TIAA Institute

Research Dialogue | Issue no. 210
April 2024

How do life events affect retirement timing? Evidence from expectations data

Aspen Gorry
Clemson University

Jonathan M. Leganza
Clemson University

Abstract

The timing of retirement is a major determinant of lifetime income and is thus a crucial factor in financial security. It is also an input in the design of retirement savings schemes, which frequently target retirement dates. Yet people face uncertainty about the timing of their retirement. Understanding how people navigate this uncertainty and how life events influence when they choose to retire is important. In this paper, we use data on retirement expectations from the Health and Retirement Study to provide new evidence on the factors that shape expected retirement timing. In a descriptive analysis, we document how expectations evolve with age and assess how demographic, economic and health characteristics relate to expectations. Retirement expectations do not fluctuate substantially with age, but older workers tend to expect to work longer than younger workers. In a causal analysis, we use an event study framework to estimate the effects of various life events on retirement expectations. We find clear evidence that many health shocks—declines in health status, cancer, lung disease and arthritis—cause decreases in the likelihood of working past common retirement ages, while there is less evidence that economic or family shocks generate such declines. Overall, our results provide a broad view of how several important events experienced by many people during working life impact the timing of later retirements.

1. Introduction

The timing of retirement is a fundamental factor for financial security. The age at which a worker retires determines their total lifetime earnings and the time horizon over which private savings—and, potentially, public pension benefits—are needed to finance consumption. Indeed, for many Americans, a dominant form of saving involves contributing to employer-sponsored retirement savings schemes. The age at which one plans to retire is an important factor for both the appropriate level of contributions to these schemes and the mix of assets to hold.

Yet, people face substantial uncertainty about the date of their retirement. For instance, Caliendo et al. (2023) show that the difference between the age workers expect to retire and their actual retirement age has a standard deviation of three to six years for a sample of older workers. Workers face many risks and uncertainties related to work capacity and employment, especially at older ages. This uncertainty about the timing of retirement could result from many risks that workers face, such as how their health will evolve with age, the labor market opportunities available to them in the future, and how their family circumstances may change with time.

Therefore, it is important to understand how people navigate these risks and how these factors—health, economic characteristics and family situations—shape retirement timing. However, providing causal evidence that quantifies how these factors during working life influence the timing of retirement is difficult. There are several empirical challenges. One is identification. For example, consider the effect of health on retirement. Correlations between health and retirement ages may not reflect causal relationships because of concerns like omitted variable bias. Unobserved variables, such as the preference for leisure, influence both health and retirement age. Another challenge relates to measurement. Retirements are often not realized until years after a person experiences circumstances that can impact the timing of their retirement. For example, a person could experience a health shock in their early 50s that causes them to retire in their 60s instead of in their 70s. Even with a long panel dataset that tracks individuals for decades, it's difficult to know when they would have retired without the health shock and to isolate the effects of that specific shock from other factors that change over time. It is especially difficult to causally link factors experienced earlier in life to retirement decisions that happen later.

In this paper, we use an event study framework and data on expectations to overcome these challenges and provide new causal evidence on how life events impact retirement timing. We leverage detailed survey data from the Health and Retirement Study (HRS) to explore how people's expectations about their retirement change after many frequently experienced events or "shocks." We study three types of events: health shocks, economic shocks and family shocks. We overcome the identification challenges by exploiting the timing of shocks within an event study framework. To overcome the measurement challenge that retirement is not necessarily contemporaneous with the shock, we study outcome variables that capture retirement expectations following Dwyer and Mitchell (1999) and McGarry (2004). Because prior work finds that people in the HRS data have rational expectations about retirement timing (Benitez-Silva & Dwyer, 2005) and their expectations strongly predict future retirements (Haider & Stephens, 2007), we can interpret our findings about how events change retirement expectations as providing evidence on how they influence the ultimate timing of retirement.

Our analysis proceeds in two steps. In the first step, we document patterns in retirement expectations as workers age from 50 to 60, and we investigate how expectations and trends in expectations relate to demographic characteristics, economic circumstances and health. Expectations don't fluctuate substantially with age—but, on average, older workers expect to retire later than younger workers. We also document how expectations vary across groups in both levels and trends, finding similar overall patterns among most groups of workers.

In the second step, we estimate the causal effects of shocks on retirement expectations. We define an analysis sample of people who experience the shock for each event. We then track the evolution of retirement expectation outcome variables around the time of the shock using the regression frameworks from Dobkin et al. (2018). Our primary outcome variables are the self-assessed probabilities of working past ages 62 and 65. These outcomes have policy-relevant interpretations because they correspond to eligibility ages for Social Security and Medicare. Crucially, these outcomes exist for workers and nonworkers in recent survey waves of the HRS, allowing us to avoid sample selection issues that would threaten the identification of causal effects.

We begin our casual analysis by investigating the impact of adverse health shocks. We first study the effect of sudden changes in self-reported health status and find that declines in health status cause people to significantly revise downward their probability of working at older ages. The event study estimates show flat pre-shock trends in outcome variables before the health decline, followed by large declines in the likelihood of working past ages 62 and 65 after the shock. Our preferred specification indicates that a decline in health status reduces the self-assessed likelihood of working past 62 by 4.0 percentage points, an

8.9% decrease from the baseline mean of 45%. Similarly, we find a reduction in the likelihood of working past age 65 of 4.5 percentage points, a 14.1% decrease from a mean of 32%.

We then study other health shocks, including hospitalizations and the diagnoses of eight different health conditions (heart attacks, strokes, cancer, lung disease, arthritis, diabetes, high blood pressure and psychiatric problems). Some, but not all, of these health shocks significantly impact retirement expectations. New diagnoses of cancer, lung disease and arthritis generate clear declines in the probabilities of working past ages 62 and 65. For instance, the point estimates for these three shocks indicate reductions in the probabilities of working past age 62 that amount to 7.2 percentage points (16.4%), 7.4 percentage points (26.4%) and 4.7 percentage points (10.2%), respectively. While these three conditions cause people to expect to retire earlier than previously expected, we find little evidence of a change in the probability of working past key ages for hospitalizations or new diagnoses of heart attacks, diabetes, high blood pressure and psychiatric problems. This finding could reflect that some conditions are more debilitating and lead to sustained decreases in work capacity, while others are more manageable.

Next, we study economic events. We study unemployment as a negative shock and large earnings gains as a positive shock. We find no statistically significant evidence that unemployment between ages 50 and 60 leads people to adjust their retirement expectations. However, the point estimates for the probabilities of working past ages 62 and 65 are negative. For earnings gains, we find some evidence of declines in the probabilities of working later, although the graphical evidence is less clear than for the health shocks.

Finally, we study two distinct family events. We first analyze the birth of a grandchild, finding no statistically significant evidence that people update their expectations after this event. We next analyze divorce or marital separation. The graphical evidence suggests an increase in the probability of working later after divorce, but our preferred point estimates for this shock are not statistically significant.

Taken together, our results provide a broad view of how several important life events that many Americans experience impact retirement timing. Some health shocks cause decreases in the likelihood of working past older ages, which we interpret as inducing earlier retirement. In contrast, there is less evidence that economic and family events significantly impact retirement expectations.

Our paper relates broadly to the large literature that studies determinants of retirement. One set of related papers studies the effect of health on current retirement and labor supply measures (e.g., McClellan, 1998; Coile, 2004; Disney et al., 2006; García-Gómez et al., 2013).[1] Another set studies how unemployment at older ages impacts retirement (e.g., Chan & Stevens, 1999; Chan & Stevens, 2001; Coile & Levine, 2007), and an emerging line of work estimates the effects of grandchildren on employment measures (e.g., Rupert & Zanella, 2018; Backhaus & Barslund, 2021; Frimmel et al., 2022; Gørtz et al., 2023; Karademir et al., 2023). Yet papers that study how health, unemployment and family events experienced earlier impact the timing of later retirements are comparatively scarce.

The most closely related papers also study how these earlier factors impact the timing of future retirements by estimating how they impact retirement expectations.[2] Influential papers by Dwyer and Mitchell (1999) and McGarry (2004) were the first to use HRS expectations data to study the relationship between health and retirement; the former used the first wave of the survey and the latter the first two waves to establish important links between health and expectations. But strictly causal interpretations of the estimates in these earlier studies could be limited by correlated unobservables—i.e., "bad controls" (Angrist & Psichke, 2009) because the regressions control for endogenous outcomes like earnings, and sample selection because the expectations variables only exist for workers in the earlier survey waves. McGarry (2004) summarizes the selection issue by noting that the "drawback of this methodology is that because the expected probability of full-time work is available only for those still in the labor force, the sample is a selected one." Gupta and Larsen (2010) build on these papers by using two waves of Danish survey data linked to register health data to study how the relationship changes between health and retirement expectations when using administrative versus survey health measures. They

1 A separate literature studies the relationship between health and retirement in the opposite direction and investigates the impact of retirement on health. For a few examples, see Coe and Zamarro (2011), Coe et al. (2012), Eibich (2015), Gorry et al. (2018) and Nielsen (2019).

2 Other related papers take different approaches to studying retirement expectations. Two recent papers study expected retirements compared to actual retirements; Munnell et al. (2018) estimate regression models that describe how a broad set of key explanatory variables relate to earlier-than-planned retirements, and Caliendo et al. (2023) show how a rich set of factors relate to the difference between actual and expected retirements. Four other papers analyze how factors our paper doesn't study influence retirement expectations: Goda et al. (2011) and McFall (2011) study stock market fluctuations and the Great Recession, Ayyagari (2019) studies the Affordable Care Act, and Hudomiet et al. (2021) study hypothetical changes to job characteristics.

also study only workers and estimate regressions similar to McGarry (2004) but use a more structural approach to correct for sample selection. Finally, Chan and Stevens (1999) estimate the relationship between unemployment and retirement expectations using three early waves of HRS data, but they also note that these estimates are limited by sample selection and are conditional on reemployment.

Our main contribution is to use a quasi-experimental framework to produce new causal estimates of shocks on retirement expectations. We use an event study approach to leverage the quasi-random timing of shocks, to cleanly isolate the effects of different shocks, and to provide graphical evidence that allows for more transparent assessments of key assumptions. Importantly, we avoid sample selection problems by using more recent waves of HRS data that include well-defined outcomes even for nonworkers. In addition to these methodological updates, we provide an analysis of additional outcome variables and produce estimates for birth cohorts approaching retirement during more recent times. The effects of life events on expectations are likely to depend on the setting; for instance, increasing capacity to work at older ages and changes to the retirement policy landscape (e.g., the rise of defined contribution retirement accounts, increasing Social Security ages) can influence how people respond to shocks. Our analysis offers an improved and broader understanding of how different life events impact retirement timing.

Our findings have implications for policymakers and practitioners concerned with the retirement income security of Americans and retirement plan design. Employer-sponsored retirement savings schemes, such as 401(k)s, are important savings vehicles in the U.S., and expected retirement dates are crucial for implementing these plans. Indeed, there has been a rapid increase in target date retirement funds, which explicitly anchor investment options to expected dates of retirement and now make up about 24% of total assets in 401(k)s (Shoven & Walton, 2021). Therefore, understanding how retirement expectations evolve with age and how life events causally influence expectations is crucial to designing retirement plans appropriately. Our analysis can help practitioners looking to account for retirement timing risk in designing retirement savings schemes. In general, understanding how life events that older people face induce changes in expectations is a first step toward understanding how individuals should adjust their retirement savings and how retirement savings schemes could help workers adjust savings after various shocks.

The rest of this paper proceeds as follows. Section 2 discusses the HRS data used in our analysis. Section 3 presents the descriptive analysis that documents the evolution of retirement expectations as workers age. Section 4 presents the casual analysis that estimates the effects of health, economic and family events on retirement expectations. Section 5 concludes.

2. Data

We use data from the Health and Retirement Study (HRS), a longitudinal, biennial survey dataset in the United States covering individuals over age 50 and their spouses. It consists of seven sample cohorts based on the date of their first interview. The first of these HRS cohorts was initially interviewed in 1992; the most recent cohort was initially interviewed in 2016. To access the data, we primarily use the RAND HRS Longitudinal File 2020 (v1) dataset (Bugliari et al., 2023a), which we supplement with the RAND HRS Family Data File 2018 (v2) (Bugliari et al., 2023b). The RAND Longitudinal File is a cleaned and streamlined product from the RAND Center for the Study of Aging that includes key information on all survey cohorts and every person interviewed. The Family Data File contains additional information on family characteristics.

The HRS data are well suited for our analysis for two reasons. First, the breadth of the survey allows for a thorough analysis. Crucially, the data contain information on retirement expectations. The data also contain detailed information on demographics, family relationships, health, income and assets, which allows us to analyze several factors that may influence retirement expectations. Second, the survey's focus on older households produces a sizable sample of individuals approaching retirement age.

2.1. Key variables capturing retirement expectations

The outcome variables in our analysis capture expectations about retirement timing. We study four outcome variables. First is the expected retirement age. The HRS contains a variable that is the calendar year an individual expects to retire. We define expected retirement age as the calendar year of the expected retirement minus the individual's birth year. The other three outcome variables are probabilities of working past key ages associated with the Social Security program. Specifically, the HRS contains variables for the probability of working past ages 62, 65, and 70. These ages are particularly relevant as they correspond with eligibility for important policies. Age 62 corresponds to the Social Security Early Eligibility Age, when individuals first become eligible to claim (reduced) old-age benefits from the Social Security program. Age 65 corresponds to when individuals are eligible to claim Medicare and to the original Full Retirement Age for Social Security when individuals were eligible for full benefits. Delaying claiming past the Full Retirement Age increases benefits until age 70.

The expected retirement age outcome and the probability-based outcomes provide complementary information and together allow for an in-depth exploration of how retirement expectations evolve. A potential advantage of the expected retirement age outcome is that it provides a clearer and more precise picture of when an individual expects to retire. For example, consider two people who expect to retire at different times: one who expects to retire at age 71 and one who expects to retire at age 75. The expected retirement age variable should accurately reflect the differences in expectations, whereas the other outcome variables may not. For instance, the two people could report the same probability of working past age 62.

On the other hand, a potential advantage of the probability-based outcomes is that they can capture more nuanced changes in expectations that the expected retirement age variable might miss. For example, consider two people who each expect to retire at age 65 (a commonly reported expected retirement age). These two people have the same expected retirement age. Still, they could have different probabilities of working past age 65 due to differences in job characteristics, health insurance or spousal retirement plans, among many other potential factors. Measuring the probabilities of working past specific ages would accurately reflect these differences in expectations, whereas the expected retirement age variable would not. Additionally, these probability-based measures have policy-relevant interpretations, as retirement before ages 62 and 65 correspond to early retirements where individuals are not yet eligible for Social Security and Medicare, respectively.

2.2. Key variables capturing demographics, family relationships, health, income and assets

The data also contain a rich set of other variables on demographic characteristics, family relationships, health and economic circumstances. Specifically, we use demographic and family variables that capture birth year, age, gender, education, marital status and the number of grandchildren. We use economic variables that capture labor market status, earnings and assets in individual retirement accounts (IRAs). Finally, we use several health-related variables that capture health status (where health is poor, fair, good, very good or excellent), overnight hospitalizations and diagnosed health conditions. The conditions that we study include heart attacks, strokes, cancer, lung disease, diabetes, high blood pressure, arthritis and psychiatric problems. The variables for the diagnoses are binary variables that take the value of one if the respondent says a doctor has ever told them that they have been diagnosed with the condition. Almost all of these variables are available in all waves of the data and come from the RAND Longitudinal File dataset. The exception is the number of grandchildren. We obtain this variable from the RAND HRS Family Data File 2018 (v2), which does not include survey wave 15, so we only have information on grandchildren in waves 1 through 14.

2.3. Constructing analysis samples

We're interested in two sets of analysis samples, one corresponding to our descriptive analysis and one corresponding to our causal analysis. In both cases, our sample of interest includes people approaching retirement age. Therefore, we start with the raw dataset of people in all available survey waves, 1 through 15, and make two sample restrictions. First, we keep only those alive and who responded to the survey. Second, we keep only observations of people between the ages of 50 (the youngest age of respondents in the HRS data) and 60 (two years before people reach key retirement ages).

Next, we confront an important data limitation: The retirement expectation outcome variables have missing values for many observations. The missing data concerns are particularly stark for the expected retirement age variable. The RAND dataset contains two variables that capture expected retirement age. Both variables come from the questionnaire corresponding to the "current job" section of the survey. The first variable is not very useful for our purposes because after imposing our two basic sample restrictions, only about 13% of observations contain nonmissing values.[3] The second variable is more useful. It combines the first variable with information from an additional question about expected retirement timing. This additional question comes right after asking respondents whether they are completely retired. After imposing our two basic sample restrictions, about 32% of observations have nonmissing values.

We use the second variable on expected retirement age, but there are still meaningful amounts of missing values.

3 The variable is based on a sequence of two questions regarding retirement plans. First, a question is asked about retirement plans in general; to that the respondent can reply with many possible answers. Then, if one of the answers the respondent gives indicates they plan to stop working altogether at some point in the future, a follow-up question asks when. The variable based on this follow-up question is sparsely populated in the data.

The most important contributing factor is that the expected retirement age is missing for people not currently working (about 25% of observations) because the underlying survey questions are only asked to those currently working. Another factor is that the additional question used to make the variable is not asked in survey wave 2, so all values are missing in that wave, which accounts for about 9% of the observations. Finally, the values are missing for people who respond by saying they will never retire, accounting for about 23% of observations.

The missing data problem is much less severe for the probability-based outcome variables. These variables come from the questionnaire corresponding to the "expectations" section of the survey. The questions about working past ages 62 and 65 are missing for people who are not currently working, but only during the earlier survey waves. Missing data is not much of a concern for later survey waves because the questions are asked regardless of working status starting in wave 8. Indeed, after imposing our two sample restrictions, about 94% of values are nonmissing in waves 8 through 15. The question about working past age 70 only exists from wave 11 onward, which somewhat limits its use, but it is similarly reliable: about 93% of values are nonmissing.

How do we handle the missing data? In our descriptive analysis, we essentially sidestep the main issues by studying a selected sample of workers who expect to retire in the future. Our goal is not to advance causal interpretations of the patterns we document but rather to describe important trends in retirement expectations for a particular group of interest: people who are working and approaching retirement. Therefore, we further restrict our attention to people who work full time or part time and who report an expected retirement date or age. We also include people who never expect to retire by studying an additional binary outcome variable for responding to the relevant survey questions with an expected retirement date of "never."

In our causal analysis, we must be more careful when addressing missing data concerns and defining our analysis samples. Here, we're interested in documenting the causal effects of various life events on retirement expectations. However, some life events we study can affect work status, thus influencing the likelihood that an observation is missing data on retirement expectations. Hence, ignoring the missing data issue will lead to biased estimates, because we'd effectively be conditioning on an endogenous outcome variable (work status).

Consider a concrete example. Suppose a health shock causes people to expect to retire earlier and to downgrade the probability that they will work past age 62. Further, suppose the health shock also causes people to drop out of the labor force (even if only temporarily). Then, when we use our event study design to compare outcomes just before and after the shock, our post-event sample will include only those with nonmissing observations, a selected group of people who didn't stop working because of the health shock. Our estimates based on this selected sample would thus be biased downward since they wouldn't detect the effects on expectations for those who drop out of the sample because of the shock.

To overcome this problem, we conduct our causal analysis using only the probability-based outcome variables and only survey waves 8 through 15. The probability-based questions are asked regardless of work status in these later waves. This approach thus crucially allows us to advance causal estimates of the shocks on expectations and to include in our analysis people who are both working and not working. Unfortunately, the expected retirement age variable is missing for nonworkers even in the later waves of the survey. This data limitation means we can't produce causal estimates for this outcome.

To summarize, nonrandom missing data means we must be careful when constructing analysis samples. Our descriptive analysis focuses on a selected sample of workers, allowing us to study people with well-defined retirement expectations. This approach highlights the general patterns of expectations for workers approaching retirement age. In contrast, our causal analysis focuses on probability-based outcomes and uses only later survey waves, avoiding selection issues that would generate biased estimates.

3. The evolution of retirement expectations as workers age

In the first phase of our analysis, we analyze the evolution of retirement expectations as workers age. Our descriptive analysis sample contains 45,143 observations on 18,572 individuals who are working either full time or part time, who are between ages 50 and 60 and who have nonmissing values for the probability-based outcomes and nonmissing values for the expected retirement age outcome. (We include people who say that they will never retire.) To provide more context for our sample restriction that focuses on workers, Figure 1 plots the fraction of people working by age. The underlying sample includes observations of people from all available birth cohorts in all survey waves 1 through 15. Each panel plots the fraction working either full time or part time. Panel (a) plots the fraction working for the full sample, and panel (b) plots the fraction working for two sets of birth cohorts because the HRS spans a relatively wide range of cohorts. The gray line corresponds to people born between 1931 and 1947 and the black line corresponds to people born between 1948 and 1971. Both graphs show that the fraction working hovers around 70% for people in their early 50s before declining for people in their mid-to-late 50s. The fraction working is just above 50% for people at 60, the oldest age we consider. While there is a notable decline, it occurs before the sizable retirement hazards at ages 62 and 65.

Table 1 reports summary statistics for this sample of workers. Since individuals differ in the number of survey waves in which they appear, the table reports means and standard deviations of variables from the first survey wave where an individual appears in the sample. Due to the structure of the HRS, survey cohorts are added to the data in 1992, 1998, 2004, 2010, and 2016. Panel A presents demographic information. The average year of first appearance in our sample is about 2004. The average age is 54. Our sample is 47% male and 70% white. The fraction married is 74%. Panel B presents information on individual-level earnings and household-level balances in IRAs to assess the economic circumstances of these older workers. Average earnings are about $45,000, and people have IRA balances of about $40,000, measured in 2010 dollars.

Panel C presents information on our outcome variables. On average, workers in our sample expect to retire at 64, but this average does not include workers who report they'll never retire. Therefore, we create an additional outcome variable that takes the value 1 if a worker responds to the expected retirement age question by saying they'll never retire and takes the value 0 if they respond with a specific retirement date. Notably, 37% report they expect to never retire. The average probability of working past age 62 is 51%, the average probability of working past 65 is 32%, and the average probability of working past 70 is 16%.

3.1. Trends in retirement expectations as workers age

We begin by analyzing age profiles of worker retirement expectations. Figure 2 documents how average outcomes evolve with age. The left-hand-side graphs, panels (a), (c), and (e), plot the age profiles of retirement expectations using the full sample, which includes people born in different birth cohorts. Panel (a) plots average expected retirement ages, panel (c) plots the average likelihood of reporting an expected retirement age of "never," and panel (e) plots the average probabilities of working past key ages. The plots reveal steady increases in the expected retirement age, the likelihood of expecting to never retire, and the probability of working past ages 62 and 65.

On average, workers in their early 50s expect to retire around age 64, but the average expected retirement age increases to 65 for workers in their late 50s. About 35% of workers at age 50 expect to never retire, and this fraction increases to just under 45% as workers approach age 60. The trend in the probability of working past age 62 is also meaningful. From age 50 to 60, the average likelihood of working past age 62 increases by 13 percentage points (from 48% to 61%), representing a 20% increase. Similarly, the average probability of working past 65 increases from 31% to 37%, a 6-percentage-point (19%) increase. In contrast, expectations about working past age 70 remain relatively constant as workers age. The average likelihood of working past age 70 hovers around 15% at all ages.

The right-hand-side graphs, panels (b), (d), and (f), show separate age profiles for different birth cohorts. The earliest cohort in our sample was born in 1931, whereas the latest cohort was born in 1971 (although only one person in the sample is from the 1971 birth cohort). Birth cohorts could have different trends in retirement expectations for many reasons. For example, they age through their 50s in different years, facing different labor market opportunities during their lives. They may also differ in the type of retirement savings scheme that they plan to use to finance consumption during retirement, with earlier birth cohorts being more likely to have defined benefit pension plans and later cohorts being more likely to have defined contribution retirement savings accounts.

Broadly speaking, while the level differences in the graphs indicate that earlier birth cohorts expect to retire at younger ages than later birth cohorts, these by-cohort graphs show similar increasing trends in retirement expectations across cohorts. For both the earlier and later cohorts, workers at older ages expect to retire later than workers at younger ages. However, panel (d) illustrates perhaps the most noticeable difference across cohorts. The earlier birth cohorts experience a sharper increase in the fraction of workers who expect to never retire.

These trends indicate that retirement expectations do not fluctuate substantially at older ages. Still, older workers tend to report later expected retirement ages and greater probabilities of working past ages associated with the Social Security and Medicare programs than workers in their early 50s. What explains these patterns? One possibility is that people update their expectations as they age, and that the aging process, as well as realizations or the lack of realizations of life events between ages 50 and 60, leads people to expect to work longer. Another possibility is that selection dynamics play a role. By studying workers, we exclude from our sample people who are retired, disabled, unemployed or not in the labor force. Some of the increasing trends could thus be explained by the changing composition of the sample, even if individuals do not update their expectations with age. For example, it could be that people who expect to retire relatively early drop out of the sample during the age window that we study and that the workers who remain in the sample at older ages are those who at all ages expect to retire later.

How important could selection be for explaining the trends? Recall from Figure 1 that the fraction of people working is roughly constant until they reach their mid-50s, at which point there is a decline. This graph thus supports the idea that the selection dynamics may be the most relevant at the oldest ages of our analysis window and less relevant at younger ages. To better understand the potential role of selection, we document trends for a subsample of consistent workers. Specifically, Figure 3 plots trends in the expected retirement age and the probability of working past key ages for people who we observe in multiple waves and are working (either full time or part time) during each wave they are observed. By focusing on this subsample of consistent workers, we avoid studying people within the main sample who retire or stop working during our observation window.

Trends in the outcome variables are less stark for this subsample, but there is still evidence of an upward trend in the expected retirement age and the probabilities of working past ages 62 and 65. The average expected retirement age of consistent workers in their early 50s is 64.5, rising to around 65 by their mid-50s; at age 60, the average expected retirement age is 65.4. Similarly, the average probability of working past age 62 increases by about 9 percentage points, from 54% to 63%, and the average probability of working past age 65 increases from about 35% to about 39%.

Overall, the patterns we document are consistent with the idea that workers update their expectations toward working longer as they age through their 50s. A simple explanation is that the individual faces many possible life events that could induce retirement (e.g., health shocks) and updates their probability of working as each additional age passes without experiencing a retirement-inducing shock. In such an environment, the conditional probability of working longer rises as individuals work each year. Beyond understanding these trends, it is also of interest to understand how various life events influence retirement expectations. In Section 4, we identify the causal effects of life events on retirement expectations, but first we turn to documenting differences in levels and trends of retirement expectations across groups.

3.2. Differences in expectations across groups

Using a simple regression framework, we analyze heterogeneity to explore how demographic, economic, and health factors relate to expectations and their evolution with age. For each dimension of heterogeneity, we split our analysis sample into two subsamples and regress the retirement expectation outcome variables on age. That is, we estimate

$$y_{it} = \alpha + \beta age_{it} + \varepsilon_{it}, \qquad (1)$$

where y_{it} is the outcome variable (e.g., expected retirement age) for individual i in survey wage t, age_{it} is the age of the individual, and ε_{it} is an error term. The coefficient of interest in this simple regression is β, which captures the change in retirement expectations associated with a one-year increase in age.

Table 2 presents the results for the expected retirement age outcome variable. Tables 3, 4 and 5 present results for the probability of working past ages 62, 65 and 70, respectively. We construct each table as follows. Panel A corresponds to the full sample, whereas panels B through H indicate the various dimensions of heterogeneity that we explore. For each group we analyze, column (1) displays the average expected retirement age in levels at age 50, column (2) displays the estimate of β, the slope parameter in the simple regression model, and column (3) displays the number of observations. The full sample estimates in panel A of each table capture the key information from the graphical analysis above in regression form. For example, Table 2 shows that, on average, people at age 50 expect to retire just before turning 64, and aging one additional year is associated with an increase in the expected retirement age by about 0.16 years, or 1.92 months.

Panels B, C and D split the sample according to demographic characteristics. The results on gender show that females expect to retire earlier than males; they report earlier expected retirement ages on average (63.66 compared with 64.31) and lower likelihoods of working past ages 62, 65 and 70. The slope parameters for expected retirement ages indicate that the increase in expected retirement dates associated with aging one additional year for females (0.170 years, or about two months) is somewhat larger than that for males (0.143, or about 1.7 months). However, the slope parameters for the probability-based outcomes are more similar. The results on race reveal similar slope parameters but differences in levels; white workers at age 50 expect to retire more than three months later than nonwhite workers and report being 10 percentage points more likely to work past age 62. Finally, results on marital status indicate that nonmarried workers expect to retire later than their married counterparts at age 50 and experience smaller increases in expected retirement ages and probabilities of working later with age. The difference in levels for married compared to nonmarried workers is sizable; the average expected retirement age for married workers is 63.72, whereas the average for nonmarried workers is 64.51, translating into over nine additional expected months of work.

Panels E and F split the sample according to health. In panel E of each table, we divide the sample based on self-reported health status, which can be either excellent, very good, good, fair or poor. We define good health status as those who report that they are in good, very good or excellent health. In panel F of each table, we divide the sample based on reports of having ever been diagnosed with one of several health conditions. For this categorization, we include four major health conditions: heart disease or heart attack, stroke, cancer and lung disease. The results on health indicate that better health is associated with expecting to work longer at age 50, and the slope parameters tend to be larger for workers with better health measures, consistent with the idea that those who experience health declines are likely to retire earlier. Again, it is important to emphasize that we are studying a sample of workers. The observation counts are substantially lower for those with worse measures of health. It's likely that some workers who experience adverse health shocks or worsening health status stop working and thus drop out of the sample, but our analysis here focuses on those who remain working.

Finally, panels G and H split the sample according to economic characteristics. In panel G of each table, we study high- and low-income workers. We first compute median earnings by age for all workers in our full sample. We then define a person-wave observation as high income if the worker has above-median nominal earnings for their age group in that wave. We define a person-wave observation as low income if the worker has below-median nominal earnings for their age group in that wave. In panel H of each table, we study high-wealth and low-wealth workers by analogously comparing the (nonhousing financial) wealth of each person-wave observation with the by-age medians.

The results on economic characteristics are perhaps more nuanced than those for health. First, consider income. When comparing means at age 50, we see that high-income workers expect to retire earlier, report greater probabilities of working past 62 and 65, and report smaller probabilities of working past age 70 than their low-income counterparts. When comparing the slope parameters, high-income workers tend to experience larger increases in their expected retirement ages and their probabilities of working past key ages. These patterns likely reflect that higher-income workers may choose to buy more leisure in the form of earlier retirement and have a larger opportunity cost of retirement, seen as foregone earnings. Next, consider wealth. Here, low-wealth workers expect to retire later than high-wealth workers. However, the slope parameters indicate that high-wealth workers experience greater changes in expectations as they age. For instance, Table 2 shows that high-wealth people at age 50 expect to retire 1.3 years earlier than low-wealth people. However, the slope parameter indicates a 2.7-month increase in expected retirement age associated with each additional year of aging. In contrast, the slope parameter for low-wealth workers indicates a smaller 1.2-month increase in expected retirement age with each additional year.

4. The effects of life events on retirement expectations

In the second phase of our analysis, we analyze the causal effects of various life events, or "shocks," on retirement expectations. We consider three broad categories of shocks: adverse health events (such as sudden declines in health status), economic events (such as significant changes in income) and family events (such as changes in marital status). We use an event study approach to identify causal effects for each specific shock within a category.

To avoid concerns with nonrandom missing data discussed in Section 2, we first restrict our data to survey waves 8 through 15. We then (i) define a given shock of interest, (ii) focus our analysis only on people that we observe experiencing the shock, and (iii) exploit the timing of shocks by comparing the evolution of retirement expectation outcomes before and after the shock occurs.

One advantage of this event study approach is that we don't rely on comparisons of people who experience the shocks of interest to those who don't. Instead, we focus on groups where everyone experiences the same shock but at different times. This approach avoids issues that might present problems for identification, like the concern that people who experience a given shock may differ in unobservable ways from people who do not experience the shock.

To implement our event study analysis, we use two regression models that follow Dobkin et al. (2018), who use administrative and HRS data to study the effects of hospitalizations on financial outcomes. First, we use a nonparametric event study framework to analyze the patterns of outcome variables around the timing of the shocks. Second, we use a parametric event study framework to quantify key magnitudes of interest.[4]

4 Other researchers also follow the useful framework laid out by Dobkin et al. (2018). For instance, Mommaerts et al. (2020) borrow their event study approach, as we do, to study the effects of hospitalization events on financial and economic outcomes across several countries.

4.1 Nonparametric event study framework

We begin our investigation of each shock by estimating equations of the following form:

$$y_{ict} = \alpha + \sum_{\tau=-3}^{-2} \delta_\tau + \sum_{\tau=0}^{1} \delta_\tau + \lambda_{ct} + \varepsilon_{ict}, \qquad (2)$$

where y_{ict} is an outcome variable (such as the probability of working past age 62) for individual i in HRS sample cohort c, during survey way t, λ_{ct} is a cohort-by-wave fixed effect, and δ_τ is a coefficient on an event time indicator for a survey wave relative to the wave during which the shock occurs. As in Dobkin et al. (2018), we include HRS cohort-by-wave fixed effects instead of simply survey wave fixed effects to account for changes in the composition of the HRS sample as cohorts are added to the survey. We analyze an event window spanning three survey waves before with one after the event of interest. The δ_τs are the parameters of interest. They capture the average difference in the outcome at event time τ relative to the omitted period, τ = –1, the survey wave before the event occurs.

The identification assumption underlying regression model (2) is that conditional on experiencing a given shock, the timing of the shock is uncorrelated with the outcome variables. For each shock that we study, the threat to identification is thus that there is some unobserved factor that influences both the timing of the shock and retirement expectations. One key concern might relate to changes in job characteristics or employment opportunities. For instance, a threat to our design is if increases in physical labor at work or increases in stress about future employment (i) cause people to update their retirement expectations but also (ii) cause people to be more likely to experience the shocks that we study, such as adverse health events.

Naturally, the plausibility of the identification assumption may well vary across the different shocks. For example, consider the event of having a heart attack or a stroke. The timing of these sudden adverse health events may be unpredictable. Indeed, several papers exploit the unpredictable nature of these emergencies as a part of their research designs (see, e.g., Chandra & Staiger, 2007; Doyle, 2011; Fadlon & Nielsen, 2021). When we study these shocks, this feature generates relatively high ex ante confidence in the identification assumption. Alternately, consider divorce, a family-related shock that we also study. People choose whether and when to divorce. Given the nature of this choice, we have lower ex ante confidence in the identification assumption for this shock. In general, we can provide an assessment of the validity of the identifying assumption by analyzing the estimated δ_τs for τ < 0 for each of the shocks that we study. The patterns of these "preperiod" estimates demonstrate if outcomes were trending prior to the shock of interest.

While we often find little to no evidence of pretrends in outcomes, there is evidence of preperiod trends in some cases. To be consistent and account for the possibility of the underlying trends when analyzing the effects of a life event, we move away from our nonparametric event study framework and instead use a parametric framework.

4.2 Parametric event study framework

To quantify magnitudes, we estimate equations of the form:

$$y_{ict} = \alpha + \beta\tau + \sum_{\tau=0}^{1} \delta_\tau + \lambda_{ct} + \varepsilon_{ict}. \qquad (3)$$

The key difference here is the inclusion of τ, a linear time trend. The coefficient β corresponds to the preshock linear trend in the outcome variables. The parameters of interest are still the δ_τs, which now capture the average difference in the outcome at event time τ compared to its linear pretrend. The identification assumption underlying regression model (3) is that, conditional on experiencing a given shock, the timing of the shock is uncorrelated with deviations in the outcome variable from its preperiod linear trend.

When estimating the parametric event study regressions, we focus on δ_0, the effect of the shock on the outcome in the survey wave of the shock. We focus on this point estimate instead of δ_1, which corresponds to the subsequent survey wave, because our definition of each shock—discussed in more detail below—requires that we observe the individual in the data during the wave of the shock (and the two previous waves) but not afterward.

4.3 The effect of health events on retirement expectations

We separately study ten different types of health events: (i) sudden declines in self-reported health status, (ii) hospitalizations, (iii) heart attacks, (iv) strokes, (v) cancer diagnoses, (vi) lung disease diagnoses, (vii) arthritis diagnoses, (viii) diabetes diagnoses, (ix) high blood pressure diagnoses and (x) diagnoses of psychiatric problems.[5]

5 The survey question for heart attack also refers to coronary heart disease, angina, congestive heart failure or other heart problems. The question for cancer refers to cancer or a malignant tumor of any kind except minor skin cancer. The question for lung disease refers to chronic lung disease, such as chronic bronchitis or emphysema, but not asthma. The question for arthritis refers to arthritis or rheumatism. The question for diabetes refers to diabetes or high blood sugar. The question for high blood pressure refers to high blood pressure or hypertension. The question for psychiatric problems also refers to any emotional or nervous problems.

How might we expect expectations to change? On one hand, some adverse health events might lead workers to update their expectations toward retiring earlier. Existing research shows across contexts that many health shocks and disabilities lead to declines in labor-supply measures, such as earnings or employment (e.g., Gertler & Gruber, 2002; Dobkin et al., 2018; Meyer & Mok, 2019). If health shocks lead to decreased work capacity, then they can cause people to be less likely to work at older ages. On the other hand, some adverse health events might lead workers to update their expectations toward retiring later. This could be the case if a health shock creates increased medical expenses, and people respond by planning to work longer or increase their labor supply later in life, perhaps even years after the shock. In that case, adverse health events might cause people to be more likely to work at older ages. Ultimately, the direction and magnitude of the potential effects are empirical questions.

We begin with changes in health status. For each person, we define a declining health status event as occurring in survey wave w if the person reports being in fair or poor health in survey wave w and being in excellent, very good or good health in both survey waves $w-1$ and $w-2$. For this health event and every other event we study, we consider only the first shock if a person experiences multiple shocks during our sample period.

Figure 4 displays the nonparametric event study results for declines in health status. Each graph corresponds to a different outcome variable and plots the δ_τ coefficients from estimating equation (2). For each graph, the pattern of the preshock period point estimates indicates a lack of significant trends in the probabilities of working past key ages before the decline in health status. Panels (a) and (b) show a clear decline in the average probability of working past ages 62 and 65, respectively. Declines in health status thus cause workers to update their expectations toward retiring earlier. In contrast, panel (c) shows no evidence of a change in the probability of working past age 70. This lack of decline could be because the average probability of working past age 70 was relatively low (12%) before the shock. Moreover, our analysis of this outcome variable is limited due to a small sample since it only exists for waves 11 onward.[6]

Panel A of Table 6 displays the parametric event study results that summarize the effects of health status declines on retirement expectations by comparing the evolution of the outcome variables to their preperiod trends (which were relatively flat in this case). The table reports the δ_0 coefficients from estimating equation (3) for each outcome variable, the mean of the outcome variable in the period before the shock, the estimated decline in percent terms (the point estimate divided by the mean), the number of individuals who experienced the shock, and the total number of observations underlying each regression.

On average, a sudden decline in health status leads to a statistically significant 4-percentage-point decline in the self-assessed probability of working past age 62 and a statistically significant 4.5-percentage-point decline in the probability of working past 65. These declines translate to a meaningful fall in the likelihood of working by 8.9% and 14.1%, respectively, compared to the baseline means. In short, declines in health status lead people to significantly reduce their expectations about working at older ages in the future.

Next, we study hospitalizations. For each person, we define a hospitalization event as occurring in survey wave w if the person reports an overnight hospitalization since his or her last interview in wave w and no such hospitalizations in both wave $w-1$ and $w-2$. We view this definition as making it more likely that we are studying unanticipated hospitalization events.

Figure 5 presents the nonparametric event study results and panel B of Table 6 presents the coefficients from the parametric event study regressions. Unlike declining health status events, we find no evidence that hospitalizations impact expectations about working at later ages. The point estimates are smaller in magnitude than for health status declines and are not statistically significant.

Finally, we study eight different diagnosis events: heart attacks, strokes, cancer, lung disease, arthritis, diabetes, high blood pressure and psychiatric problems. For each one of these health conditions, we define people as experiencing a shock in wave w if they report that a doctor has ever told them that they have the condition in wave w, and if they report that a doctor has never told them that they have the condition in both wave $w-1$ and wave $w-2$.

Figures 6 through 13 and Tables 7 and 8 report the event study results. We find that some health conditions impact expectations about future work, whereas others appear to have little to no effect. We find no statistically significant evidence that heart attacks, diabetes, high blood pressure

6 The implication of this limitation is that some person-waves observations become unusable because of the missing data. For instance, consider a person who experiences the health decline event during wave 10. We would observe the probability of working past age 70 for this person in the survey wave after their shock (wave 11), but we would not observe their probability in the survey wave of their shock, nor would we observe their probabilities in the survey waves before their shock.

or psychiatric problems impact the probability of working past age 62, 65, or 70. Note that while graphs (a) and (b) in Figure 13 for psychiatric problems show evidence of a post-shock decline in the probabilities of working past ages 62 and 65, the preshock trends for these outcomes in this case are more noticeable, and the parametric event study estimates that account for these preperiod trends reveal no statistically significant evidence of a decline. We find some evidence that strokes affect expectations about working past age 65 (see Figure 7 and panel B of Table 7), but the sample sizes underlying the regressions for stroke diagnoses are small.

We find the clearest evidence of declines in the probabilities of working later after cancer, lung disease and arthritis diagnoses. Figure 8 and panel C of Table 7 present the results for cancer. Graphs (a) and (b) show some evidence of preperiod trends, but also clear declines in the self-assessed probabilities of working past ages 62 and 65 after the shock, especially in the survey wave during which the diagnosis is initially reported. The parametric event study results indicate that the probability of working past age 62 declines by 7.2 percentage points, which translates to a 16.3% decline when compared with the baseline mean, and that the probability of working past age 65 declines by 5.8 percentage points, an even larger 19.3% decline. Cancer diagnosis events cause people to meaningfully downgrade their probability of working at later ages.

Figure 9 and panel D of Table 7 present the results for lung disease. Like cancer diagnoses, lung disease diagnoses lead to sizable decreases in self-assessed likelihoods of working later. The graphs reveal sharp decreases in the probabilities of working past ages 62 and 65, and the corresponding point estimates in the table indicate declines of 7.4 percentage points (26%) and 6.5 percentage points (34%), respectively. Figure 10 and panel A of Table 8 present the results for arthritis. The patterns of the point estimates in graphs (a) and (b) reveal some evidence of decreases in the key outcome variables. The parametric event study point estimates indicate a statistically significant 4.7 percentage point (10.2%) decrease in the probability of working past age 62, indicating that arthritis diagnoses can cause people to expect to retire earlier than they otherwise would have.

Overall, our results on health show that the effects of adverse shocks can be nuanced; some events meaningfully impact retirement expectations whereas others do not. Of course, different shocks will likely impact current work capacity, future work capacity, medical expenses, savings and other factors differently. Our findings provide evidence on which shocks impact expectations. We find that cancer, lung disease, and changes in health status can lead to major declines in the likelihood of working in the future. In contrast, shocks like diabetes and high blood pressure appear less likely to influence expectations about future retirements.

4.4 The effect of economic events on retirement expectations

Next, we study the effects of economic events. We consider two economic events related to employment. First, we study unemployment as an adverse shock. We define people experiencing unemployment if they report being unemployed as their labor force status in wave w and if they report being employed either full time or part time in wave $w-1$ and $w-2$. Next, we study earnings increases as a positive shock. We define people experiencing an earnings-increase event if their nominal earnings in wave w are between 25% and 100% greater than their nominal earnings in wave $w-1$ and, to be consistent with our definitions of other shocks that require two look-back waves, if we also observe them in the data in wave $w-2$.

Like the health shocks that we study, these economic shocks could shift expectations about future retirement in both directions. Consider unemployment. On the one hand, existing research indicates that older displaced workers might have trouble returning to work and that labor market conditions affect retirement transitions (e.g., Chan & Stevens, 2001; Coile & Levine, 2007), which might lead unemployed individuals to downgrade their probabilities of working at future ages. On the other hand, with the decline in earnings that accompanies unemployment, workers could respond by choosing to work longer after they become reemployed. Now consider an increase in earnings. If these increases are due to changes in wages, then there are both income and substitution effects. The income effect would lead workers to consume more leisure and less work, so we might expect a decrease in the probability of working at later ages. However, the substitution effect would lead workers to consume less leisure and work more, so we might expect an increase in the probability of working at later ages if the increase in wages is permanent. Importantly, we note that while some people could respond more strongly to income effects and others to substitution effects, we are unable to separate these types of responses and our design estimates an average effect for the full sample of workers experiencing an earnings shock.

Figure 14 and panel A of Table 9 present the results for unemployment shocks. The graphs reveal a modest upward preshock trend that seems to continue through the post-shock periods, and the parametric event study results indicate no statistically significant effect on the probability-based outcome variables. We find no evidence that the unemployment events we study, when people are in their 50s, impact their self-assessed probabilities of working at later ages.

Figure 15 and panel B of Table 9 present the results for earnings increases. Graphs (a) and (b) reveal some evidence of a decrease in the likelihood of working in the future, although the graphical evidence is not as strong as that for some of the health events. The parametric point estimates, which measure deviations from the preperiod trends (and there does appear to be a modest increasing trend in graph (b), for instance) indicate decreases in the probabilities of working past ages 62 and 65 that are statistically significant at the 10% level. The estimates are smaller in magnitude than those for the health events that lead to a decreased likelihood of working at later ages. Still, they amount to meaningful declines of 4.9% and 7.7%, compared to the baseline means.

4.5 The effect of family events on retirement expectations

Finally, we study the effects of family events. We study two distinct events: the birth of a grandchild and divorce or separation. We use the number of grandchildren to define grandchild birth events. We define people experiencing a grandchild birth if they report a number of grandchildren in wave w that is greater than the number of grandchildren they report in wave $w-1$ and if we also observe the person in wave $w-2$. We define people experiencing a divorce event if they report being divorced or separated in wave w and if they report being either married or partnered in both wave $w-1$ and $w-2$.

We find no evidence that having an additional grandchild affects the probability of working later. Figure 16 shows no graphical evidence of a change in expectations around the timing of the grandchild births, and panel A of Table 10 displays point estimates that are relatively small and not statistically significant. We do find some weak evidence that a divorce or separation may increase the likelihood of working at later ages. Figure 17 presents the graphical evidence. The patterns of the point estimates suggest increases in the probabilities of working past ages 62, 65 and perhaps 70, although there is some evidence of a pretrend. The leading point estimates in panel B of Table 10, which focus on the wave of the reported shock, are not statistically significant. However, individuals may also have some advanced knowledge of their divorce or separation that could be consistent with them revising their probability of working later in advance of the measured shock. Overall, we conclude that there is little evidence that the family events we study impact retirement expectations.

5. Conclusion

In this paper, we advance our understanding of retirement timing by analyzing expectations data. We first conduct a descriptive analysis that documents how retirement expectations change as workers age and assesses how demographic, financial and health factors are associated with expectations. We then use an event study framework to estimate the causal effects of common life events on retirement expectations. Our findings underscore the important role of health in shaping retirement timing. We find large declines in the probabilities of working later for people who experience declines in health status and for people diagnosed with cancer, lung disease and arthritis.

Our study has implications for assessing individuals' financial security and financial wellness as they approach retirement. The timing of retirement is a critical choice that plays a key role in determining financial security in retirement, but it is subject to considerable uncertainty. By providing new evidence on retirement expectations and how life events shape them, we help clarify the factors contributing to this uncertainty. Understanding how retirement expectations evolve is paramount to understanding how people anticipate, experience and cope with events that ultimately affect retirement timing.

Our findings also inform the design of retirement savings schemes, since the expected retirement date often determines optimal strategies within a savings plan. However, evidence shows that retirement savings often evolve passively (e.g., Madrian & Shea, 2001; Chetty et al., 2014) and that inertia within retirement savings can shape responses to pension reforms that influence retirement timing (Garcia-Miralles & Leganza, forthcoming). This passivity raises concerns that workers may not optimally update their savings in response to events that change their expected retirement date. Studying the evolution of expectations is an essential first step in understanding whether individuals adjust savings appropriately and how financially secure they will be in retirement.

References

Angrist, J.D., and Pischke, J.S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Ayyagari, P. (2019). Health insurance and early retirement plans: Evidence from the affordable care act. *American Journal of Health Economics*, 5(4), 533–560.

Backhaus, A., and Barslund, M. (2021). The effect of grandchildren on grandparental labor supply: evidence from Europe. *European Economic Review*, 137, 103817.

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

Bugliari, D., Carroll, J., Hayden, O., Hayes, J., Hurd, M.D., Lee, S., Main, R., McCullough, C.M., Meijer, E., Pantoja, P., and Rohwedder, S. (2023). RAND HRS Longitudinal File 2020 (V1) Documentation. *RAND Center for the Study of Aging*.

Bugliari, D., Carroll, J., Hurd, M.D., Lee, S., Main, R., and Rohwedder, S. (2023). RAND HRS Family Data File 2018 (V2) Documentation. *RAND Center for the Study of Aging*.

Caliendo, F.N., Casanova, M., Gorry, A., and Slavov, S. (2023). Retirement timing uncertainty: Empirical evidence and quantitative evaluation. *Review of Economic Dynamics*.

Chan, S., and Stevens, A.H. (1999). Employment and retirement following a late-career job loss. *American Economic Review*, 89(2), 211–216.

Chan, S., and Huff Stevens, A. (2001). Job loss and employment patterns of older workers. *Journal of Labor Economics*, 19(2), 484–521.

Chandra, A., and Staiger, D.O. (2007). Productivity spillovers in health care: evidence from the treatment of heart attacks. *Journal of Political Economy*, 115(1), 103–140.

Chetty, R., Friedman, J.N., Leth-Petersen, S., Nielsen, T.H., and Olsen, T. (2014). Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark. *Quarterly Journal of Economics*, 129(3), 1141–1219.

Coe, N.B., von Gaudecker, H.M., Lindeboom, M., and Maurer, J. (2012). The effect of retirement on cognitive functioning. *Health Economics*, 21(8), 913–927.

Coe, N.B., and Zamarro, G. (2011). Retirement effects on health in Europe. *Journal of Health Economics*, 30(1), 77–86.

Coile, C. (2004). Health shocks and couples' labor supply decisions. Working Paper 10810, National Bureau of Economic Research.

Coile, C.C., and Levine, P.B. (2007). Labor market shocks and retirement: Do government programs matter? *Journal of Public Economics*, 91(10), 1902–1919.

Disney, R., Emmerson, C., and Wakefield, M. (2006). Ill health and retirement in Britain: A panel data-based analysis. *Journal of Health Economics*, 25(4), 621–649.

Dobkin, C., Finkelstein, A., Kluender, R., and Notowidigdo, M.J. (2018). The economic consequences of hospital admissions. *American Economic Review*, 108(2), 308–352.

Doyle Jr, J.J. (2011). Returns to local-area health care spending: evidence from health shocks to patients far from home. *American Economic Journal: Applied Economics*, 3(3), 221–243.

Eibich, P. (2015). Understanding the effect of retirement on health: Mechanisms and heterogeneity. *Journal of Health Economics*, 43, 1–12.

Fadlon, I., and Nielsen, T.H. (2021). Family labor supply responses to severe health shocks: Evidence from Danish administrative records. *American Economic Journal: Applied Economics*, 13(3), 1–30.

Frimmel, W., Halla, M., Schmidpeter, B., and Winter-Ebmer, R. (2022). Grandmothers' labor supply. *Journal of Human Resources*, 57(5), 1645–1689.

García-Gómez, P., Van Kippersluis, H., O'Donnell, O., and Van Doorslaer, E. (2013). Long-term and spillover effects of health shocks on employment and income. *Journal of Human Resources*, 48(4), 873–909.

García-Miralles, E., and Leganza, J.M. (forthcoming). Public Pensions and Private Savings. *American Economic Journal: Economic Policy*.

Gertler, P., and Gruber, J. (2002). Insuring consumption against illness. *American Economic Review*, 92(1), 51–70.

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

Gørtz, M., Sander, S., and Sevilla, A. (2023). Does the Child Penalty Strike Twice, and If So Why? IZA Discussion Paper No. 16557, IZA Institute of Labor Economics.

Gorry, A., Gorry, D., and Slavov, S.N. (2018). Does retirement improve health and life satisfaction? *Health Economics*, 27(12), 2067–2086.

Haider, S.J., and Stephens Jr., M., (2007). Is there a retirement-consumption puzzle? Evidence using subjective retirement expectations. *The Review of Economics and Statistics*, 89(2), 247–264.

Hudomiet, P., Hurd, M.D., Parker, A.M., and Rohwedder, S. (2021). The effects of job characteristics on retirement. *Journal of Pension Economics & Finance*, 20(3), 357–373.

Karademir, S., Laliberté, J.W., and Staubli, S. (2023). The Multigenerational Impact of Children and Childcare Policies. IZA Discussion Paper No. 15894, IZA Institute of Labor Economics.

Madrian, B.C., and Shea, D.F. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics*, 116(4), 1149–1187.

McClellan, M.B. (1998). Health events, health insurance, and labor supply: Evidence from the health and retirement survey. In David A. Wise (Ed.), *Frontiers in the Economics of Aging* (pp. 301–350). University of Chicago Press.

McFall, B.H. (2011). Crash and wait? The impact of the Great Recession on the retirement plans of older Americans. *American Economic Review*, 101(3), 40–44.

Meyer, B.D.. and Mok, W.K. (2019). Disability, earnings, income and consumption. *Journal of Public Economics*, 171, 51–69.

Munnell, A.H., Sanzenbacher, G.T., and Rutledge, M.S. (2018). What causes workers to retire before they plan? *The Journal of Retirement*, 6(2), 35–52.

Nielsen, N.F. (2019). Sick of retirement? *Journal of Health Economics*, 65, 133–152.

Rupert, P., and Zanella, G. (2018). Grandchildren and their grandparents' labor supply. *Journal of Public Economics*, 159, 89–103.

Shoven, J.B., and Walton, D.B. (2021). An analysis of the performance of target date funds. *The Journal of Retirement*, 8(4), 43–65.

FIGURE 1. TRENDS IN THE LIKELIHOOD OF WORKING

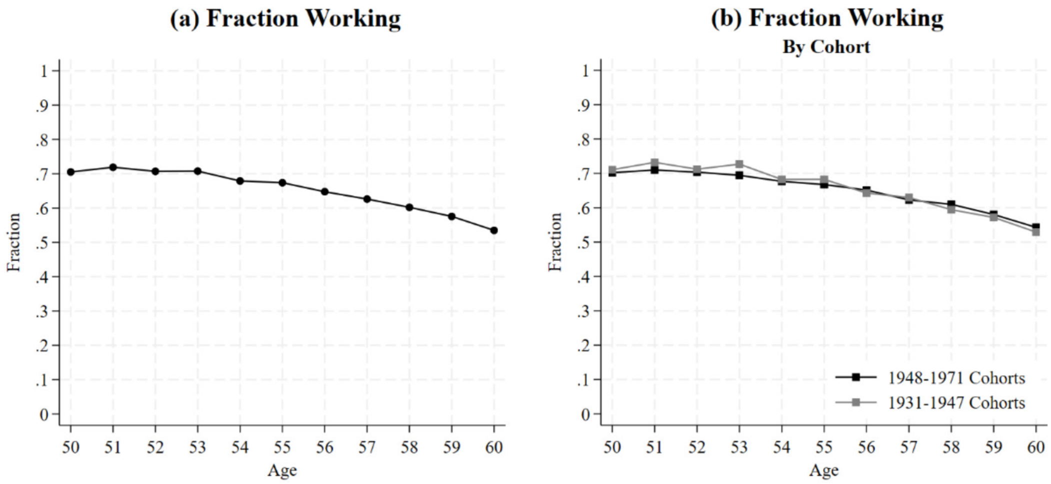

Notes: This figure plots the fraction of people working (either full time or part time) by age. Panel (a) plots the fraction working for the full sample of people in waves 1 through 15 of the Health and Retirement Study (HRS) between ages 50 and 60. Panel (b) plots the fraction working for early and late birth cohorts.

FIGURE 2. TRENDS IN RETIREMENT EXPECTATIONS

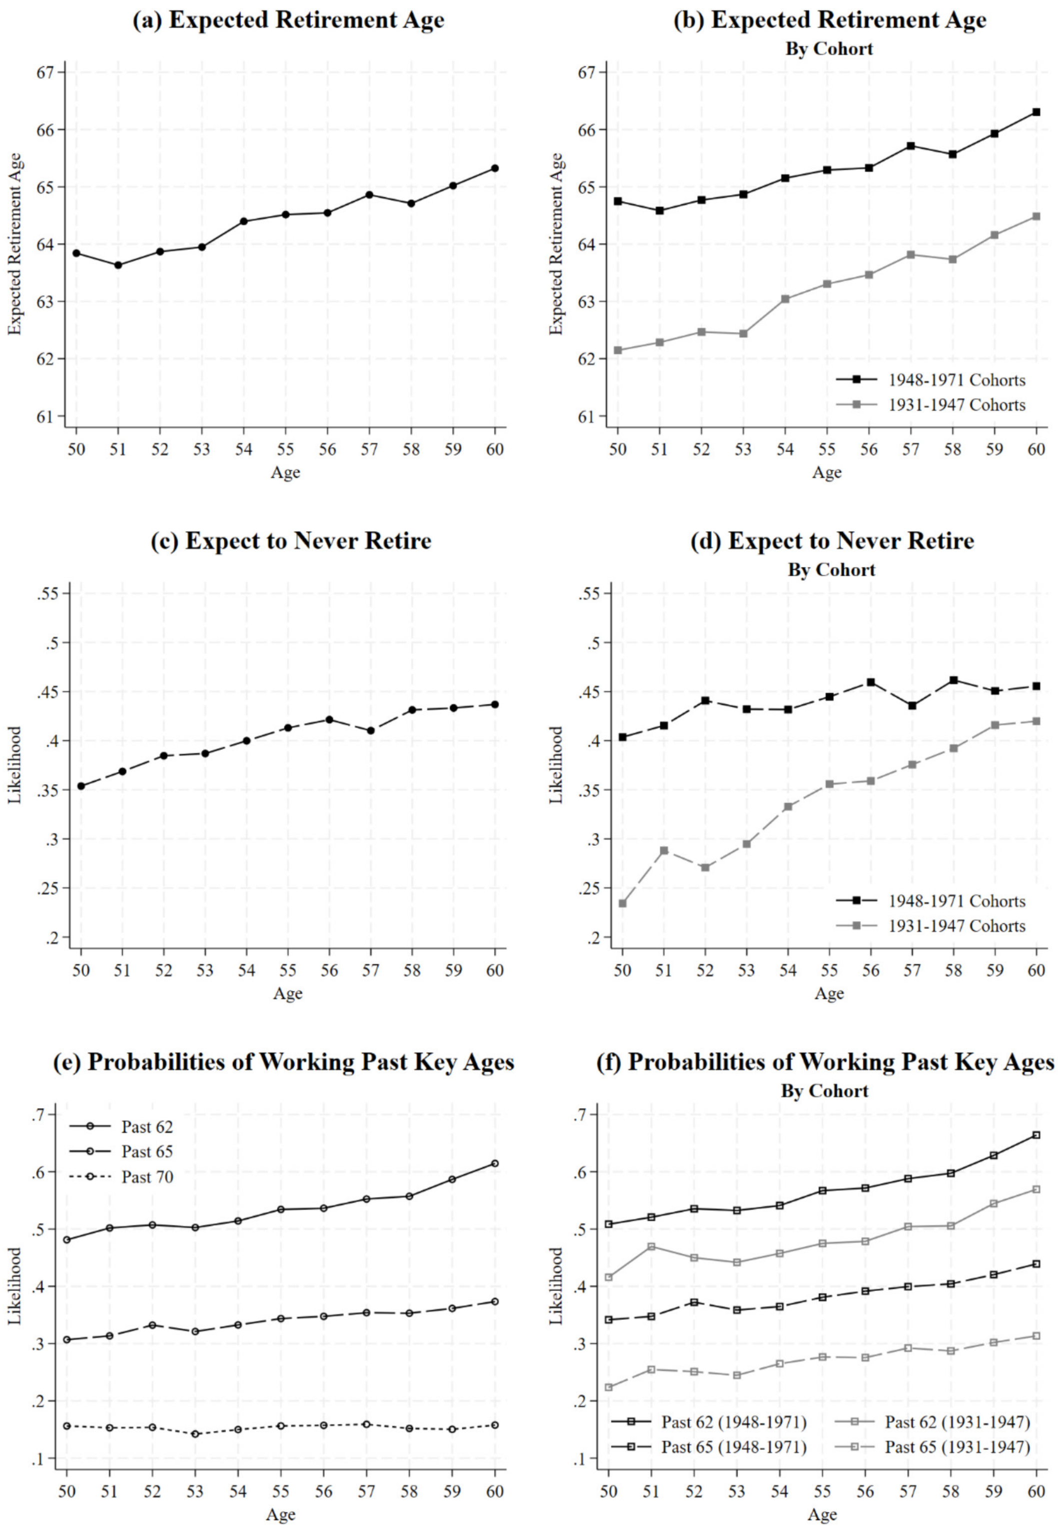

Notes: This figure plots the evolution of retirement expectations. Each panel plots averages of an outcome variable against age. Panels (a), (c) and (e) use the full analysis sample. Panels (b), (d) and (f) split the sample by birth cohort.

FIGURE 3. TRENDS IN RETIREMENT EXPECTATIONS FOR CONSISTENT WORKERS

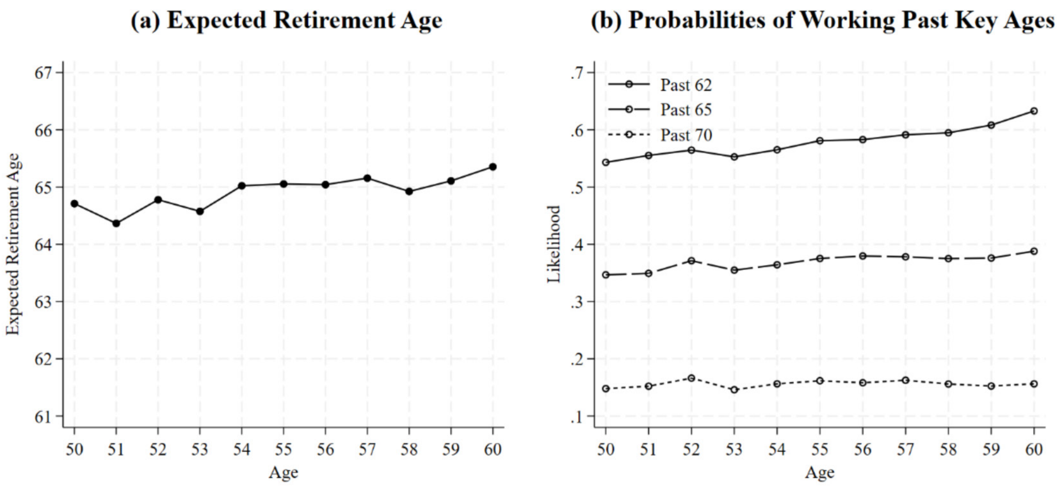

Notes: This figure illustrates how retirement expectations evolve for a subsample of consistent workers, defined as people we observe in multiple waves and who are working during each wave they are observed. Each panel plots averages of outcome variables against age.

FIGURE 4. EFFECTS OF CHANGES IN HEALTH STATUS ON RETIREMENT EXPECTATIONS

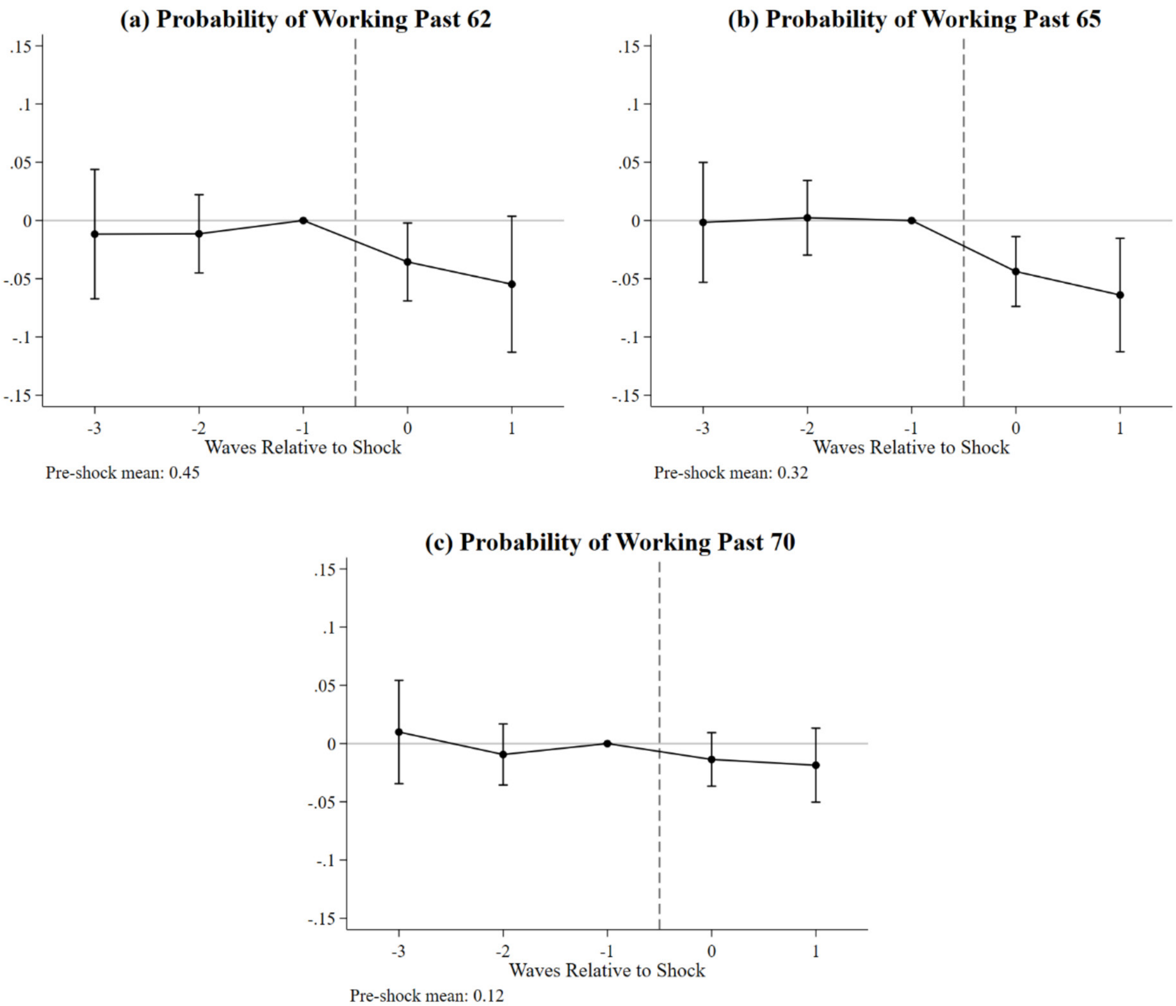

Notes: This figure presents the nonparametric event study estimates for changes in health status. Each graph plots point estimates and confidence intervals for the δ_τ coefficients in equation (2).

FIGURE 5. EFFECTS OF HOSPITALIZATION ON RETIREMENT EXPECTATIONS

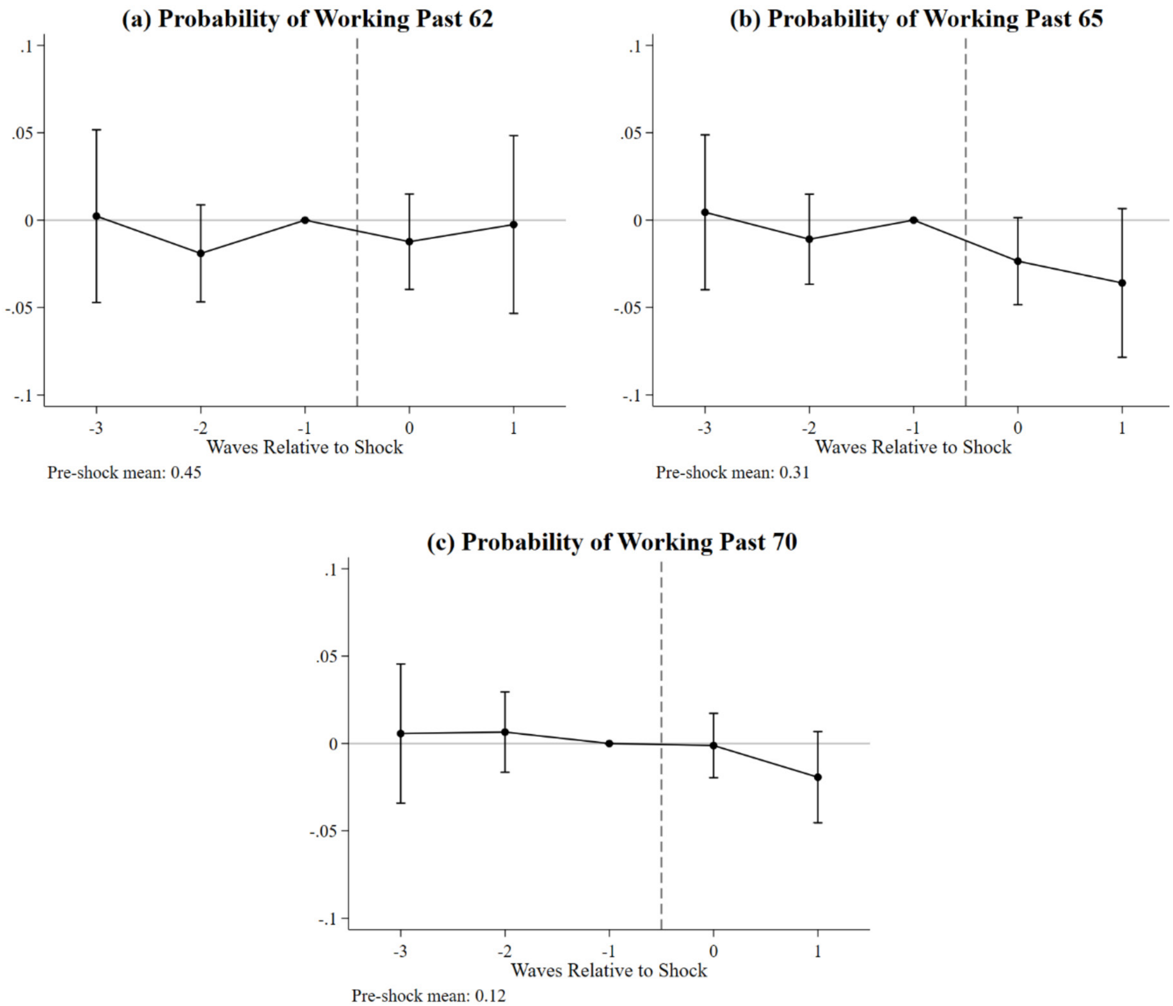

Notes: This figure presents the nonparametric event study estimates for hospitalization events. Each graph plots point estimates and confidence intervals for the δ_{τ} coefficients in equation (2).

FIGURE 6. EFFECTS OF HEART ATTACKS ON RETIREMENT EXPECTATIONS

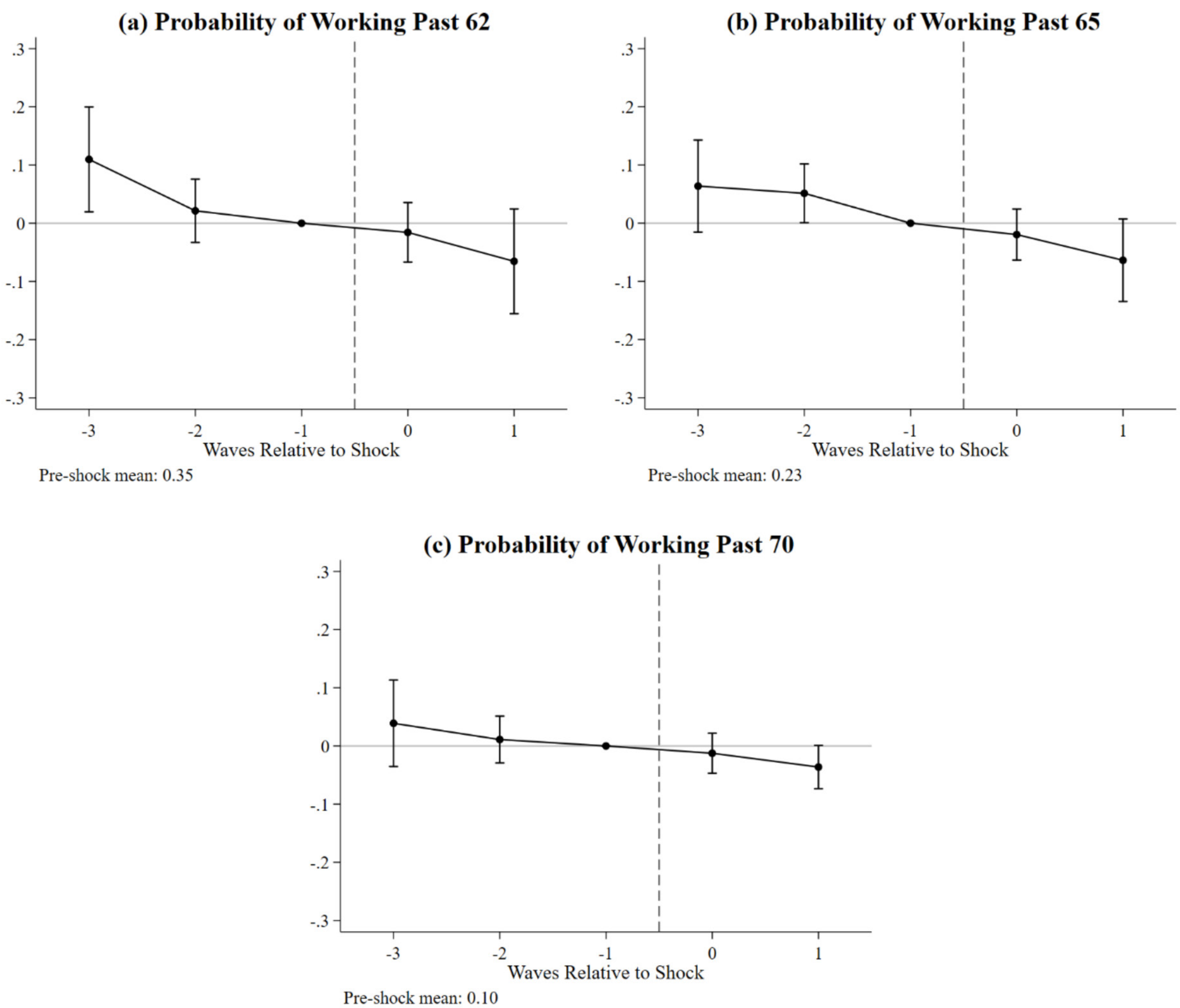

Notes: This figure presents the nonparametric event study estimates for heart attacks. Each graph plots point estimates and confidence intervals for the δ_τ coefficients in equation (2).

FIGURE 7. EFFECTS OF STROKES ON RETIREMENT EXPECTATIONS

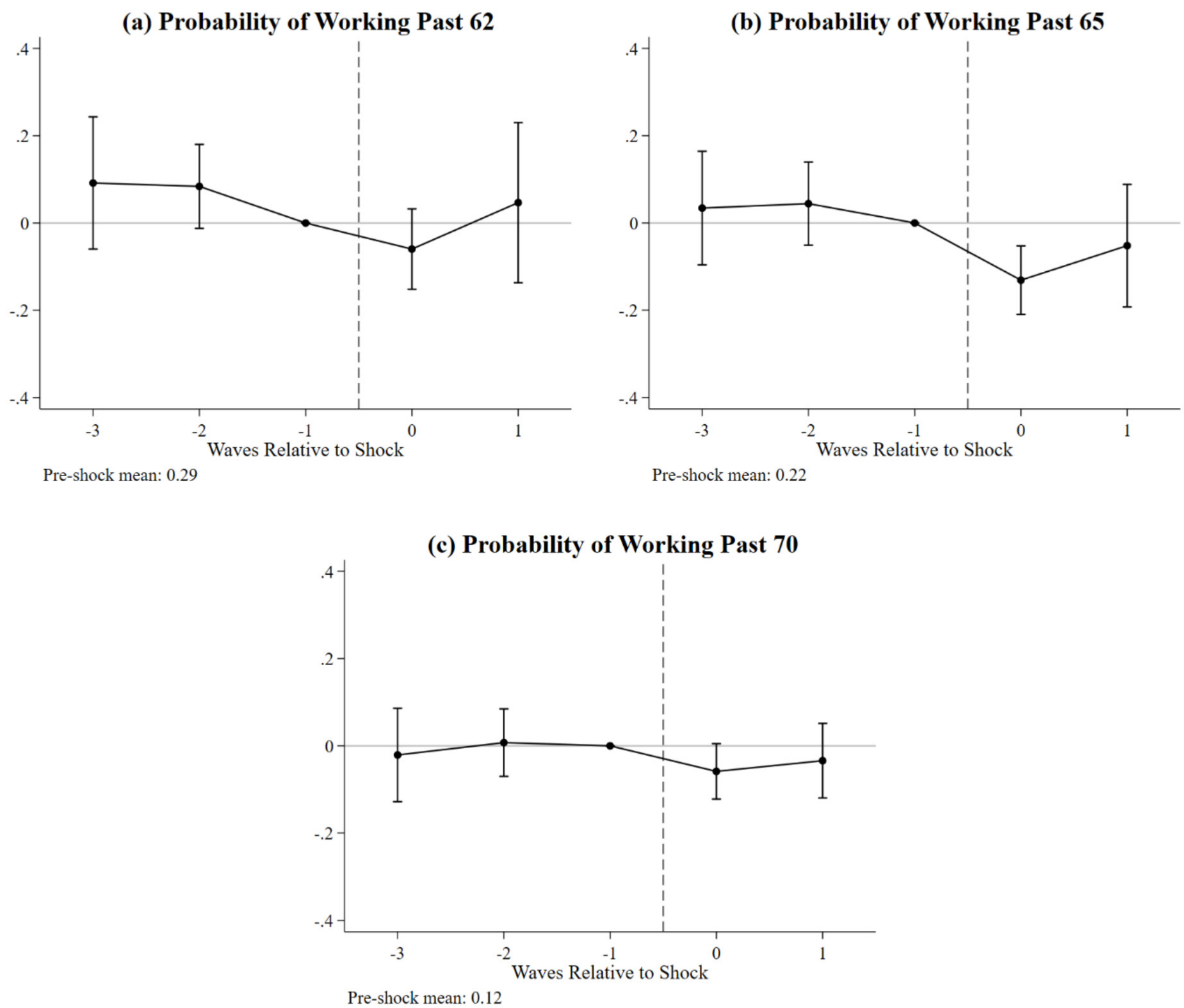

Notes: This figure presents the nonparametric event study estimates for strokes. Each graph plots point estimates and confidence intervals for the δ_τ coefficients in equation (2).

FIGURE 8. EFFECTS OF CANCER ON RETIREMENT EXPECTATIONS

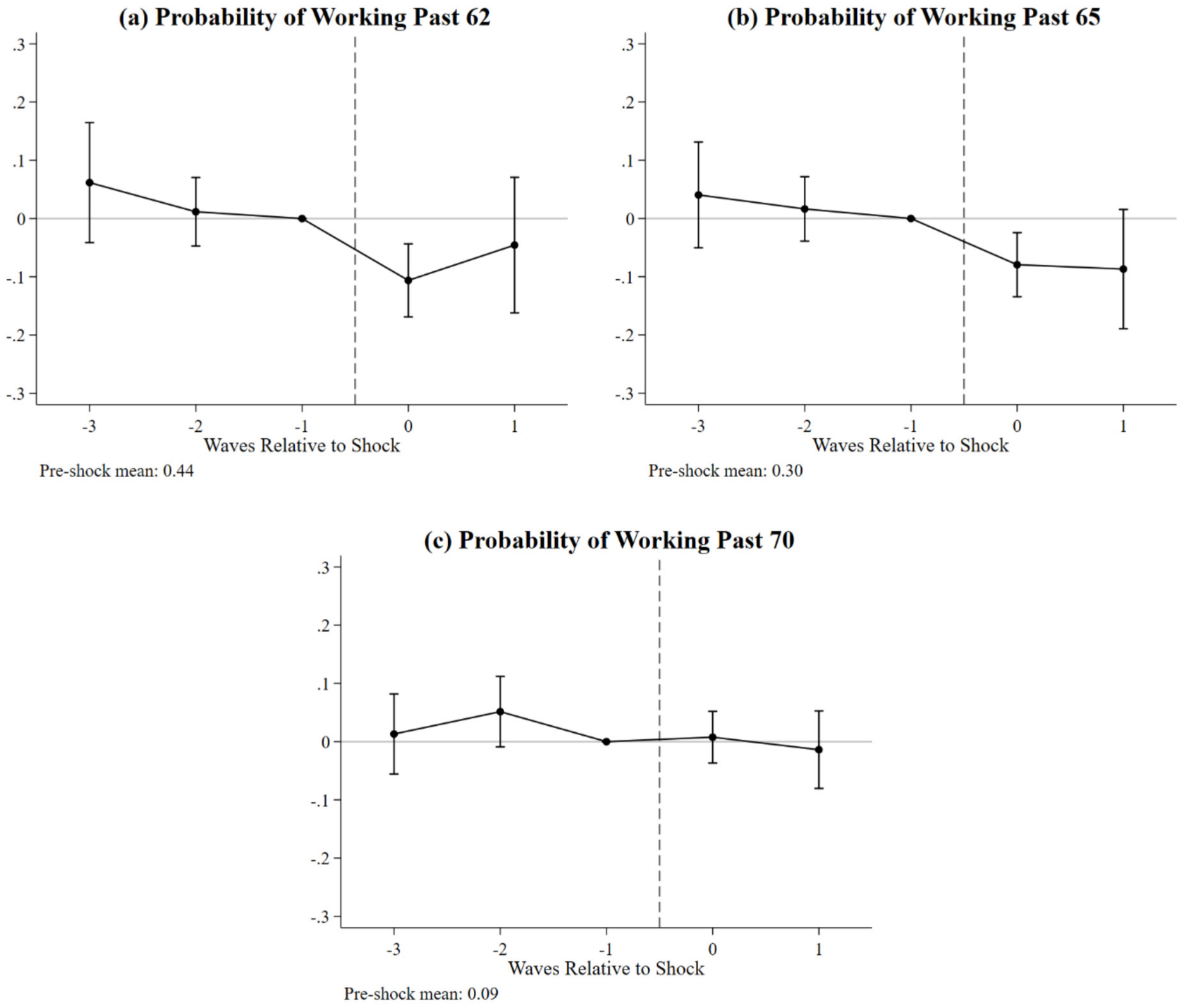

Notes: This figure presents the nonparametric event study estimates for cancer diagnoses. Each graph plots point estimates and confidence intervals for the δ_τ coefficients in equation (2).

FIGURE 9. EFFECTS OF LUNG DISEASE ON RETIREMENT EXPECTATIONS

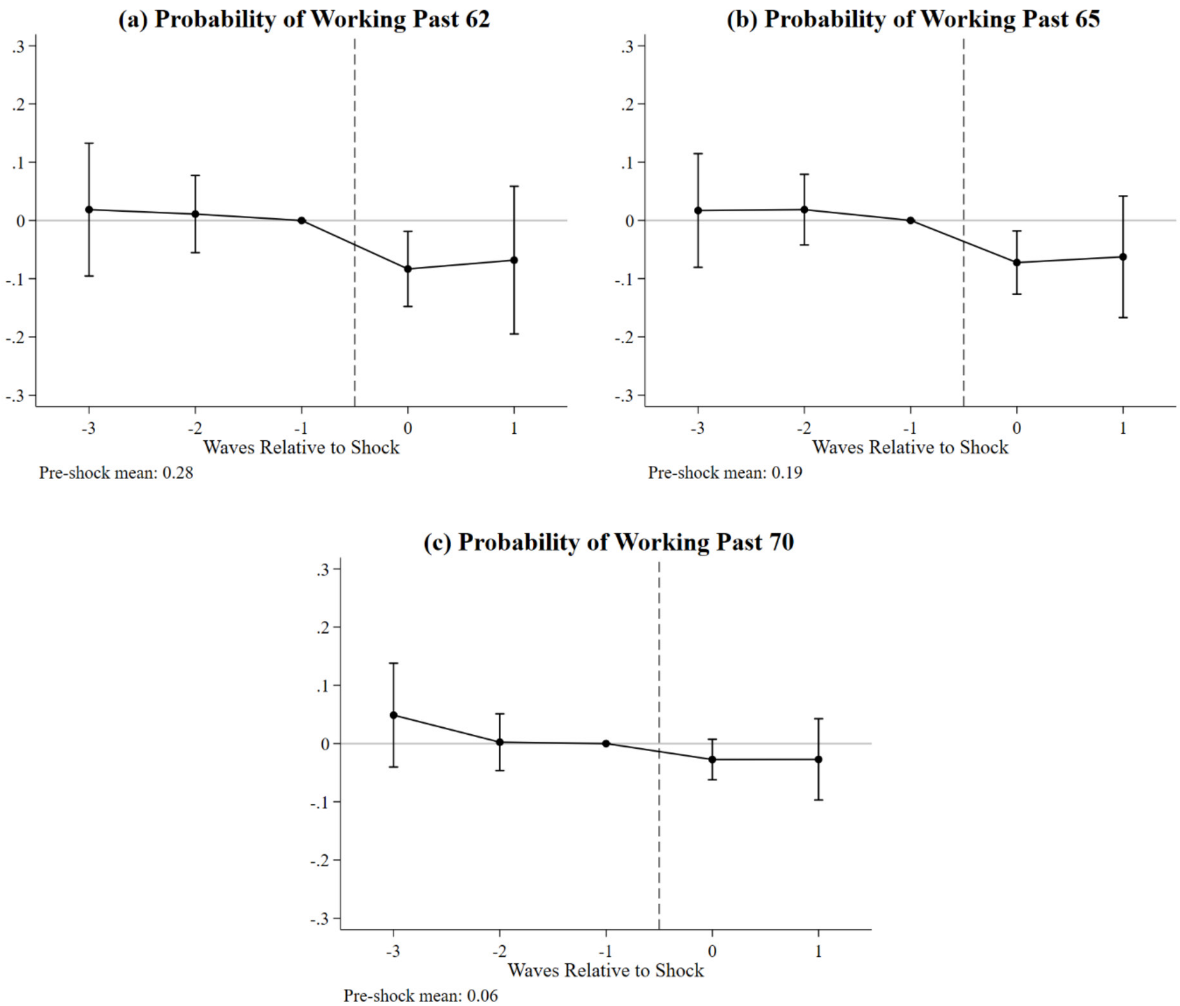

Notes: This figure presents the nonparametric event study estimates for lung disease. Each graph plots point estimates and confidence intervals for the δ_τ coefficients in equation (2).

FIGURE 10. EFFECTS OF ARTHRITIS ON RETIREMENT EXPECTATIONS

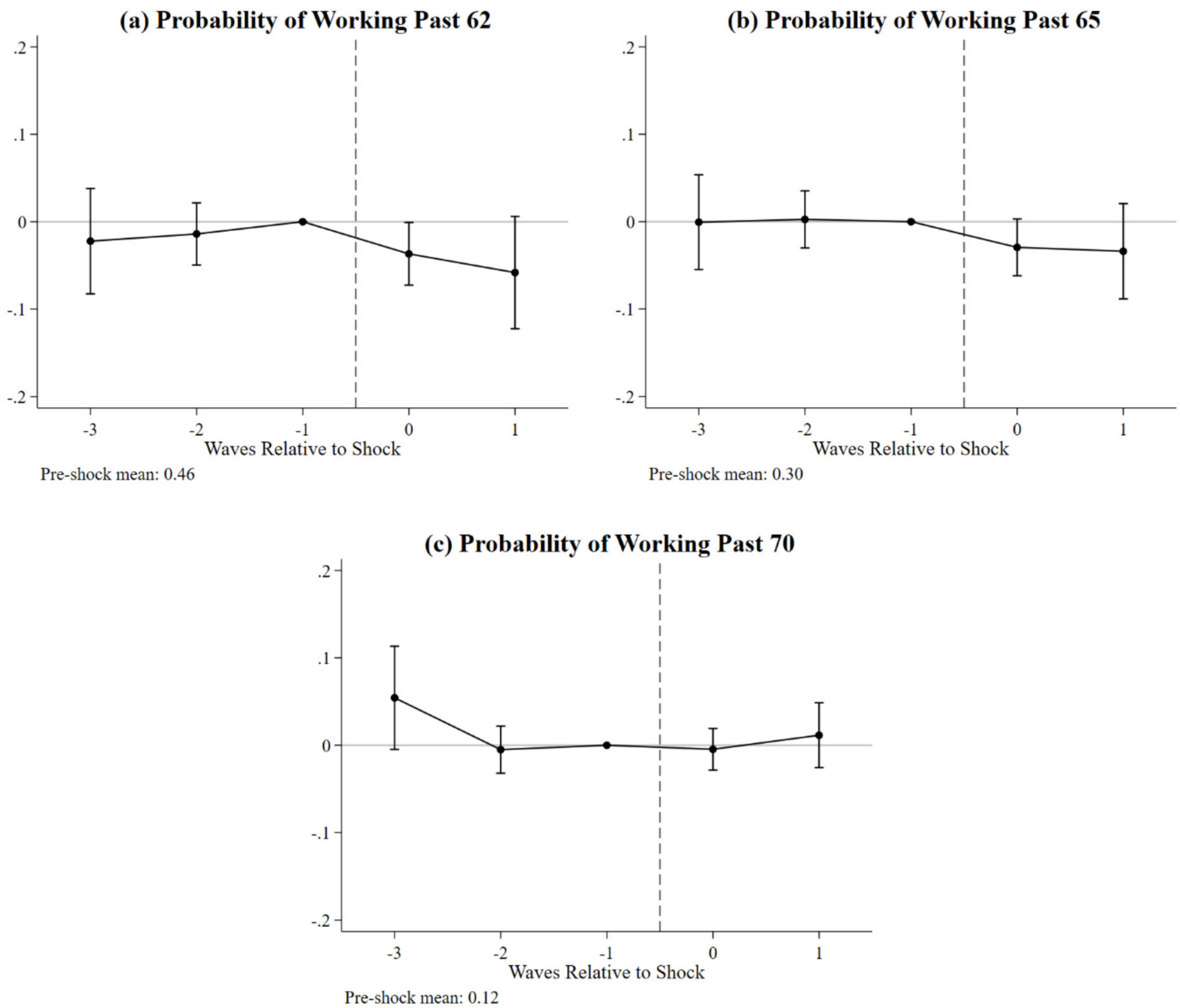

Notes: This figure presents the nonparametric event study estimates for arthritis. Each graph plots point estimates and confidence intervals for the δ_τ coefficients in equation (2).

FIGURE 11. EFFECTS OF DIABETES ON RETIREMENT EXPECTATIONS

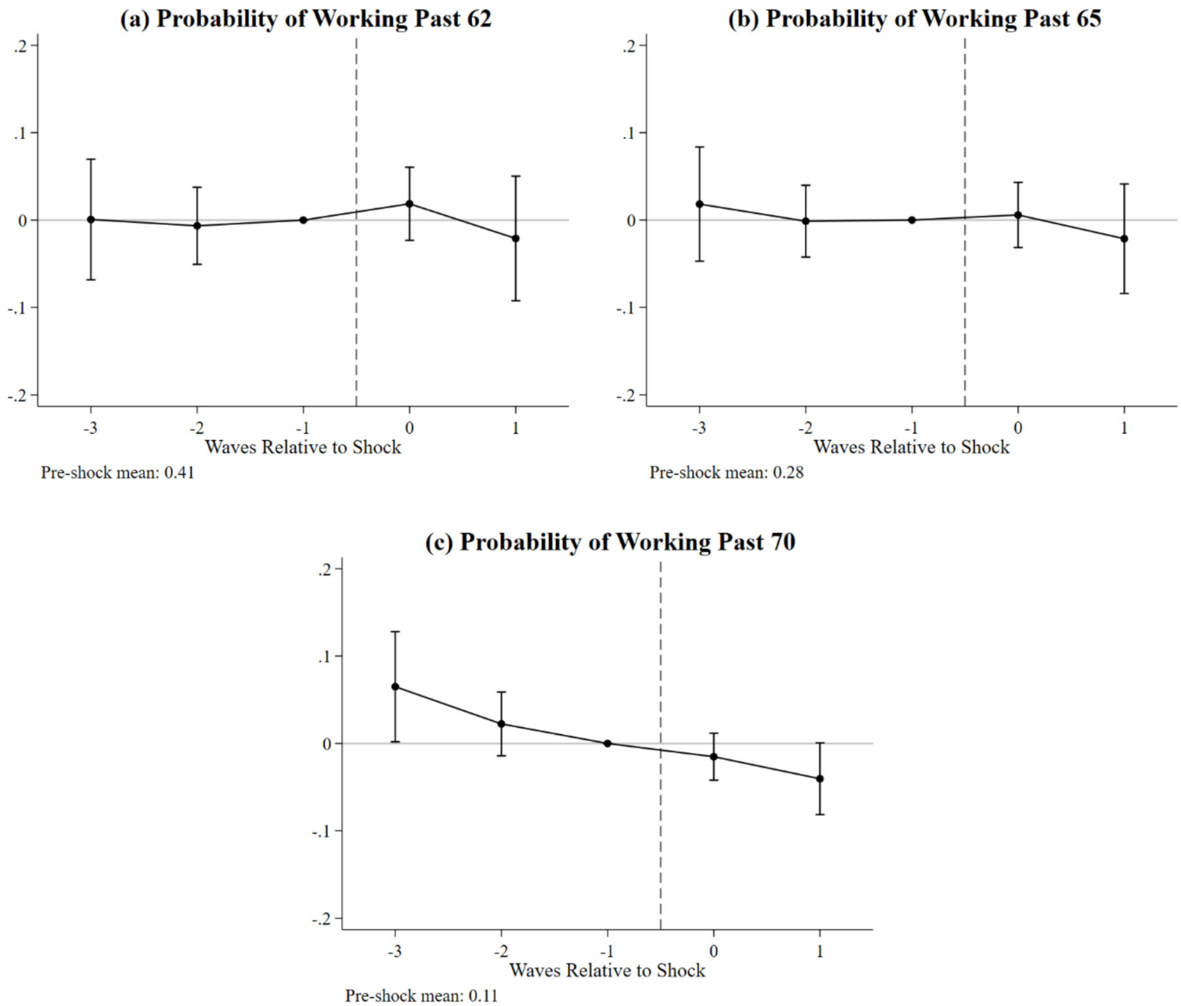

Notes: This figure presents the nonparametric event study estimates for diabetes. Each graph plots point estimates and confidence intervals for the δ_τ coefficients in equation (2).

FIGURE 12. EFFECTS OF HIGH BLOOD PRESSURE ON RETIREMENT EXPECTATIONS

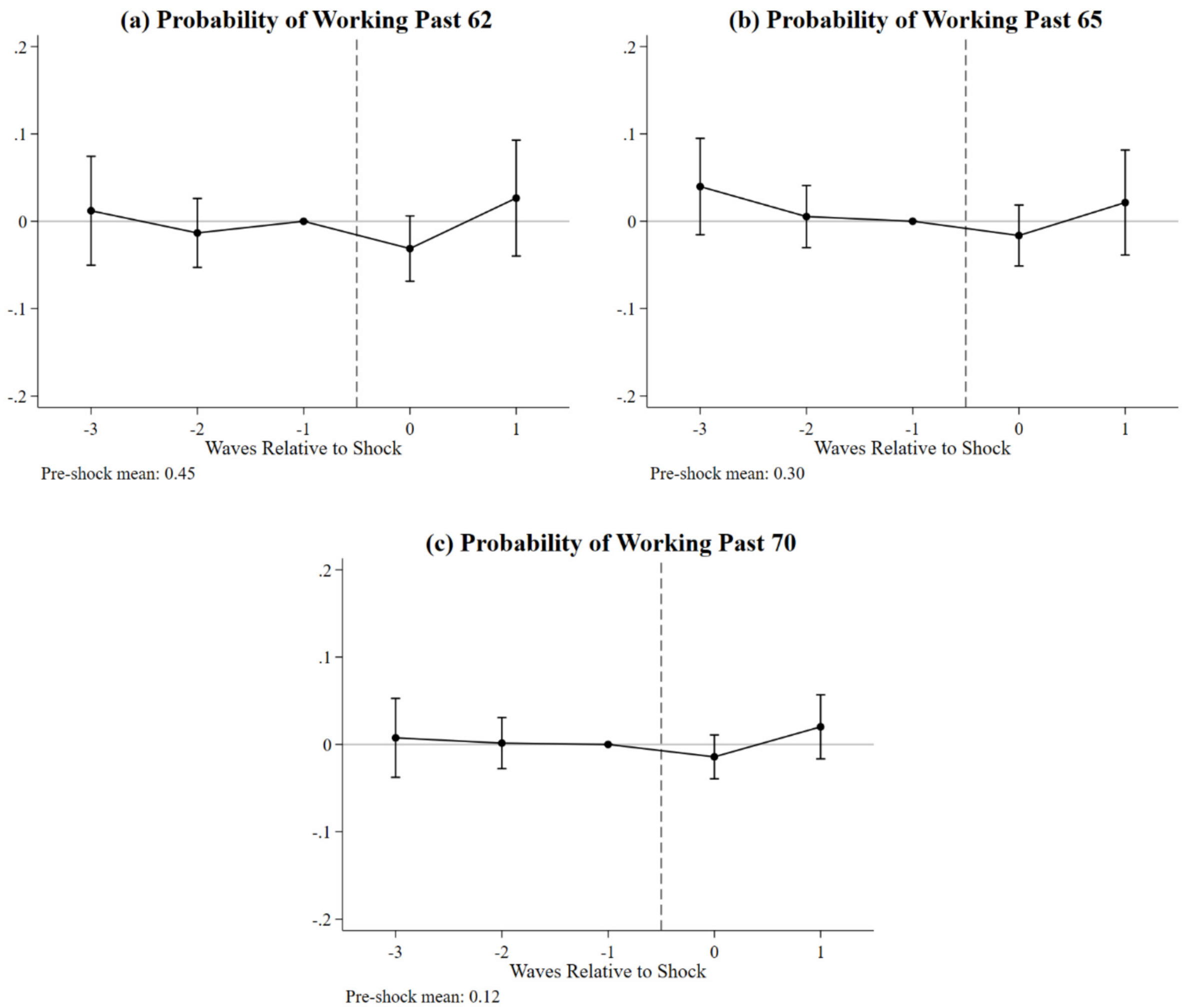

Notes: This figure presents the nonparametric event study estimates for high blood pressure. Each graph plots point estimates and confidence intervals for the δ_τ coefficients in equation (2).

FIGURE 13. EFFECTS OF PSYCHIATRIC PROBLEMS ON RETIREMENT EXPECTATIONS

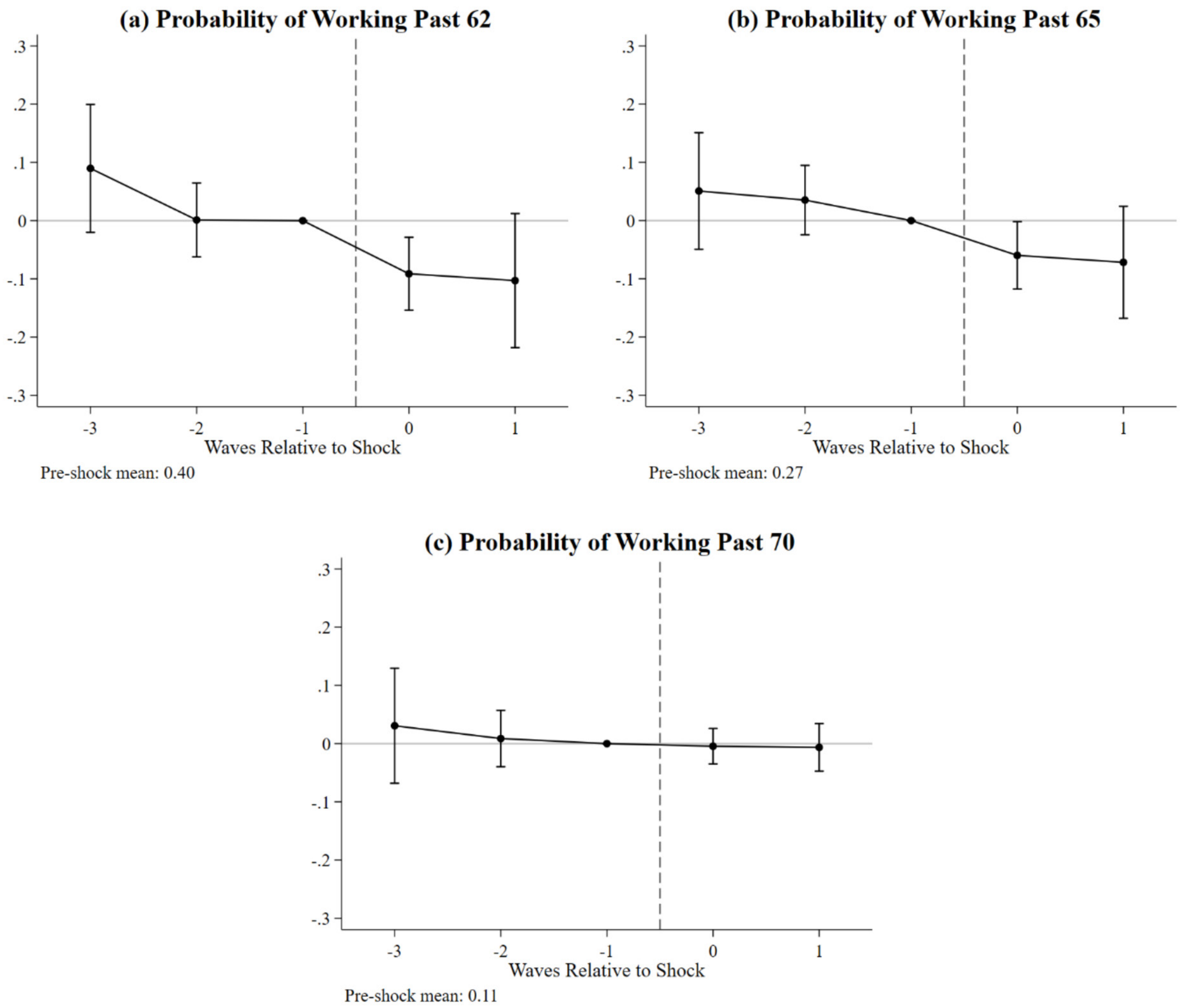

Notes: This figure presents the nonparametric event study estimates for psychiatric problems. Each graph plots point estimates and confidence intervals for the δ_τ coefficients in equation (2).

FIGURE 14. EFFECTS OF UNEMPLOYMENT EVENTS ON RETIREMENT EXPECTATIONS

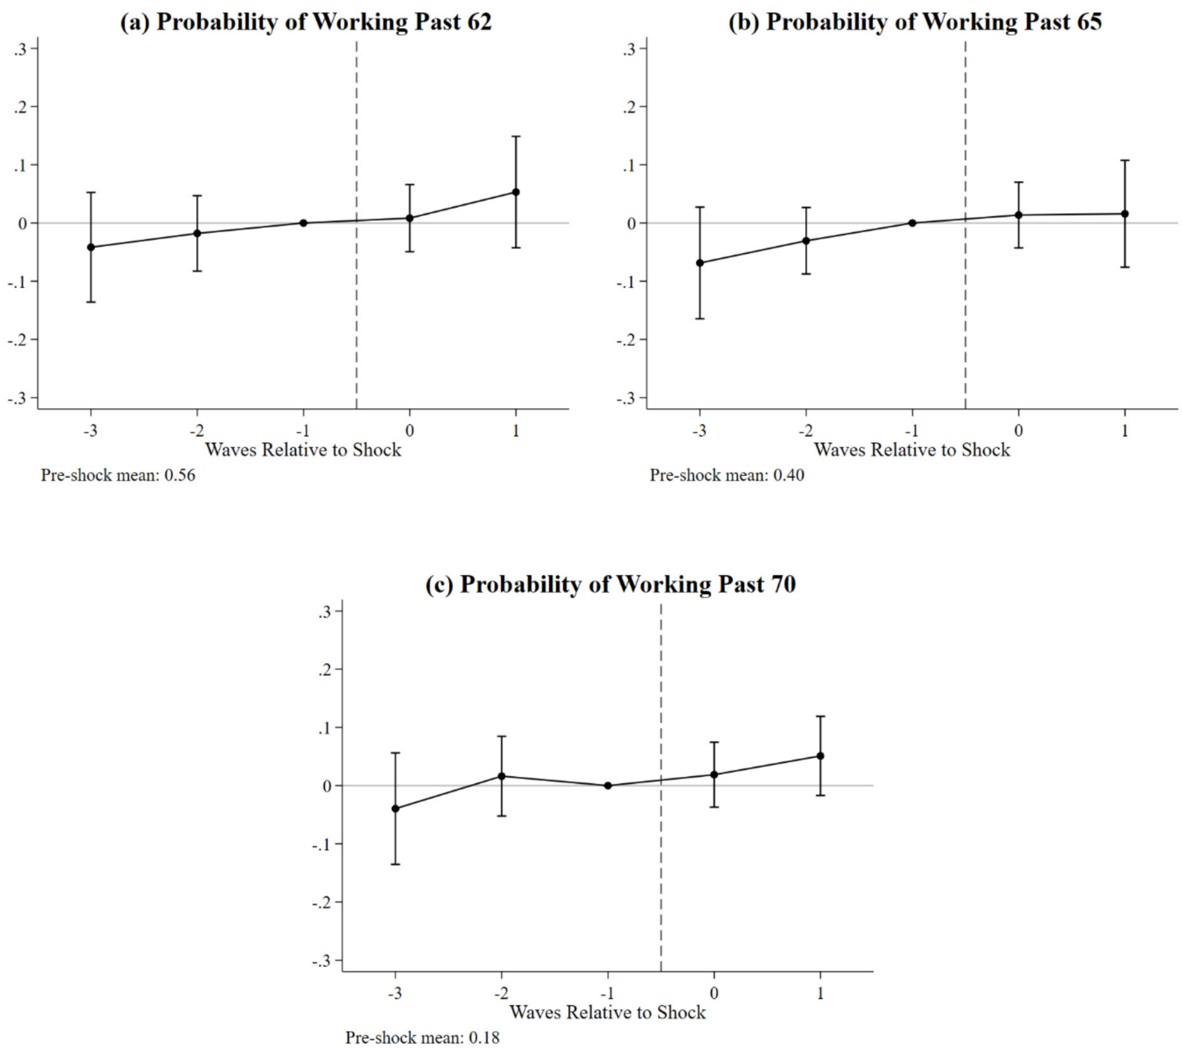

Notes: This figure presents the nonparametric event study estimates for unemployment events. Each graph plots point estimates and confidence intervals for the δ_τ coefficients in equation (2).

FIGURE 15. EFFECTS OF EARNINGS INCREASES ON RETIREMENT EXPECTATIONS

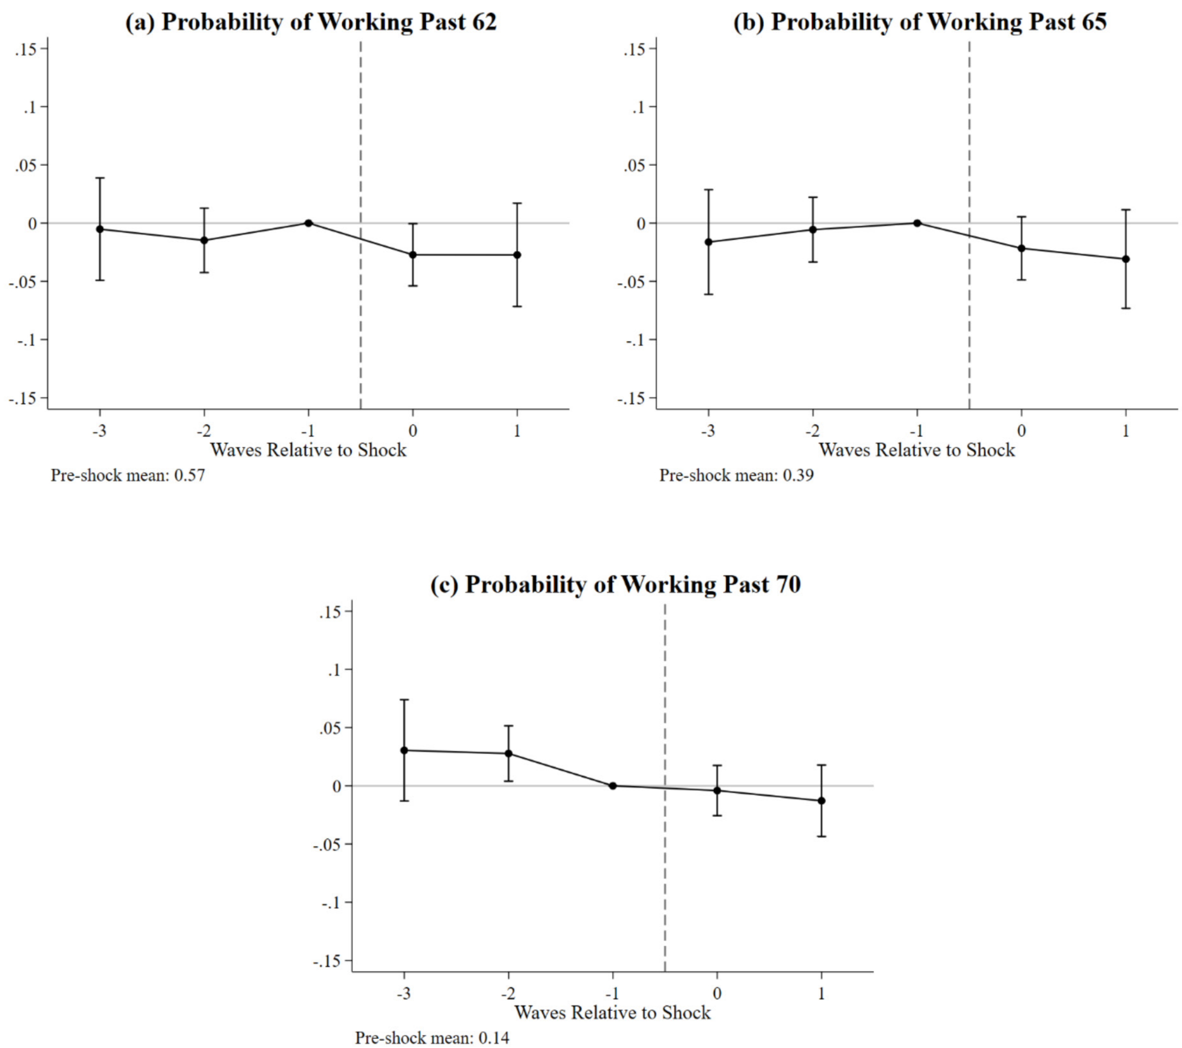

Notes: This figure presents the nonparametric event study estimates for increases in earnings. Each graph plots point estimates and confidence intervals for the δ_τ coefficients in equation (2).

FIGURE 16. EFFECTS OF GRANDCHILD BIRTHS ON RETIREMENT EXPECTATIONS

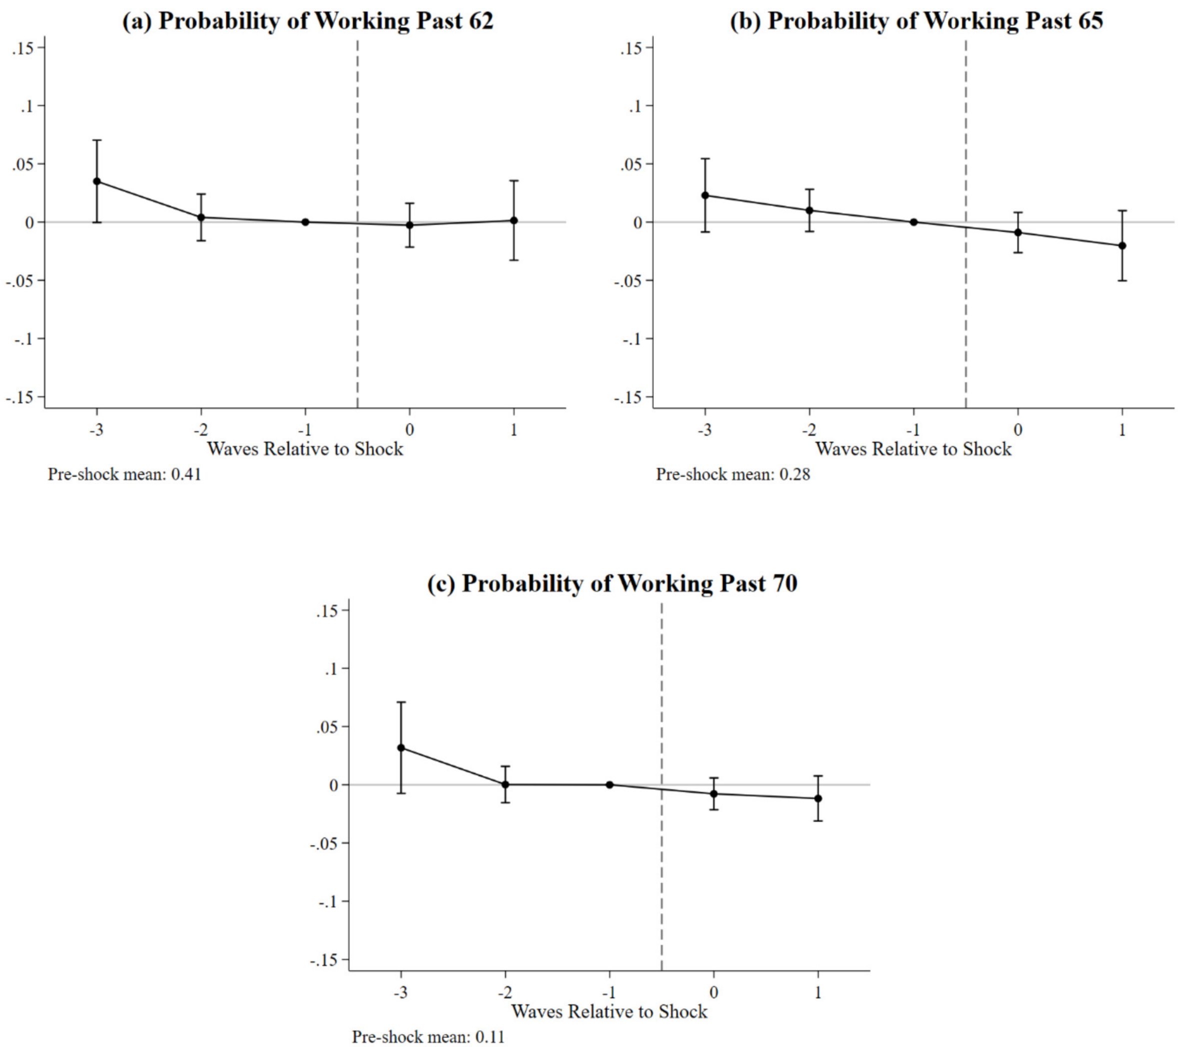

Notes: This figure presents the nonparametric event study estimates for grandchildren births. Each graph plots point estimates and confidence intervals for the δ_τ coefficients in equation (2).

FIGURE 17. EFFECTS OF DIVORCE ON RETIREMENT EXPECTATIONS

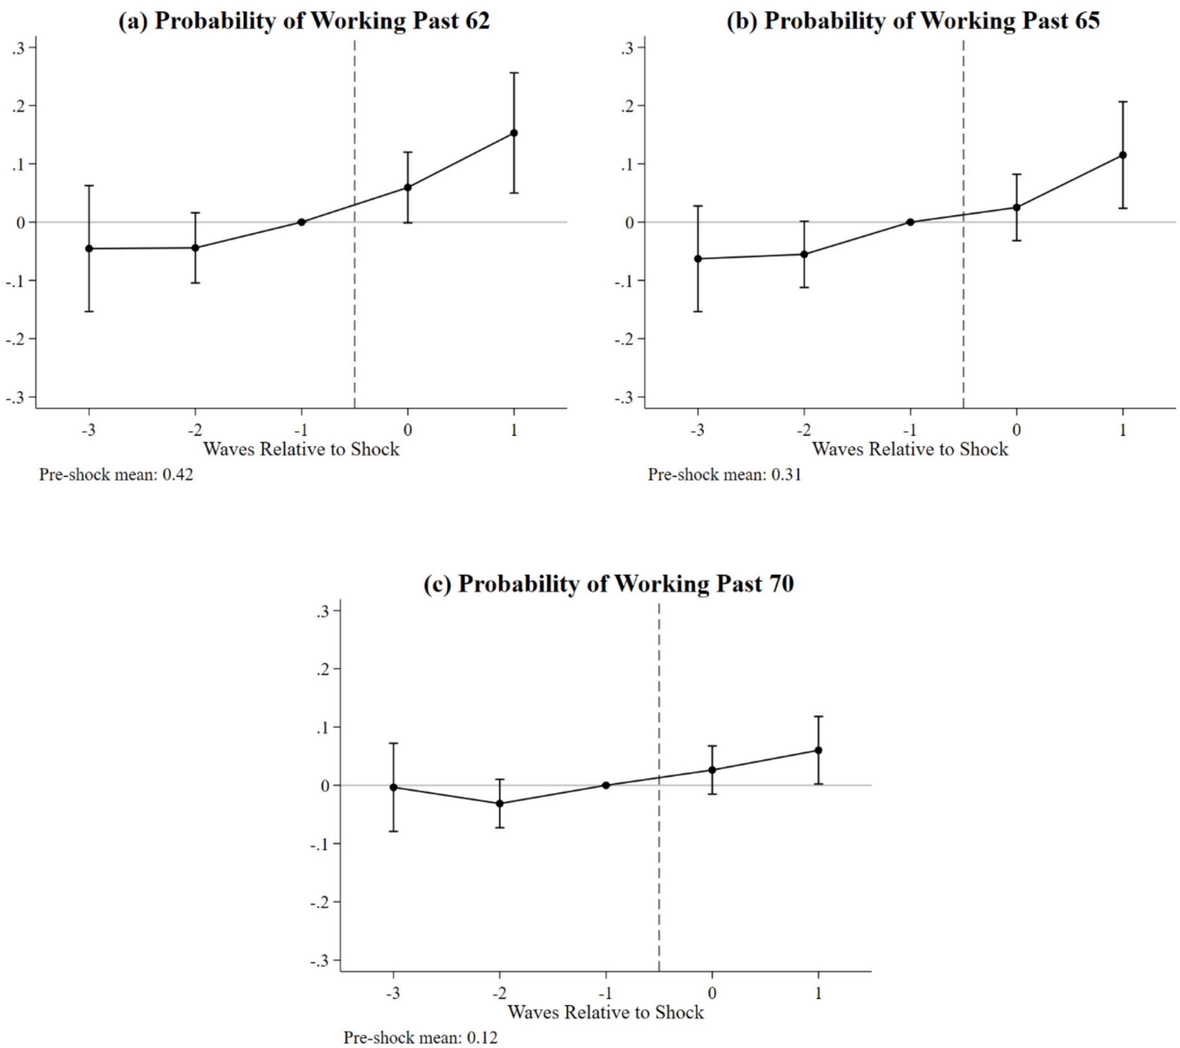

Notes: This figure presents the nonparametric event study estimates for divorce or separation events. Each graph plots point estimates and confidence intervals for the δ_τ coefficients in equation (2).

TABLE 1. SUMMARY STATISTICS FOR DESCRIPTIVE ANALYSIS SAMPLE

	Mean (1)	Std. Dev. (2)	Observations (3)
Panel A: Demographics			
Age	53.92	2.75	18,572
Birth Year	1949.26	10.07	18,572
Survey Year	2003.6	87.7	18,572
Male	0.47	0.50	18,572
White	0.70	0.46	18,572
Married	0.74	0.44	18,572
Attended Some College	0.53	0.50	18,572
Panel B: Economic Characteristics			
Earnings ($)	44,939	56,862	18,572
IRA Balances ($)	39,859	142,457	18,572
Panel C: Retirement Expectations			
Probability of Working Past Age 62	0.51	0.37	18,572
Probability of Working Past Age 65	0.32	0.34	18,572
Probability of Working Past Age 70	0.16	0.24	4,328
Expected Retirement Age	64.04	4.57	18,572
Fraction Expecting to Never Retire	0.37	0.48	11,725

Notes: This table reports summary statistics for our descriptive analysis sample of 18,572 people who are working either full time or part time, who are between ages 50 and 60, and who have nonmissing values for the probability-based outcome variables and nonmissing values for the expected retirement age outcome variable (we include people who say that they will never retire). Since individuals differ in the number of waves in which they appear, we report means and standard deviations of variables from the first survey wave in which an individual appears in the sample. Monetary values are expressed in 2010 dollars.

TABLE 2. DIFFERENCES IN EXPECTED RETIREMENT AGE ACROSS GROUPS

	Average Level at Age 50 (1)	Slope Parameter (2)	Observations (3)
Panel A: Full Sample			
	63.84	0.163***	26,689
		(0.010)	
Panel B: Gender			
Male	64.31	0.143***	11,255
		(0.016)	
Female	63.66	0.170***	15,434
		(0.012)	
Panel C: Race			
White	63.91	0.161***	19,083
		(0.011)	
Not white	63.65	0.166***	7,606
		(0.019)	
Panel D: Marital Status			
Married	63.72	0.173***	19,998
		(0.011)	
Not married	64.51	0.101***	6,691
		(0.021)	
Panel E: Health Status			
Good health	63.90	0.167***	22,744
		(0.010)	
Not good health	63.44	0.142***	3,945
		(0.027)	
Panel F: Health Diagnosis			
No health condition	63.89	0.166***	22,240
		(0.011)	
Diagnosed with a health condition	63.49	0.142***	4,449
		(0.025)	
Panel G: Income			
High income	63.42	0.180***	13,704
		(0.013)	
Low income	64.25	0.148***	12,985
		(0.014)	
Panel H: Wealth			
High wealth	63.18	0.223***	13,671
		(0.013)	
Low wealth	64.48	0.102***	13,018
		(0.014)	

Notes: This table displays the average expected retirement age at age 50 and the results from estimating equation (1) for several different groups. We report robust standard errors. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

TABLE 3. DIFFERENCES IN PROBABILITIES OF WORKING PAGE AGE 62 ACROSS GROUPS

	Average Level at Age 50 (1)	Slope Parameter (2)	Observations (3)
Panel A: Full Sample			
	0.48	0.012^{***}	45,143
		(0.001)	
Panel B: Gender			
Male	0.53	0.011^{***}	20,481
		(0.001)	
Female	0.46	0.012^{***}	24,662
		(0.001)	
Panel C: Race			
White	0.51	0.012^{***}	32,516
		(0.001)	
Not white	0.41	0.010^{***}	12,627
		(0.001)	
Panel D: Marital Status			
Married	0.47	0.012^{***}	33,396
		(0.001)	
Not married	0.54	0.009^{***}	11,747
		(0.001)	
Panel E: Health Status			
Good health	0.49	0.013^{***}	38,467
		(0.001)	
Not good health	0.41	0.007^{***}	6,676
		(0.002)	
Panel F: Health Diagnosis			
No health condition	0.48	0.012^{***}	37,581
		(0.001)	
Diagnosed with a health condition	0.47	0.013^{***}	7,562
		(0.002)	
Panel G: Income			
High income	0.51	0.014^{***}	23,070
		(0.001)	
Low income	0.45	0.010^{***}	22,073
		(0.001)	
Panel H: Wealth			
High wealth	0.46	0.015^{***}	22,917
		(0.001)	
Low wealth	0.50	0.009^{***}	22,226
		(0.001)	

Notes: This table displays the average probability of working past age 62 at age 50 and the results from estimating equation (1) for several different groups. We report robust standard errors. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

TABLE 4. DIFFERENCES IN PROBABILITIES OF WORKING PAGE AGE 65 ACROSS GROUPS

	Average Level at Age 50 (1)	Slope Parameter (2)	Observations (3)
Panel A: Full Sample			
	0.31	0.006***	45,143
		(0.001)	
Panel B: Gender			
Male	0.34	0.005***	20,481
		(0.001)	
Female	0.29	0.005***	24,662
		(0.001)	
Panel C: Race			
White	0.32	0.006***	32,516
		(0.001)	
Not white	0.27	0.005***	12,627
		(0.001)	
Panel D: Marital Status			
Married	0.30	0.006***	33,396
		(0.001)	
Not married	0.36	0.003**	11,747
		(0.001)	
Panel E: Health Status			
Good health	0.32	0.007***	38,467
		(0.001)	
Not good health	0.24	0.002	6,676
		(0.001)	
Panel F: Health Diagnosis			
No health condition	0.31	0.007***	37,581
		(0.001)	
Diagnosed with a health condition	0.31	0.004***	7,562
		(0.001)	
Panel G: Income			
High income	0.32	0.007***	23,070
		(0.001)	
Low income	0.30	0.005***	22,073
		(0.001)	
Panel H: Wealth			
High wealth	0.29	0.007***	22,917
		(0.001)	
Low wealth	0.33	0.005***	22,226
		(0.001)	

Notes: This table displays the average probability of working past age 65 at age 50 and the results from estimating equation (1) for several different groups. We report robust standard errors. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

TABLE 5. DIFFERENCES IN PROBABILITIES OF WORKING PAGE AGE 70 ACROSS GROUPS

	Average Level at Age 50 (1)	Slope Parameter (2)	Observations (3)
Panel A: Full Sample			
	0.16	0.001	15,987
		(0.001)	
Panel B: Gender			
Male	0.20	-0.002	7,249
		(0.001)	
Female	0.14	0.001*	8,738
		(0.001)	
Panel C: Race			
White	0.17	-0.000	9,658
		(0.001)	
Not white	0.14	0.002	6,329
		(0.001)	
Panel D: Marital Status			
Married	0.16	-0.000	11,276
		(0.001)	
Not married	0.15	0.001	4,711
		(0.001)	
Panel E: Health Status			
Good health	0.17	0.001	13,268
		(0.001)	
Not good health	0.10	-0.000	2,719
		(0.002)	
Panel F: Health Diagnosis			
No health condition	0.16	0.001	13,022
		(0.001)	
Diagnosed with a health condition	0.12	-0.001	2,965
		(0.002)	
Panel G: Income			
High income	0.15	-0.000	9,380
		(0.001)	
Low income	0.17	0.002*	6,607
		(0.001)	
Panel H: Wealth			
High wealth	0.15	-0.001	6,803
		(0.001)	
Low wealth	0.16	0.002*	9,184
		(0.001)	

Notes: This table displays the average probability of working past age 70 at age 50 and the results from estimating equation (1) for several different groups. We report robust standard errors. *** p<0.01, ** p<0.05, * p<0.1.

TABLE 6. EFFECTS OF HEALTH DECLINES AND HOSPITALIZATIONS ON RETIREMENT EXPECTATIONS

	Probability of Working Past 62 (1)	Probability of Working Past 65 (2)	Probability of Working Past 70 (3)
Panel A: Change in Health Status			
Parametric Event Study Estimate	-0.040**	-0.045**	-0.010
	(0.019)	(0.018)	(0.015)
Mean	0.45	0.32	0.12
Estimate in Percent Terms	-8.9%	-14.1%	8.3%
Clusters	855	854	792
Observations	3,153	3,116	2,170
Panel B: Hospitalization			
Parametric Event Study Estimate	-0.008	-0.019	-0.002
	(0.016)	(0.015)	(0.013)
Mean	0.45	0.31	0.12
Estimate in Percent Terms	-1.8%	-6.1%	-1.7%
Clusters	1,210	1,210	1,087
Observations	4,463	4,431	2,914

Notes: This table presents the parametric event study estimates for declines in health status and hospitalization events. The table displays estimates of δ_0 from estimating equation (3) for each regression. We report robust standard errors clustered at the individual level. *** p<0.01, ** p<0.05, * p<0.1.

TABLE 7. EFFECTS OF HEART ATTACKS, STROKES, CANCER, AND LUNG DISEASE ON RETIREMENT EXPECTATIONS

	Probability of Working Past 62 (1)	Probability of Working Past 65 (2)	Probability of Working Past 70 (3)
Panel A: Heart Attacks			
Parametric Event Study Estimate	0.046	0.009	0.006
	(0.030)	(0.029)	(0.027)
Mean	0.35	0.23	0.10
Estimate in Percent Terms	13.1%	3.9%	6.0%
Clusters	351	351	315
Observations	1,275	1,266	890
Panel B: Strokes			
Parametric Event Study Estimate	-0.021	-0.119***	-0.068*
	(0.051)	(0.045)	(0.037)
Mean	0.29	0.22	0.12
Estimate in Percent Terms	-7.2%	-54.1%	-56.7%
Clusters	127	127	118
Observations	436	431	305
Panel C: Cancer			
Parametric Event Study Estimate	-0.072**	-0.058*	0.018
	(0.034)	(0.032)	(0.029)
Mean	0.44	0.30	0.09
Estimate in Percent Terms	-16.4%	-19.5%	20.0%
Clusters	233	233	214
Observations	854	849	580
Panel D: Lung Disease			
Parametric Event Study Estimate	-0.074**	-0.065**	-0.008
	(0.037)	(0.030)	(0.025)
Mean	0.28	0.19	0.06
Estimate in Percent Terms	-26.4%	-34.2%	-13.3%
Clusters	208	208	193
Observations	751	744	519

Notes: This table presents the parametric event study estimates for heart attacks, strokes, cancer diagnoses, and lung disease diagnoses. The table displays estimates of δ_0 from estimating equation (3) for each regression. We report robust standard errors clustered at the individual level. *** p<0.01, ** p<0.05, * p<0.1.

TABLE 8. EFFECTS OF ARTHRITIS, DIABETES, HIGH BLOOD PRESSURE, AND PSYCHIATRIC PROBLEMS ON RETIREMENT EXPECTATIONS

	Probability of Working Past 62 (1)	Probability of Working Past 65 (2)	Probability of Working Past 70 (3)
Panel A: Arthritis			
Parametric Event Study Estimate	-0.047**	-0.030	0.018
	(0.021)	(0.019)	(0.016)
Mean	0.46	0.30	0.12
Estimate in Percent Terms	-10.2%	-10.0%	15.0%
Clusters	815	816	754
Observations	2,988	2,959	2,082
Panel B: Diabetes			
Parametric Event Study Estimate	0.020	0.017	0.016
	(0.023)	(0.021)	(0.020)
Mean	0.41	0.28	0.11
Estimate in Percent Terms	4.9%	6.1%	14.5%
Clusters	521	521	493
Observations	1,932	1,905	1,406
Panel C: High Blood Pressure			
Parametric Event Study Estimate	-0.022	0.006	-0.011
	(0.022)	(0.020)	(0.017)
Mean	0.45	0.30	0.12
Estimate in Percent Terms	-4.9%	2.0%	-9.2%
Clusters	692	692	620
Observations	2,496	2,470	1,656
Panel D: Psychiatric Problems			
Parametric Event Study Estimate	-0.041	-0.036	0.009
	(0.038)	(0.034)	(0.025)
Mean	0.40	0.27	0.11
Estimate in Percent Terms	-10.3%	-13.3%	8.2%
Clusters	316	316	294
Observations	1,144	1,139	786

Notes: This table presents the parametric event study estimates for the diagnosis of arthritis, diabetes, high blood pressure and psychiatric problems. The table displays estimates of δ_0 from estimating equation (3) for each regression. We report robust standard errors clustered at the individual level. *** p<0.01, ** p<0.05, * p<0.1.

TABLE 9. EFFECTS OF ECONOMIC EVENTS ON RETIREMENT EXPECTATIONS

	Probability of Working Past 62 (1)	Probability of Working Past 65 (2)	Probability of Working Past 70 (3)
Panel A: Unemployment			
Parametric Event Study Estimate	-0.013	-0.021	0.005
	(0.037)	(0.037)	(0.039)
Mean	0.56	0.40	0.18
Estimate in Percent Terms	-2.3%	-5.3%	2.8%
Clusters	317	317	268
Observations	1,129	1,121	681
Panel B: Earnings Increases			
Parametric Event Study Estimate	-0.028*	-0.030*	0.013
	(0.016)	(0.016)	(0.013)
Mean	0.57	0.39	0.14
Estimate in Percent Terms	-4.9%	-7.7%	9.3%
Clusters	1,327	1,328	1,246
Observations	4,901	4,846	3,319

Notes: This table presents the parametric event study estimates for unemployment and earnings increases. The table displays estimates of δ_0 from estimating equation (3) for each regression. We report robust standard errors clustered at the individual level. *** p<0.01, ** p<0.05, * p<0.1.

TABLE 10. EFFECTS OF FAMILY EVENTS ON RETIREMENT EXPECTATIONS

	Probability of Working Past 62 (1)	Probability of Working Past 65 (2)	Probability of Working Past 70 (3)
Panel A: Grandchild Births			
Parametric Event Study Estimate	0.016	0.003	0.002
	(0.011)	(0.010)	(0.009)
Mean	0.41	0.28	0.11
Estimate in Percent Terms	3.9%	1.1%	1.8%
Clusters	2,814	2,815	2,497
Observations	10,467	10,420	6,198
Panel B: Divorce			
Parametric Event Study Estimate	0.039	-0.004	0.017
	(0.036)	(0.035)	(0.025)
Mean	0.42	0.31	0.12
Estimate in Percent Terms	9.3%	-1.3%	14.2%
Clusters	267	267	256
Observations	984	978	689

Notes: This table presents the parametric event study estimates for grandchild births and divorces or separations. The table displays estimates of δ_0 from estimating equation (3) for each regression. We report robust standard errors clustered at the individual level. *** p<0.01, ** p<0.05, * p<0.1.

About the authors

Aspen Gorry is an associate professor in the John E. Walker Department of Economics at Clemson University. Prior to joining Clemson, Professor Gorry was an assistant professor at University of California, Santa Cruz and Utah State University. He was also a research fellow at the American Enterprise Institute. Dr. Gorry's research focuses on macroeconomics, labor economics and public economics. In particular, he studies the effect of policies such as taxes, retirement programs and minimum wages on labor market outcomes. His research has been published in journals including American Economic Journal: Macroeconomics, Journal of Public Economics, Quantitative Economics and Review of Economic Dynamics. His policy writing has appeared in Tax Notes, the Los Angeles Times, and Real Clear Markets. Dr. Gorry received his MA and PhD in economics from the University of Chicago and a BS in mathematics and economics from Arizona State University.

Jonathan M. Leganza is an assistant professor in the John E. Walker Department of Economics at Clemson University. His primary fields of interest are public finance, labor economics, and health economics. He is especially interested in the economics of aging, and he studies topics related to retirement, social security policy, and Medicare policy. He received his PhD in economics from the University of California San Diego in 2021, his BA in economics from Indiana University in 2015, and his BS in mathematics from Indiana University in 2015.

About the TIAA Institute

The TIAA Institute helps advance the ways individuals and institutions plan for financial security and organizational effectiveness. The Institute conducts in-depth research, provides access to a network of thought leaders, and enables those it serves to anticipate trends, plan future strategies, and maximize opportunities for success.

To learn more, visit tiaainstitute.org.

Join the conversation online:
@TIAAInstitute

This research was supported by funding from the TIAA Institute. The content, findings and conclusions are the responsibility of the authors and do not necessarily represent the views of TIAA or the TIAA Institute. Gorry: John E. Walker Department of Economics. Clemson University (email: aspen.gorry@gmail.com). Leganza: John E. Walker Department of Economics. Clemson University (email: jleganz@clemson.edu).

GRE-3506930PR-O0424WX

TIAA Institute

(04/24)