Essays in Labor Economics

by

Brooke Helppie McFall

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2011

Doctoral Committee:

Professor Robert J. Willis, Chair
Professor Trivellore E. Raghunathan
Professor Matthew D. Shapiro
Professor Jeffrey Smith

For my family.

ACKNOWLEDGEMENTS

I would like to express my gratitude to family, friends, colleagues and faculty members for the help and support that have allowed me to complete this dissertation. There are many people whose efforts combined to get me to the finish line, and I will be forever grateful for their help.

First and foremost, I thank my committee. Bob Willis was originally assigned to be my third-year paper advisor, but his roles through the years have included valued mentor, kind teacher, brilliant collaborator, and supportive supervisor. His contagious enthusiasm and incredible mind have inspired me; his anecdotes, advice and unflagging belief in my abilities have sustained me. I am also indebted to Matthew Shapiro for his help and guidance in the latter stages of the dissertation process. His attention to detail and determination that I should get more out of my dissertation has greatly improved this final product, and I look forward to his continued mentorship and sound advice. Thanks, also, to Jeff Smith, for the friendly smile, helpful suggestions and expertise in applied econometrics; and to Trivellore Raghunathan for his time and advice, especially at the early stages of the Job Seekers' Project.

I am also grateful to my friend, colleague and collaborator, Marta Murray-Close. She has devoted immense time and effort to making the Job Seekers' Project a success. Thanks also to Carol Bowen and Aneesa Buageila for help with the grant proposals and budgets that kept the Job Seekers and Cognitive Economics Projects moving forward. I gratefully acknowledge the Sloan Foundation, the A. Regula Herzog Young Investigators Fund, and the Survey Research Center for support of the Job Seekers'

Project, as well as the Social Security Administration and the National Institute on Aging for funding my assistantships. For their excellent work keeping the Job Seekers and Cognitive Economics Projects on track over the past year, I thank Uniko Chen and Mike Gideon.

Greg Clark and Christian Weller's support before graduate school allowed me to embark on this adventure. Charlie Brown, Mick Couper, Miles Kimball, David Lam, and David Weir offered encouragement and helpful comments at the early stages of my survey and research work. For crucial early reads of my work, I thank Pamela Giustinelli, Joanne Hsu, Jo Ann McFall and Matt Rutledge. For positive feedback when I needed it most, I thank Stephan Lindner, Zoë McLaren, Mary Beth Ofstedal and Purvi Sevak. For advice on some important details, I thank Peter Hudomiet, Brian Kovak, Jon Lanning and Shanthi Ramnath. For the little slip of paper taped to my monitor, I thank Jessica Faul. For her ever-positive attitude and unwavering support, I thank Gwen Fisher. For her kind words, and for helping me figure out the graduate school process, I thank Mary Braun.

For the twice-annual flights to Brazil that got me through my first Michigan winters, as well as so much other support, I thank my mother and Abel. For standing by me through it all, I thank my husband, Curtis. For their love and support– and their understanding when I couldn't be there– I thank all of my friends, my siblings, the Michigan Helppies, the McFalls, the Swarm and Ellie.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

ABSTRACT

Essays in Labor Economics

by

Brooke Helppie McFall

Chair: Robert J. Willis

This dissertation is composed of two essays, both which use data from original survey projects to examine issues related to work choice.

The first essay 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 essay 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 consis-

tent 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.

The second essay is a product of the Job Seekers Study. The essay extends Mincer's seminal theory of family migration to allow couples to adjust to migration constraints by living apart, and examine the ways in which new PhD economists adjust to migration constraints imposed on them by their spouses or partners. Both the impact of migration constraints on job outcomes and the impact of job considerations on relationship outcomes are analyzed. The essay finds that migration constraints result in small costs in terms of job outcomes, relative to many existing studies, and that adjustment through living apart is common. These results imply that existing studies may overestimate the job impact of migration constraints.

CHAPTER I

Introduction

This dissertation is composed of two essays. Although each essay falls within the broad area of labor economics, the greatest commonality between them is that both use data from original survey projects in which I have played a role. Because the essays make use of data from surveys that were designed to answer the very research questions addressed in my essays, the analyses I conduct are quite simple, yet have direct relationships to economic theory.

In the case of the first essay, my involvement in the Cognitive Economics Study over the recent economic crisis both kindled my interest in the impact of the crisis on retirement age and allowed me to help gather the data I would later use to study this issue. In this essay, I examine the labor supply effects of the wealth losses during the stock market crash of 2008 and 2009. A simple life-cycle model incorporating both consumption and retirement timing implies that exogenous wealth losses should delay optimal retirement timing. The short time-period over which massive amounts of wealth were lost presents a natural experiment with which to study the implications of this model.

Using pre- and post-crash data from the Cognitive Economics Study and the Health and Retirement Study, I quantify 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 analysis and Tobit regressions 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. I find that the 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 broadly consistent with the results of other recent studies and life-cycle model, but stand in contrast to other examinations of actual wealth shocks on the general population of older Americans, which have tended to find little relationship between changes in wealth and retirement timing.

The Job Seekers Study, a web-based survey project undertaken with Marta Murray-Close and Robert J. Willis, was developed with the express goal of examining the job market decisions of economists. The second essay in this dissertation, co-written with Marta Murray-Close, is the product of more than four years of original data collection work. In this essay, we examine the ways in which new Ph.D. economists adjust to migration constraints imposed on them by their spouses or partners.

Mincer's (1978) unitary model of family migration predicts that a couple will either move together to a location that maximizes the members' joint utility or break up. If an individual moves to a location that is not the best option for her, individually, she is said to be a "tied" migrant. In this essay, we adapt Mincer's theory to allow for a third margin of adjustment, the choice to live apart, while maintaining a relationship.

By using data from a survey explicitly designed to examine migration decisions, as

well as job outcomes, break-ups and long-distance relationships, we analyze both the impact of migration constraints on job outcomes and the impact of job considerations on relationship outcomes. Using a combination of statistical tests and linear regression analyses, we find that migration constraints result in small costs in terms of job outcomes, relative to many existing studies. We also find that adjustment along the relationship margin, through decisions to maintain long-distance relationships, are quite common. Our results imply that, by excluding individuals who adjust along this margin, existing studies may have resulted in overestimates of the job impact of migration constraints.

Together, the essays in this dissertation underscore the advantages of using data gathered to answer specific research questions. Because the data used in these essays contain direct measures of variables relevant to my research questions, relatively simple analyses yielded substantively interesting results about how individuals make life decisions, including how, when and where to supply their labor.

CHAPTER II

The Impact of the Great Recession on the Retirement Plans of Older Americans

2.1 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 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.

2.2 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.

2.3 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 2.1 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 A.1, I present a simplified version of a life-cycle model of consumption and labor supply developed by Kimball and Shapiro (2008; 2003). Figure 2.2 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 2.2 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 2.3. 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 2.4 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.

2.4 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

$$\triangle retirement\ timing_j = \alpha + \beta_1 \left(R_{\overline{sc}} - R_0\right)_j + Z_j'\gamma + \varepsilon_j \tag{2.1}$$

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 2.5. 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]

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:

$$\triangle retirement\ timing_j^* = \alpha + \beta_1 \left(R_{\overline{sc}} - R_0\right)_j + Z_j'\gamma + \varepsilon_j \tag{2.2}$$

$$\triangle retirement\ timing_j = max\left(0, \triangle retirement\ timing_j^*\right) \tag{2.3}$$

[1] Indeed, evidence of heaping, seen in my descriptive analysis (Section 2.6.1), suggests that fixed costs and other discontinuities are important in predicting the adjustment of retirement plans to wealth losses in 2008 and 2009.

$$\varepsilon_j | \left(R_{\overline{sc}} - R_0\right)_j, Z_j \sim N(0, \sigma^2) \tag{2.4}$$

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.

$$\triangle retirement\ timing_j = \alpha + \beta_1 \left(R_{\overline{sc}} - R_0\right)_j + \beta_2 \left(R_{\overline{sc}} - R_0\right)_j^2 + Z_j'\gamma + \varepsilon_j \tag{2.5}$$

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.

2.5 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

[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.

advantage of the strengths of both studies, I conduct analyses using both datasets.

2.5.1 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 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]

[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 2.6. 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.

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.

2.5.2 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 expecta-

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

[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.

tions 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. Table 2 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 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

[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.

be different from the average person of comparable age.[12]

2.5.3 Measurement

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

$$\triangle retirement\ timing_j = \alpha + \beta_1 \left(R_{\overline{sc}} - R_0\right)_j + Z_j'\gamma + \varepsilon_j \tag{2.6}$$

and

$$\triangle retirement\ timing_j = \alpha + \beta_1 \left(R_{\overline{sc}} - R_0\right)_j + \beta_2 \left(R_{\overline{sc}} - R_0\right)_j^2 + Z_{ij}'\gamma + \varepsilon_j \tag{2.7}$$

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.

2.5.3.1 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

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

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 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

[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."

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.

2.5.3.2 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 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

[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.

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 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

[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.

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:

$$EPDV(earnings)_j = \sum_{s=\tau}^{R_0} \left((1 - UE\ rate_j) \times \frac{earn}{(1+r)^s} \right) \tag{2.8}$$

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

[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.

by county-level FIPS code.[19] In calculations of post-crash wealth, Equation 2.8 uses county unemployment rates from May 2009.[20] 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.

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

[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 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.

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 A.2 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, $EPDV\ earnings_{08,j}$ 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

[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.

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 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

[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.

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 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

[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.

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 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,

[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.

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:

$$EPDV(earnings)_j = \sum_{s=\tau}^{R_0} \left((1 - UE\ rate_j) \times \frac{earn}{(1+r)^s} \right) \tag{2.9}$$

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

[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.

$EPDV(earnings)$) 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 information from several variables. The imputation of this variable is described in detail in Appendix A.3. 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 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

[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.

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.

2.6 Results

2.6.1 Descriptive analysis

To provide a sense of the material impact of the shock on sustainable consumption levels, Table 2.2 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

[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

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 2.3 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 2.7 displays the reported changes in retirement age since July 2008. The mean change reported by all respondents in

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.

this sample was 1.6 years, with median of 0 years and a change of 3 years at the 75th percentile. In Figure 2.8, 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 for the subjective probability of full-time work after age 62 and 65. Table 2.5 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 ($\triangle retirement\, timing$) compare with the extra years respondents would need to work to attain their pre-crash sustainable consumption paths ($R_{\overline{sc}} - R_0$). Table 2.4 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 2.9 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 2.10 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 2.4.

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.

2.6.2 Regression analysis

As discussed in Section 2.3, 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 2.4, 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 2.4 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 2.4), 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 2.6.1 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 2.6.2.4 do not provide significant cause for concern about the Tobit specification, most analyses presented in this paper use the Tobit specification.

2.6.2.1 CogEcon base regressions

Table 2.6 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 2.6 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 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 $R_{\overline{sc}} - R_0$

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

[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.

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 2.11 and 2.12. In Figure 2.11, the predicted the probability of an increase in retirement age (based on the Tobit regression) and the proportion of respondents actually reporting 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 2.12, 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.

2.6.2.2 HRS base regressions

Tables 2.7 and 2.8 present the results from Tobit regressions like those in Columns 3 and 4 of Table 2.6, 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 2.7, 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

[32]See Appendix A.4 for the results from these regressions.

about six days in the planned age of retirement.

The first column of Table 2.7 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 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 2.13 and 2.14. Figure 2.13 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 2.14 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 under-prediction 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 sub-sample, 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 2.8, 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 2.7, 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 2.15 and 2.16. Figure 2.15 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 2.16 illustrates the predicted and observed adjustments along the intensive margin. Again, consistent with the attenuation bias discussion with respect to Figures 2.12 and 2.14, it can be seen that the Tobit regression under-predicts the size of the reported changes.

Columns 3 and 4 of Table 2.8 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.

2.6.2.3 Comparison of CogEcon and HRS findings

Table 2.9 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 2.9 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 2.17 through 2.22 allow for a visual comparison of the results from the different datasets. Figures 2.17 and 2.18 present the results from the preferred CogEcon specification presented in Column 4 of Table 2.6, 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 2.18 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 retire-

ment 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 2.19 and 2.20 are parallel graphs for the preferred full-sample HRS regressions using ${}_{08}\triangle_{09}Pr(FT62)$ as the dependent variable, and Figures 2.21 and 2.22 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.

2.6.2.4 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 2.10 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 2.11 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 A.5.

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 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.

2.6.2.5 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*

[33] See Wooldridge (2002) page 539.

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 2.12 and 2.13, 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 2.12, I present results based on exclusion of primary home equity.

In Table 2.13, 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 A.6.

2.6.2.6 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 2.14 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 2.14 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 2.15, 2.16 and 2.17 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 2.15, 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 2.6. 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 2.16, 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 2.6, 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.

[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.

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 2.16 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 2.17 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 extremely 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

[37] Because county-level unemployment data are not seasonally-adjusted, I have used unemployment rates from exactly one year apart, with the end date coinciding with the CogEcon 2009 survey fielding.

to the base specification. Indicators for this variable are not statistically significant, and do not greatly change the results from the base specification (Table 2.6, 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

[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.

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 required 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 2.18 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 2.6, 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 2.19 and 2.20, 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 2.19, 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 2.7, 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 2.20 repeats the analyses in Table 2.19 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 2.19. 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.

2.7 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 retire-

ment 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.

2.8 Figures and tables

Figure 2.1: Life-cycle saving and consumption

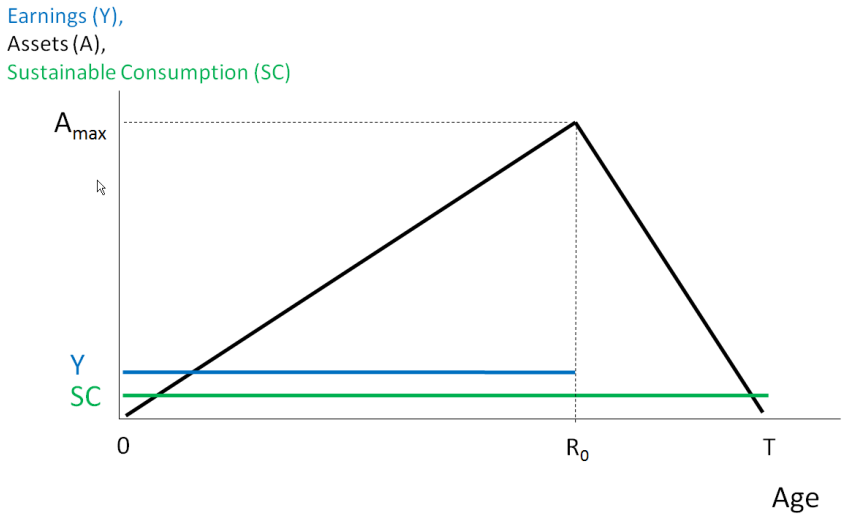

Figure 2.2: Optimal retirement choice

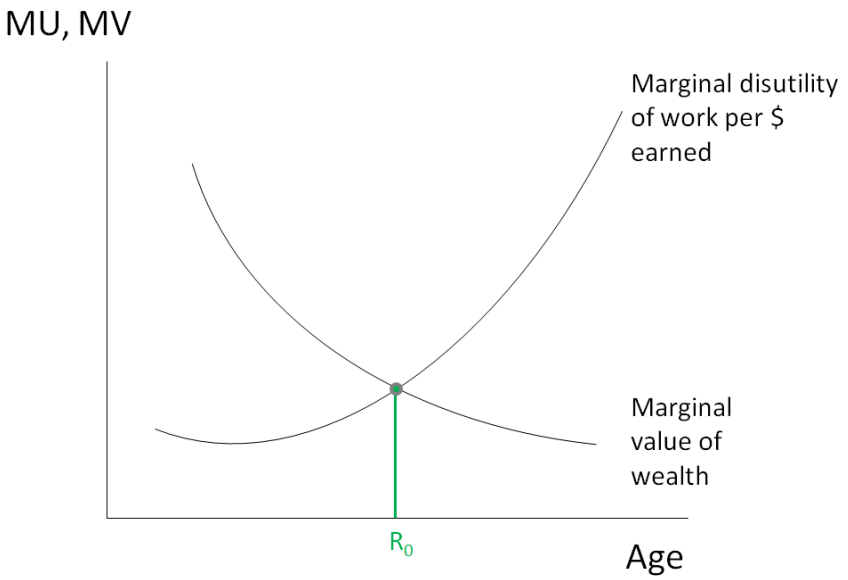

Figure 2.3: Optimal retirement choice after a wealth shock

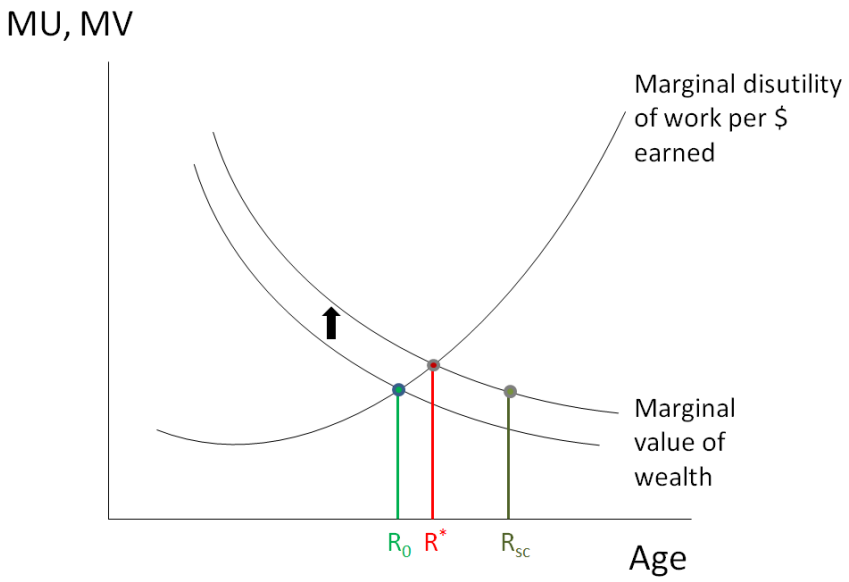

Figure 2.4: Life-cycle saving and consumption with variable retirement timing

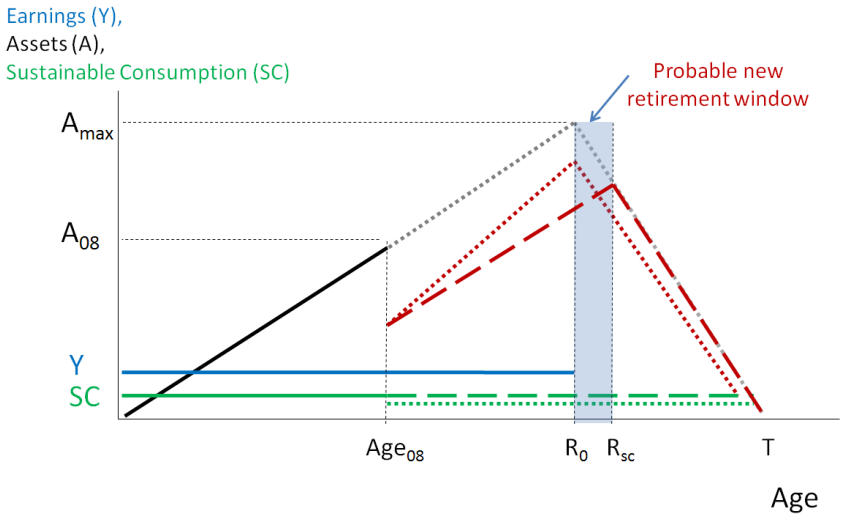

Figure 2.5: Dow Jones Industrial Average closing values

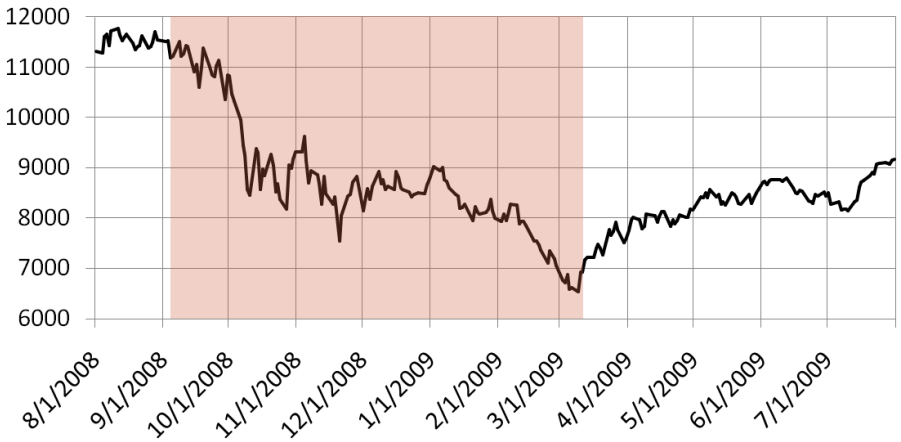

Source: Yahoo! Finance

Figure 2.6: Timeline of surveys and the Dow Jones Industrial Average

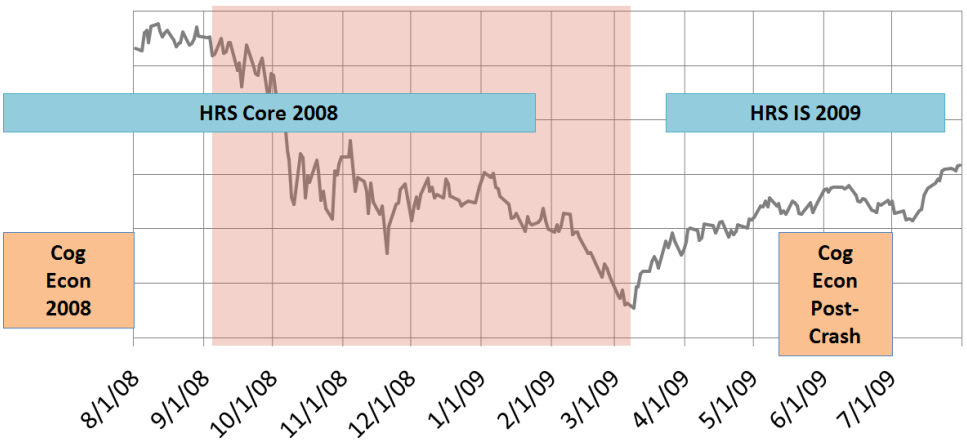

Figure 2.7: Changes in retirement age owing to crash (CogEcon sample)

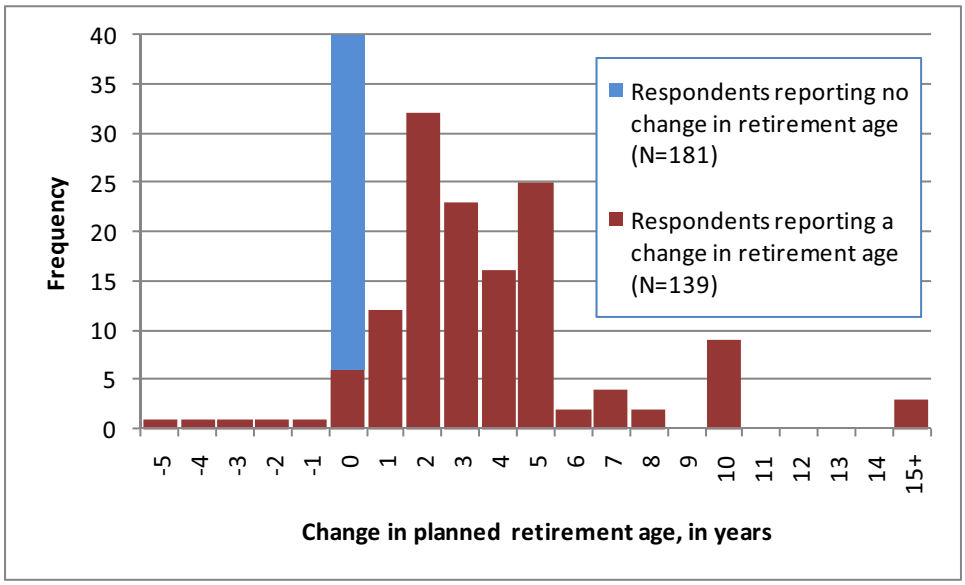

Figure 2.8: Cumulative distribution of expected retirement ages (CogEcon sample)

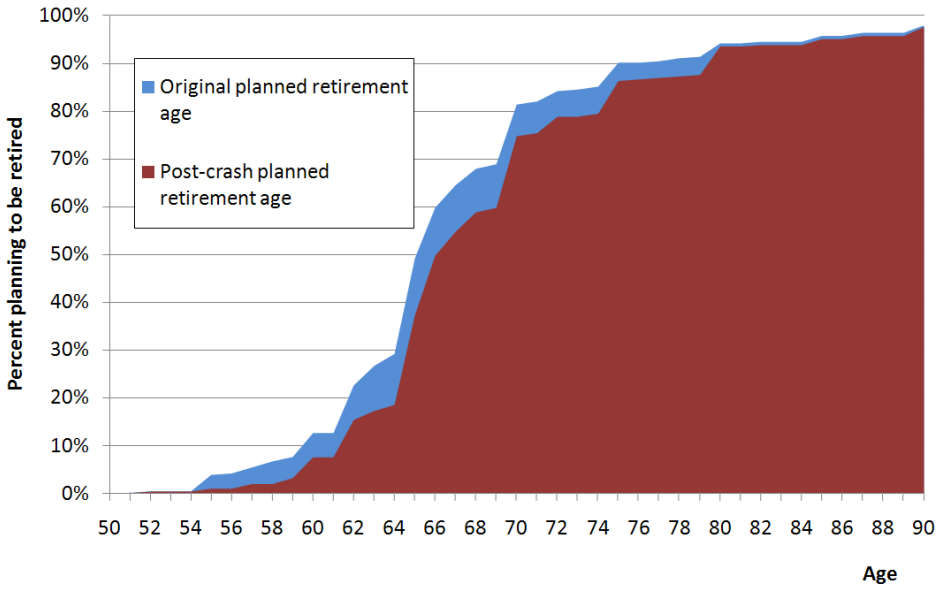

Figure 2.9: Comparison of retirement age needed to attain pre-crash consumption path and planned post-crash retirement age (CogEcon sample)

Figure 2.10: Comparison of changes in retirement age needed to attain pre-crash consumption path and reported changes (CogEcon sample)

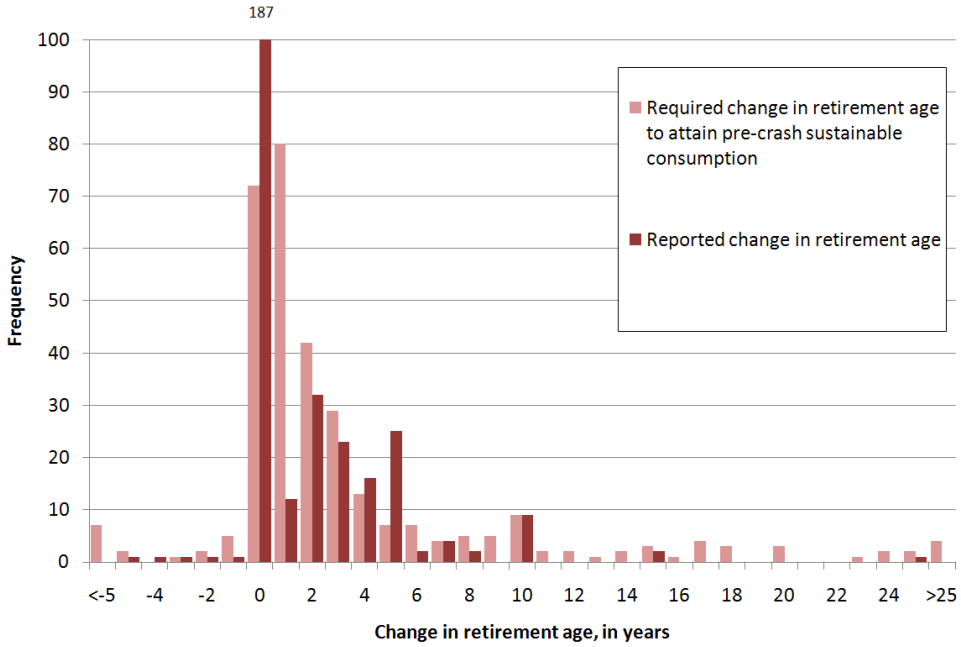

Figure 2.11: Extensive margin: Tobit prediction versus observed probability of increase in planned retirement age (CogEcon sample)

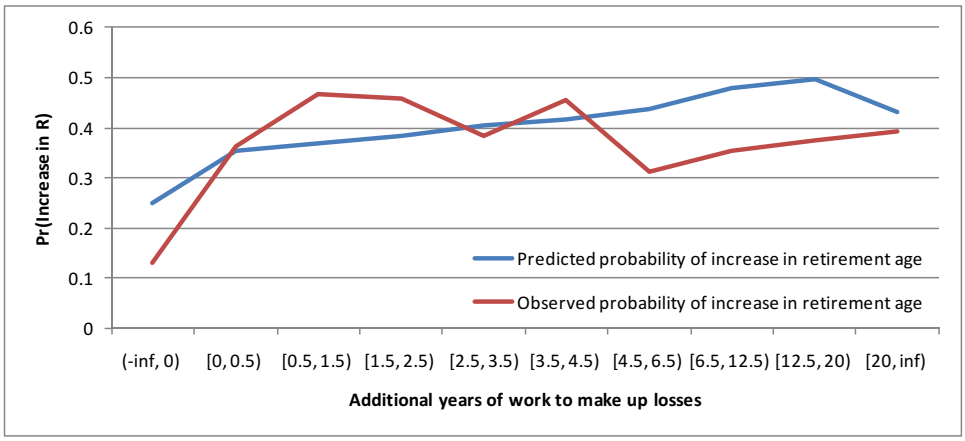

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is probability of an increase in the planned age of retirement. Lines are plotted by connecting the average for each bin.

Figure 2.12: Intensive margin: Tobit prediction versus observed increase in planned retirement age (CogEcon sample)

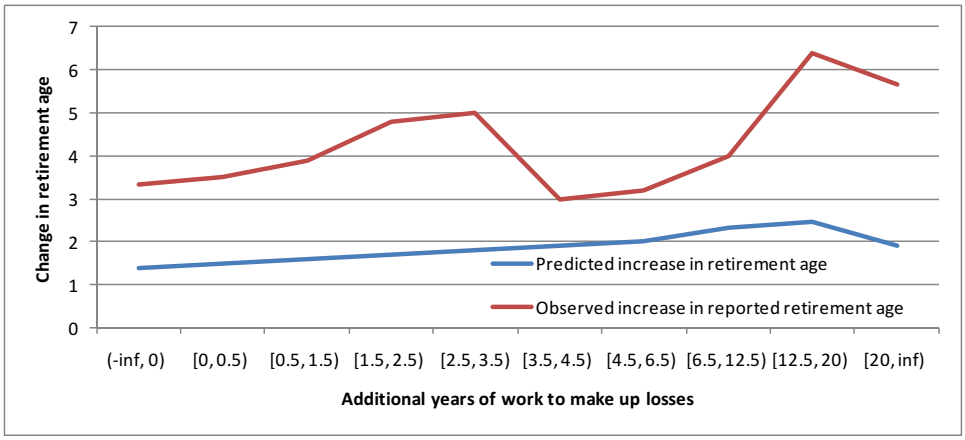

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis represents years of increase in the planned age of retirement. Lines are plotted by connecting the average for each bin.

Figure 2.13: Extensive margin: Tobit prediction versus observed probability of increase in Pr(FT62) (HRS <62 sample)

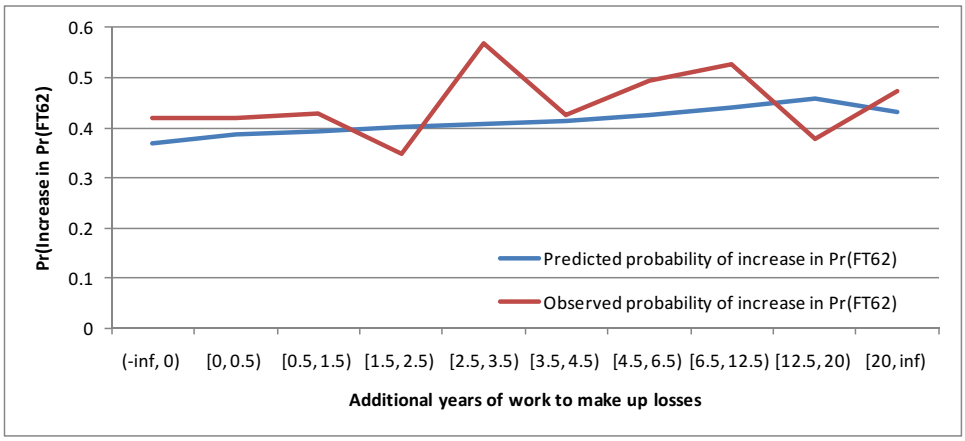

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is probability of an increase in the subjective probability of full-time work after age 62 per year of $R_{\overline{sc}} - R_0$. Lines are plotted by connecting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

Figure 2.14: Intensive margin: Tobit prediction versus observed increase in Pr(FT62) (HRS <62 sample)

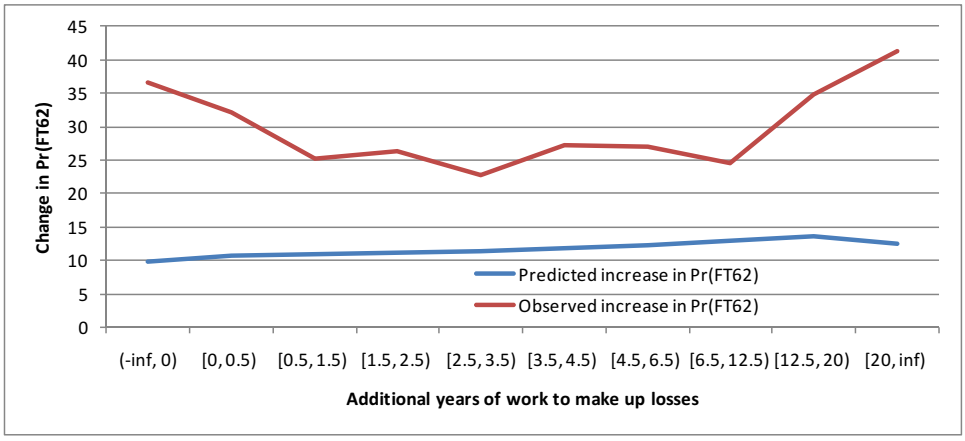

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is the increase in the subjective probability of full-time work after age 62 per year of $R_{\overline{sc}} - R_0$. Lines are plotted by connecting the average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

Figure 2.15: Extensive margin: Tobit prediction versus observed probability of increase in Pr(FT65) (HRS <65 sample)

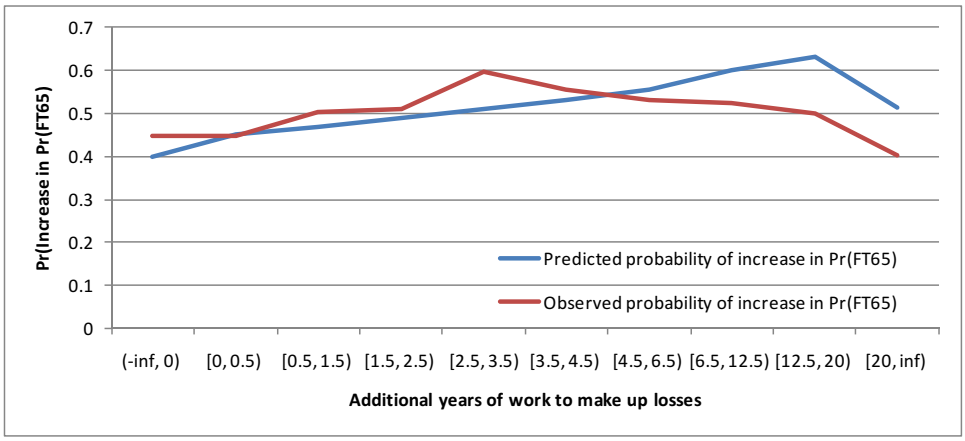

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is probability of an increase in the subjective probability of full-time work after age 65 per year of $R_{\overline{sc}} - R_0$. Lines are plotted by connecting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

Figure 2.16: Intensive margin: Tobit prediction versus observed increase in Pr(FT65) (HRS <65 sample)

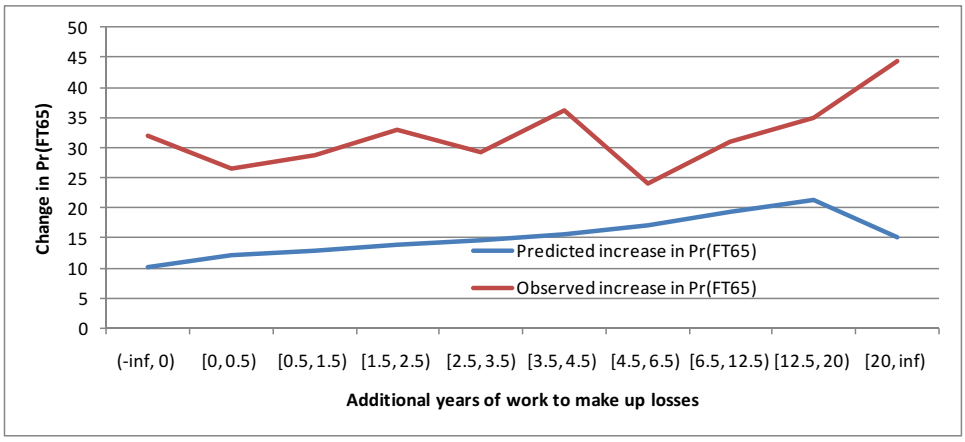

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is the increase in the subjective probability of full-time work after age 65 per year of $R_{\overline{sc}} - R_0$. Lines are plotted by connecting the average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

Figure 2.17: Average marginal effects by $R_{\overline{sc}} - R_0$ group (CogEcon sample)

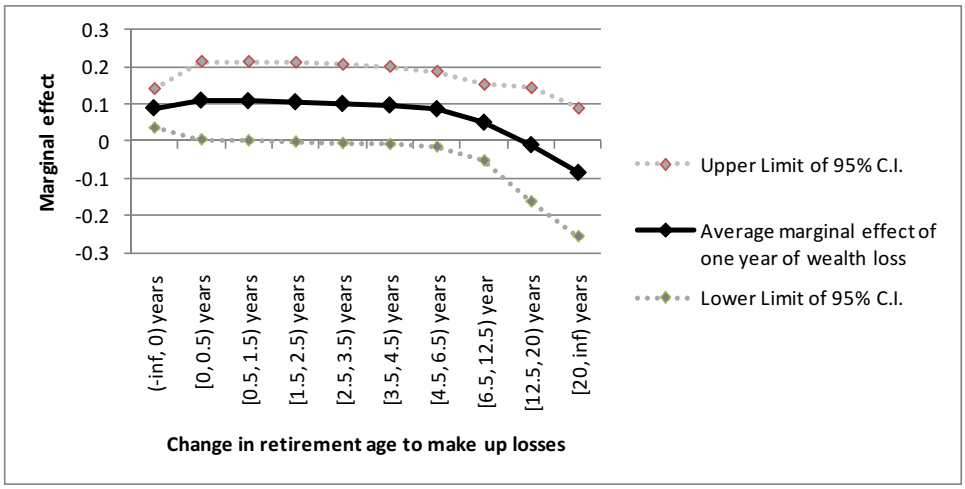

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Estimates based on results from regression shown in Column 4 of Table 2.6.

Figure 2.18: Average predicted change in retirement age by $R_{\overline{sc}} - R_0$ group (CogEcon sample)

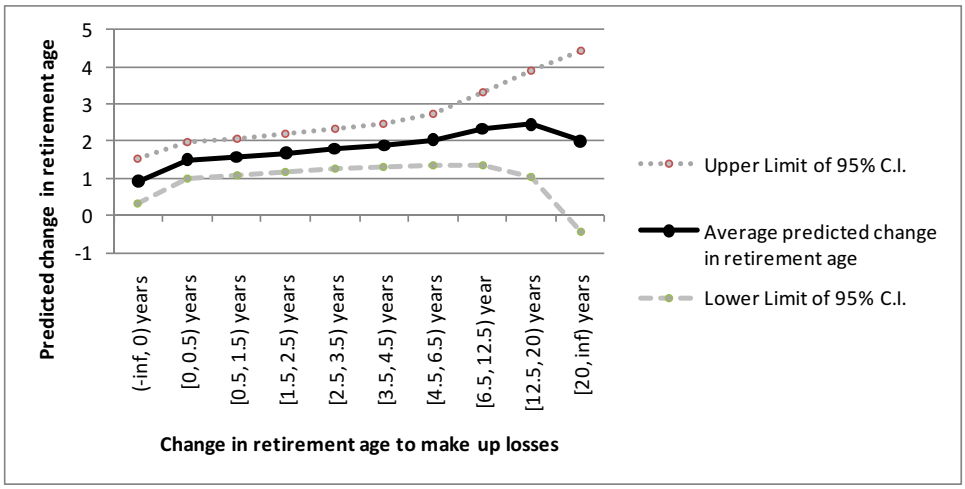

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Estimates based on results from regression shown in Column 4 of Table 2.6.

Figure 2.19: Average marginal effects by $R_{\overline{sc}} - R_0$ group (HRS <62 sample)

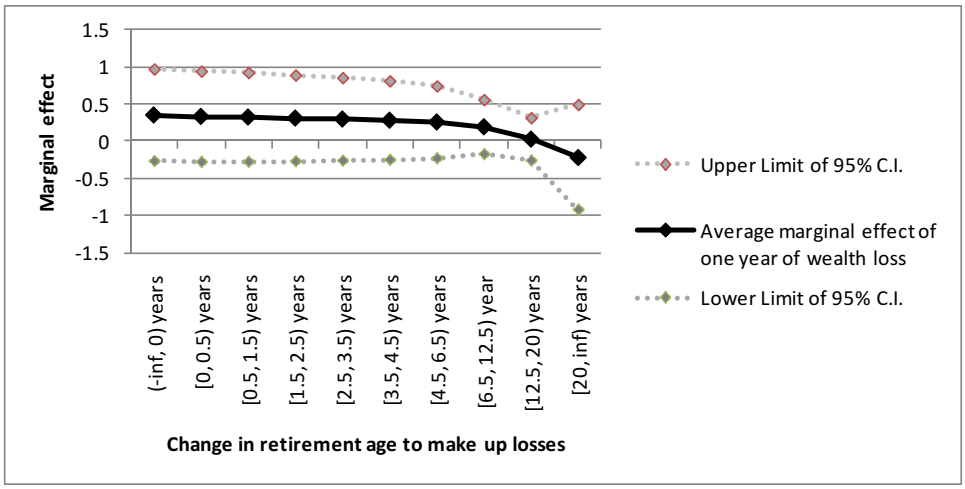

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is percentage point change in the probability of full-time work after age 62 per year of $R_{\overline{sc}} - R_0$. Estimates based on results from regression shown in Column 2 of Table 2.7.

Figure 2.20: Average predicted change in probability of full-time work by $R_{\overline{sc}} - R_0$ group (HRS <62 sample)

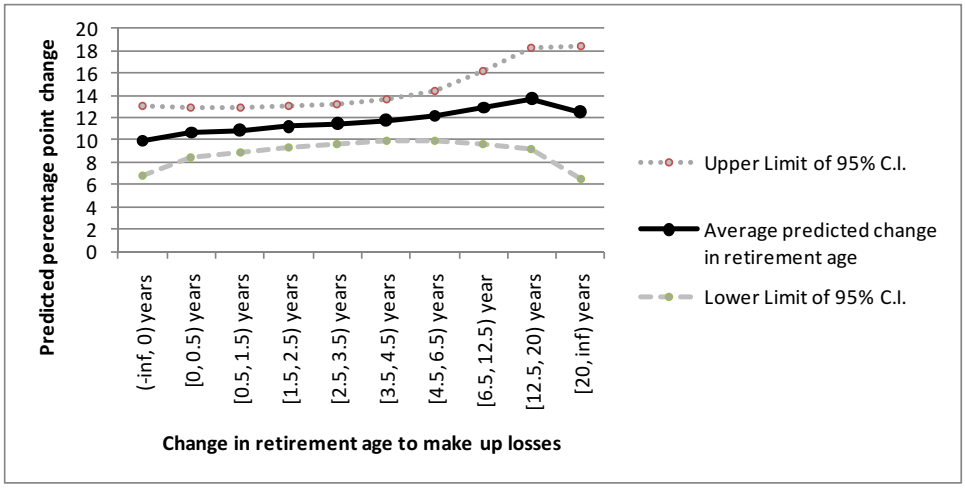

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is percentage point change in the probability of full-time work after age 62 per year of $R_{\overline{sc}} - R_0$. Estimates based on results from regression shown in Column 2 of Table 2.7.

Figure 2.21: Average marginal effects by $R_{\overline{sc}} - R_0$ group (HRS <65 sample)

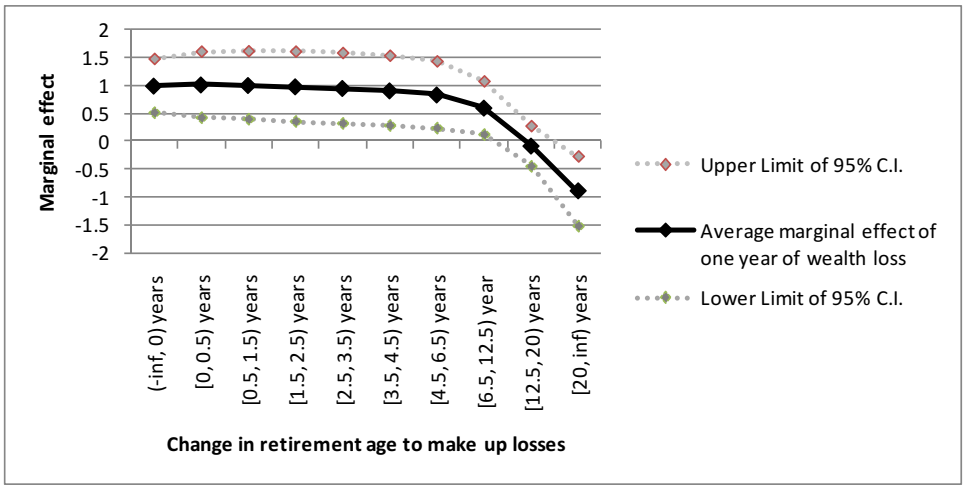

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is percentage point change in the probability of full-time work after age 62 per year of $R_{\overline{sc}} - R_0$. Estimates based on results from regression shown in Column 2 of Table 2.8.

Figure 2.22: Average predicted change in probability of full-time work by $R_{\overline{sc}} - R_0$ group (HRS <65 sample)

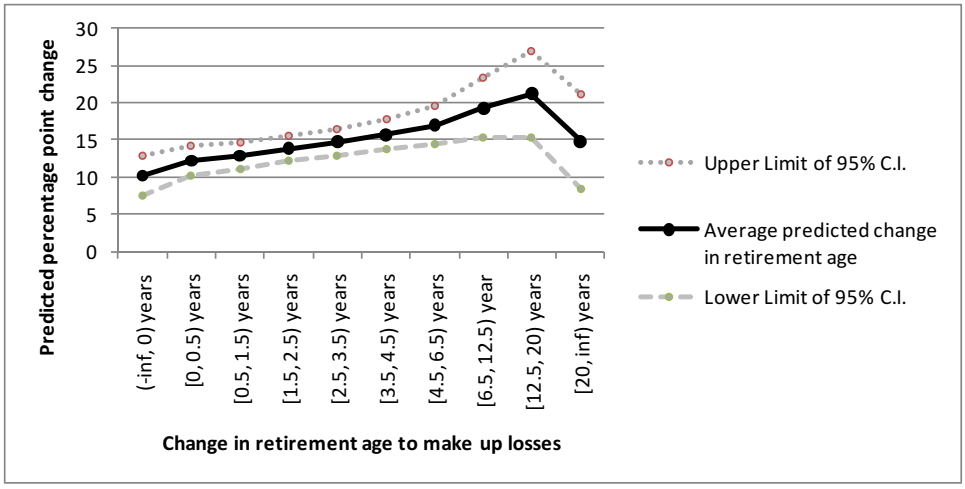

Notes: Horizontal axis categories are bins representing different wealth effect sizes (in terms of $R_{\overline{sc}} - R_0$). Vertical axis is percentage point change in the probability of full-time work after age 62 per year of $R_{\overline{sc}} - R_0$. Estimates based on results from regression shown in Column 2 of Table 2.8.

Table 2.1: 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 2.2: 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 2.3: 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 2.4: 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 2.5: 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 2.6: 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 2.7: Impact of wealth losses on probability of full-time work after age 62 (HRS <62 sample)

	(1)	(2)	(3)	(4)
Specification:	Tobit	Tobit	Tobit	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 2.8: 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 2.9: 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 2.10: 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 2.11: 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 2.12: Robustness check excluding primary home from total wealth

Sample	(1) CogEcon sample	(2) HRS <62 sample	(3) 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 2.13: Robustness check excluding expected bequests from total wealth

Sample	(1) HRS <62 sample	(2) 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 2.14: 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 2.15: 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 2.16: 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

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

Optimism Indicators:	(1) All	(2) Stock Market	(3) Labor Market	(4) Housing Market
Pseudo-R^2	-380.1	-379	-433	-438.9
Log-Likelihood	0.0134	0.0219	0.0243	0.00574
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 2.17: 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 2.18: 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)*

Tercile Indicators Used:	(1) Financial Knowledge (FK)	(2) Financial Knowledge (FK)	(3) Fluid Intelligence (FI)	(4) 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 2.19: 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 2.20: 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.

CHAPTER III

Moving Out to Move Up? New Economists Sacrifice Job Opportunities for Proximity to Significant Others– and Vice Versa

With Marta Murray-Close

3.1 Introduction

In the past half century, historic increases in women's labor-force participation have prompted growing interest in the migration decisions of dual-career couples. In 1970, 41 percent of married women in the United States were in the labor force. By 2009, 61 percent of married women were in the labor force, and married couples with two earners outnumbered married couples with a single earner (*United States Census Bureau*, 2010). Unlike couples in which only the husband or, less frequently, only the wife works for pay, couples in which both partners work must balance the potentially competing demands of two careers. Career-related migration opportunities, which can arise for partners at different times and in different locations, may be a source of conflict for these couples.

How couples respond to conflicting locational preferences has implications for their well-being. On one hand, living together may harm the career prospects of one or

both partners (*Sandell*, 1977; *Mincer*, 1978); on the other hand, living apart may harm their relationship. Empirical studies suggest that living together constrains the location choices of married workers (*Costa and Kahn*, 2000; *Gemici*, 2008; *McKinnish*, 2008; *Mincer*, 1978) and lowers their earnings relative to what they could obtain in their individually optimal locations (*Sandell*, 1977; *Lichter*, 1983; *Jacobson and Levin*, 1997; *Gemici*, 2008; *Boyle et al.*, 2001). Because many studies have found that the negative impact of living together falls disproportionately on women (*Mincer*, 1978; *Sandell*, 1977; *Jacobson and Levin*, 1997; *Compton and Pollak*, 2007; *Cooke*, 2008; *Boyle et al.*, 2001), some researchers have also suggested that the migration decisions of couples contribute to the gender gap in earnings and career attainment (*Bielby and Bielby*, 1992). On the relationship side of the trade-off, one study found that career-motivated migration is associated with higher divorce rates (*Gemici*, 2008).

Implicit in early theories of family migration was an assumption that couples who remain together live together (*Sandell*, 1977; *Mincer*, 1978). Subsequent empirical work has shown that this assumption is unwarranted. A number of qualitative studies outside of economics have identified long-distance relationships as an alternative to career sacrifices or relationship dissolution for couples with conflicting locational preferences (*Gerstel and Gross*, 1982; *Magnuson and Norem*, 1999; *Gross*, 1980; *Rhodes*, 2002). These studies have explored the circumstances under which dual-career couples live apart and have assessed the implications of living apart for the careers and relationships of couples who choose the arrangement. Because they have relied on non-representative samples, however, studies of long-distance relationships have not estimated the prevalence of living apart in the population of dual-career couples. Nor can we be confident that the results of these studies generalize to the population.

This paper uses data from original surveys of new entrants to the junior PhD job market in economics – all of whom have invested heavily in their human capital, most of whom will move for their first job in thier field of training, and many of whom have

highly educated partners – to assess the impact of conflicting locational preferences on a group for whom the problem is likely to be severe. The surveys combine questions about the job-market decisions and outcomes of new entrants to the junior PhD job market in economics, hereafter referred to as "new economists," with questions about their partners and living arrangements. In addition, the surveys contain direct counterfactual questions about the job-market outcomes new economists think they would have had if they had responded differently to conflicts over location. Using data from the surveys, we are able to characterize the impact of conflicting locational preferences on the career outcomes of new economists who live with their partner. We are also able to estimate, for the first time, the prevalence and predictors of long-distance relationships in a known sub-population of dual-career couples.

Our results indicate that the impact of conflicting locational preferences on choices new economists make between job offers is modest. Just 14 percent of the partnered economists we surveyed had rejected their first-choice job offer for the benefit of their relationship and, among those who had rejected their first-choice job offer, the sacrifices entailed were not large. At the same time, 16 percent of the economists expected to be living apart from their partner in the year after they entered the job market. Economists who faced large career costs of living with their partner were the most likely to live apart. In light of these patterns, we argue that long-distance relationships attenuate the impact of conflicting locational preferences on the career outcomes of new economists.

Our results corroborate several findings from the qualitative literature on long-distance relationships. Dual-career couples are motivated to live apart when the benefits of the arrangement to their careers are large, but they are not primarily concerned with financial compensation (*Gerstel and Gross*, 1982; *Gross*, 1980; *Magnuson and Norem*, 1999). The economists we surveyed were more likely to live apart when they believed that the arrangement would increase their likelihood of publishing in

top journals and of having opportunities to move to different kinds of jobs; they were not more likely to live apart when they believed that the arrangement would increase their lifetime income. Our results also corroborate the finding that couples are less likely to live apart when they are parents or expect to become parents (*Gerstel and Gross*, 1982; *Gross*, 1980).

While this paper assesses the impact of conflicting locational preferences on the careers and relationships of new economists, our results have relevance for specialized professionals more generally. In the contemporary United States, highly educated men and women tend to marry highly educated partners (*Schwartz and Mare*, 2005). Because educational attainment is positively associated with occupational mobility (*McKinnish*, 2008), the pairing of highly educated partners is likely to complicate migration decisions for dual-career couples across the professions.

3.2 Theoretical predictors of living apart

The seminal theoretical work of Jacob *Mincer* (1978) is the point of departure for our analysis. Although he ignored the possibility that couples with conflicting locational preferences live apart, Mincer characterized the circumstances under which they live together and the circumstances under which they break up. Couples in the Mincer model solve

$$\underset{x_A, x_B}{\text{maximize}} \quad G_A^{x_A} + G_B^{x_B} + I_{x_A = x_B} \cdot (M_A + M_B)\,,$$

where x_i is the location of partner i; $G_i^{x_i}$ is the utility gain of partner i from locational amenities and career opportunities in location x_i, net of the cost of moving to that location; M_i is the utility gain of partner i from the couple's relationship; and $I_{x^A = x^B}$ is an indicator variable for the relationship.

Let G_A^I and G_B^I be the net utility gains of the partners from locational amenities

and career opportunities in the locations that solve

$$\underset{x_i}{\text{maximize}} \quad G_i^{x_i}$$

for $i = A, B$. Let G_A^F and G_B^F be the utility gains of the partners in the location that solves

$$\underset{x_A = x_B}{\text{maximize}} \quad G_A^{x_A} + G_B^{x_B}.$$

Mincer defines the *migration tie* of partner i as the difference between the net utility gain of partner i from locational amenities and career opportunities in the location he or she would choose as a single person, G_i^I, and the utility gain of partner i in the location that maximizes the joint utility of the couple, G_i^F:

$$T_i = G_i^I - G_i^F.$$

He predicts that couples live together when the sum of their gains from their relationship exceeds the sum of their migration ties:

$$M_A + M_B > T_A + T_B. \tag{3.1}$$

When the sum of their migration ties exceeds the sum of their gains from their relationship, Mincer predicts that couples break up.

As discussed above, evidence from qualitative studies suggests that some couples neither move together nor break up. Instead, these couples reconcile conflicts between relationship commitments and career opportunities by maintaining long-distance relationships. To explore the implications of long-distance relationships for relationship and career outcomes, we develop a simple extension of the Mincer model. We decom-

pose the gain from a relationship, M_i, into a component that accrues to every person in a relationship, K_i and a component that accrues only to people who live with their partners, H_i:

$$M_i = K_i + H_i.$$

Because we are interested in relationship stability only to the extent that it is influenced by conflicting migration opportunities, we assume that couples maintain their relationships in the absence of conflicting migration opportunities:

$$M_A + M_B > 0.$$

We also assume that couples prefer cohabiting relationships to long-distance relationships. Given a choice between living together and living apart in separate but otherwise identical locations, couples choose to live together:

$$H_A + H_B > 0.$$

Finally, we allow for the possibility that some couples would rather live apart than break up and for the possibility that some would rather break up:

$$-H_A - H_B < K_A + K_B < M_A + M_B.$$

In our extension of the Mincer model, couples solve

$$\underset{x_A, x_B}{\text{Maximize}} \quad G_A^{x_A} + G_B^{x_B} + I_{r=1}\left(K_A + K_B\right) + I_{r=1} I_{x_A = x_B} \cdot \left(H_A + H_B\right),$$

where $I_{r=1}$ is an indicator variable for a relationship, either long-distance or cohabiting; $I_{x^A = x^B}$ is an indicator variable for cohabitation; and the other variables are defined as above.

For couples who would rather break up than live apart, the predictions of the extended model coincide with the predictions of the Mincer model. These couples live together when the sum of their gains from their relationship exceeds the sum of their migration ties and break up when the reverse is true. Formally, couples with $K_A + K_B < 0$ live together when Equation (3.1) holds, and break up otherwise.

In contrast, for couples who would rather live apart than break up, the predictions of the extended model and the predictions of the Mincer model diverge. By maintaining long-distance relationships, these couples can enjoy utility from from their relationships without sacrificing utility to migration ties. Consequently, it is never optimal for them to break up. Instead, these couples live together when the sum of their gains from *cohabitation* exceeds the sum of their migration ties and *live apart* when the reverse is true. Formally, couples with $K_A + K_B > 0$ live together when

$$H_A + H_B > T_A + T_B \tag{3.2}$$

and live apart otherwise.

A comparison of Equations (3.1) and (3.2) indicates that, if there are couples who would rather live apart than break up, the Mincer model makes incomplete predictions about their responses to conflicting migration opportunities. While the extended model predicts that couples with

$$M_A + M_B > T_A + T_B > H_A + H_B \tag{3.3}$$

and

$$T_A + T_B > M_A + M_B > H_A + H_B \tag{3.4}$$

live apart, the Mincer model predicts that the former live together and the latter break up. Thus, relative to the extended model, the Mincer model posits a stark

trade-off between personal relationships and professional success. As a result of this simplification, the Mincer model overstates the negative impact of migration ties on careers. Depending on the distribution of couples between the circumstances described in Equations (3.3) and (3.4), the Mincer model may overstate or understate the negative impact of migration ties on relationships.

3.3 Data

This paper is a product of the Job Seekers Project, a longitudinal survey project that tracks the professional and personal trajectories of recent entrants to the junior PhD job market in economics.[1] The project combines information from original web surveys with information from job-placement and professional websites to create a uniquely rich dataset for the study of work-family trade-offs. Since the 2007-08 academic year, the project has contacted three graduating cohorts of economists as they enter the job market and has followed up with them several months later to learn about their professional and personal circumstances. At the same time, the project has gathered detailed background information about the economists, including their demographic characteristics, educational credentials, and professional accomplishments, from the CVs they post on the job-placement websites of their graduate departments.

The sampling frame for the Job Seekers Project is comprehensive. We use publicly available information to compile a list sample of a clearly defined population: job candidates whose names and contact information appear on job-placement websites linked by the National Bureau of Economic Research.[2] Between the 2007-08 and 2009-10 academic years, the three job market cohorts included in the Job Seekers sample

[1]This project is a joint effort between Marta Murray-Close, Robert J. Willis and myself. We gratefully acknowledge grant funding from the Sloan Foundation.

[2]The National Bureau of Economic Research posts links to job-placement websites of graduate departments on their own job-market website: http://www.nber.org/candidates/.

included 2,756 job candidates from 134 job-placement websites. A large majority of the job candidates in the sample posted information on websites maintained by departments of economics (90 percent) and departments in the United States (91 percent); a minority posted information on websites maintained by departments of business, public policy, or other fields closely related to economics, or departments in Canada or Europe.[3] We believe that the Job Seekers sample is nearly the universe of job candidates who expected to participate in first-round job interviews at the annual meetings of the American Economic Association between 2008 and 2010.

We invite job candidates to participate in the pre-market survey in late December, just before most begin their first-round job interviews. The survey is available for job candidates to complete online during the period leading up to the annual meetings of the American Economic Association and remains available for several months after the meetings. While the fielding window for the pre-market survey is long, most respondents complete the survey in a timely manner. Across all three cohorts, 63 percent of the job candidates who completed the pre-market survey submitted their responses before the meetings. Eighty-eight percent submitted their responses within one month of receiving the invitation to participate.

We invite job candidates to participate in the post-market survey approximately six months after the job market closes, when most have concluded their job search and know whether and where they will be working in the coming year. For the 2007-08 and 2009-10 job market cohorts, we sent the invitation to the post-market survey in August; for the 2008-09 cohort, we sent the invitation in November. Like the pre-market survey, the post-market survey is available for job candidates to complete online over a period of several months, and like pre-market respondents, most post-market respondents complete the survey in a timely manner. Across all three cohorts,

[3]The sample did not include job candidates from European departments until 2009-10. Prior to 2009-10, 95 percent of the job candidates in the sample were from departments in the United States; in 2009-10, 84 percent were from departments in the United States.

79 percent of the job candidates who completed the post-market survey submitted their responses within one month of receiving the invitation to participate.

The data sources compiled by the Job Seekers project have a number of desirable features. First, the project's combination of data from web surveys and job placement websites allows for a detailed analysis of sample selectivity. We are confident that the estimates derived from this data are reasonably representative of the population of new entrants to the new PhD job market in economics. Response rates for the Job Seekers surveys are comparable or superior to response the response rate for a typical web survey. Over all cohorts, the response rate for the pre-market survey was 53 percent, and the response rate for the post-market survey was 39 percent. By way of comparison, a meta-analysis of response rates to web and internet surveys used in academic studies of well-definied populations since 1994 found a mean response rate of 40 percent (*Cook et al.*, 2000).

We supplement the Job Seekers survey with publicly available information from the job placement websites of graduate departments. Using the ranking of graduate departments in economics from the US News and World Report, we ranked the departments of 70 percent of the job candidates in the Job Seekers sample.[4] Using the photographs and CVs that job candidates posted on the job placement websites, and supplementing with coding based on first names, we identified the gender of 97 percent of the job candidates. From the CVs, we identified the countries in which 93 percent of the job candidates obtained their undergraduate degrees.[5] Finally, also from the CVs, we obtained information about the doctoral training of job candidates,

[4]The US News rankings are available at http://grad-schools.usnews.rankingsandreviews.com/best-graduate-schools/top-humanities-schools/economics-rankings. Programs are only ranked through place 79; only programs in the United States are ranked. Most (66 percent) of job candidates from unranked departments were from lower-ranked economics departments in the United States. A sizable minority (30 percent) were from departments outside the United States, and a small number (4 percent) were from departments in fields closely related to economics, such as business or public policy.

[5]Country of origin is not listed on most CVs, so we use country of undergraduate education as a proxy for this variable. Data from the web surveys indicate that country of undergraduate education is a good proxy for the job candidate's country of origin.

including their research fields, teaching experience, and research productivity, as well as information about their previous education.

A second feature of the Job Seekers data is the availability of information about a wide range of job-market outcomes and expectations regarding future career trajectories. Previous studies of family migration have focused on the current employment status and earnings of couples likely to face migration ties. This narrow focus is largely attributable to data limitations. The large-scale datasets most studies have used do not contain other measures of career attainment. The narrow focus of previous studies is also a serious limitation. It is likely that specialized professionals – the group most vulnerable to migration ties – care a great deal about career outcomes beyond their employment and earnings. Especially at the beginning of their careers, they are likely to value less tangible aspects of their jobs, such as prestige, and more forward-looking aspects, such as access to career ladders. To the extent that previous studies have neglected these other outcomes, they may have misrepresented or understated the impact of migration ties on highly educated workers. In contrast, the Job Seekers surveys contain detailed questions about a comprehensive list of career outcomes and expectations.

A third feature of the Job Seekers data is the combination of information about relationship status with information about household composition. Previous studies of family migration have assessed its impact on couples who live together but have largely ignored couples who live apart. Again, the narrow focus of previous studies is understandable but unfortunate. Many large-scale datasets contain information about household composition but not about family members who live outside the household. Other datasets contain information about spouses or partners but assume that couples live in the same household. Despite the possibility that dual-career concerns induce couples to live apart, few datasets contain the information that would be necessary to study living arrangements as a margin of adjustment to conflicting

migration opportunities. In contrast, the Job Seekers surveys contain questions both about the partners of new economists and about the living arrangements of couples in the year after the job market.

A fourth feature of the Job Seekers data, and a key innovation of the project relative to other studies of family migration, is our use of individual-specific measures of job candidates' counterfactual job-market outcomes. The structure of the PhD job market in economics, where most job candidates submit applications, complete interviews, travel to fly-outs, and receive job offers during narrow, pre-determined windows of time, provides job candidates with well-defined choice sets, including well-defined counterfactual outcomes. The structure of the job market also allows us to survey job candidates about their choices while their memories of the job market are fresh. To this end, the post-market survey includes a series of questions about the outcomes respondents actually had on the job market and the outcomes they think they would have had under counterfactual scenarios where their responses to migration ties were different.

To determine whether respondents made individually optimal choices on the job market, and to assess the impact of migration ties on the job placements of those who did not make individually preferred choices, the post-market survey asks respondents to consider the following counterfactual scenario: "Suppose your [husband/wife/significant other] could have an equally satisfying professional and personal life in any location – that is, suppose it would not be a sacrifice for [him/her] to move with you anywhere."[6] The survey then asks respondents to describe the decisions they would have made and the outcomes they think would have had at each stage of the job market under this scenario. For the remainder of the paper, we refer to options the respondent would have chosen in the absence of migration ties as

[6]This is the text that introduced the counterfactual questions in the 2009-10 post-market survey. The wording of the counterfactual questions has varied slightly over time, but the changes do not appear to have affected response patterns. The text of the questions from other survey years is available in the appendix.

individually preferred, or *IP*, options.

To determine whether respondents had forgone living with their partner to accept their individually preferred job, the post-market survey asked respondents who expected to be living apart the following March to consider a second counterfactual scenario: "Please imagine the life you would have had if you and your [husband\wife\significant other] were constrained to live together next March (i.e., share your primary residence or live close enough to each other that you could see each other after work on weeknights)." To assess the impact migration ties *would have had* on the job placements of respondents who rejected their jointly preferred option, had they instead decided to accept it and live with their partner, the survey asks respondents to describe the outcomes they think they would have had if living apart were not an option. We refer to options the respondent would have chosen if living apart were not an option as *jointly preferred*, or *JP*, options.[7]

Each analysis presented in this paper uses slightly different sample restriction criteria. Because changes over time in the survey questions and skip logic mean that some variables are not available for all respondents in all years, the sample restriction criteria are formulated to ensure comparability while maintaining as large a sample as possible. Column headers and footnotes in each table describe the samples used in each analysis, while the Appendix provides detail about changes in question wording and response scales between cohorts.

3.4 Results

The Mincer model provides strong reasons to believe that family migration harms the career prospects of new economists. Like other specialized professionals, economists participate in national and international labor markets. Their career opportunities

[7] We added questions about jointly preferred options to the post-market survey in 2009-10. We do not have information about the jointly preferred options of job candidates in the 2007-08 and 2008-09 job-market cohorts.

are geographically dispersed. Most economists move for their first job, and many move for subsequent jobs.

At the same time, the Job Seekers data indicate that many new economists are in relationships with highly educated partners. A majority (73 percent) of the economists who responded to the Job Seekers surveys were in a relationship at the beginning of their job-market year (Table 3.1). Almost half (48 percent) were married; another fifth (20 percent) characterized their relationship as marriage-like or committed. The partners of the economists, like the economists themselves, had strong educational credentials. More than three quarters (76 percent) of the partners had earned or were pursuing a graduate degree, and more than one third (40 percent) had earned or were pursuing a PhD (Table 3.2).

While both male and female economists reported personal circumstances that made them vulnerable to migration ties, we observe gender differences in two domains. On one hand, male economists were more likely than female economists to be in a relationship during their job-market year. At the time they entered the job market, 76 percent of the men who responded to the Job Seekers surveys were in a relationship, and 51 percent were married. The comparable figures for women were 67 percent and 40 percent. On the other hand, conditional on being in a relationship, female economists were more likely than male economists to have a partner whose educational attainment equaled their own. More than half (57 percent) of the women who were in a relationship during their job-market year had a partner who had earned or or was pursuing a PhD. The same was true of less than one third (32 percent) of the men.

3.4.1 Impact of migration ties on career outcomes

The Job Seekers' data show that the impact of migration ties on economists' choices at the final stage of the job market was surprisingly small. Of 631 respondents

who received at least one job offer and had a relationship that spanned the job-market months, a large majority (85 percent) reported that they had accepted their individually preferred job from their final choice set. A small number (1 percent) reported that their individually preferred option had been to reject all of their job offers and that they had, in fact, rejected all of their offers. Unexpectedly, in a population theoretically vulnerable to severe migration ties, just 14 percent of these respondents reported that their job choice would have been different in the absence of relationship-related constraints. These results are summarized in Table 3.3.

Both the Mincer model and the extended model predict that job candidates are more likely to forgo their individually preferred job for the benefit of their relationship when the career sacrifices involved are small. The 2009-10 post-market survey assessed the career costs of forgoing an individually preferred job with respect to eight long-term career outcomes a typical economist might value: earning tenure at a research university, earning tenure at a four-year college, publishing regularly in top journals, having opportunities to move to more prestigious jobs, having opportunities to move to different kinds of jobs, having a lifetime income higher than average for their field, finding their everyday work satisfying, and having plenty of time for life outside of work. Respondents rated their likelihood of realizing each outcome, in light of the job they accepted, using a six-point scale that ranged from 1 (*extremely unlikely*) to 6 (*extremely likely*).

When respondents faced a trade-off between their individually preferred job and their jointly preferred job, the survey also asked them to rate their their likelihood of realizing the long-term career outcomes under a counterfactual scenario where their response to the trade-off was different. Respondents who rejected their individually preferred job rated their likelihood of realizing each outcome under the counterfactual scenario where they accepted it. Respondents who accepted their individually preferred job rated their likelihood of realizing each outcome under the counterfactual

scenario where they rejected it in favor of their jointly preferred job.

Finally, to assess the overall impact of rejecting an individually preferred job on the career outcomes of new economists, the survey asked respondents who faced a tradeoff between their individually preferred job and their jointly preferred job to compare their overall career prospects with the prospects they would have had under the relevant counterfactual scenario. Respondents who rejected their individually preferred job rated their career prospects at that job, relative to the job they accepted, using a five-point scale that ranged from 1 (*much better*) to 5 (*much worse*). Respondents who accepted their individually preferred job rated their career prospects at that job, relative to their jointly preferred job, using the same scale.

Table 3.4 presents mean ratings for the eight specific outcome measures and the overall outcome measure, for both the individually preferred job and the jointly preferred job.[8] For each measure, results are presented separately for each of three groups: partnered job candidates for whom the individually preferred job coincided with the jointly preferred job (*not constrained*), partnered job candidates who were constrained by their partners and accepted the individually preferred job over the jointly preferred alternative (*constrained, accepted IP job*), and partnered job candidates who were constrained by their partners and accepted the jointly preferred job over the individually preferred alternative (*constrained, rejected IP job*).

On the whole, job candidates in all three groups believed that both their individually preferred jobs and their jointly preferred jobs would position them to succeed in their careers. Mean ratings for all of the specific outcome measures exceed the scale value corresponding to *somewhat unlikely* for each group. For job candidates in the unconstrained group, ratings for the individually preferred job and the jointly preferred job are, by definition, identical. For job candidates in the constrained groups,

[8]The ratings we summarize in this table are ordinal data. We present means rather than ordinal measures of central tendency, such as medians or modes, because the ratings for most outcomes cluster at the high end of the scale, and the ordinal measures obscure important variation within that range.

the results in Table 3.4 are consistent with the definition of *individually preferred job* as the job that maximizes career-related utility: on average, job candidates believed that their individually preferred jobs were more likely than their jointly preferred jobs to produce most of the specific outcomes we assessed. Job candidates also believed that their individually preferred jobs offered better overall career prospects than their jointly preferred jobs. Mean ratings of career prospects with the individually preferred job, relative to career prospects with the jointly preferred job, fell between the scale value corresponding to *about the same* and the scale value corresponding to *somewhat better* for both constrained groups.

Tables 3.5 and 3.6 apply sign tests to assess the statistical significance of differences between the individually preferred and jointly preferred jobs of job candidates in the constrained groups. Table 3.5 indicates that, among job candidates who rejected their individually preferred job, the number who reported that their career prospects had suffered from the decision statistically exceeds the number who reported that their career prospects had improved. Differences with respect to the specific career outcomes are not statistically significant for this group.

Table 3.6 indicates that, among job candidates who accepted their individually preferred job, the number who reported that their career prospects would have suffered from rejecting it statistically exceeds the number who reported that their career prospects would have improved. Job candidates who accepted their individually preferred job also reported a statistically significant preponderance of differences favoring that job in the likelihood of realizing five specific career outcomes: earning tenure at a research university, publishing regularly in top journals, having opportunities to move to more prestigious jobs, having opportunities to move to different kinds of jobs, and finding everyday work satisfying. Taken together, the results in Tables 3.5 and 3.6 show that migration ties are a salient issue for some new economists.

Consistent with the predictions of the Mincer model and the extended model,

job candidates who rejected their individually preferred job – and who therefore endured the career sacrifices associated with their migration ties – described the decision as less costly than did respondents who accepted their individually preferred job. While job candidates in both constrained groups believed that their individually preferred jobs were more likely than their jointly preferred jobs to produce most of the long-term career outcomes we assessed, the differences tend to be smaller among job candidates who rejected their individually preferred job. Returning to Table 3.4, for most career outcomes, job candidates who rejected their individually preferred job reported that their individually and jointly preferred jobs were more similar, on average, than did job candidates who accepted their individually preferred job. Returning to Tables 3.5 and 3.6, results from the sign tests tell a similar story. Six of the tests show a statistically significant difference between the individually preferred and jointly preferred jobs for job candidates who accepted their individually preferred job. Despite a larger number of observations and correspondingly greater power, just one test shows a statistically significant difference for job candidates who rejected their individually preferred job.

While the evidence we have presented to this point suggests that the impact of migration ties on the career prospects of new economists is modest, two caveats are in order. First, even if migration ties do not shape outcomes substantially in the final stage of the job market, they may shape outcomes at earlier stages, when job candidates make decisions about which applications to submit and which interviews and fly-outs to accept. To the extent that job candidates alter their application, interview, or fly-outs sets in response to relationship commitments, the job offers from which they choose in the final stage of the job market may differ from the offers they would have obtained in the absence of migration ties.

Table 3.7 presents evidence that migration ties do, in fact, influence the decisions of job candidates with respect to applications, interviews, and fly-outs. A sizable

minority (44 percent) of Job Seekers respondents with partners reported that they would have applied to a different set of jobs if they had not been constrained by relationship commitments. Twelve percent reported that they would have accepted a different set of interviews, and 14 percent reported that they would have accepted a different set of fly-outs. By the time they reached the offer stage of the job market, almost half (49 percent) of job candidates with partners had altered their choices in some way in response to migration ties.

The results in Table 3.7 suggest that our measure of counterfactual job outcomes is not accurate in every case. The magnitude of the results, however, suggests that our measure is accurate in most cases. For example, even among the minority of respondents who altered their application set in response to migration ties, the median change in the size of the set was just five applications withheld. Given that the the median application set contained 100 applications, changes of this magnitude are small, and seem unlikely to have shaped the offer sets of respondents in dramatic ways. Furthermore, at the interviews stage, the median change reported by the 12 percent of respondents who indicated that their interview decisions were influenced by their partners was an *increase* of two interviews accepted.

The second caveat is that the migration ties of new economists may influence their job-market outcomes through another indirect channel: the behavior of employers in the job market. Even if the choices of job candidates are unaffected by their relationship commitments, employers may consider family circumstances when deciding which candidates to interview, invite for fly-outs, or hire. Employers may learn about the relationships of job candidates in at least two ways. First, job candidates may tell employers about their relationships when they meet for interviews or fly-outs. Second, academic advisors and other members of the academic community may discuss the relationships of job candidates in an attempt to facilitate good job matches.

Table 3.7 indicates that, whatever the source of the information, employers are

likely to learn about the relationships of job candidates before they extend job offers. Seventy-two percent of Job Seekers respondents with partners said that some or all of their prospective employers knew about their relationship by the time they completed their interviews. Eighty-four percent said that some or all of their employers knew about their relationship by the time they completed their fly-outs. To the extent that employers learn about the relationships of job candidates from third parties, without the knowledge of the job candidates, these results may understate the amount of information available to employers. Because the Job Seekers surveys focus on the supply side of the job market, we do not know how employers incorporate information about relationships into their decision making and cannot rule out the possibility that employer responses shape the offer sets of job candidates in meaningful ways. In particular, we cannot rule out the possibility that the individually and jointly preferred jobs we observe are more similar – and the impact of migration ties we infer less pronounced – than they would be in the true counterfactual situation.

3.4.1.1 Comparing the Job Seekers approach to existing work

Previous studies of family migration have not had access to direct measures of migration ties. Consequently, previous tests of the Mincer model have relied on proxies. Most often, studies have assumed the migration ties are more severe among married men and women, and among men and women with highly educated partners. These studies have shown that career outcomes theoretically related to being a tied migrant (for example, reduced earnings or labor supply after a move) are more likely for married couples – especially married women – and for men and women whose partners have college or graduate degrees.

The Job Seekers dataset contains uniquely detailed information about the responses of new economists to migration ties. This information allows us to examine the association between the proxies previous studies have used and direct measures of

migration-induced career sacrifices. Both the Mincer model and the extended model predict that job candidates are more likely to reject their individually preferred job when their migration ties are large and when the value of their relationship is low. Because previous studies have suggested that migration ties increase with educational attainment, we expect that new economists whose partners have graduate degrees are more likely to alter their job-market choices in response to migration than new economists whose partners have lower levels of education. Because we hypothesize that the value of relationships usually increases with commitment, we also expect that new economists in less committed relationships are more likely to alter their job-market choices than new economists in more committed relationships.

Table 3.8 presents results from probit regressions examining the association between the probability that job candidates altered their job-market choices in response to migration ties and the proxies for migration ties that previous studies have used. In Column 1, the dependent variable takes a value of one for respondents who rejected their individually preferred job and a value of zero otherwise. In Column 2, the dependent variable takes a value of one for respondents who altered their choice set at the application, interview, or fly-out stage of the job market and a value of zero otherwise. In Column 3, the dependent variable takes a value of one for respondents who rejected their individually preferred job, altered their choice set at an intermediate stage of the job market, or did both, and a value of zero otherwise.

Consistent with our expectations, the results in Table 3.8 suggest that new economists whose partners have graduate degrees are more likely to reject their individually preferred job than new economists whose partners have college degrees or less. The estimates in Columns 2 and 3 are similar in sign and magnitude to the results in Column 1 but are not statistically significant from zero. Also consistent with our expectations, the negative coefficients on the relationship status indicators for committed and dating relationships suggest that new economists who are in less formal

relationship are less likely than new economists who are married to reject their individually preferred job or to alter their application, interview, or fly-outs sets in response to migration ties.

In contrast to the large number of studies finding that family migration imposes larger costs on women than on men, we observe only minor gender differences in the impact of migration ties on new economists. Of 631 job candidates who provided information about their counterfactual job choices, 17 percent of women and 12 percent of men that they had rejected their individually preferred and had chosen an option more favorable to their partner. This difference is not statistically significant, $\chi^2\,(1, N = 631) = 2.07$, $p = 0.15$.

On the other hand, women were somewhat more likely than men to report that migration ties had influenced their decisions at intermediate stages of the job market. Fifty-seven percent of women, but just 45 percent of men, reported that migration ties had shaped the set of applications they submitted, or the set of interviews or fly-outs they accepted, $\chi^2\,(1, N = 631) = 5.51$, $p = 0.02$. Results from the probit regressions suggest that gender differences in relationship status and the educational attainment of respondents' partners do not explain women's greater likelihood of altering job-market choices in response to migration ties. Controlling for these characteristics, women were still 9 percentage points more likely than men to report that migration ties had shaped their application, interview, or fly-out sets (Table 3.8).

3.4.2 Living apart to avoid career sacrifices

Results from the Job Seekers surveys indicate that living arrangements are an important margin of adjustment for couples facing migration ties. Of 454 respondents who described their expectations for their relationship in the year after the job market, 16 percent reported that a long-distance relationship was the most likely outcome. The prevalence of long-distance relationships among Job Seekers respondents equaled

or exceeded the prevalence rejecting individually preferred offers: we saw, in the previous section, that 14 percent of respondents rejected their individually preferred job for the benefit of their relationship, and just 7 percent of respondents reported that a break-up was the most likely outcome for their relationship.

Table 3.10 presents summary results from adjustment along both the relationship and career outcome margins. Of 360 respondents who were still in their relationships at the time of the post-market survey and who provided information about both their counterfactual job outcomes and their expected relationship outcomes, 73 percent accepted their individually preferred job and expected to be living with their partner in the year after the job market. (See right panel of Table 3.10.) That economists were likely to obtain optimal outcomes in both their careers and their relationships suggests that many couples did not face migration ties or, more likely, that the migration ties of the economists dominated the migration ties of their partners. In the latter situation, it would be the partners rather than the economists who rejected their individually preferred job.

Table 3.11 presents evidence that, as suggested in the previous section, the economists most likely to live apart from their partners are those whose careers would suffer the most if they lived together. Specifically, the table presents results from ordinary least squares regressions of the subjective probability of living apart in the year after the job market on a series of indicator variables for benefits of the individually preferred job over the jointly preferred job. We examine the association between living apart and the belief that the individually preferred job is more likely than the jointly preferred job to produce each of the long-term career outcomes enumerated in the previous section: better overall career prospects, tenure at a research university, tenure at a four-year college, regular publication in top journals, opportunities to move to more prestigious jobs, opportunities to move to different kinds of jobs, a lifetime income higher than average for their field, everyday work that is satisfying, and plenty of

time for life outside of work.

Column 1 of Table 3.11 presents results from regressions in which each of the career outcomes enters as the sole regressor. Consistent with the extended model, the coefficient on overall career prospects, as well as a majority of the coefficients on other specific career outcomes, are positive. Respondents who believed that their individually preferred job was more likely to produce a desirable outcome were more likely to live apart than respondents who believed that their jointly preferred job was as likely or more likely to produce the outcome. The two exceptions to this pattern were a higher-than-average lifetime income and time for life outside of work. Respondents who believed that their individually preferred job was more likely than their jointly preferred job to produce these outcomes were no more likely than other respondents to live apart. The finding that responses to migration ties were not sensitive to changes in expected lifetime income is consistent with evidence from previous studies that couples who choose long-distance relationships pursue career opportunities not primarily as a source of income, but rather as a "central life interest" (*Gerstel and Gross*, 1982).

Column 2 of Table 3.11 presents results from a regression in which all of the specific career outcomes enter together. In contrast to the coefficients estimated for these variables in Column 1, most of the estimates in Column 2 are statistically indistinguishable from zero. The estimates that remain statistically significant in the combined regression model suggest that superiority of the individually preferred job with respect to two specific career outcomes, publishing regularly in top journals and having opportunities to move to different kinds of jobs, is associated with a substantially greater likelihood of living apart.

The change in the pattern of estimates between Columns 1 and 2 probably indicates that the eight outcome measures we use tap a smaller number of underlying job characteristics. Chi-square tests show that, among respondents whose individually

and jointly preferred jobs were different, those who believed that their individually preferred job was more likely to help them publish regularly in top journals also tended to believe that it was more likely to help them earn tenure at a research university, to facilitate moving to more prestigious jobs, and to offer satisfying work (Table 3.12). Similarly, respondents who believed that their individually preferred job was more likely to facilitate moving to different kinds of jobs also tended to believe that it was more likely to facilitate moving to more prestigious jobs and to offer satisfying work.

In Table 3.13, we present results from a "horserace" between the dummy variable indicating that the overall career prospects are better at the individually preferred job, relative to the jointly preferred job, and each of the specific career outcomes. The first row contains estimates from the base case, a linear regression of the probability of living apart on the dummy variable representing overall career prospects. In each additional row, we also include one specific career outcome dummy variable. Column 1 displays the estimated coefficient and standard error from the overall measure, while Column 2 displays the estimated coefficient and standard error from the specific career outcome. For five of the eight specific career outcomes, the estimated coefficient on the specific outcome is small and statistically insignificant, and the coefficient on the overall career prospects measure is virtually unchanged from the base case in which the overall career prospects measure is the sole regressor. In these cases, the adjusted R-squared statistic is very close to, or less than, the adjusted R-squared statistic in the base regression. In the row containing the results from the regression with the dummy variable for "have opportunities to move to more prestigious jobs," the overall career outcomes coefficient is reduced, and neither coefficient is statistically significant, implying that this variable may be highly correlated with the overall career outcomes variable. In the row including "have opportunities to move to different kinds of jobs," the coefficient on the overall career prospects is similarly reduced, but the coefficient on the specific outcome is quite large and statistically significant.

The adjusted R-squared statistic also indicates that the explanatory power of this regression is also higher. We interpret these results to imply that respondents choosing to live apart from their partners may particularly value jobs that will give them more flexibility in employment options in the future. This is consistent with our belief that living apart is a temporary but important solution to the two-body problem.[9]

3.4.3 When does it pay to live apart?

Our extension of the Mincer model, which allows for the possibility that couples with conflicting locational preferences live apart, carries predictions that are testable using data from the Job Seekers project. In the remainder of this section, we assess the extent to which circumstances in which Job Seekers respondents live apart correspond to circumstances in which the model predicts that couples live apart. In particular we estimate ordinary least squares regression models of the form

$$SPAPART_i = X_i'\beta + \epsilon_i,$$

where $SPAPART_i$ is the job candidate's subjective probability of living apart in the year after the job market and X_i contains characteristics of the job candidate, his or her partner, and their relationship that are theoretically likely to influence the couple's response to migration ties. Our focus in these analyses is the choice of living arrangements by couples who expect to maintain their relationship. Accordingly, we focus on predictions from Equation (3.2) and exclude from the regression sample respondents who reported that they were more likely than not to break up with their partner in the year after the job market.

[9] Indeed, while discussion of the dynamics of cohabitation and career decision-making are beyond the scope of this paper, our study has begun to gather data from follow-up surveys to examine these dynamics.

Role of partner's career-related utility

Our extended Mincer model assumes that couples consider the migration ties of both partners when deciding where to live and work. Accordingly, the model predicts that economists are more likely to live apart from their partner not just when their own migration tie is large, but also when their partner's migration tie is large. This prediction is consistent with findings from qualitative studies of long-distance relationships, which indicate that the ability of dual-career couples to preserve their relationship while pursuing desirable career opportunities in separate locations is their primary impetus for living apart (*Gerstel and Gross*, 1982).

The Job Seekers data contain several measures of partner migration ties. First, like most studies of family migration, we use the educational attainment of the partner as a proxy for labor-force attachment and the possession of specialized human capital. Second, we ask about the school enrollment of the partner in the year after the respondent was on the job market. Third, we ask respondents what they think is the percent chance that their partner will work at least 20, 40, and 60 hours per week over most of the next ten years. Fourth, we ask respondents how good they thought their partner's job opportunities would be in the location of their new job, at the time they accepted the job. Finally, we assess agreement with the following statement: "My [husband's/wife's/significant other's] career will *not* suffer if we move to the places that are best for my career."

Table 3.14 presents results from regressions of the subjective probability of living apart on these measures of partner migration ties. Column 1 indicates that, as expected, higher levels of educational attainment are associated with greater likelihoods of living apart. Economists whose partner had or was pursuing a master's degree believed they were 9 percentage points more likely to live apart than those whose partner had no more than a college degree. Economists whose partner had or was pursuing a PhD believed they were 16 percentage points more likely to live

apart. School enrollment was also associated with living apart. Economists whose partner expected to be in school in the year after the job market believed they were 9 percentage points more likely to live apart than economists whose partner did not expect to be in school.

The results in Column 2 suggest that labor force attachment has little impact on the likelihood of living apart. The estimated coefficients on the percent chance of working at least 20 and at least 40 hours per week are small and are not statistically different from zero. The estimated coefficient on the percent chance of working at least 60 hours per week is also small, but is statistically significant. In particular, the estimate indicates that an increase of 10 percentage points in the percent chance that the partner will work long hours over most of the next ten years is associated with an increase of 3 percentage points in the percent chance that the couple will live apart in the year after the job market.

The quality of the partner's job prospects in the location of the economist's job has a substantial and statistically significant association with the percent chance that the couple lives apart. The results in Column 3 suggest that, compared with economists whose partner had good job prospects in the location of their job, economists whose partner had fair prospects believed they were 11 percentage points more likely to live apart, and economists whose partner had poor prospects believed they were 33 percentage points more likely to live apart. The results in Column 4 are consistent with the results in Column 3. Economists who agreed that the career attainment of their partner would *not* be harmed by following them to their individually preferred location believed they were 18 percentage points less likely to live apart than economists who thought that the career attainment of their partner would be harmed.

On the whole, the estimates from Columns 1 through 4 are robust to the inclusion of additional variables in the regression model. Column 5 presents estimates from a regression model that includes all of the measures of partner migration ties. While

some estimates that were statistically significant in the partial models lose significance in the full model, the signs and magnitudes of most estimates change only modestly. Interestingly, the association between educational attainment and living apart is not statistically significant in the model that includes direct measures of migration ties. This result suggests that, while education is a reasonable proxy for migration ties in studies that lack direct measures, it is not itself responsible for the decision to live apart.

Role of relationship-related utility

The extended model predicts that economists are more likely to live apart when their value of living together is lower. A reasonable hypothesis is that couples have a stronger preference for living together when they are more committed to each other and more satisfied with their relationship. On the other hand, qualitative studies of long-distance relationships have found that couples who live apart are not motivated to do so by problems in their relationship, and do not expect to break up (*Gerstel and Gross*, 1982).

The Job Seekers data contain two measures of relationship satisfaction and commitment. First, we ask respondents to classify their relationship as "married," "marriage-like," "committed," or "dating." Second, we ask them to rate their satisfaction with the relationship in the months leading up to the job market. In addition, the surveys contain six items designed to assess the relationship-related costs of living apart. Specifically, we ask respondents how upset they would be if they were living apart from their partner in the year after the job market, and we ask them to rate their agreement with the following statements: (1) "It would be possible for me to have a fulfilling relationship while living apart from my [husband/wife/significant other]," (2) "I would never consider living apart from my [husband/wife/significant other]," (3) "I would be willing to make a large career sacrifice so that I could live with my

[husband/wife/significant other]," (4) "I would consider jobs that require me to live apart from my [husband/wife/significant other] for up to one year," and (5) "I would consider jobs that require me to live apart from my [husband/wife/significant other] for up to five years."

Table 3.15 presents results from regressions of the subjective probability of living apart on relationship commitment, relationship satisfaction, and other relationship-related cost measures. Column 1 indicates that lower levels of relationship commitment are associated with greater likelihood of living apart. Compared with economists who were married, those in marriage-like and committed relationships believed they were between 11 and 33 percentage points more likely to live apart. Column 2 suggests that relationship satisfaction plays a moderately important role in decisions about living arrangements. Economists who were extremely satisfied with their relationship in the months leading up to the job market believed they were 15 percentage points less likely to live apart than economists who were less satisfied.

Column 3 presents estimates from a regression model that includes dummy variables indicating that the respondent would be very upset or extremely upset to be living apart in the year after the job market, along with dummy variables indicating that the respondent agreed with the statements enumerated above. Surprisingly, just two estimates in Column 3 are statistically significant. Economists who would be extremely upset to be living apart believed they were 19 percentage points less likely to live apart than economists who would be less upset. Those who agreed that they would consider jobs that required them to live apart from their partner for up to five years believed they were 19 percentage points more likely to live apart than those who disagreed that they would consider such jobs.

Column 4 presents estimates from a model that includes all of the relationship-related cost measures. The estimates in this column are comparable to estimates in the preceding columns. Economists who are less committed to their relationships are

more likely to live apart than those who are more committed, and economists who report a strong preference for living together are less likely to live apart than those who report a weaker preference for living together. Taken together, these results contrast with, but do not contradict, the findings of the qualitative studies. Previous studies of long-distance relationships have identified participants using non-random sampling methods. Consequently, they have not estimated predictors of living apart for the population – or a known sub-population – of dual-career couples. To the extent that the snowball sampling methods employed by these studies identified the most successful and enduring long-distance relationships, they may have overrepresented couples who were committed and happy in their relationships.

Role of children

Couples are likely to place a higher value on living together when they have children. Qualitative studies of long-distance marriages have found that couples with young children find living apart more stressful than couples without children. In addition, couples who anticipate having children report that they will not continue living apart when they become parents (*Gerstel and Gross*, 1982).

The Job Seekers data includes measures of current parental status and expectations of future fertility. We identify respondents who already had children by the time of the job-market year, and we ask all respondents what they think is the percent chance that they will have a child in the next year and the next five years. In addition to these measures, the surveys contain several items designed to assess the child-related costs of living apart. Specifically, we ask respondents to rate their agreement with the following statements: (1) "I would consider having a child while living apart from my [husband/wife/significant other]," (2) "Living apart from my [husband/wife/significant other] over the next year would prevent us from having as large a family as we would like," (3) "Living apart from my [husband/wife/significant

other] over the next five years would prevent us from having as large a family as we would like," (4) "My children would live with me if my [husband/wife/significant other] and I were living apart," and (5) "I could have a very good relationship with my children even if they were not living with me."

Table 3.16 presents results from regressions of the subjective probability of living apart on parental status, fertility expectations, and other child-related costs of living apart. Column 1 shows that, as expected, parenthood is associated with a lower likelihood of living apart. Parents believed they were 15 percentage points less likely to live apart than non-parents.

Column 2 shows that the expectation of having children is also associated with a lower likelihood of living apart. An inspection of the distributions of subjective probabilities of having children indicated that they were tri-modal, with responses clustering near 0, 0.5, and 1. To assess the role of fertility expectations on decisions about living arrangements, we regressed the expected probability of living apart on dummy variables corresponding to subjective probabilities of less than 0.25, probabilities between 0.25 and 0.75, and probabilities greater than 0.75. Estimates from this regression model are small in magnitude and, with one exception, are not statistically different from zero. The exception is a strong expectation of having children within five years. Economists who thought that they would probably have a child in the next five years believed they were 10 percentage points less likely to live apart than economists who were less sure that they would have children.

Results from Column 3 suggest that, consistent with the findings of qualitative studies of long-distance marriages, some couples will not consider living apart while growing their families. Economists who disagreed that they would consider having a child while living apart from their partner believed they were 18 percentage points less likely to live apart than economists who viewed long-distance relationships and parenting as compatible. With one exception, the estimated coefficients for the re-

maining child-related cost measures are of the expected signs. None, however, are statistically different from zero. Notably, while we find suggestive evidence that respondents who would bear the burden of daily caretaking for their children are less likely to live apart, and that respondents who could maintain good relationships with their children are more likely to live apart, these results are not statistically significant.

Column 4 presents estimates from a regression model that includes all of the child-related cost measures. The estimates in this column are comparable to the estimates in Columns 1 through 3. Like the qualitative studies of long-distance relationships, we find consistent evidence that children increase the cost of living apart and deter couples from adopting the arrangement.

3.5 Conclusions

The job market for PhD economists may be more robust than other job markets for new PhDs. It seems likely, for example, that job candidates in less robust academic markets may be more affected by migration ties than our study of economists might imply. While the results from this study may not be informative about the prevalence and magnitude of migration ties in weaker academic markets, our results may generalize well to other strong national markets for specialized workers, such as lawyers and business executives.

Findings from the Job Seekers project show that some of the new economists we surveyed rejected desirable career outcomes in order to live with their partner. Surprisingly, however, given that many economists are members of highly educated dual-career couples – precisely the sort of couples most vulnerable to severe migration ties – the career sacrifices they described were not large. Just 14 percent of Job Seekers respondents rejected their individually preferred job for the benefit of their relationship. Among respondents who rejected their individually preferred job, the

differences between that job and the job they accepted were moderate.

We argue that the gap between the substantial career sacrifices we expected based on the Mincer model and the relatively minor career sacrifices we observe is explained, in part, by the availability of an option Mincer never considered: living apart. Studies using non-probability samples have found that living apart allows some couples with severe migration ties to avoid both career sacrifices and relationship dissolution. Our study corroborates that finding using representative data from a known sub-population of dual-career couples. Sixteen percent of the new economists we surveyed expected to be living apart from their partner in the year after they entered the job market, and economists whose careers stood to gain most from living apart were the most likely to adopt the arrangement.

Previous research on the migration decisions of dual-career couples has assessed the impact of migration ties on their employment status and earnings. Our results suggest that this focus is too narrow. Of the eight specific career outcomes we considered as likely components of new economists' assessments of their overall career prospects, expected lifetime income was one of just two outcomes that did *not* significantly influence the living arrangements of new economists. Instead, the economists we surveyed were motivated to live apart when they believed that the arrangement would improve their research productivity and facilitate their future career mobility. To the extent that these findings are representative of dual-career couples more generally, studies that focus exclusively on earnings and employment status neglect important costs that migration ties impose on highly educated workers.

Finally, results from the Job Seekers project suggest that living apart is a more viable option for some couples than others. While severe migration ties can induce even happy couples to live apart, relationship commitment and satisfaction are deterrents to long-distance relationships. Children and the expectation of having children are also deterrents. In light of these findings, we posit that the impact of migration ties

on the professional outcomes of dual-career couples is conditioned by their personal circumstances. Couples who are deeply engaged in family life find it more difficult than other couples to protect their careers when their locational preferences diverge.

3.6 Tables

Table 3.1: Relationship status

	All respondents	Male respondents	Female respondents
In relationship	73%	76%	67%
Married	48%	51%	40%
Marriage-like	8%	8%	10%
Committed	12%	12%	12%
Dating	5%	5%	5%
Not in relationship	27%	24%	33%
Observations	1,503	707	503

Notes: Table includes respondents from the 2007-08, 2008-09 and 2009-10 job-market cohorts who gave relationship information in the post-market survey. "In relationship" and "Not in relationship" indicate whether the respondent was partnered in November of the job-market year. Relationship status is the most committed status the respondent ever reported with respect to that relationship.

Table 3.2: Partner education

	All respondents	Male respondents	Female respondents
Bachelor's degree or less	24%	27%	18%
Master's or professional degree	36%	41%	25%
PhD	40%	32%	57%
Observations	1,057	730	327

Notes: Table includes data from 2007-08, 2008-09 and 2009-10 cohort respondents who were in a relationship in November of the job-market year and responded to questions about the educational attainment of their partner. Educational attainment is the highest degree the partner had earned or was pursuing during the job-market year.

Table 3.3: Did respondents choose their individually-preferred job outcomes?

	Percent
Chose preferred outcome	86.4%
Rejected preferred outcome	13.6%
Observations	631

Notes: The individually preferred job outcome is the outcome the respondent would have preferred in the absence of constraints imposed by thee partner's preferences or career. It may refer to a particular job offer or to a preference to reject all job offers. Sample includes respondents from 2007-10 who were in relationships at least from November until March of the job-market year (2007-09 cohorts), or through the post-market survey (2009-10 cohort).

Table 3.4: Expected career outcome ratings of individually preferred (IP) and jointly preferred (JP) jobs (7-point scale)

Expected career outcome	Mean[b]		Obs
Constraint group	IP	JP	
Tenure at a research university			
Not constrained	3.76		143
Constrained, accepted IP job	4.43	3.57	7
Constrained, rejected IP job	4.10	3.90	21
Tenure at a four-year college			
Not constrained	3.14		139
Constrained, accepted IP job	3.29	3.14	7
Constrained, rejected IP job	3.35	3.15	20
Regular publication in top journals			
Not constrained	3.76		144
Constrained, accepted IP job	4.14	3.14	7
Constrained, rejected IP job	4.10	3.95	21
Opportunities to move to more prestigious jobs			
Not constrained	4.04		143
Constrained, accepted IP job	4.20	3.20	5
Constrained, rejected IP job	4.38	4.29	21
Opportunities to move to different kinds of jobs			
Not constrained	4.26		143
Constrained, accepted IP job	4.57	3.71	7
Constrained, rejected IP job	4.24	4.52	21
Higher-than-average lifetime income for field			
Not constrained	3.93		144
Constrained, accepted IP job	3.71	3.57	7
Constrained, rejected IP job	4.19	4.00	21
Satisfying everyday work			
Not constrained	4.95		144
Constrained, accepted IP job	4.83	3.83	6
Constrained, rejected IP job	5.00	4.67	21

Expected career outcome	Mean[b]		Obs
Constraint group	IP	JP	
Plenty of time for life outside of work			
Not constrained	4.48		143
Constrained, accepted IP job	3.71	3.57	7
Constrained, rejected IP job	4.43	4.52	21
Rating of overall career prospects at IP job, relative to JP job[a]			
Not constrained	–		–
Constrained, accepted IP job	2.00		10
Constrained, rejected IP job	2.39		28

Notes: Observations are from coupled respondents from the 2009-10 cohort only, since these questions were not asked before 2009-10. Respondents rated the likelihood of each outcome for both the individually preferred and jointly preferred jobs on a six-point scale, where 1 is "extremely unlikely," 2 is "very unlikely," 3 is "somewhat unlikely," 4 is "somewhat likely," 5 is "very likely," and 6 is "extremely likely." "Not constrained" indicates that respondent's individually preferred choice coincided with couple's jointly preferred choice. For those who were constrained, "accepted individually preferred job" indicates that the respondent accepted the individually preferred job and "rejected individually preferred job" indicates that the respondent accepted the jointly preferred job.

[a] Respondents rated the overall quality of the individually preferred job relative to the jointly preferred job on a five-point scale, where 1 is "much better," 2 is "somewhat better," 3 is "about the same," 4 is "somewhat worse," and 5 is "much worse." Mean ratings below 3 indicate that, on average, respondents felt that the individually preferred job was more likely to yield better long-term career prospects than the jointly preferred job.

[b] The ratings we summarize in this table are ordinal data. We present means rather than ordinal measures of central tendency, such as medians or modes, because the ratings for most outcomes cluster at the high end of the scale, and the ordinal measures obscure important variation within that range.

Table 3.5: Career costs experienced by respondents rejecting individually preferred job offer

Chance respondent will...	Number of respondents rating ... IP > AJ	IP < AJ	Sign test (p)
Earn tenure at a research university	6	3	0.254
Earn tenure at a four-year college	5	5	0.623
Publish regularly in top journals	5	2	0.226
Have opportunities to move to more prestigious jobs	6	2	0.145
Have opportunities to move to different kinds of jobs	3	4	0.773
Have higher-than-average lifetime income for field	6	4	0.377
Have everyday work that is satisfying	8	3	0.113
Have plenty of time for life outside of work	5	5	0.623
Observations: 21			
Overall career prospects[a]	14	2	0.002
Observations: 28			

Notes: Observations represent partnered respondents from the 2009-10 job-market cohort who rejected the individually preferred job (IP) and accepted an alternative job (AJ). The middle two columns in the table present the frequencies of ratings indicating that the individually preferred job is more likely than the accepted job to yield a particular outcome (IP > AJ) or that the individually preferred job is less likely to yield the outcome (IP < AJ). The excluded category is IP = AJ. The sign test column presents one-sided p-values based on the probability of observing the given frequencies of positive signs on the differences between ratings of the individually preferred job and the accepted job from a binomial distribution with mean 0.5. For the eight specific outcomes, respondents rated the likelihood of realizing the outcome on a six-point scale, where 1 is "extremely unlikely," and 6 is "extremely likely."

[a] Respondents rated their overall career prospects with the individually preferred job relative to their prospects with the accepted job on a five-point scale, where 1 is "much better," 2 is "somewhat better," 3 is "about the same," 4 is "somewhat worse," and 5 is "much worse." Ratings below 3 indicate that the individually preferred job is better than the accepted job, while ratings above 3 that the accepted job is better than the individually preferred job.

Table 3.6: Career costs avoided by respondents accepting individually preferred job over jointly preferred job

Chance respondent will...	Number of respondents rating ... IP > JP	IP < JP	Sign test (p)
Earn tenure at a research university	5	0	0.031
Earn tenure at a four-year college	3	2	0.500
Publish regularly in top journals	5	0	0.031
Have opportunities to move to more prestigious jobs	4	0	0.063
Have opportunities to move to different kinds of jobs	5	0	0.031
Have higher-than-average lifetime income for field	2	1	0.500
Have everyday work that is satisfying	4	0	0.063
Have plenty of time for life outside of work	1	2	0.875
Observations: 7			
Overall career prospects[a]	8	0	0.004
Observations: 10			

Notes: Observations represent partnered respondents from the 2009-10 job-market cohort who faced a choice between an individually preferred job (IP) and a jointly preferred alternative (JP), and who chose to accept the individually preferred job. The middle two columns in the table present the frequencies of ratings indicating that the individually preferred job is more likely than the jointly preferred job to yield a particular outcome (IP > JP) or that the individually preferred job is less likely to yield the outcome (IP < JP). The excluded category is IP = JP. The sign test column presents one-sided p-values based on the probability of observing the above frequencies of positive signs on the differences between ratings of the individually preferred job and the jointly preferred job from a binomial distribution with mean 0.5. For the eight specific outcomes, respondents rated the likelihood of realizing the outcome on a six-point response scale, where 1 is "extremely unlikely," and 6 is "extremely likely."

[a] Respondents rated their overall career prospects with the individually preferred job relative to their prospects with the jointly preferred job on a five-point scale, where 1 is "much better," 2 is "somewhat better," 3 is "about the same," 4 is "somewhat worse," and 5 is "much worse." Ratings below 3 indicate that the individually preferred job is better than the jointly preferred job, while ratings above 3 that the jointly preferred job is better than the individually preferred job.

Table 3.7: Impact of migration ties on job choice set by stage of job search

	Stage of job search							
	Applications		Interviews		Fly-outs		Job offers	
	Obs	Statistic	Obs	Statistic	Obs	Statistic	Obs	Statistic
By stage								
Median size of choice set[a]	620[c]	100	766	17	761	6	751	3
Percent reporting changed choice set due to migration ties[b]	710	44%	740	12%	707	14%	628	14%
Median change in size of choice set due to migration ties, conditional on having changed choice set[d]	103	-5	33	2	35	-1		
Cumulative to this stage								
Percent reporting changed choice set due to migration ties[e]	544	43%	544	46%	544	49%	544	54%
Percent reporting some or all employers knew about partner			747	72%	604	84%		

Notes: Statistics are computed for respondents who were in a relationship from November through March of the job market year. Sample includes observations from all cohorts for which data are available (see footnotes for more detail). Questions about applications were asked in the pre-market survey for the 2007-08 cohort and in both the post-market and pre-market surveys for later cohorts. Questions about interviews, fly-outs and job offers were only asked in the post-market surveys, so fewer observations are available compared to the applications stage.

[a] At the application stage, *choice set* refers to the number of applications respondents submitted. At the interview and fly-out stages, *choice set* refers to the number of invitations extended by prospective employers for interviews and fly-outs. At the job-offer stage, *choice set* refers to the sum of job offers extended by prospective employers and job offers not officially extended but that respondents are sure would have been extended if they had not indicated to the employers that they would not accept.

[b] At the application stage, *changed choice set* indicates that the respondent applied to a different set of jobs (not necessarily fewer) because of migration ties. At the interview and fly-out stages, *changed choice set* indicates that the respondent accepted a different set of interviews or fly-outs (again, not necessarily fewer). At the job-offer stage, *changed choice set* indicates that respondents rejected their individually preferred job.

[c] In 2007-08, total number of applications is missing for job candidates who did not complete the post-market survey.

[d] Due to changes in the surveys over time, this information is available at the application stage only for the 2008-09 job-market cohort, and at the interview and fly-out stages only for the 2007-08 job-market cohort.

[e] Sample is restricted to the subset of respondents who provided information about choices at every stage. Requires that respondents had non-zero applications, interviews, fly-outs and offers, and provided answers to questions about outcomes and migration ties at each stage.

Table 3.8: Probit regressions: Impact of migration ties and gender on job-market choices

Dependent variable:	(1) Altered job choice	(2) Altered earlier choice	(3) Altered any choice
Female	0.041	0.095**	0.057
	(0.030)	(0.048)	(0.049)
Relationship status			
Married	–	–	–
	–	–	–
Marriage-like	−0.035	−0.084	−0.062
	(0.040)	(0.067)	(0.068)
Committed	−0.074**	−0.094	−0.109*
	(0.033)	(0.064)	(0.064)
Dating	0.006	−0.346***	−0.254**
	(0.072)	(0.085)	(0.102)
Partner education			
Bachelor's degree or less	–	–	–
	–	–	–
Master's or professional degree	0.072**	0.017	0.042
	(0.033)	(0.056)	(0.056)
PhD	0.059*	0.069	0.074
	(0.032)	(0.056)	(0.057)
Observations	613	537	537

Notes: Sample includes partnered respondents from all cohorts. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Regression samples include observations from partnered respondents from all cohorts. Column 1 requires at least one job offer, while Columns 2 and 3 also require that respondents provided information about the impact of migration ties at application, interview and fly-out stages. Dependent variables are dummy variables for rejecting the individually preferred job (Column 1), for altering the application, interview, or fly-outs set in response to migration ties (Column 2), or for doing either of these (Column 3). Partner education is highest degree completed or in progress.

Table 3.9: Relationship outcomes in March after the job-market year

Likely relationship outcome	Percent
Live together	77.8%
Live apart	16.1%
Break up	6.2%
Observations	454

Notes: Sample includes respondents from 2008-10 who were in relationships from November until March of the job-market year. Expected relationship outcomes are the expected outcomes in March of the year after the job-market year, asked at the time of the post-market survey. *Live together* refers to sharing a primary residence or living close enough that the partners can see each other on weeknights; *live apart* refers to not living close enough that the partners can see each other on weeknights. Outcomes are coded as the most likely outcome based on respondents' subjective probabilities of each outcome; respondents were also coded as *break up* if their relationship had already ended by the time of the post-market survey.

Table 3.10: Career and relationship trade-offs

	Expected relationship outcome			Expected relationship outcome		
Sample includes past break-ups:	Yes			No		
Cohorts:	2008-09			2008-10		
Chose individually preferred outcome?	Live together	Live apart	Break up	Live together	Live apart	Break up
Yes	138	32	13	261	54	0
	68.7%	15.9%	6.5%	72.5%	15.0%	0%
No	15	1	2	41	4	0
	7.5%	0.5%	1.0%	11.4%	1.1%	0%
Observations		201			360	

Notes: Sample is restricted to respondents who received at least one job offer and were in a relationship from November until March of the job-market year. Expected relationship outcomes are the expected outcomes in March of the year after the job-market year, asked at the time of the post-market survey. The individually preferred job outcome is the outcome the respondent would have preferred in the absence of constraints imposed by thee partner's preferences or career. It may refer to a particular job offer or to a preference to reject all job offers. In the first panel, respondents whose relationships had ended between March of the job-market year and the time of the post-market survey are included under the *break up* category. In the second panel, respondents whose relationships had ended by the time of the post-market survey are excluded because job preference data for these respondents is missing in 2009-10.

Table 3.11: Determinants of living apart: Career outcome trade-offs

	(1)			(2)
	Coef	Obs	Adj R^2	Coef
IP > JP with respect to chance respondent will...				
Regressors:	Single regressor			All
Have better job prospects overall	0.246***	176	0.089	–
	(0.060)			–
Earn tenure at a research university	0.228***	171	0.049	−0.135
	(0.073)			(0.124)
Earn tenure at a four-year college	0.163*	166	0.016	0.087
	(0.084)			(0.130)
Publish regularly in top journals	0.326***	172	0.083	0.423***
	(0.080)			(0.157)
Have opportunities to move to more prestigious jobs	0.261***	169	0.057	−0.177
	(0.078)			(0.148)
Have opportunities to move to different kinds of jobs	0.446***	171	0.117	0.346**
	(0.092)			(0.141)
Have higher than average income for field	0.139	172	0.006	0.068
	(0.097)			(0.167)

Table 3.11: Determinants of living apart: Career outcome trade-offs *(continued)*

	(1)			(2)
	Coef	Obs	Adj R^2	Coef
IP > JP with respect to chance respondent will...				
Have satisfying everyday work	0.169**	171	0.025	−0.124
	(0.072)			(0.136)
Have plenty of time for life outside of work	0.041	171	−0.005	−0.128
	(0.112)			(0.129)
Observations	–	–	–	160
Adjusted R^2	–	–	–	0.085

Notes: Sample is restricted to 2008-09 and 2009-10 cohorts. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Dependent variable is the subjective probability of living apart in the year after the job market. For overall career prospects, counts are based on a direct counterfactual question with a 1 to 5 response scale, where 1 indicates that "IP is much better than AJ" and 5 indicates that "IP is much worse than AJ." For each of the eight specific career outcomes, respondents rated the likelihood of each outcome on a six-point scale where 1 was "extremely unlikely," 2 was "very unlikely," 3 was "somewhat unlikely," 4 was "somewhat likely," 5 was "very likely," and 6 was "extremely likely." Regressors for all regressions are dummy variables indicating that respondent thought the individually preferred job was better than the jointly preferred job. Column 1 reports estimates from regressions in which the dummy variable for each outcome enters as the sole regressor. Column 2 reports estimates from a regression in which all of the dummy variables enter together.

Table 3.12: Association between relative likelihoods of achieving long-term career outcomes at individually and jointly preferred job

	"You will publish regularly in top journals."				"You will have opportunities to move to different kinds of jobs."			
% IP > JP	IP > JP	IP <= JP	Chi-sq	p	IP > JP	IP <= JP	Chi-sq	p
"You will earn tenure at a research university."	82%	29%	7.337	0.007	63%	45%	0.700	0.403
"You will earn tenure at a four-year college."	45%	38%	0.171	0.679	38%	42%	0.050	0.824
"You will have opportunities to move to more prestigious jobs."	70%	25%	5.106	0.024	86%	26%	7.394	0.007
"Your lifetime income will be higher than average for your field."	36%	24%	0.539	0.463	50%	20%	2.520	0.112
"Your everyday work will be satisfying to you."	73%	38%	3.240	0.072	86%	40%	4.340	0.037
"You will have plenty of time for your life outside of work."	18%	24%	0.113	0.736	25%	20%	0.085	0.771

Notes: Sample is restricted to respondents from the 2008-09 and 2009-10 cohorts who reported that their individually preferred job was not the same as the jointly preferred job. Number of observations ranges from 26 to 28. IP is individually preferred job. JP is jointly preferred job.

Table 3.13: Determinants of living apart: Which specific career outcome trade-offs matter?

IP > JP with respect to chance respondent will...	(1) Overall prospects	(2) Specific outcome	Obs	R^2	*Adj* R^2
Have better job prospects overall	0.246***	–	176	0.089	0.0833
	(0.060)	–			
Earn tenure at a research university	0.222***	0.081	171	0.096	0.085
	(0.080)	(0.089)			
Earn tenure at a four-year college	0.259***	0.016	166	0.092	0.081
	(0.073)	(0.092)			
Publish regularly in top journals	0.164*	0.185*	172	0.107	0.096
	(0.088)	(0.110)			
Have opportunities to move to more prestigious jobs	0.090	0.182	169	0.067	0.056
	(0.097)	(0.115)			
Have opportunities to move to different kinds of jobs	0.115	0.338***	171	0.133	0.122
	(0.082)	(0.120)			
Have higher than average income for field	0.282***	-0.052	172	0.093	0.0824
	(0.073)	(0.105)			

Table 3.13: Determinants of living apart:Which specific career outcome trade-offs matter? *(continued)*

IP > JP with respect to chance respondent will...	(1) Overall prospects	(2) Specific outcome	Obs	R^2	*Adj* R^2
Have satisfying everyday work	0.241***	-0.011	171	0.071	0.060
	(0.089)	(0.098)			
Have plenty of time for life outside of work	0.276***	-0.072	171	0.094	0.083
	(0.066)	(0.111)			

Notes: Sample is restricted to 2008-09 and 2009-10 cohorts. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Dependent variable is the subjective probability of living apart in the year after the job market. For overall career prospects, counts are based on a direct counterfactual question with a 1 to 5 response scale, where 1 indicates that "IP is much better than AJ" and 5 indicates that "IP is much worse than AJ." For each of the eight specific career outcomes, respondents rated the likelihood of each outcome on a six-point scale where 1 was "extremely unlikely," 2 was "very unlikely," 3 was "somewhat unlikely," 4 was "somewhat likely," 5 was "very likely," and 6 was "extremely likely." Regressors for all regressions are dummy variables indicating that respondent thought the individually preferred job was better than the jointly preferred job. Column 1 reports estimates from regressions in which the dummy variable for each outcome enters as the sole regressor. Column 2 reports estimates from a regression in which all of the dummy variables enter together.

Table 3.14: Determinants of living apart: Partner's career-related variables

Cohorts:	(1) 2008-10	(2) 2009-10	(3) 2008-10	(4) 2009-10	(5) 2009-10
Highest degree completed or in progress					
Bachelor's degree or less	–				–
	–				–
Master's or professional degree	0.087**				0.034
	(0.042)				(0.081)
PhD	0.155***				0.096
	(0.043)				(0.076)
Expects to be in school next year	0.092**				0.245***
	(0.041)				(0.085)
"What do you think is the percent chance that your husband/wife/significant other will work at least 20/40/60 hours per week over most of the next ten years?"					
At least 20 hours per week		0.003			0.002
		(0.002)			(0.003)
At least 40 hours per week		−0.001			0.001
		(0.002)			(0.002)

Table 3.14: Determinants of living apart: Partner's career-related variables *(continued)*

	(1)	(2)	(3)	(4)	(5)
Cohorts:	2008-10	2009-10	2008-10	2009-10	2009-10
At least 60 hours per week		0.003***			0.001
		(0.001)			(0.001)
"At the time you accepted your job, how would you have rated your husband's/wife's/significant other's job prospects in this location?"					
Good			–		–
			–		–
Fair			0.113***		0.171***
			(0.040)		(0.060)
Poor			0.327***		0.276***
			(0.061)		(0.103)
Agrees with: "My husband's/wife's/significant other's career will *not* suffer if we move to the places that are best for my career."				−0.178***	−0.138**
				(0.052)	(0.059)
Observations	428	205	344	162	130
R-squared	0.047	0.068	0.086	0.068	0.251

Note: Sample in each column includes partnered respondents from cohorts listed in column header. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Dependent variable is the subjective probability of living apart in the year after the job market. Wording of labor-force attachment questions changed slightly over time. Question text for all years is available in the appendix.

Table 3.15: Determinants of living apart: Relationship-related variables

	(1)	(2)	(3)	(4)
Cohorts:	2008-10	2008-09	2009-10	2009-10
Relationship status				
Married	–			–
	–			–
Marriage-like	0.107**			0.230***
	(0.047)			(0.072)
Committed or dating	0.330***			0.188***
	(0.044)			(0.070)
"In the past few months, how satisfied have you been with your romantic relationship?"				
Somewhat satisfied or dissatisfied		–		
		–		
Very satisfied		−0.042		
		(0.087)		
Extremely satisfied		−0.152*		
		(0.083)		

Table 3.15: Determinants of living apart: Relationship-related variables *(continued)*

	(1)	(2)	(3)	(4)
Cohorts:	2008-10	2008-09	2009-10	2009-10
Agrees with: "It would be possible for me to have a fulfilling relationship while living apart from my husband/wife/significant other."			−0.082	−0.051
			(0.052)	(0.051)
Agrees with: "I would never consider living apart from my husband/wife/significant other."			−0.008	−0.001
			(0.055)	(0.053)
Agrees with: "I would be willing to make a large career sacrifice so that I could live with my husband/wife/significant other."			−0.071	−0.015
			(0.058)	(0.058)
Agrees with: "I would consider jobs that require me to live apart from my husband/wife/significant other for up to one year."			0.059	0.027
			(0.058)	(0.057)
Agrees with: "I would consider jobs that require me to live apart from my husband/wife/significant other for up to five years."			0.194***	0.178**
			(0.071)	(0.069)

Table 3.15: Determinants of living apart: Relationship-related variables *(continued)*

	(1)	(2)	(3)	(4)
Cohorts:	2008-10	2008-09	2009-10	2009-10
"How upset would you be if you were in your current relationship, but living in a different area from your husband/wife/significant other one year from now?"				
Somewhat upset or not upset			–	–
			–	–
Very upset			−0.075	−0.081
			(0.065)	(0.062)
Extremely upset			−0.190***	−0.163**
			(0.066)	(0.064)
Observations	430	144	178	178
R-squared	0.120	0.037	0.186	0.250

Note: Sample in each column includes partnered respondents from cohorts listed in column header. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Dependent variable is the subjective probability of living apart in the year after the job market.

Table 3.16: Determinants of living apart: Child-related variables

	(1)	(2)	(3)	(4)
Cohorts:	2008-10	2008-10	2009-10	2009-10
Has at least one child	−0.149***			−0.136**
	(0.036)			(0.053)
"What do you think is the percent chance that you will have or adopt a/another child in the next year?"				
Low ($0 \leq p < .25$)		–		–
		–		–
Medium ($.25 \leq p < .75$)		−0.020		0.015
		(0.052)		(0.067)
High ($.75 \leq p \leq 1$)		−0.058		−0.016
		(0.058)		(0.079)
"What do you think is the percent chance that you will have or adopt a/another child in the five years?"				
Low ($0 \leq p < .25$)		–		–
		–		–
Medium ($.25 \leq p < .75$)		0.025		−0.123
		(0.053)		(0.087)
High ($.75 \leq p \leq 1$)		−0.102**		−0.213**
		(0.051)		(0.093)

Table 3.16: Determinants of living apart: Child-related variables *(continued)*

	(1)	(2)	(3)	(4)
Cohorts:	2008-10	2008-10	2009-10	2009-10
Agrees with: "I would consider having a child while living apart from my [husband/wife/significant other]."			0.183**	0.192**
			(0.081)	(0.081)
Agrees with: "Living apart from my [husband/wife/significant other] over the next year would prevent us from having as large a family as we would like."			−0.072	−0.034
			(0.056)	(0.057)
Agrees with: "Living apart from my [husband/wife/significant other] over the next five years would prevent us from having as large a family as we would like."			0.015	0.076
			(0.067)	(0.077)
Agrees with: "My children would live with me if my [husband/wife/significant other] and I were living apart."			−0.036	−0.046
			(0.051)	(0.050)
Agrees with: "I could have a very good relationship with my children even if they were not living with me."			0.043	0.018
			(0.053)	(0.053)
Observations	428	397	197	197
R-squared	0.038	0.040	0.040	0.106

Note: Sample in each column includes partnered respondents from cohorts listed in column header. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Dependent variable is the subjective probability of living apart in the year after the job market.

APPENDICES

APPENDIX A

Appendix to Chapter 2: The Impact of the Great Recession on the Retirement Plans of Older Americans

A.1 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

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

subject to

$$\dot{A} = rA_t + \omega_t \chi_t - C_t \tag{A.2}$$

where

$$\chi_t = \begin{cases} 0 \, if \; working \; at \; time \; t \\ \\ 1 \, if \; not \; working \; at \; time \; t \end{cases} \tag{A.3}$$

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

$$\mathcal{H} = \frac{C_t^{1-\frac{1}{\theta}}}{1-\frac{1}{\theta}} - e^{\alpha-\zeta t}\chi_t + \lambda_t \left[rA_t + \omega_t \chi_t - C_t \right] \tag{A.4}$$

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

$$h_c = 0 \Leftrightarrow C_t^{-1/\theta} = \lambda_t \tag{A.5}$$

$$h_A = \rho\lambda_t - \dot{\lambda}_t \Leftrightarrow r\lambda_t = \rho\lambda_t - \dot{\lambda}_t \tag{A.6}$$

$$\dot{A} = rA_t + \omega_t\chi_t - C_t \tag{A.7}$$

Letting χ, the decision to work, be characterized by

$$\chi_t = \begin{cases} 0\, if \quad \lambda_t\omega_t \geq e^{\alpha+\zeta t} \\ 1\, if \quad \lambda_t\omega_t \leq e^{\alpha+\zeta t} \end{cases} \tag{A.8}$$

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

$$\omega_R\lambda_R = e^{\alpha+\zeta t} \tag{A.9}$$

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

$$\lambda_R = \lambda_t e^{(\rho-r)(R-t)} \tag{A.10}$$

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

$$\omega_R\lambda_t e^{(\rho-r)(R-t)} = e^{\alpha+\zeta t} \tag{A.11}$$

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,

$$R = \frac{ln(\lambda_t) + ln(\omega_R) - (\rho - r)t - \alpha}{\varsigma - \rho + r} \tag{A.12}$$

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.

A.2 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.

A.3 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.

A.4 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 A.1: 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.

A.5 Specification comparison results without sampling weights

The following tables present the regression results on which the likelihood ratio tests discussed in Section 2.6.2.4 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 A.2: 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).

Table A.3: 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).

A.6 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 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.

APPENDIX B

Appendix to Chapter 3: Moving Out to Move Up? New Economists Sacrifice Job Opportunities for Proximity to Significant Others– and Vice Versa

B.1 Text of key survey questions

Job preferences: Ranking of job types

Cohort	Survey	Question	Responses	Universe
2007-08	Pre-market	Please rank the following types of job in order of your preference, with "1" being your most preferred and "5" being your least preferred. Research university	Whole number between 1 and 5	All respondents
2007-08	Pre-market	Please rank the following types of job in order of your preference, with "1" being your most preferred and "5" being your least preferred. Four-year college	Whole number between 1 and 5	All respondents

Job preferences: Ranking of job types *(continued)*

Cohort	Survey	Question	Responses	Universe
2007-08	Pre-market	Please rank the following types of job in order of your preference, with "1" being your most preferred and "5" being your least preferred. Postdoctoral fellowship	Whole number between 1 and 5	All respondents
2007-08	Pre-market	Please rank the following types of job in order of your preference, with "1" being your most preferred and "5" being your least preferred. Non-academic, research	Whole number between 1 and 5	All respondents

Job preferences: Ranking of job types *(continued)*

Cohort	Survey	Question	Responses	Universe
2007-08	Pre-market	Please rank the following types of job in order of your preference, with "1" being your most preferred and "5" being your least preferred. Non-academic, non-research	Whole number between 1 and 5	All respondents
2008-09	Pre-market	Please rank the following types of job in order of your preference, with "1" being your most preferred and "5" being your least preferred. Assistant professor at university	Whole number between 1 and 5	All respondents

Job preferences: Ranking of job types *(continued)*

Cohort	Survey	Question	Responses	Universe
2008-09	Pre-market	Please rank the following types of job in order of your preference, with "1" being your most preferred and "5" being your least preferred. Assistant professor at four-year college	Whole number between 1 and 5	All respondents
2008-09	Pre-market	Please rank the following types of job in order of your preference, with "1" being your most preferred and "5" being your least preferred. Postdoctoral fellow	Whole number between 1 and 5	All respondents

Job preferences: Ranking of job types *(continued)*

Cohort	Survey	Question	Responses	Universe
2008-09	Pre-market	Please rank the following types of job in order of your preference, with "1" being your most preferred and "5" being your least preferred. Non-academic researcher (e.g., researcher at a think tank, government research unit, central bank, or international financial organization)	Whole number between 1 and 5	All respondents

Job preferences: Ranking of job types *(continued)*

Cohort	Survey	Question	Responses	Universe
2008-09	Pre-market	Please rank the following types of job in order of your preference, with "1" being your most preferred and "5" being your least preferred. Private sector researcher	Whole number between 1 and 5	All respondents
2009-10	Pre-market	Please rank the following jobs in order of your preference, with "1" being your most preferred and "5" being your least preferred. Assistant professor at a university	Whole number between 1 and 5	All respondents

Job preferences: Ranking of job types *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Pre-market	Please rank the following jobs in order of your preference, with "1" being your most preferred and "5" being your least preferred. Assistant professor at a four-year college	Whole number between 1 and 5	All respondents
2009-10	Pre-market	Please rank the following jobs in order of your preference, with "1" being your most preferred and "5" being your least preferred. Postdoctoral fellow	Whole number between 1 and 5	All respondents

Job preferences: Ranking of job types *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Pre-market	Please rank the following jobs in order of your preference, with "1" being your most preferred and "5" being your least preferred. Researcher at a non-profit, governmental, or quasi-governmental organization	Whole number between 1 and 5	All respondents
2009-10	Pre-market	Please rank the following jobs in order of your preference, with "1" being your most preferred and "5" being your least preferred. Researcher at a business or industry establishment	Whole number between 1 and 5	All respondents

Expected career outcomes

Cohort	Survey	Question	Responses	Universe
2009-10	Pre-market	Imagine that, next year, the faculty in your department compile a list of the job placements of their graduates over the last five years. They put the placements they consider to be most impressive at the top of the list. Thinking about the kind of job you expect to obtain, where do you think you would fall in this list?	Select one: (1) 1st decile (most impressive), (2) 2nd decile, (3) 3rd decile, ..., (9) 9th decile, (10) 10th decile	All respondents

Expected career outcomes *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	How likely do you think it is that, over the course of your life, you will earn tenure at a research university?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who accepted a job
2009-10	Post-market	How likely do you think it is that, over the course of your life, you will earn tenure at a four-year college?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who accepted a job

Expected career outcomes *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	How likely do you think it is that, over the course of your life, you will publish regularly in top journals?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who accepted a job
2009-10	Post-market	How likely do you think it is that, over the course of your life, you will have opportunities to move to more prestigious jobs?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who accepted a job

Expected career outcomes *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	How likely do you think it is that, over the course of your life, you will have opportunities to move to different kinds of jobs?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who accepted a job
2009-10	Post-market	How likely do you think it is that, over the course of your life, your lifetime income will be higher than average for your field?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who accepted a job

Expected career outcomes *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	How likely do you think it is that, over the course of your life, your everyday work will be satisfying to you?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who accepted a job
2009-10	Post-market	How likely do you think it is that, over the course of your life, you will have plenty of time for your life outside of work?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who accepted a job

Individual counterfactual: Applications

Cohort	Survey	Question	Responses	Universe
2007-08	Pre-market	Would the set of all jobs for which you applied have been different in any way if you were not in a relationship?	Select one: (1) Yes, (2) No	Respondents who were in a relationship
2008-09	Pre-market	Are there jobs for which you applied, but for which you would not have applied if you were not in a relationship?	Select one: (1) Yes, (2) No	Respondents who (1) were in a relationship and (2) applied for at least one job
2008-09	Pre-market	Are there jobs for which you did not apply, but for which you would have applied if you were not in a relationship?	Select one: (1) Yes, (2) No	Respondents who were in a relationship

Individual counterfactual: Applications *(continued)*

Cohort	Survey	Question	Responses	Universe
2008-09	Post-market	Imagine that your [husband/wife/significant other] would have been be equally happy with any outcome of your job search — no matter what location or job you eventually chose. In that situation, would you have applied to a different set of jobs?	Select one: (1) Yes, (2) No	Respondents who (1) were in a relationship and (2) applied for at least one job

Individual counterfactual: Applications *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Pre-market	Imagine that your [husband/wife/significant other] would be equally happy with any outcome of your job search — no matter what location or job you eventually chose. In that situation, would you have applied for a different set of jobs?	Select one: (1) Yes, (2) No	Respondents who were in a relationship

Individual counterfactual: Applications *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	Suppose your [husband/wife/significant other] could have an equally satisfying professional and personal life in any location — that is, suppose it would not be a sacrifice for [him/her] to move with you anywhere. In that situation, would you have applied to a different set of jobs?	Select one: (1) Yes, (2) No	Respondents who (1) were in a relationship and (2) applied for at least one job

Individual counterfactual: Interviews

Cohort	Survey	Question	Responses	Universe
2007-08	Post-market	Did you refuse any interviews you would have accepted if you were not in a relationship?	Select one: (1) Yes, (2) No	Respondents who (1) were in a relationship and (2) refused at least one interview invitation
2007-08	Post-market	Did you accept any interviews you would have refused if you were not in a relationship?	Select one: (1) Yes, (2) No	Respondents who (1) were in a relationship and (2) accepted at least one interview invitation

Individual counterfactual: Interviews *(continued)*

Cohort	Survey	Question	Responses	Universe
2008-09	Post-market	Imagine that your [husband/wife/significant other] would have been be equally happy with any outcome of your job search – no matter what location or job you eventually chose. In that situation, would you have accepted a different set of interviews?	Select one: (1) Yes, (2) No	Respondents who (1) were in a relationship and (2) received at least one interview invitation

Individual counterfactual: Interviews *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	Suppose your [husband/wife/significant other] could have an equally satisfying professional and personal life in any location — that is, suppose it would not be a sacrifice for [him/her] to move with you anywhere. In that situation, would you have accepted a different set of interviews?	Select one: (1) Yes, (2) No	Respondents who (1) were in a relationship and (2) received at least one interview invitation

Individual counterfactual: Fly-outs

Cohort	Survey	Question	Responses	Universe
2007-08	Post-market	Did you refuse any fly-outs you would have accepted if you were not in a relationship?	Select one: (1) Yes, (2) No	Respondents who (1) were in a relationship and (2) refused at least one fly-out invitation
2007-08	Post-market	Did you accept any fly-outs you would have refused if you were not in a relationship?	Select one: (1) Yes, (2) No	Respondents who (1) were in a relationship and (2) accepted at least one fly-out invitation

Individual counterfactual: Fly-outs *(continued)*

Cohort	Survey	Question	Responses	Universe
2008-09	Post-market	Imagine that your [husband/wife/significant other] would have been be equally happy with any outcome of your job search – no matter what location or job you eventually chose. In that situation, would you have accepted a different set of fly-outs?	Select one: (1) Yes, (2) No	Respondents who (1) were in a relationship and (2) received at least one fly-out invitation

Individual counterfactual: Fly-outs *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	Suppose your [husband/wife/significant other] could have an equally satisfying professional and personal life in any location — that is, suppose it would not be a sacrifice for [him/her] to move with you anywhere. In that situation, would you have accepted a different set of fly-outs?	Select one: (1) Yes, (2) No	Respondents who (1) were in a relationship and (2) received at least one fly-out invitation

Individual counterfactual: Job choice

Cohort	Survey	Question	Responses	Universe
2007-08	Post-market	If you had considered only your own preferences and had ignored the preferences of your [husband/wife/significant other], would you have accepted the same job or would you have accepted a different job?	Select one: (1) I would have accepted the same job, (2) I would have accepted a different job	Respondents who (1) were in a relationship and (2) accepted a job
2008-09	Post-market	If your [husband/wife/significant other] would have been equally happy with any outcome of your job search, would you have accepted the same job, or would you have accepted a different job?	Select one: (1) Same job, (2) Different job	Respondents who (1) were in a relationship and (2) accepted a job

Individual counterfactual: Job choice *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	Suppose your [husband/wife/significant other] could have an equally satisfying professional and personal life in any location — that is, suppose it would not be a sacrifice for [him/her] to move with you anywhere. In that situation, what do you think the outcome of your job search would have been?	Select one: (1) I would have accepted the same job, (2) I would have accepted a different job, (3) I would not have accepted a job	Respondents who (1) were in a relationship and (2) accepted a job

Individual counterfactual: Job choice *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	Suppose your [husband/wife/significant other] could have an equally satisfying professional and personal life in any location — that is, suppose it would not be a sacrifice for [him/her] to move with you anywhere. In that situation, would you have accepted a job?	Select one: (1) Yes, (2) No	Respondents who (1) were in a relationship, (2) received at least one job offer, and (3) did not accept a job

Individual counterfactual: Expected career outcomes

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	If you had accepted your individually preferred job, how likely do you think it is that, over the course of your life, you would have earned tenure at a research university?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who rejected their individually preferred job
2009-10	Post-market	If you had accepted your individually preferred job, how likely do you think it is that, over the course of your life, you would have earned tenure at a four-year college?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who rejected their individually preferred job

Individual counterfactual: Expected career outcomes *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	If you had accepted your individually preferred job, how likely do you think it is that, over the course of your life, you would have published regularly in top journals?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who rejected their individually preferred job
2009-10	Post-market	If you had accepted your individually preferred job, how likely do you think it is that, over the course of your life, you would have had opportunities to move to more prestigious jobs?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who rejected their individually preferred job

Individual counterfactual: Expected career outcomes *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	If you had accepted your individually preferred job, how likely do you think it is that, over the course of your life, you would have had opportunities to move to different kinds of jobs?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who rejected their individually preferred job
2009-10	Post-market	If you had accepted your individually preferred job, how likely do you think it is that, over the course of your life, your lifetime income would have been higher than average for your field?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who rejected their individually preferred job

Individual counterfactual: Expected career outcomes *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	If you had accepted your individually preferred job, how likely do you think it is that, over the course of your life, your everyday work would have been be satisfying to you?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who rejected their individually preferred job
2009-10	Post-market	If you had accepted your individually preferred job, how likely do you think it is that, over the course of your life, you would have had plenty of time for your life outside of work?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who rejected their individually preferred job

Joint counterfactual: Job choice

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	Where do you think you would have been living next March if you and your were constrained to live together?	Select one: (1) The area where you actually expect to be living next March, (2) The area where your [husband/wife/significant other] actually expects to be living next March, (3) Some other area	Respondents who (1) were in a relationship and (2) did not expect to be living with their partner the following March

Joint counterfactual: Job choice *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	What do you think you would have been doing as your primary activity next March if you and your [husband/wife/significant other] were constrained to live together?	Select all that apply: (1) Working for pay, (2) Looking for a paid job, (3) Going to school, (4) Stay-at-home parenting or other unpaid work, (5) Other	Respondents who, if constrained to live with their partner the following March, would not have lived in the area where they actually expected to be living

Joint counterfactual: Job choice *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	Would you have been working at a job from your job-market choice set?	Select one: (1) Yes, (2) No	Respondents who, if constrained to live with their partner the following March, (1) would not have lived in the area where they actually expected to be living and (2) would be working for pay

Joint counterfactual: Expected career outcomes

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	If you and your [husband/wife/significant other] were constrained to live together next March, how likely do you think it is that, over the course of your life, you would have earned tenure at a research university?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who, if constrained to live with their partner the following March, would not have lived in the area where they actually expected to be living

Joint counterfactual: Expected career outcomes *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	If you and your [husband/wife/significant other] were constrained to live together next March, how likely do you think it is that, over the course of your life, you would have earned tenure at a four-year college?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who, if constrained to live with their partner the following March, would not have lived in the area where they actually expected to be living

Joint counterfactual: Expected career outcomes *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	If you and your [husband/wife/significant other] were constrained to live together next March, how likely do you think it is that, over the course of your life, you would have published regularly in top journals?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who, if constrained to live with their partner the following March, would not have lived in the area where they actually expected to be living

Joint counterfactual: Expected career outcomes *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	If you and your [husband/wife/significant other] were constrained to live together next March, how likely do you think it is that, over the course of your life, you would have had opportunities to move to more prestigious jobs?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who, if constrained to live with their partner the following March, would not have lived in the area where they actually expected to be living

Joint counterfactual: Expected career outcomes *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	If you and your [husband/wife/significant other] were constrained to live together next March, how likely do you think it is that, over the course of your life, you would have had opportunities to move to different kinds of jobs?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who, if constrained to live with their partner the following March, would not have lived in the area where they actually expected to be living

Joint counterfactual: Expected career outcomes *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	If you and your [husband/wife/significant other] were constrained to live together next March, how likely do you think it is that, over the course of your life, your lifetime income would have been higher than average for your field?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who, if constrained to live with their partner the following March, would not have lived in the area where they actually expected to be living

Joint counterfactual: Expected career outcomes *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	If you and your [husband/wife/significant other] were constrained to live together next March, how likely do you think it is that, over the course of your life, your everyday work would have been be satisfying to you?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who, if constrained to live with their partner the following March, would not have lived in the area where they actually expected to be living

Joint counterfactual: Expected career outcomes *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	If you and your [husband/wife/significant other] were constrained to live together next March, how likely do you think it is that, over the course of your life, you would have had plenty of time for your life outside of work?	Select one: (1) Extremely unlikely, (2) Very unlikely, (3) Somewhat unlikely, (4) Somewhat likely, (5) Very likely, (6) Extremely likely	Respondents who, if constrained to live with their partner the following March, would not have lived in the area where they actually expected to be living

Expected partner work hours

Cohort	Survey	Question	Responses	Universe
2008-09	Post-market	What do you think is the percent chance that your [husband/wife/significant other] will work the following schedules most weeks for most of the next ten years? At least [20/40/60] hours per week.	Whole number between 1 and 100	All respondents
2009-10	Post-market	What do you think is the percent chance that your [husband/wife/significant other] will work at least [20/40/60] hours per week over most of the next ten years?	Whole number between 1 and 100	All respondents

Partner job prospects

Cohort	Survey	Question	Responses	Universe
2007-08	Post-market	At the time you accepted your job, how would you have rated your [husband's/wife's/significant other's] job prospects in this location?	Select one: (1) Good, (2) Fair, (3) Poor, (4) Not applicable	Respondents who (1) were in a relationship and (2) accepted a job
2008-09	Post-market	At the time you accepted your job, how would you have rated your [husband's/wife's/significant other's] job prospects in this location?	Select one: (1) Good, (2) Fair, (3) Poor, (4) Not applicable	Respondents who (1) were in a relationship and (2) accepted a job

Partner job prospects *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	At the time you accepted your job, how would you have rated your [husband's/wife's/significant other's] job prospects in this location?	Select one: (1) Good, (2) Fair, (3) Poor, (4) Does not plan to work	Respondents who (1) were in a relationship and (2) accepted a job

Relationship satisfaction

Cohort	Survey	Question	Responses	Universe
2007-08	Pre-market	In the past few months, how satisfied have you been with your romantic relationship?	Select one: (1) Extremely dissatisfied, (2) Very dissatisfied, (3) Somewhat dissatisfied, (4) Somewhat satisfied, (5) Very satisfied, (6) Extremely satisfied	Respondents who were in a relationship

Relationship satisfaction *(continued)*

Cohort	Survey	Question	Responses	Universe
2008-09	Pre-market	In the past few months, how satisfied have you been with your romantic relationship?	Select one: (1) Extremely dissatisfied, (2) Very dissatisfied, (3) Somewhat dissatisfied, (4) Somewhat satisfied, (5) Very satisfied, (6) Extremely satisfied	Respondents who were in a relationship

Attitudes toward living apart: Relationship

Cohort	Survey	Question	Responses	Universe
2009-10	Pre-market	How much do you agree or disagree with the following statements? It would be possible for me to have a fulfilling relationship while living apart from my [husband/wife/significant other].	Select one: (1) Strongly disagree, (2) Disagree, (3) Somewhat disagree, (4) Somewhat agree, (5) Agree, (6) Strongly agree	Respondents who were in a relationship
2009-10	Pre-market	How much do you agree or disagree with the following statements? I would never consider living apart from my [husband/wife/significant other].	Select one: (1) Strongly disagree, (2) Disagree, (3) Somewhat disagree, (4) Somewhat agree, (5) Agree, (6) Strongly agree	Respondents who were in a relationship

Attitudes toward living apart: Relationship *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Pre-market	How much do you agree or disagree with the following statements? I would be willing to make a large career sacrifice so that I could live with my [husband/wife/significant other].	Select one: (1) Strongly disagree, (2) Disagree, (3) Somewhat disagree, (4) Somewhat agree, (5) Agree, (6) Strongly agree	Respondents who were in a relationship
2009-10	Pre-market	How much do you agree or disagree with the following statements? I would consider jobs that require me to live apart from my [husband/wife/significant other] for up to one year.	Select one: (1) Strongly disagree, (2) Disagree, (3) Somewhat disagree, (4) Somewhat agree, (5) Agree, (6) Strongly agree	Respondents who were in a relationship

Attitudes toward living apart: Relationship *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Pre-market	How much do you agree or disagree with the following statements? I would consider jobs that require me to live apart from my [husband/wife/significant other] for up to five years.	Select one: (1) Strongly disagree, (2) Disagree, (3) Somewhat disagree, (4) Somewhat agree, (5) Agree, (6) Strongly agree	Respondents who were in a relationship
2009-10	Pre-market	How upset would you be in you were in your current relationship, but living in a different area from your [husband/wife/significant other] one year from now?	Select one: (1) Not upset, (2) Somewhat upset, (3) Very upset, (4) Extremely upset	Respondents who were in a relationship

Expected fertility

Cohort	Survey	Question	Responses	Universe
2008-09	Post-market	What do you think is the percent chance that you will or adopt [a/another] child in the next [year/five years/ten years]?	Whole number between 1 and 100	All respondents
2009-10	Pre-market	What do you think is the percent chance that you will or adopt [a/another] child in the next [year/five years/ten years]?	Whole number between 1 and 100	All respondents
2009-10	Post-market	What do you think is the percent chance that you will or adopt [a/another] child in the next [year/five years/ten years]?	Whole number between 1 and 100	All respondents

Attitudes toward living apart: Children

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	How much do you agree or disagree with the following statements? I would consider having a child while living apart from my [husband/wife/significant other].	Select one: (1) Strongly disagree, (2) Disagree, (3) Somewhat disagree, (4) Somewhat agree, (5) Agree, (6) Strongly agree	Respondents who were in a relationship
2009-10	Post-market	How much do you agree or disagree with the following statements? Living apart from my [husband/wife/significant other] over the next year would prevent us from having as large a family as we would like.	Select one: (1) Strongly disagree, (2) Disagree, (3) Somewhat disagree, (4) Somewhat agree, (5) Agree, (6) Strongly agree	Respondents who were in a relationship

Attitudes toward living apart: Children *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	How much do you agree or disagree with the following statements? Living apart from my [husband/wife/significant other] over the next five years would prevent us from having as large a family as we would like.	Select one: (1) Strongly disagree, (2) Disagree, (3) Somewhat disagree, (4) Somewhat agree, (5) Agree, (6) Strongly agree	Respondents who were in a relationship
2009-10	Post-market	How much do you agree or disagree with the following statements? My children would live with me if my [husband/wife/significant other] and I were living apart.	Select one: (1) Strongly disagree, (2) Disagree, (3) Somewhat disagree, (4) Somewhat agree, (5) Agree, (6) Strongly agree	Respondents who were in a relationship

Attitudes toward living apart: Children *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	How much do you agree or disagree with the following statements? I could have a very good relationship with my children even if they were not living with me.	Select one: (1) Strongly disagree, (2) Disagree, (3) Somewhat disagree, (4) Somewhat agree, (5) Agree, (6) Strongly agree	Respondents who were in a relationship

Composition of application set

Cohort	Survey	Question	Responses	Universe
2007-08	Pre-market	Of the jobs for which you have applied in the United Stats and Canada, how many fall into the following categories? Research university	Whole number	Respondents who applied to at least one job in the United States or Canada
2007-08	Pre-market	Of the jobs for which you have applied in the United Stats and Canada, how many fall into the following categories? Four-year college	Whole number	Respondents who applied to at least one job in the United States or Canada

Composition of application set *(continued)*

Cohort	Survey	Question	Responses	Universe
2007-08	Pre-market	Of the jobs for which you have applied in the United Stats and Canada, how many fall into the following categories? Postdoctoral fellowship	Whole number	Respondents who applied to at least one job in the United States or Canada
2007-08	Pre-market	Of the jobs for which you have applied in the United Stats and Canada, how many fall into the following categories? Non-academic research (for example, the Federal Reserve Board, think tanks, policy research organizations)	Whole number	Respondents who applied to at least one job in the United States or Canada

Composition of application set *(continued)*

Cohort	Survey	Question	Responses	Universe
2007-08	Pre-market	You indicated that you have applied to jobs in non-academic settings like the Federal Reserve Board, think tanks, and policy research organizations. Have you applied for any non-academic jobs in other settings (for example, consulting, banking, management)?	Whole number	Respondents who applied to at least one job in the United States or Canada

Composition of application set *(continued)*

Cohort	Survey	Question	Responses	Universe
2007-08	Pre-market	For how many such jobs have you applied?	Whole number	Respondents who applied to non-academic jobs in non-academic settings like consulting, banking, or management
2008-09	Pre-market	Of the jobs for which you have applied, how many are of the following types? Assistant professor at university	Whole number	Respondents who applied to at least one job

Composition of application set *(continued)*

Cohort	Survey	Question	Responses	Universe
2008-09	Pre-market	Of the jobs for which you have applied, how many are of the following types? Assistant professor at four-year college	Whole number	Respondents who applied to at least one job
2008-09	Pre-market	Of the jobs for which you have applied, how many are of the following types? Postdoctoral fellow	Whole number	Respondents who applied to at least one job 2008-09
2008-09	Pre-market	Of the jobs for which you have applied, how many are of the following types? Private sector researcher	Whole number	Respondents who applied to at least one job

Composition of application set *(continued)*

Cohort	Survey	Question	Responses	Universe
2008-09	Pre-market	Of the jobs for which you have applied, how many are of the following types? Other	Whole number	Respondents who applied to at least one job
2009-10	Pre-market	Of the applications you have submitted, how many (or about how many) are for jobs of the following types? Assistant professor at a university	Whole number	Respondents who applied to at least one job
2009-10	Pre-market	Of the applications you have submitted, how many (or about how many) are for jobs of the following types? Assistant professor at a four-year college	Whole number	Respondents who applied to at least one job

Composition of application set *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Pre-market	Of the applications you have submitted, how many (or about how many) are for jobs of the following types? Post-doctoral fellow	Whole number	Respondents who applied to at least one job
2009-10	Pre-market	Of the applications you have submitted, how many (or about how many) are for jobs of the following types? Researcher at a non-profit, governmental, or quasi-governmental organization	Whole number	Respondents who applied to at least one job

Composition of application set *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Pre-market	Of the applications you have submitted, how many (or about how many) are for jobs of the following types? Researcher at a business or industry establishment	Whole number	Respondents who applied to at least one job
2009-10	Pre-market	Of the applications you have submitted, how many (or about how many) are for jobs of the following types? Other	Whole number	Respondents who applied to at least one job

Composition of interview and fly-out sets

Cohort	Survey	Question	Responses	Universe
2007-08	Post-market	Of the interviews/fly-outs you were offered, how many were for the following types of jobs? Assistant professor at university	Whole number	Respondents who received at least one interview/fly-out offer
2007-08	Post-market	Of the interviews/fly-outs you were offered, how many were for the following types of jobs? Assistant professor at four-year college	Whole number	Respondents who received at least one interview/fly-out offer
2007-08	Post-market	Of the interviews/fly-outs you were offered, how many were for the following types of jobs? Postdoctoral fellow	Whole number	Respondents who received at least one interview/fly-out offer

Composition of interview and fly-out sets *(continued)*

Cohort	Survey	Question	Responses	Universe
2007-08	Post-market	Of the interviews/fly-outs you were offered, how many were for the following types of jobs? Non-academic researcher (e.g., researcher at a think tank, government research unit, central bank, or international financial organization)	Whole number	Respondents who received at least one interview/fly-out offer
2007-08	Post-market	Of the interviews/fly-outs you were offered, how many were for the following types of jobs? Private sector researcher	Whole number	Respondents who received at least one interview/fly-out offer

Composition of interview and fly-out sets *(continued)*

Cohort	Survey	Question	Responses	Universe
2007-08	Post-market	Of the interviews/fly-outs you were offered, how many were for the following types of jobs? Other	Whole number	Respondents who received at least one interview/fly-out offer
2008-09	Post-market	Of the interviews/fly-outs you were offered, how many (or about how many) were for the following types of job? Assistant professor at a university	Whole number	Respondents who received at least one interview/fly-out offer

Composition of interview and fly-out sets *(continued)*

Cohort	Survey	Question	Responses	Universe
2008-09	Post-market	Of the interviews/fly-outs you were offered, how many (or about how many) were for the following types of job? Assistant professor at a four-year college	Whole number	Respondents who received at least one interview/fly-out offer
2008-09	Post-market	Of the interviews/fly-outs you were offered, how many (or about how many) were for the following types of job? Postdoctoral fellow	Whole number	Respondents who received at least one interview/fly-out offer

Composition of interview and fly-out sets *(continued)*

Cohort	Survey	Question	Responses	Universe
2008-09	Post-market	Of the interviews/fly-outs you were offered, how many (or about how many) were for the following types of job? Researcher at a non-profit, governmental, or quasi-governmental organization	Whole number	Respondents who received at least one interview/fly-out offer
2008-09	Post-market	Of the interviews/fly-outs you were offered, how many (or about how many) were for the following types of job? Researcher at a business or industry establishment	Whole number	Respondents who received at least one interview/fly-out offer

Composition of interview and fly-out sets *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	Of the interviews/fly-outs you were offered, about how many were for the following types of jobs? Assistant professor at a university	Whole number	Respondents who received at least one interview/fly-out offer
2009-10	Post-market	Of the interviews/fly-outs you were offered, about how many were for the following types of jobs? Assistant professor at a four-year college	Whole number	Respondents who received at least one interview/fly-out offer
2009-10	Post-market	Of the interviews/fly-outs you were offered, about how many were for the following types of jobs? Postdoctoral fellow	Whole number	Respondents who received at least one interview/fly-out offer

Composition of interview and fly-out sets *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	Of the interviews/fly-outs you were offered, about how many were for the following types of jobs? Researcher at a non-profit, governmental, or quasi-governmental organization	Whole number	Respondents who received at least one interview/fly-out offer
2009-10	Post-market	Of the interviews/fly-outs you were offered, about how many were for the following types of jobs? Researcher at a business or industry establishment	Whole number	Respondents who received at least one interview/fly-out offer

Composition of interview and fly-out sets *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	Of the interviews/fly-outs you were offered, about how many were for the following types of jobs? Other	Whole number	Respondents who received at least one interview/fly-out offer

Composition of job-offer set

Cohort	Survey	Question	Responses	Universe
2007-08	Post-market	How many offers did you receive for the following types of job? Assistant professor at university	Whole number	Respondents who received at least one job offer
2007-08	Post-market	How many offers did you receive for the following types of job? Assistant professor at four-year college	Whole number	Respondents who received at least one job offer
2007-08	Post-market	How many offers did you receive for the following types of job? Postdoctoral fellow	Whole number	Respondents who received at least one job offer

Composition of job-offer set *(continued)*

Cohort	Survey	Question	Responses	Universe
2007-08	Post-market	How many offers did you receive for the following types of job? Non-academic researcher (e.g., researcher at a think tank, government research unit, central bank, or international financial organization)	Whole number	Respondents who received at least one job offer
2007-08	Post-market	How many offers did you receive for the following types of job? Private sector researcher	Whole number	Respondents who received at least one job offer
2007-08	Post-market	How many offers did you receive for the following types of job? Other	Whole number	Respondents who received at least one job offer

Composition of job-offer set *(continued)*

Cohort	Survey	Question	Responses	Universe
2008-09	Post-market	How many offers did you receive for the following types of job? Assistant professor at a university	Whole number	Respondents who received at least one job offer
2008-09	Post-market	How many offers did you receive for the following types of job? Assistant professor at a four-year college	Whole number	Respondents who received at least one job offer
2008-09	Post-market	How many offers did you receive for the following types of job? Postdoctoral fellow	Whole number	Respondents who received at least one job offer

Composition of job-offer set *(continued)*

Cohort	Survey	Question	Responses	Universe
2008-09	Post-market	How many offers did you receive for the following types of job? Researcher at a non-profit, governmental, or quasi-governmental organization	Whole number	Respondents who received at least one job offer
2008-09	Post-market	How many offers did you receive for the following types of job? Researcher at a business or industry establishment	Whole number	Respondents who received at least one job offer
2009-10	Post-market	How many offers did you receive for the following types of jobs? Assistant professor at a university	Whole number	Respondents who received at least one job offer

Composition of job-offer set *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	How many offers did you receive for the following types of jobs? Assistant professor at a four-year college	Whole number	Respondents who received at least one job offer
2009-10	Post-market	How many offers did you receive for the following types of jobs? Postdoctoral fellow	Whole number	Respondents who received at least one job offer
2009-10	Post-market	How many offers did you receive for the following types of jobs? Researcher at a non-profit, governmental, or quasi-governmental organization	Whole number	Respondents who received at least one job offer

Composition of job-offer set *(continued)*

Cohort	Survey	Question	Responses	Universe
2009-10	Post-market	How many offers did you receive for the following types of jobs? Researcher at a business or industry establishment	Whole number	Respondents who received at least one job offer
2009-10	Post-market	How many offers did you receive for the following types of jobs? Other	Whole number	Respondents who received at least one job offer

BIBLIOGRAPHY

BIBLIOGRAPHY

Benitez-Silva, H., and D. S. Dwyer (2005), The rationality of retirement expectations and the role of new information, *Review of Economics and Statistics*, *87*(3), 587–592.

Bielby, W. T., and D. D. Bielby (1992), I will follow him: Family ties, Gender-Role beliefs, and reluctance to relocate for a better job, *American Journal of Sociology*, *97*(5), 1241–1267.

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

Blundell, R., T. Macurdy, O. Ashenfelter, and D. Card (1999), Chapter 27: Labor supply: A review of alternative approaches, in *Handbook of Labor Economics*, vol. Volume 3, Part 1, pp. 1559–1695, Elsevier.

Bosworth, B., and G. Burtless (2010), Recessions, wealth destruction, and the timing of retirement, *Retirement Research Center at Boston College Working Paper Series*, *WP 2010-22*.

Boyle, P. J., T. J. T. J. Cooke, K. Halfacree, and D. Smith (2001), A Cross-National comparison of the impact of family migration on women's employment status, *Demography*, *38*(2), 201–213, doi:10.1353/dem.2001.0012.

Chan, S., and A. H. Stevens (2004), Do changes in pension incentives affect retirement? A longitudinal study of subjective retirement expectations, *Journal of Public Economics*, *88*(7-8), 1307–1333.

Coile, C. C., and P. B. Levine (2005), Bulls, bears, and retirement behavior, *Industrial and Labor Relations Review*, *59*.

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

Compton, J., and R. A. Pollak (2007), Why are power couples increasingly concentrated in large metropolitan areas?, *Journal of Labor Economics*, *25*(3), 475–512, doi:10.1086/512706.

Cook, C., F. Heath, and R. L. Thompson (2000), A meta-analysis of response rates in web- or internet-based surveys, *Educational and Psychological Measurement*, *60*(6), 821–836, doi:10.1177/00131640021970934.

Cooke, T. J. (2008), Migration in a family way, *Population, Space and Place*, *14*(4), 255–265, doi:10.1002/psp.500.

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

Costa, D. L., and M. E. Kahn (2000), Power couples: Changes in the locational choice of the college educated, 1940-1990, *The Quarterly Journal of Economics*, *115*(4), 1287–1315.

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

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

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

Gemici, A. (2008), Family migration and labor market outcomes, *Working paper*, New York University Department of Economics.

Gerstel, N., and H. E. Gross (1982), Commuter marriages – a review, *Marriage & Family Review*, *5*(2), 71–93, doi:10.1300/J002v05n02_05.

Goda, G. S., J. B. Shoven, and S. N. Slavov (2010), Does stock market performance influence retirement expectations?, *National Bureau of Economic Research Working Paper Series*, *No. 16211*.

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

Gross, H. E. (1980), Dual-Career couples who live apart: Two types, *Journal of Marriage and Family*, *42*(3), 567–576, doi:10.2307/351900.

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

Gustman, A. L., T. L. Steinmeier, and N. Tabatabai (2009), 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 (2010a), *Pensions in the Health and Retirement Study*, Harvard University Press.

Gustman, A. L., N. Tabatabai, and T. L. Steinmeier (2010b), Table data: Pensions in the Health and Retirement Study (Final version 1.0).

Health and Retirement Study (2009), 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.

Health and Retirement Study (2010a), 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.

Health and Retirement Study (2010b), 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.

Health and Retirement Study (2010c), 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.

Health and Retirement Study (2010d), 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.

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

Holtz-Eakin, D., D. Joulfaian, and H. S. Rosen (1993), The Carnegie conjecture: Some empirical evidence, *The Quarterly Journal of Economics*, *108*(2), 413–435.

Hurd, M., and K. McGarry (1995), Evaluation of the subjective probabilities of survival in the health and retirement study, *Journal of Human Resources*, *30*, S268–S292.

Hurd, M., and M. Reti (2001), 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*.

Hurd, M., and J. P. Smith (2002), Expected bequests and their distribution, *National Bureau of Economic Research Working Paper Series*, *No. 9142*.

Hurd, M., M. Reti, and S. Rohwedder (2009), The effect of large capital gains or losses on retirement, in *Developments in the Economics of Aging*, edited by D. Wise, National Bureau of Economic Research Conference Report, pp. 127–172, University of Chicago Press, Chicago, IL.

Hurd, M. D. (2009), Subjective probabilities in household surveys, *Annual Review of Economics*, *1*(1), 543–564.

Hurd, M. D., and S. Rohwedder (2010a), 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*.

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

Imbens, G. W., D. B. Rubin, and B. I. Sacerdote (2001), 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.

Jacobson, J. P., and L. M. Levin (1997), Marriage and migration: Comparing gains and losses from migration for couples and singles, *Social Science Quarterly, 78*(3), 688–709.

Kapinos, K. (2010), 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.

Kapinos, K., C. Brown, M. Nolte, H. Stolyarova, and D. Weir (2011), Prospective Social Security Wealth Measures of Pre-retirees, Public release, Version 4.0.

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

Kimball, M. S., and M. D. Shapiro (2003), 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), Labor supply: Are the income and substitution effects both large or both small?, *National Bureau of Economic Research Working Paper Series, No. 14208*.

Lichter, D. T. (1983), Socioeconomic returns to migration among married women, *Social Forces, 62*(2), 487–503, doi:10.2307/2578318.

Low, H., C. Meghir, and L. Pistaferri (2010), Wage risk and employment risk over the life cycle, *American Economic Review, 100*(4), 1432–67.

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

Magnuson, S., and K. Norem (1999), Challenges for higher education couples in commuter marriages: Insights for couples and counselors who work with them, *The Family Journal, 7*(2), 125 –134, doi:10.1177/1066480799072005.

Manski, C. F. (2004), Measuring expectations, *Econometrica*, *72*(5), 1329–1376.

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

McKinnish, T. (2008), Spousal mobility and earnings, *Demography*, *45*(4), 829–849.

Mincer, J. (1978), Family migration decisions, *The Journal of Political Economy*, *86*(5), 749–773.

Mitchell, O. S., J. M. Poterba, M. J. Warshawsky, and J. R. Brown (1999), New evidence on the money's worth of individual annuities, *The American Economic Review*, *89*(5), 1299–1318.

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

RAND Corporation (2010), 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.

Rhodes, A. R. (2002), Long-Distance relationships in Dual-Career commuter couples: A review of counseling issues, *The Family Journal*, *10*(4), 398 –404, doi:10.1177/106648002236758.

Sandell, S. H. (1977), Women and the economics of family migration, *The Review of Economics and Statistics*, *59*(4), 406–414, doi:10.2307/1928705.

Sass, S. A., C. Monk, and K. Haverstick (2010), Workers' response to the market crash: Save more, work more?, *Center for Retirement Research at Boston College, Issue Brief 10-3*.

Schwartz, C. R., and R. D. Mare (2005), Trends in educational assortative marriage from 1940 to 2003, *Demography*, *42*(4), 621–646.

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

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

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

S&P/Case-Shiller (2010), Home Price Index Data, seasonally-adjusted.

Stock, J. H., and D. A. Wise (1990), Pensions, the option value of work, and retirement, *Econometrica*, *58*(5), 1151–1180.

United States Census Bureau (2010), *Statistical Abstract of the United States: 2011*, chap. Section 12: Labor Force, Employment, and Earnings, 130th ed., Washington, D.C.

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

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

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

U.S. Census Bureau (n.d.), American community survey 5-Year estimates, 2005-2009.

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

Yaari, M. E. (1965), Uncertain lifetime, life insurance, and the theory of the consumer, *The Review of Economic Studies*, *32*(2), 137–150.

Yahoo Finance (n.d.), Dow Jones Industrial Average Historical Price Data.