the hypothetical gain sample is identical to the hypothetical loss sample.

future together"). Mean, median and standard deviations of annual earnings, age and age at which respondents expect to "retire completely" are also similar across the samples. The CogEcon sub-samples are slightly older and have somewhat lower levels of mean earnings, though the median earnings match up well with the ALP sub-samples.

This is a somewhat higher-earning group than the average population in this age range: the CPS general population estimate for weekly earnings for full-time workers between 55 and 64 is $904, which works out to approximately $45,000 for the slightly older group. However, our sample is necessarily somewhat select, in that it includes only respondents with tax-advantaged retirement account holdings; low-earning individuals are less likely than higher-earning individuals to have retirement accounts.

Table 2 displays summary statistics about the value of tax-advantaged retirement savings reported by respondents, and the size of the hypothetical shocks asked about in the survey experiments. The median holdings for both the CogEcon and the ALP are approximately $110,000, with the 25th percentile at about one third that, and the 75th percentile at about three times that. By construction, the hypothetical wealth shocks shown in the next column are 30% of the retirement savings values for the ALP, and a bit more for CogEcon (due to the way the question was asked in the mail mode portion of the sample). The "shock" amounts are substantial, but are generally smaller than respondents' annual earnings.[11]

Annuitizing all financial assets and the present discounted value of future earnings yields the distribution of sustainable consumption estimates illustrated in the top panel of Table 3. Note that CogEcon estimates of sustainable consumption are higher than that of the ALP due to the more comprehensive measures of wealth (defined benefit pensions and Social Security, for example) available in the survey. In the case of hypothetical losses, the sustainable consumption distributions shift downward, with more compression at the top of the wealth distribution than the bottom. In the case of

[11] In the ALP data, a linear regression of the shock size on earnings yields a coefficient of 0.523 (s.e. 0.018); the R-squared statistic for this regression is 0.671.

hypothetical gains, the both sustainable consumption distributions shift upward, with more expansion at the top of the wealth distribution than the bottom.

The within-person changes in sustainable consumption levels implied by the hypothetical shocks are summarized in Table 4. The interquartile range of changes runs from just over $400 to approximately $4,000 in both ALP samples.[12] The interpretation of the median here is that the median ALP respondent would lose $1,439 in consumption every year until death under the hypothetical loss scenario, holding retirement timing constant; the median respondent would gain $1,582 in consumption every year until death under the hypothetical gain scenario, holding retirement timing constant. While likely not enough to move many individuals into or out of poverty or the "top 1%," these changes would be a meaningful proportion of individuals' annual spending *for the rest of their lives*. The patterns are qualitatively similar for both CogEcon waves, though the magnitudes of changes are slightly larger.

By how much would an individual's retirement age need to shift to hold sustainable consumption at pre-shock levels? Table 5 presents summary statistics for this measure. The median number of years by which retirement age would need to shift to offset the hypothetical shock is around a year in the case of either losses or gains in the ALP. Only about a quarter of the samples would need to shift their retirement ages by approximately two years or more. For CogEcon, the hypothetical losses (gains) would require modestly longer (shorter) durations of work to recover the original level of sustainable consumption.

While Table 5 presents the shifts in retirement age that would be needed to offset the impact of wealth shocks, Table 6 presents summary statistics that shed light on how individuals actually *think* they would react to such a wealth loss or gain. Columns (1) and (4) display summary statistics about question F15, which asked about respondents' subjective probability of working full-time after reaching age 65. The average ALP response (column 1) indicated about a 60 percent chance that a respondent would be

[12] These are not perfectly symmetric because the gain and loss sub-samples include a different set of respondents in both surveys.

working full-time after age 65; the median was also 60 percent. One quarter of respondents thought there was a 20 percent chance or lower of working full-time beyond age 65, while one quarter of respondents thought there was a 90 percent chance or greater of working full-time beyond 65. For CogEcon (column 4), the mean baseline probability is 41 percent, potentially due to the older age of the sample. Again in this sample, observations cover the full distribution of probabilities.

Columns (2) and (5) display differences between the response to question F15 and the subjective probability of full-time work after age 65 in the case of the hypothetical wealth loss. For both datasets, while half of both samples report no changes, more than a quarter of the sample reported a change, with the top quartiles reporting changes of at least 20 percentage points. The mean changes are about 10 percentage points in both datasets. Overall, 47 percent of ALP respondents and 49 percent of CogEcon respondents gave answers indicating at least a 10 percentage point increase in the probability of full-time work in reaction to the hypothetical loss. In contrast, just under 6 percent of ALP respondents and 9 percent of CogEcon respondents gave answers indicating a decrease in the probability of full-time work.

Columns (3) and (6) display the change in subjective probability of full-time work after age 65 in the case of the hypothetical wealth gain. While the median changes are zero in both datasets, more than a quarter of the sample reported a change, with the bottom quartile of the distribution reporting at least a 10 percentage point change. The mean change is 3.7 (6.4) percentage points reductions in the probability of working at age 65 for ALP (CogEcon) respondents. Overall, approximately thirty percent of respondents on both surveys gave answers indicating at least a 10 percentage point decrease in the probability of full-time work in reaction to the hypothetical gain. In contrast, about 11 percent of ALP respondents and 13 percent of CogEcon respondents gave answers indicating an increase in the probability of full-time work with a gain.

As expected, the impact of a hypothetical gain on the probability of full-time work after age 65 is less pronounced than the impact of a hypothetical loss. Fewer respondents indicated changes in the theoretically expected directions in the case of a gain than a loss (and only a small minority indicated changes counter to expectations, in both the case of a loss and a gain). In the case of the wealth gain, the distribution of changes was also more compressed toward zero.

Figure 2 uses a different approach to characterize respondents' reports of the probability of full-time work, using ALP data. The bins along the x-axis of the histogram represent the answer options for the questions about the subjective probability of full-time work. The first bar in each bin displays the frequency of responses under the respondents' actual circumstances. The second bar in each bin displays the frequency of responses in the case of a hypothetical wealth loss. The third bar in each bin displays the frequency of responses in the case of a hypothetical wealth gain. In eight out of the eleven bins, the first bar is between the heights of the other two columns, suggesting that the hypothetical gain and loss move the probability of full-time work in opposite directions. In the leftmost several bins, the hypothetical gain bar tends to be the highest, indicating that respondents are moving to lower probabilities of full-time work after 65 after a gain. Similarly, in the rightmost bins, the hypothetical loss bar tends to be the highest, suggesting that respondents move to higher probabilities of full-time work after a loss.

5. Regression Analyses

We have examined variables characterizing the hypothetical wealth shocks, in terms of the years by which retirement age would need to shift to offset the impact of the losses or gains on sustainable consumption levels. These variables, $(R_{SC}^{loss} - R_0)$ and $(R_{SC}^{gain} - R_0)$, enter into the regressions below as the independent variables of interest. As described above, the dependent variable in each regression is the reported change in the subjective probability of full-time work from the respondent's actual situation to the case in which the respondent has experienced a hypothetical loss or gain of approximately 30 percent of his or her tax-advantaged retirement savings. In the case of the loss, the dependent variable is $\Delta \text{Pr}(Work\ FT\ after\ 65)_{loss}$; in the case of the gain, it is $\Delta \text{Pr}(Work\ FT\ after\ 65)_{gain}$. A graphical scatterplot of these relationships in the CogEcon samples can be seen in Figure 3.

As described in the empirical framework, we use Tobit regressions due to the heaping at zero (likely caused by underlying fixed cost of adjustment and/or non-linearities in the underlying optimization problem). For each sample, we present results from a Tobit specification including only the relevant $(R_{SC} - R_0)$ variable and the square of that regressor.

Table 7 presents the regression results considering the impact of a hypothetical wealth loss on the probability of working full-time after reaching age 65. In Column (1), the estimated coefficient on the linear term is positive, as expected, and precisely measured at the 10 percent level. The squared term is slightly negative and not statistically significant. The marginal effect is positive and significantly different from zero at the 5 percent level of significance.

Columns (2) and (3) of Table 7 present results from similar analysis using the CogEcon data. These specifications are identical to the ALP analysis in column (1), but pool together two years of data and adjust the standard errors for clustering at the individual level. Column (2) presents the results omitting a dummy variable for year, while column (3) includes it. The marginal effects are nearly identical in magnitude to those measured in the ALP in spite of being measured in a completely different sample of respondents.

Table 8 presents a comparison of the wealth loss results from our ALP analysis (Col 1) and HRS analysis (Col 2) with those from the Great Recession estimated by McFall (2011b) (Col 3). The similarities between the sizes of the losses and the changes in the probability of full-time work past 65 show that the three analyses examined shocks of relatively comparable magnitudes, and that respondents reported similar changes in the probabilities of working full-time after 65. Strikingly, the estimated marginal effects are also of similar magnitudes. To convert these results into something resembling change in retirement age, the last row in the table uses estimates from McFall (2011b) that equate a 1 percentage point change in the probability of working full-time past age 65 to at least 6.7 days' delay of retirement (per McFall, this estimate is likely a lower bound). The ALP analysis implies that for an individual experiencing a loss that would take 1.78 extra years of work to make up (the mean in the sample), we expect that individual to retire about 14 days earlier (6.7days/yr x 1.146p.p. x 1.78yrs=13.7 days). Using the same conversion method, the CogEcon analysis implies a delay of 22 days, and McFall (2011b) estimated a delay of approximately 25 days due to losses in the Great Recession. It is interesting that the magnitudes of the effects are similar between the analyses using hypothetical losses and the analysis examining the Great Recession. We take this as evidence that respondents may be able to answer questions about the impact of hypothetical losses on expectations questions, and subjective probabilities, in particular, in a way that is reflective of the way they would react to an actual loss.

Table 9 presents the regression results considering the impact of a hypothetical wealth gain on the probability of working full-time after reaching age 65. Column 1 shows the results of the ALP analysis, in which the estimated coefficient on the linear term is negative and significant at the 1 percent level. The squared term is smaller, also negative and statistically significant at the 1 percent level. The marginal effect is also negative, and significantly different from zero at the 1 percent level of significance. Using the rule-of-thumb of 6.7 days retirement change per percentage point effect (McFall (2011b)), this estimate implies that a wealth gain that would take one fewer year of work than planned to offset results in a shift forward of retirement date by of just under 10 days. At the mean, the impact of a sudden, 30% increase in one's tax-advantaged retirement account holdings works out to approximately 18 days, or just over half a month, of a shift forward in retirement.

In Column 2, we repeat the specification from Columns 1, but exclude respondents who would have had to retire in the past to end up with the pre-gain level of sustainable consumption. The results are very similar to Column 1, with an implied effect on retirement timing of 9 days at the mean. In Column 3, by contrast, the CogEcon estimates imply a smaller change in retirement age, but are also quite imprecisely measured. It reveals virtually no measurable response in retirement to a hypothetical positive wealth shock. The coefficient on the size of the shock remains negative, meaning that larger positive shocks are measured to lead to lower probabilities of full time work past 65, though this small coefficient is very imprecisely measured.

Comparing this null effect in the CogEcon sample to the positive effect of a *loss* on the probability of full time work may imply asymmetric responses to shocks that are gains rather than losses, as we expected from theory. Surprisingly, though, we do not see evidence of asymmetry in the impact of wealth losses versus gains on the probability of full-time work after reaching age 65 in the regression results for the ALP. Based on the summary statistics about the survey questions used to create the dependent variables (Table 6), we expected to see evidence of less adjustment in the case of a gain than a loss. At this stage, we are still the differences between the “gain” results between samples, and the differences between the “gain” and “loss” analyses.

6. Conclusion

This paper makes use of unique data from survey experiments on the American Life Panel as well as the Cognitive Economics Study on the impact of hypothetical wealth shocks on retirement timing. Specifically, in this survey respondents were asked about their subjective probability of full-time work after reaching age 65. They were then asked to answer the same question, considering the impact of a hypothetical loss of approximately 30 percent of their current tax-advantaged retirement savings balance, followed by the same question about a 30 percent gain. The data are drawn from two distinct nationwide samples of Americans between the ages of 50 and 65. While both are slightly wealthier than the average population of this age range, likely at least partially due to the requirement that respondents had tax-advantaged retirement savings accounts, results from this survey may nonetheless provide some insight into the potential policy impact of changes in the Social Security benefit structure or changes in the taxation of retirement account disbursements.

McFall (2011a,b) found a relationship between wealth lost in the Great Recession and retirement expectations in the Health and Retirement Study and the Cognitive Economics survey. Past empirical work on the impact of wealth shocks on retirement expectations has yielded mixed results, with some researchers finding no impact, and others finding measureable impacts in the expected direction. This work applies McFall's approach to examine the impact of *hypothetical* retirement wealth gains and losses on the subjective probability of working full-time past age 65. This analytical approach specifically takes into account potential fixed costs and non-linearities in the optimization problem the life-cycle hypothesis implies that individuals may solve when deciding on retirement timing. Additionally, using data from survey questions specifically-designed to test the impact of hypothetical wealth shocks enables us to calculate the variables most needed to examine the research question using empirical strategies grounded by theory.

Analyses show that the labor supply reactions to hypothetical wealth shocks are in the expected directions, in that losses result in increased probabilities of full-time work beyond age 65, while gains result in decreased probabilities of full-time work beyond age 65, relative to baseline. Attenuation bias due to measurement error in component data needed to calculate total wealth or sustainable consumption is likely a significant factor in the regression results; the estimated marginal effects might be best interpreted as lower bounds.

In the analyses examining the impact of a hypothetical loss (Table 7), the marginal effect at the mean value of the regressor is 1.146 percentage points (s.e. 0.552) for the ALP sample and very similar for the CogEcon sample. That is, a hypothetical wealth loss that would take an additional year of work to "make up," in terms of sustainable consumption, increases the subjective probability of full-time work after reaching age 65 by just over one percentage point. The ALP result is statistically significant at the 5 percent level. At the mean value of the regressor, 1.78 years delay in time to retirement that would be needed to make up lost sustainable consumption power, this works out to just over 2 percentage points' change in the probability of full-time work after reaching age 65. Using the approximate conversion metric estimated by McFall (2011), this works out to approximately 14 days, or just under half a month, of retirement delay at the mean. The estimates using the much smaller CogEcon sample are not statistically significant, but are of very similar magnitude.

In the analyses of the impact of a hypothetical gain for the ALP sample (Table 9, Column 1), the marginal effect at the mean value of the regressor is 1.443 percentage points (s.e. 0.429). That is, a hypothetical wealth gain that would take one fewer year of work to hold sustainable consumption constant at the pre-gain level decreases the subjective probability of full-time work after reaching age 65 by almost one and a half percentage points. This result is statistically significant at the 1 percent level. At the mean value of the regressor, 1.84 years shift in time to retirement, this works out to 2.66 percentage points' change in the probability of full-time work after reaching 65. Using the approximate conversion metric estimated by McFall (2011b), this works out to approximately 18 days, or just over half a month, shift forward of retirement at the mean. The CogEcon result is of the same sign but smaller, and again not statistically significant.

We also compare our results examining the impact of a hypothetical loss with McFall's (2011b) findings about the impact of the Great Recession on the probability of full-time work after age 65. The overall effect of the Great Recession wealth losses on the reported changes in retirement age is estimated to be very similar to the estimates from our experimental samples. We interpret this as suggesting that respondents may be able to answer questions about the impact of hypothetical losses on expectations questions, and subjective probabilities, in particular, in a way that is reflective of the way they would react to an actual loss.

This may have important policy implications, since using the framework used in this paper to analyze results from a survey question about a hypothetical change may provide valid insight into the likely effect of a proposed tax or Social Security policy change. Furthermore, it implies that changes to Social Security benefit amounts or tax treatment of retirement savings may be effectual policy levers to affect retirement timing, a conclusion that is not obvious from the existing literature examining the impact of wealth shocks on labor force decisions. Given the likely significant impact of attenuation bias in our results, the true impact of policy changes resulting in wealth changes of the order analyzed in this paper may be much greater.

7. References

American Life Panel. (2014). MS 307: Decision Making. Version March 2014. Produced and distributed by the Rand Corporation. Module MS 307 produced with funding from the National Institute on Aging (grant number NIA P01 AG026571).

Cognitive Economics Study. (2014). 2013 Public Use Dataset. Version 1.2. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA P01 AG026571), Ann Arbor, MI.

Cognitive Economics Study. (2014). 2011 Public Use Dataset. Version 1.1. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA P01 AG026571), Ann Arbor, MI.

Goda, G. S., J. B. Shoven, and S. N. Slavov (2011, May). What explains changes in retirement plans during the Great Recession? *American Economic Review 101 (3)*, 29-34.

Gustman, A. L., T. L. Steinmeier, and N. Tabatabai (2009, October). 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*.

Gustman, A. L., T. L. Steinmeier, and N. Tabatabai (2015). Retirement and the Great Recession. *The Journal of Retirement 3 (1),* 87-106.

McFall, B. H. (2011). Crash and Wait? The Impact of the Great Recession on the Retirement Plans of Older Americans. *American Economic Review 101 (3).*

McFall, B. H. (2011). Essays in Labor Economics. Doctoral dissertation, University of Michigan. Available at http://hdl.handle.net/2027.42/89744.

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

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

Modigliani, F. and R. Brumberg (2005). Utility analysis and the consumption function: An interpretation of Cross-Section data. In F. Franco (Ed.), The Collected Papers of Franco Modigliani, Volume 6, pp. 3-45. Cambridge, Massachusetts: MIT Press. (Originally published in 1954).

Ondrich, J. and A. Falevich (2016). The Great Recession, housing wealth and the retirement decisions of older workers. *Public Finance Review 44 (1)*, 109-131.

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

Social Security Administration. Period Life Table 2013, Actuarial Life Tables. http://www.ssa.gov/OACT/STATS/table4c6.html. Downloaded June 20, 2017.

8. Tables and figures

Table 1: Sample Descriptions

	Mean	Median	St. Dev.
Panel A: ALP hypothetical loss sample (N=378)			
Female	0.51	--	--
Single	0.28	--	--
Annual Earnings	$108,959	$52,000	$430,801
Age	56.86	56.68	4.06
Planned Retirement Age	66.94	66	5.04
Panel B: ALP hypothetical gain sample (N=402)			
Female	0.52	--	--
Single	0.27	--	--
Annual Earnings	$104,076	$50,000	$418,253
Age	56.89	56.65	4.07
Planned Retirement Age	67.55	66	6.57
Panel C: CogEcon hypothetical loss sample (weighted, N=254)			
Female	0.46	--	--
Single	0.16	--	--
Annual Earnings	$69,936	$54,000	$102,522
Age	59.69	59.55	2.62
Planned Retirement Age	66.61	66	3.67
Panel D: CogEcon 2013 (weighted, N=264)			
Female	0.48	--	--
Single	0.168	--	--
Annual Earnings	$68,393	$50,000	$101,855
Age	59.67	59.55	2.64
Planned Retirement Age	66.69	66	3.96

Table 2: Size of tax-advantaged retirement savings and hypothetical shocks to accounts from survey experiments (ALP and CogEcon)

	Tax-Advantaged Retirement Savings (ALP)	Hypothetical Loss/Gain (ALP)	Tax-Advantaged Retirement Savings (CogEcon)	Hypothetical Loss/Gain (CogEcon)
Mean	\$300,507	\$90,152	\$305,895	\$88,705
25th percentile	\$30,000	\$9,000	\$41,613	\$16,005
Median	\$110,000	\$33,000	\$106,699	\$50,148
75th percentile	\$300,000	\$90,000	\$412,105	\$117,369
St. Dev.	\$889,017	\$266,705	\$436,507	\$125,754
Observations	402	402	254	254

Note: Data source is ALP and CogEcon. Statistics shown are for samples used in "loss" analyses, but, within study, are nearly identical for "gain" samples.

Table 3: Sustainable Consumption Levels, Actual and After Hypothetical Shock (ALP and CogEcon)

	Hypothetical Loss Sample (ALP)	Hypothetical Gain Sample (ALP)	Hypothetical Loss Sample (CogEcon)	Hypothetical Gain Sample (CogEcon)
Panel A: Actual Estimated Sustainable Consumption				
Mean	\$81,920	\$82,073	\$102,183	\$102,705
25th percentile	\$22,201	\$22,591	\$51,782	\$52,235
Median	\$40,174	\$41,920	\$82,109	\$82,695
75th percentile	\$73,533	\$74,108	\$112,288	\$112,695
Panel B: Sustainable Consumption After Hypothetical Shock				
Mean	\$77,959	\$86,225	\$98,281	\$107,213
25th percentile	\$21,310	\$23,427	\$47,280	\$55,889
Median	\$38,510	\$44,059	\$77,299	\$85,885
75th percentile	\$69,835	\$78,663	\$107,903	\$122,005
Observations	378	402	254	264

Note: Data source is ALP and CogEcon. For ALP, calculation is based on DC pension holdings, other financial accounts, businesses and other assets, plus net real estate wealth. It does not include defined benefit pension or Social Security wealth, nor does it subtract out non-mortgage debt. For CogEcon, calculation is based on DC holdings, other financial accounts, businesses and other assets plus real estate wealth, less credit card and other debt. In both datasets, sustainable consumption is computed as the ratio of wealth to annuity price, where annuity price takes into account gender, life expectancy and assumed a load factor of 0.18, and a real interest rate of 2.9 percent. Sample size differs because, in some cases, loss is too large relative to earnings to ever make up; these observations are excluded from the regressions and, therefore, other analyses in this paper.

Table 4: Changes in sustainable consumption levels generated by hypothetical gains/losses to retirement wealth

	Hypothetical Loss (ALP)	Hypothetical Gain (ALP)	Hypothetical Loss (CogEcon)	Hypothetical Gain (CogEcon)
Mean	-$3,961	$4,153	-$4,219	$4,194
25th percentile	-$3,794	$431	-$5,380	$778
Median	-$1,439	$1,582	-$2,311	$2,311
75th percentile	-$422	$4,061	-$771	$5,491
Observations	378	402	252	262

Table 5: Change in years-to-retirement required to hold sustainable consumption constant under a hypothetical change to retirement wealth

	Hypothetical Loss (ALP)	Hypothetical Gain (ALP)	Hypothetical Loss (CogEcon)	Hypothetical Gain (CogEcon)
Mean	1.78	-1.84	2.99	-0.92
10th percentile	0.08	-3.98	0.51	-3.06
25th percentile	0.37	-1.94	1.14	-1.76
Median	1.11	-1.08	1.98	-0.48
75th percentile	1.98	-0.37	3.43	0.18
90th percentile	3.71	-0.08	6.53	0.75
Standard deviation	2.64	2.66	3.32	2.00
Observations	378	402	254	264

Table 6: Subjective probabilities of full-time work after age 65 at baseline, and changes to probabilities under a hypothetical changes to retirement wealth

	Pr(FT after 65) ALP	ΔPr(FT after 65) ALP		Pr(FT after 65) CogEcon	ΔPr(FT after 65) CogEcon	
	(1) Baseline Probability	(2) With Loss	(3) With Gain	(4) Baseline Probability	(5) With Loss	(6) With Gain
Mean	59.4 percent	10 p.p.	-3.7 p.p.	41.0 percent	12.3 pp	-6.4 pp
25th percentile	20 percent	0 p.p.	-10 p.p.	10 percent	0 pp	-10 pp
Median	60 percent	0 p.p.	0 p.p.	30 percent	0 pp	0 pp
75th percentile	90 percent	20 p.p.	0 p.p.	80 percent	20 pp	0 pp
Observations	402	378	402	254	254	254

Note: Based on responses from 11-point probability scales. Fox example, "10 pp" change means respondent chose checkboxes that were 10 percentage points apart for the subjective probability of full-time work after 65 given current conditions, and after a hypothetical wealth loss (or gain).

Table 7: Tobit estimates of the effect of a hypothetical loss on the probability of working full time at age 65

	(1)	(2)	(3)
Sample:	**ALP**	**CogEcon**	**CogEcon**
Marginal effect at mean	1.146	1.098	1.034
	(0.552)**	(0.719)	(0.727)
$(R_{sc}-R_0)$	2.927	2.750	2.528
	(1.505)*	(2.284)	(2.306)
$(R_{sc}-R_0)^2$	-0.122	-0.075	-0.061
	(0.090)	(0.122)	(0.123)
Constant	-6.458	-7.711	-5.836
	(2.936)**	(6.890)	(7.660)
Year indicator	n/a	No	Yes
Sigma	31.837	36.176	36.116
	(1.897)***	(4.671)***	(4.703)***
Observations	378	254	254
Number uncensored obs.	177	124	124
Pseudo R-squared	0.002	0.0034	0.0036
Log-Likelihood	-995.8	-572.5	-572.4
Chi-squared	4.474		
P-value of chi-squared	0.034		

Note: Data source is the American Life Panel 2012 and the Cognitive Economics Study 2011 and 2013. Results reported for Tobit regressions where the dependent variable is difference between hypothetical change in subjective probability of working full-time after reaching age 65 and the lower censoring point is 0. The mean of the explanatory variable $(R_{sc}-R_0)$ *is 1.78 years for the ALP and 2.99 years for CogEcon. Columns (2) and (3) report standard errors that are adjusted for clustering at the individual level. *** p<0.01, ** p<0.05, * p<0.1.*

Table 8: Comparison of Impact of Hypothetical Loss with Realized Losses from the Great Recession

	Hypothetical Loss		Great Recession
Source	Current paper	Current paper	McFall (2011)
Sample	ALP	CogEcon	Health and Retirement Study
Median loss in sustainable consumption	4.14%	2.52%	4.96%
Median value of (R_{sc}-R_0)	1.11 years	1.98 years	1.88 years
Change in Pr(FT at 65) – Mean	10 p.p.	12 p.p	8.1 p.p.
Change in Pr(FT at 65) – Median	0 p.p.	0 p.p.	2 p.p.
Change in Pr(FT at 65) - 75th percentile	20 p.p.	20 p.p.	25 p.p.
Estimated Marginal Effect (s.e.)	1.146 p.p. (0.552)***	1.098p.p (0.719)	0.736 p.p. (0.235)***
Marginal Effect x Mean	2.04 p.p.	3.28 p.p.	3.67 p.p.
Interpretation: Effect on Retirement Timing	14 days[ᵻ]	22 days[ᵻ]	25 days[ᵻ]

*Note: Data source is the American Life Panel, Cognitive Economics Study, and the Health and Retirement Study. The ALP version of the Pr(FT at 65) questions offered multiple choice response options of 0%, 10%, 20%... 80%, 90%, 100%. In the Health and Retirement Study, respondents could give an integer answer ranging from 0 to 100 (percent). [ᵻ] denotes calculations based on figures presented in McFall (2011). These numbers may be subject to rounding error. CogEcon sustainable consumption measure is based on wealth that includes the expected present value of Social Security and defined benefit pensions. CogEcon results presented in this table use regression results from column (2) of table 7, which were generated in an analysis without a year indicator variable.*** p<0.01, ** p<0.05, * p<0.1.*

Table 9: Tobit Regressions of the Effect of a Hypothetical Gain on the probability of working full time at age 65 (ALP and CogEcon)

	(1)	(2)	(3)
Sample:	ALP	ALP	CogEcon
Marginal effect at mean	-1.443	-1.516	-0.312
	(0.429)***	(0.490)***	(0.554)
$(R_{sc}-R_0)$	-7.239	-6.913	-0.899
	(2.245)***	(2.396)***	(1.790)
$(R_{sc}-R_0)^2$	-0.644	-0.551	-0.006
	(0.216)***	(0.234)**	(0.246)
Constant	-22.972	-22.849	-12.550**
	(3.886)***	(4.036)***	(5.637)
Year indicator	n/a	n/a	Yes
Sigma	30.619	31.005	35.703***
	(2.278)***	(2.418)***	(3.345)
Observations	402	364	264
Number uncensored obs.	122	111	90
Pseudo R-squared	0.008	0.007	0.0008
Log-Likelihood	-724.6	-660.5	-434.8
Chi-squared	12.040	9.862	
P-value of chi-squared	0.001	0.002	

Note: Data source is the American Life Panel 2012 and the Cognitive Economics Study 2011 and 2013. Results reported for Tobit regressions where the dependent variable is the hypothetical change in subjective probability of working full-time after reaching age 65, with an upper censoring point of 0. The mean of the explanatory variable of interest $R_{sc}-R_0$ *is -1.84 years in the ALP and -0.92 years in CogEcon. Column (3) reports standard errors that are adjusted for clustering at the individual level. *** p<0.01, ** p<0.05, * p<0.1.*

Figure 1: Subjective Probability of Working Full-Time after 65 (ALP, N=378)

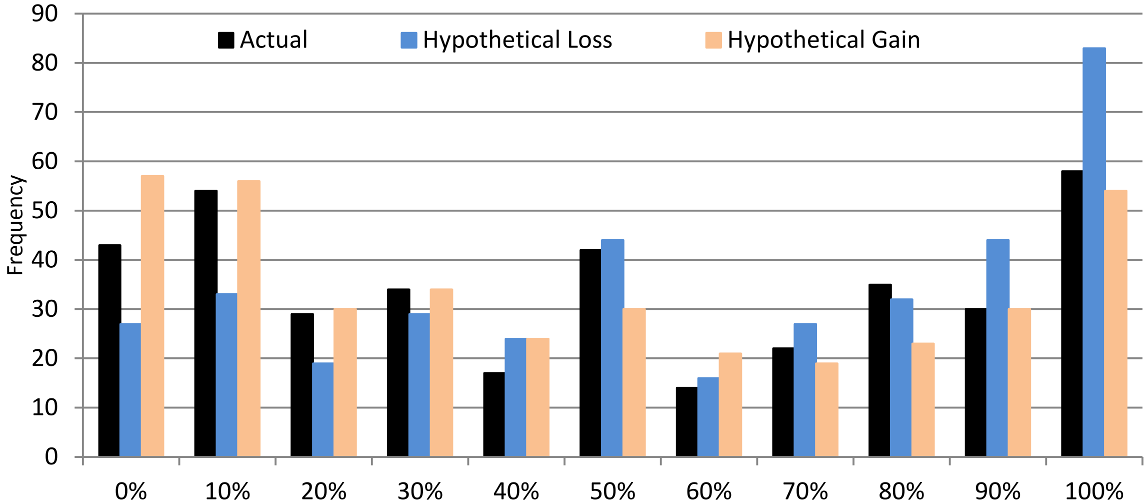

Note: For consistency of number of observations across series, this figure displays only the subset of 378 observations used in the hypothetical loss regressions.

Figure 2: : Scatterplot of the change in the probability of full time work at age 65, by magnitude of hypothetical wealth losses and gains (CogEcon)