Revisiting the Effects of Health on Retirement: Event Study Evidence from Expectations Data*

Aspen Gorry† Jonathan M. Leganza‡

October 15, 2024

Abstract

How does health affect the timing of retirement? We revisit this question using data from the Health and Retirement Study and an event study framework to estimate the causal effects of various health shocks on retirement expectations. Importantly, the more-recent survey data contain information on expectations for current workers and non-workers, which avoids sample selection issues. We find that declines in subjective health status and new diagnoses of objective health conditions decrease the probabilities of working past 62 and 65. Our findings highlight how poor health can cause people to retire before becoming eligible for Social Security and Medicare benefits.

Keywords: retirement expectations, retirement timing, health shocks
JEL codes: J26, J22, I10

*We thank Brent Davis and participants at the public economics workshop at Clemson University for helpful comments and conversations. 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.

†John E. Walker Department of Economics. Clemson University. (email: aspen.gorry@gmail.com)
‡John E. Walker Department of Economics. Clemson University. (email: jleganz@clemson.edu)

1 Introduction

The timing of retirement is a major determinant of lifetime earnings and a crucial factor in financial security at older ages. However, people face uncertainty about the timing of their retirement. One prominent risk people face relates to the evolution of their health. On the one hand, poor health may decrease work capacity and prevent people from working as long as they had planned. On the other hand, poor health may lead to greater medical expenses that induce people to work longer. Therefore, an important task for empirical work is to provide evidence on the extent to which health impacts the timing of retirement and to assess the importance of different health risks.

Yet, providing causal evidence on how health affects the timing of retirement is difficult. There are two key challenges. One is identification. Correlations between health and retirement may not reflect causal relationships because of concerns like omitted variable bias or reverse causality. Health can affect retirement, but retirement can also affect health. The other challenge is measurement. Retirements may not be realized until years after a person experiences health conditions that can impact the timing of their retirement. For example, a person could experience a health shock in their 50s that causes them to retire in their 60s instead of in their 70s. It can be difficult to isolate the causal link between that specific health shock and a later retirement decision.

In this paper, we use an event study framework and data on retirement expectations to overcome these challenges, providing new evidence on how health impacts retirement timing. To overcome the identification challenge, we leverage the quasi-random timing of various health shocks. To overcome the measurement challenge, we follow Dwyer and Mitchell (1999) and McGarry (2004) in using retirement expectations data from the Health and Retirement Study (HRS). Since people in the HRS data have rational expectations about retirement timing (Benitez-Silva and Dwyer, 2005) and their expectations strongly predict future retirements (Haider and Stephens Jr., 2007), our findings about how health shocks shape retirement expectations provide evidence on how health influences the timing of retirement.

For each health shock, we define an analysis sample of people between the ages of 50 and 60 who experience the shock and then track the evolution of retirement expectation outcome variables around the timing of the shock. Our data allow us to study health shocks defined using both subjective and objective health measures. Our primary outcome variables are the self-assessed probabilities of working past 62 and 65; we also study the probability of working past 70. The first two outcomes are especially relevant for policy because they correspond to eligibility ages for Social Security and Medicare.

We begin by studying declines in subjective health. Specifically, we estimate the effects of a sudden decline in overall health status. The event study estimates reveal flat pre-shock trends in expectations before the health status decline and large decreases in the likelihood of working past 62 and 65 after the shock. Our preferred specification indicates that a decline in health status reduces the likelihood of (i) working past 62 by 4.0 percentage points, an 8.9% decrease from the baseline mean of 45%, and (ii) working past 65 by 4.5 percentage points, a 14.1% decrease.

Next, we study declines in objective health, focusing on two main measures. First, we estimate the effects of hospitalizations. The signs of the estimates suggest that hospitalizations induce earlier retirements, but the estimates are not statistically different from zero. Second, we estimate the effects of a newly diagnosed health condition. In the spirit of Hosseini, Kopecky and Zhao (2022), we construct a summary measure equal to the total number of diagnoses a person has and define a new diagnosis shock as an increase in this measure. We find that new diagnoses lead to statistically significant decreases in the likelihood of working past 62 and 65 that amount to 2.5 and 2.6 percentage points, respectively, which are decreases of 6.0% and 9.0% when compared to the means.

We then unpack this result on new diagnoses by studying separate events for each of the eight health conditions in our summary measure. Some, but not all, of these conditions significantly impact retirement expectations. For example, we find clear evidence of decreases in the likelihood of working past older ages after diagnoses of cancer and lung disease. The point estimates for these shocks indicate sizable reductions in the probabilities of working past 65 of 19.3% and 34.2% when compared to the means. In contrast, we find little to no evidence that new diagnoses of diabetes and high blood pressure lead to changes in retirement expectations.

These findings on specific diagnoses could reflect that some conditions are more debilitating while others are more manageable, a takeaway that connects to work on health and income inequality. O'Donnell, Van Doorslaer and Van Ourti (2015) highlight that poor health influences individual income through employment and Hosseini, Kopecky and Zhao (2024) show that differences in health measured using a frailty index account for about a quarter of the variation in lifetime earnings. Our findings help demonstrate which health conditions may be most important for contributing to shorter careers and earnings inequality.

Our paper relates broadly to the large literature that asks how health impacts labor supply at older ages (e.g., McClellan, 1998; Bound et al., 1999; Blau and Gilleskie, 2001;

Coile, 2004; Disney, Emmerson and Wakefield, 2006; García-Gómez et al., 2013; Gustman and Steinmeier, 2018; Blundell et al., 2023; Hsu, Morrill and Pathak, 2024).[1] Several papers provide reviews of research on this longstanding question (e.g., Currie and Madrian, 1999; Coile, 2016; O'Donnell, Van Doorslaer and Van Ourti, 2015; Blundell, French and Tetlow, 2016; French and Jones, 2017). While the methods and data naturally vary across studies, most of the reduced form papers show that declining health leads to declines in employment.

We relate most closely to the few other papers that also use expectations data to advance our understanding of how health shapes the timing of future retirements (Dwyer and Mitchell, 1999; McGarry, 2004; Gupta and Larsen, 2010; Munnell, Sanzenbacher and Rutledge, 2018; Caliendo et al., 2023; Giustinelli and Shapiro, 2024). Two of these papers, Dwyer and Mitchell (1999) and McGarry (2004), stand out as the most similar to ours, as they also use HRS data to study the effect of realized health on retirement expectations.[2] They use the initial survey waves and regression analyses to establish important links between several health measures and retirement expectations for workers. But causal interpretations of the estimates in these earlier studies could be threatened by sample selection, because the expectations variables only exist for workers in the earlier survey waves, and "bad controls" (Angrist and Pischke, 2009), because the regressions condition on outcomes like earnings. McGarry (2004) summarizes the sample selection problem 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."

Our main contribution is to use a quasi-experimental framework and updated data to produce new causal evidence on how health shocks affect retirement expectations. Two features of our study allow us to advance the literature in this way. First, we avoid sample selection problems by using only the more recent waves of HRS data that include well-defined outcomes, even for nonworkers. Specifically, beginning with wave 8, the survey consistently asks everyone—not just workers—about the likelihood that they work past 62 and 65. The newer data thus allow us to track changes in expectations for a crucial group

[1] A separate but related literature studies the relationship between health and retirement in the opposite direction and asks how retirement impacts health (e.g., Coe and Zamarro, 2011; Coe et al., 2012; Eibich, 2015; Gorry, Gorry and Slavov, 2018; Nielsen, 2019; Gorry and Slavov, 2021, 2023).

[2] The other related papers take different approaches. Gupta and Larsen (2010) use data from Denmark to study how the relationship between health and retirement expectations varies when using administrative versus survey data. Munnell, Sanzenbacher and Rutledge (2018) study how various factors, including health, contribute to earlier-than-planned retirements, and Caliendo et al. (2023) study how health and other factors influence the difference between expected and actual retirement ages. Finally, Giustinelli and Shapiro (2024) estimate person-specific "subjective ex ante treatment effects," which compare a person's own estimate of the probability they work in poor health to their estimate of the probability they work in good health.

of people absent from data covering only workers: those who (temporarily or permanently) stop working because of their health. If health shocks cause some people to (i) stop working immediately and (ii) believe they are less likely to work past 62 and 65, then non-workers dropping out of the data after health shocks would bias estimates upwards. Consistent with this idea, when we conduct our analysis on a selected sample of workers, we find no statistically significant evidence that health shocks cause workers to expect to retire earlier. Second, our event study framework allows us to leverage the quasi-random timing of shocks, cleanly isolate the effects of different shocks, and produce graphical evidence that allows for transparent assessments of the key identification assumptions.

In addition to these methodological advancements, we provide estimates for additional outcomes (the probabilities of working past 65 and 70) and for birth cohorts approaching retirement during the more recent period, which is important because the effects of health on retirement are likely to depend on setting. For instance, increasing capacity to work at older ages and changes to the retirement policy landscape (e.g., the decline of defined benefit pensions) may influence how people respond to health shocks.

2 Data

We use data from the Health and Retirement Study (HRS). The HRS is a longitudinal, biennial survey covering Americans over 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 use the RAND HRS Longitudinal File 2020 (v1) dataset (Bugliari et al., 2023), which is a cleaned and streamlined HRS product from the RAND Center for the Study of Aging that includes key information on all survey cohorts and every person interviewed.

These 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 and also detailed information on health. Second, the survey's focus on older people produces a sizable sample of individuals approaching retirement age.

2.1 Outcome Variables: Retirement Expectations

The outcomes in our analysis capture expectations about retirement timing. Specifically, we study self-assessed probabilities of working past older ages. The HRS contains variables for the probabilities of working past 62, 65, and 70. We focus mostly on the probabilities of working past 62 and 65 because these variables are available in all survey waves. They are also especially relevant because they correspond to eligibility ages for important government programs; age 62 is when people become eligible to claim old-age benefits from Social Security, and age 65 is when people become eligible for health insurance through Medicare. In contrast, while we also study the probability of working past 70, this variable is available only from wave 11 onward, which limits its use.

An advantage of these probability-based outcomes is that they capture different types of changes in retirement expectations. For example, consider a person who originally plans to retire at 65 and who then experiences a health shock. One possibility is that this person updates their plans to retire at a different age. Another possibility is that this person continues to plan to retire at 65, but they may be less confident in their ability to do so. Our outcomes should capture each of these important types of changes.

The main disadvantage of these variables used to be that there were many missing values. The underlying survey questions were only asked to workers in earlier waves. However, starting with wave 8 (corresponding to 2006), the questions were asked regardless of work status. Appendix Figure A.1 illustrates this point by plotting the fraction of observations missing values for each of the outcomes by work status across waves. For the two main outcomes, values are mechanically missing for non-workers in waves 1 and 3 through 7. From wave 8 onward, the rates of missing values are similar for workers and non-workers.[3]

2.2 Health Variables Used to Define Health Shocks

We use several health variables to define the health shocks we study and their corresponding analysis samples, which we detail in the next subsection. To study subjective health, we use a categorical variable that captures whether the individual considers their health to be poor, fair, good, very good, or excellent. To study objective health, we use one indicator variable that captures whether the individual has been hospitalized overnight since their last interview, and another set of indicator variables that capture whether a doctor has ever told the individual that they have a specified health condition. We use variables for diagnoses of (i) arthritis, (ii) cancer, (iii) diabetes, (iv) heart attacks or heart disease, (v) high blood pressure, (vi) lung disease, (vii) strokes, and (viii) psychiatric problems.[4]

[3]The data contain another retirement expectations variable: the expected retirement age. While it would be interesting to analyze too, it is still mechanically missing for non-workers.

[4]The survey question for heart disease refers to heart attacks, 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

In our main analysis, we combine these diagnosis indicators into an index by summing them to create one measure that captures the total number of diagnosed health conditions that a person has. We use this measure, which is similar to frailty indices used in gerontology (e.g., Searle et al., 2008) and recently in economics (Hosseini, Kopecky and Zhao, 2022), to study a general shock defined as a new diagnosis. Using a summary measure like this one has advantages. It is a useful way to aggregate information, allowing us to pool data on individuals who all experience a new diagnosis. The tradeoff is that by including all conditions in the index, the sample of people who experience a new diagnosis can develop different health conditions that could generate different responses. Therefore, we also analyze each condition-specific diagnosis separately.

2.3 Analysis Samples

We construct several analysis samples based on the shocks that we study. We begin with the RAND Longitudinal File, which includes everyone interviewed in the HRS, and implement three basic sample restrictions. First, we keep only people who are alive and who responded to the survey. Second, we keep people between ages 50 and 60, who are approaching typical retirement ages and have not yet reached the specific ages referenced for the probability-based retirement expectations outcomes. Third, we only keep observations from survey waves 8 through 15. In these waves, the retirement expectations questions are asked of all people, not just workers, which allows us to sidestep the concern about sample selection.

Next, we define the analysis sample for each health event we consider. We study three main health shocks: (i) declines in health status, (ii) new diagnoses of health conditions, and (iii) hospitalizations. The health status events reflect changes in a subjective measure of health, and we define a health status event to capture new and meaningful declines in overall health. Specifically, we define a person as experiencing a health status event in survey wave w if they report their health status as either fair or poor in wave w and as either excellent, very good, or good in the two prior consecutive waves, $w-1$ and $w-2$. In contrast, diagnoses of health conditions and hospitalizations reflect changes in objective health measures. We define a person as experiencing a new diagnosis event in wave w if they report more diagnosed health conditions in wave w than in wave $w-1$ and $w-2$. Finally, we define a person as experiencing a hospitalization event in wave w if they report

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.

in wave w that they were hospitalized overnight since their last interview and report no such hospitalizations in waves $w-1$ and $w-2$.

These events constitute our main health shocks of interest, but the detailed nature of our data allows us to also study separate events for each of the eight underlying diagnosis variables. For each diagnosis, we define a person as experiencing the event in wave w if they report having the condition in wave w and not in waves $w-1$ and $w-2$.

The analysis samples for each event consist of only people experiencing the shock. If a person experiences the shock more than once, we study the first instance. To focus on the evolution of outcomes around the timing of the shock, we limit the analysis samples to observations of people up to three waves before and up to one wave after the event.

Appendix Table A.1 presents summary statistics for our main shocks. On average, people report roughly a 45% chance of working past 62, a 30% chance of working past 65, and a 12% chance of working past 70. The observation counts reflect some missing data for the outcomes. Between 3.5% and 5.3% of values are missing for the probabilities of working past 62 and 65, whereas the fraction missing is greater for the probability of working past 70 because the underlying survey question is not asked until wave 11. In our main analysis, we use all available observations because our sample sizes are relatively small, but we check the sensitivity of our estimates to missing data by limiting our sample to people who have no missing values for the outcome variables.

3 Identification Strategy

We use an event study framework to analyze the effects of health shocks on retirement expectations. For each shock of interest, we exploit the timing of the shock by limiting our attention to people who experience the shock and comparing the evolution of retirement expectations before and after the shock occurs. One advantage of this approach is that we do not rely on comparisons of people who experience the shocks of interest to those who do not. By focusing on groups in which everyone experiences the same shock, we avoid identification concerns centered on the idea 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. We begin by using a nonparametric event study framework to graphically analyze the evolution of outcomes around the timing of each shock. For each analysis sample, we estimate equations of the following form:

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

where y_{ict} is an outcome variable (such as the probability of working past 65) for individual i in HRS cohort c during survey wave t, λ_{ct} is an HRS-cohort-by-wave fixed effect, 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. Our event window spans three survey waves: before the shock, the wave of the shock, and one wave after the shock. The δ_τs are the coefficients of interest and capture the average difference in the outcome at event time τ relative to the omitted period, $\tau = -1$, the survey wave before the event occurs.

The identification assumption underlying regression model (1) is that conditional on experiencing a given shock, the timing of the shock is uncorrelated with the outcomes. 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 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 a health shock.

The ex-ante plausibility of the identification assumption may vary across shocks. For example, some health conditions may be sudden and unpredictable, whereas others may be expected based on previous conditions or gradually deteriorating health. In general, we can provide an assessment of the validity of the identifying assumption for each shock by analyzing the estimated δ_τs for $\tau < 0$. The patterns of these pre-period estimates provide evidence on whether outcomes were trending before the shock of interest.

We often find little to no evidence of problematic pre-trends in outcomes, but there is some evidence of pre-period trends in some cases. Therefore, to be consistent and account for the possibility of underlying trends when analyzing the effects of different shocks, we move away from the nonparametric regression model and instead use a parametric specification.

Specifically, we estimate

$$y_{ict} = \alpha + \beta\tau + \sum_{\tau=0}^{1} \delta_\tau + \lambda_{ct} + \varepsilon_{ict}. \tag{2}$$

The key difference here is the inclusion of τ, a linear time trend. The coefficient β corresponds to the pre-shock linear trend in the outcome variable. The parameters of interest are the δ_τs, which now capture the average difference in the outcome at event time τ compared to its linear pre-trend. The identification assumption underlying regression model (2) is that, conditional on experiencing a given shock, the timing of the shock is uncorrelated with deviations in the outcome variable from its pre-period 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 the shock is reported. We focus on this point estimate instead of δ_1, which corresponds to the subsequent wave, because our definition of each shock requires that we observe the individual in the data during the wave of the shock (and the two previous waves) but not afterward.

We rely on this baseline parametric specification to quantify magnitudes and assess the statistical significance of the results. Later, we assess the robustness of the estimates to changes in the regression specification and the composition of the analysis samples. Specifically, we report estimates from alternative specifications where we (i) include demographic control variables, (ii) use survey weights, (iii) include wave fixed effects instead of cohort-by-wave fixed effects, (iv) include individual fixed effects, (v) use only the subsample of people who are never missing values for the outcomes, and (vi) use only the subsample of observations that form a balanced panel.

4 The Effects of Health Shocks on Retirement Expectations

In this section, we present our results. For each health shock, we note that the sign of the effect is theoretically ambiguous. On the one hand, health shocks might lead workers to update their expectations towards retiring earlier if they lead to declines in work capacity. On the other hand, health shocks might lead workers to update their expectations towards retiring later if they lead to increased medical expenses or temporary earnings declines. This ambiguity is particularly important to emphasize when looking at retirement expectations rather than contemporaneous labor supply responses because some health shocks could induce immediate exits from the labor market, but people then might expect to work longer once they return to work.

4.1 The Effects of Health Status Declines, New Diagnoses, and Hospitalizations

We begin with an analysis of our three main health shocks: (i) a decline in subjective health status, (ii) a new diagnosis of a health condition, and (iii) a hospitalization. Figure 1 displays the nonparametric event study results. Each graph corresponds to one of the shocks and one of the main outcomes, either the probability of working past 62 or the probability of working past 65. The point estimates in the graphs are the δ_τ coefficients from estimating equation (1).

First, consider panels (a) and (b), which present the results for a decline in subjective health status. The pre-shock period point estimates indicate a lack of significant trends in the probabilities of working past key ages before a decline in health status, consistent with the interpretation that the underlying causes of the declines in health may occur suddenly. The post-shock period estimates show a sharp and clear decline in the average probability of working past 62 and 65. These two panels indicate that declines in subjective health status cause workers to update their expectations towards retiring earlier.

Next, consider panels (c) and (d), which present the results for a new diagnosis of a health condition. Here, we also find pre-shock period point estimates that are not statistically distinguishable from zero, although there is more of a visible downward trend before the shock. Still, the post-shock period estimates reveal a sharp and greater decline in retirement expectations after people are diagnosed with a new health condition. Finally, consider panels (e) and (f), which present the results for a hospitalization. While there is perhaps some visual evidence of a decline in the likelihood of working past 65, the two panels show point estimates for the pre- and post-shock periods that are not statistically different from zero at the 5-percent level and reveal mostly flat trends in outcomes.

To quantify these findings, Table 1 presents the results from the parametric event study analysis. The parametric approach summarizes the effects by comparing the evolution of the outcome variables to their pre-period trends (which were relatively flat for declines in health status and hospitalizations but were more noticeable for new diagnoses). Panel A of the table presents results for each of the main shocks. The first four columns present results for the probability of working past 62. Column (1) displays the point estimate for δ_0 from estimating equation (2) for this outcome, and the subsequent columns display the dependent variable mean in the survey wave preceding the shock, the number of clusters (unique individuals), and the number of observations in the analysis sample. The next four columns present the corresponding results for the probability of working past 65.

On average, a decline in health status causes a statistically significant 4.0 percentage point decline in the self-assessed probability of working past 62 and a statistically significant 4.5 percentage point decline in the probability of working past 65. These declines translate to meaningful reductions in the likelihood of working by 8.9% and 14.1%, respectively, compared to the baseline means. Similarly, new diagnoses of health conditions lead to statistically significant declines in the probabilities of working past older ages. The point estimates indicate a 2.5 percentage point decline in the likelihood of working past 62, which is a 6.0% decline compared to the baseline mean, and a 2.6 percentage point decline in the likelihood of working past 65, which is a 9.0% decline compared to the baseline mean.

In contrast, we find no evidence that hospitalizations lead to changes in the probabilities of working past 62 or 65. The estimates are smaller in magnitude than those for the other health shocks, and they are not statistically significant. We also find little to no evidence that any of the main health shocks impact the likelihood of working past 70 (see Appendix Figure A.2 and Appendix Table A.2), which could be due to the lower likelihoods of working past 70 to begin with; the means for this outcome range from 11% to 12%. However, recall that our analysis of working past 70 is also more limited due to smaller sample sizes because it is only available for waves 11 onward.[5]

4.2 The Effects of Specific Health Conditions

Next, we unpack the findings on new diagnoses by studying the underlying conditions that make up the summary measure separately. For each specific condition, we analyze the impact of the shock on the probabilities of working past 62, 65, and 70. This exercise allows us to provide evidence on which diagnoses could be driving the overall effect. A disadvantage is that the sample sizes become (sometimes substantially) smaller.

We present graphs of the nonparametric event study results in Figure 2 for the probability of working past 65. We present results for 62 and 70 in Appendix Figures A.3 and A.4. Each graph corresponds to a different diagnosis. There are almost no pre-shock period estimates that are statistically different from zero. Indeed, for many of the diagnoses, the pattern of the pre-shock period estimates reveals encouragingly flat trends. However, there are a few shocks for which the pattern of the estimates suggests that the outcomes were trending downward even before the shock. For example, the probabilities of working past 62 and 65 appear to exhibit downward trends before heart disease diagnoses. Recall that the parametric event

[5]This data limitation implies that some person-wave observations become unusable because the outcome does not exist in a relevant wave. For instance, consider someone who experiences a health shock in wave 10. We observe the probability of working past 70 for this person in the wave after their shock (wave 11), but not in the wave of the shock or earlier waves.

study regressions will account for these trends when quantifying the overall effects.

Some health conditions impact expectations about future work, whereas others appear to have little to no effect. For example, the graphs for arthritis, cancer, and lung disease in Figure 2 and Appendix Figure A.3 show visually apparent declines in the likelihood of working past older ages after the diagnoses. In contrast, the graphs for diabetes, heart disease, high blood pressure, and psychiatric problems show point estimates that appear to evolve smoothly. The results for strokes are perhaps less clear; graph (h) in Figure 2 shows a marked decline in the probability of working past 65 in the wave of the diagnosis, but graph (h) in Appendix Figure A.3 shows less evidence of a change.

Panel B of Table 1 presents the results from the parametric event study regressions. Consistent with the graphs, we find statistically significant declines in the likelihood of working past 62 for arthritis, cancer, and lung disease, and we find statistically significant declines in the likelihood of working past 65 for cancer, lung disease, and strokes. In contrast, we find no statistically significant evidence that diagnoses of heart disease, diabetes, high blood pressure, or psychiatric problems induce immediate changes in retirement expectations. Some of the estimates that are not statistically significant are relatively small (like the point estimate for working past 65 after a diagnosis of high blood pressure). In contrast, others are more sizable (like the estimates for psychiatric problems) but are not precisely estimated.

The magnitudes of the statistically significant estimates are large. The point estimate for the decline in the likelihood of working past 62 after an arthritis diagnosis is 4.7 percentage points, which corresponds to a 10.2% decrease when compared to the baseline mean. The results for cancer diagnoses indicate 7.2 and 5.8 percentage point declines in the likelihood of working past 62 and 65, which are 16.4% and 19.3% decreases when compared to the baseline means of 44% and 30%, respectively. Similarly, the analogous point estimates for lung disease indicate 26.4% and 34.2% decreases. Strikingly, the point estimate for the likelihood of working past 65 after a stroke corresponds to a 54.1% decline (although we note the sample size for this shock is small).

One qualitative takeaway from these estimates is that the effects of adverse health shocks can be nuanced. Some diagnoses clearly impact retirement expectations, whereas the evidence for others is less clear. For instance, our findings highlight how cancer and lung disease are two conditions that lead to major changes in expectations about working longer. In contrast, the evidence for high blood pressure and diabetes suggests a lack of an effect. Of course, different health conditions may impact current work capacity, future work capacity, medical expenses, savings, and other factors differently. Some conditions, like cancer and lung disease, might greatly reduce work capacity. Others, like diabetes, might have a less severe impact on work capacity while meaningfully increasing future medical expenses.

4.3 Robustness Checks

We conduct several robustness checks. For simplicity, we focus on our main health shocks and the probabilities of working past 62 and 65, although we also assess the robustness of our estimates for the specific health conditions and for the probability of working past 70 in the appendix (see Appendix Tables A.3, A.4, A.5, and A.6).

Table 2 presents the main robustness results. The columns of the table correspond to different health shocks and outcome variables. The rows of the table correspond to different robustness checks and indicate how the robustness specification differs from the baseline specification. Panel A reproduces the baseline estimates for ease of comparison.

First, we assess the robustness of our results to standard regression specification checks in rows B through E. Row B includes control variables in the regression specification. Specifically, we include a variable for age and indicator variables for being female, for being white, for having attended college, and for being married. This change has little effect on our estimates. The magnitudes of the estimates are similar, and the estimates that are statistically significant in the baseline specification remain statistically significant. Row C uses person-level analysis sample weights. The magnitudes of the point estimates are similar to our baseline specification, but most of the estimates are no longer statistically significant because the standard errors are larger; importantly, the survey weights are not yet available for the most recent survey wave of the data, which results in smaller sample sizes and less precise estimates.

Row D includes wave fixed effects instead of cohort-by-wave fixed effects. This change results in (i) estimates for health status declines that are larger than the baseline estimates and that are statistically significant at the 1-percent level, (ii) estimates for new diagnoses that are similar to the baseline estimates, and (iii) estimates for hospitalizations that are meaningfully larger than their baseline level. Row E includes individual fixed effects. To include these fixed effects, we need to address the identification problem that arises because of the collinearity between individual fixed effects (which subsume time-of-shock fixed effects), time fixed effects, and time-relative-to-shock fixed effects, our parameters of interest (Dobkin et al., 2018). To address this issue and obtain identification, we include individual fixed effects and two-wave fixed effects instead of wave fixed effects, which effectively assumes that the macroeconomic conditions controlled for by time fixed effects evolve slowly. The results for health status declines and new diagnoses are similar to their baseline estimates. In contrast, the estimates for hospitalizations from this specification are larger in magnitude than their baseline counterparts, and they are statistically significant.

Second, we assess the robustness of our results to sample construction choices in rows F and G. Because our sample sizes are not especially large, our baseline specification includes as much data as possible; we include all usable person-wave observations in the analysis samples between three waves before the shock and one wave after the shock. Row F takes a different approach by limiting the people in the analysis sample to only those who experience the shock of interest and have no missing values for any outcomes. The takeaways do not change when we look at this subsample. Row G takes an even more restrictive approach by focusing on a shorter, balanced panel. This analysis is limited to subsamples of observations that (i) are for people with non-missing values for all outcomes and (ii) are for survey waves corresponding to event times between $\tau = -2$ and $\tau = 0$, which are the waves for which everyone appears in the data by construction. The advantage of the balanced panels are that the sample composition is the same at all relative time periods. The disadvantage is that sample sizes are smaller. The magnitudes of the estimates are mostly similar to their baseline counterparts, but all of the standard errors are larger.

Overall, these results bolster our confidence in our empirical approach and the baseline estimates. The size and statistical significance of the estimates for health status declines and new diagnoses are reasonably stable across specifications and samples. The estimates for hospitalizations are more sensitive to specification choices; our baseline specification indicates a lack of evidence of responses, but some of the alternative specifications produce estimates that are larger in magnitude and statistically significant.

4.4 Assessing the Importance of Avoiding Sample Selection Issues

An important feature of our study that supports causal interpretations of the estimates is the more recent survey data. The earlier waves of the HRS contain retirement expectations outcomes only for workers. In contrast, we leverage updated data that contain these outcomes even for non-workers. We take two approaches to assess the importance of overcoming the sample selection issues that the earlier waves of the data presented.

First, we redo our analysis using only the earlier waves instead of only the later waves. Specifically, in our main analysis, one of our initial sample construction steps involves limiting the underlying data to only waves 8 through 15. Here, we limit the underlying data to only waves 1 through 7. We then define the main health shocks of interest, construct the analysis samples, and estimate regressions in the same way as our main analysis.

Appendix Figure A.5 presents the nonparametric event study results, and panel A of Appendix Table A.7 presents the parametric results. These results from the earlier sample waves starkly contrast with our main results. Overall, the patterns of the nonparametric estimates indicate either a lack of evidence of a response or, if anything, an increase in the likelihood of working at later ages. Indeed, the parametric estimates indicate no statistically significant evidence of an effect for health status declines or hospitalizations but strong evidence of an increase in the likelihood of working longer after a new diagnosis.

Second, we redo our analysis using the later survey waves but restrict the sample to include only workers. Specifically, we use the same underlying samples of people in our main analysis, but we keep only the person-wave observations for which the person reports working full-time or part-time. In contrast to the first approach, which generates a different sample of people (those who experience the health shocks of interest in earlier waves), this second approach has the advantage of keeping the people in our sample fixed while still allowing us to look at the impact that selection would have on our estimates if the survey did not begin recording responses for non-workers.

Appendix Figure A.6 presents the nonparametric event study results, and panel B of Appendix Table A.7 presents the parametric results. When we use this selected sample, we find no evidence that health shocks impact retirement expectations.

Overall, the results from both approaches are consistent with the idea that people who drop out of the selected sample because they (temporarily or permanently) stop working also believe they are less likely to work past 62 and 65. Avoiding sample selection is important for assessing the impact of health on retirement expectations, as we would have reached different conclusions if we could not use the updated data that cover workers and non-workers.

5 Discussion and Conclusion

This paper uses an event study framework and data on retirement expectations to provide new causal evidence on how health shocks affect retirement timing. The results are relevant for policymakers and practitioners concerned with assessing the retirement income security of older Americans.

First, a better understanding of how health shocks influence retirement timing can help inform public policy. Social Security and Medicare are two of the largest public programs in the U.S. and provide access to old-age benefits and health insurance starting at 62 and 65.

Therefore, our findings show how declines in health status and new diagnoses of health conditions can cause people to retire before becoming eligible for these benefits, which highlights a potentially important role for the Disability Insurance program.

Second, by studying expectations, our results are informative for retirement plan design and the administration of tax-advantaged retirement savings accounts. The expected retirement date is a key input into optimal strategies within a savings plan. Indeed, target date funds explicitly anchor investments to expected retirement dates and are an increasingly important part of savings portfolios (Shoven and Walton, 2021). Our results highlight the importance of health in shaping retirement timing uncertainty and suggest that people may value flexibility in savings schemes like the ability to update their elected expected retirement age. Additionally, our analysis points to specific shocks that may require changes to savings plans.

Finally, our analysis advances our understanding of how people manage uncertain events in preparation for retirement. People may modify their plans by, for instance, working longer, changing the amount they work while working, saving more, or consuming less. While Bronshtein et al. (2019) highlight the power of working longer for retirement readiness, our analysis implies that extending employment may not always be an available option. In these cases, savings responses may be critical. However, evidence shows that retirement savings often evolve passively (Madrian and Shea, 2001; Chetty et al., 2014) and that retirement savings decisions are strongly linked to labor supply decisions (García-Miralles and Leganza, 2024). This passivity raises concerns that many may not optimally update savings in response to events that change expected retirement dates. While assessing optimal savings responses to health shocks is beyond the goal of this paper, studying changes in retirement expectations is an essential first step in understanding whether people should adjust savings, the extent to which they should do so, and how financially secure they will be in retirement.

References

Angrist, Joshua D, and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton University Press.

Benitez-Silva, Hugo, and Debra S. Dwyer. 2005. "The Rationality of Retirement Expectations and the Role of New Information." *Review of Economics and Statistics*, 87(3): 587–592.

Blau, David M, and Donna B Gilleskie. 2001. "Retiree Health Insurance and the Labor Force Behavior of Older Men in the 1990s." *Review of Economics and Statistics*, 83(1): 64–80.

Blundell, Richard, Eric French, and Gemma Tetlow. 2016. "Retirement Incentives and Labor Supply." In *Handbook of the Economics of Population Aging.* Vol. 1, 457–566. Elsevier.

Blundell, Richard, Jack Britton, Monica Costa Dias, and Eric French. 2023. "The Impact of Health on Labor Supply Near Retirement." *Journal of Human Resources*, 58(1): 282–334.

Bound, John, Michael Schoenbaum, Todd R. Stinebrickner, and Timothy Waidmann. 1999. "The Dynamic Effects of Health on the Labor Force Transitions of Older Workers." *Labour Economics*, 6(2): 179–202.

Bronshtein, Gila, Jason Scott, John B Shoven, and Sita Nataraj Slavov. 2019. "The power of working longer." *Journal of Pension Economics & Finance*, 18(4): 623–644.

Bugliari, Delia, Joanna Carroll, Orla Hayden, Jessica Hayes, Michael D. Hurd, Stephen Lee, R. Main, C.M. McCullough, E. Meijer, P. Pantoja, et al. 2023. "RAND HRS Longitudinal File 2020 (V1) Documentation." *RAND Center for the Study of Aging.*

Caliendo, Frank N, Maria Casanova, Aspen Gorry, and Sita Slavov. 2023. "Retirement Timing Uncertainty: Empirical Evidence and Quantitative Evaluation." *Review of Economic Dynamics*, 51: 226–266.

Chetty, Raj, John N Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen. 2014. "Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark." *Quarterly Journal of Economics*, 129(3): 1141–1219.

Coe, Norma B, and Gema Zamarro. 2011. "Retirement Effects on Health in Europe." *Journal of Health Economics*, 30(1): 77–86.

Coe, Norma B, Hans-Martin von Gaudecker, Maarten Lindeboom, and Jürgen Maurer. 2012. "The Effect of Retirement on Cognitive Functioning." *Health Economics*, 21(8): 913–927.

Coile, Courtney. 2004. "Health Shocks and Couples' Labor Supply Decisions."

Coile, Courtney C. 2016. "Economic Determinants of Workers' Retirement Decisions." *A Collection of Reviews on Savings and Wealth Accumulation*, 271–296.

Currie, Janet, and Brigitte C Madrian. 1999. "Health, Health Insurance and the Labor

Market." *Handbook of Labor Economics*, 3: 3309–3416.

Disney, Richard, Carl Emmerson, and Matthew Wakefield. 2006. "Ill Health and Retirement in Britain: A Panel Data-Based Analysis." *Journal of Health Economics*, 25(4): 621–649.

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo. 2018. "The Economic Consequences of Hospital Admissions." *American Economic Review*, 108(2): 308–352.

Dwyer, Debra Sabatini, and Olivia S. Mitchell. 1999. "Health Problems as Determinants of Retirement: Are Self-Rated Measures Endogenous?" *Journal of Health Economics*, 18(2): 173–193.

Eibich, Peter. 2015. "Understanding the Effect of Retirement on Health: Mechanisms and Heterogeneity." *Journal of Health Economics*, 43: 1–12.

French, Eric, and John Bailey Jones. 2017. "Health, Health Insurance, and Retirement: A Survey." *Annual Review of Economics*, 9(1): 383–409.

García-Gómez, Pilar, Hans Van Kippersluis, Owen O'Donnell, and Eddy Van Doorslaer. 2013. "Long-term and Spillover Effects of Health Shocks on Employment and Income." *Journal of Human Resources*, 48(4): 873–909.

García-Miralles, Esteban, and Jonathan M Leganza. 2024. "Public pensions and private savings." *American Economic Journal: Economic Policy*, 16(2): 366–405.

Giustinelli, Pamela, and Matthew D. Shapiro. 2024. "SeaTE: Subjective Ex Ante Treatment Effect of Health on Retirement." *American Economic Journal: Applied Economics*, 16(2): 278–317.

Gorry, Aspen, Devon Gorry, and Sita Nataraj Slavov. 2018. "Does Retirement Improve Health and Life Satisfaction?" *Health Economics*, 27(12): 2067–2086.

Gorry, Devon, and Sita Nataraj Slavov. 2021. "The Effect of Retirement on Health Biomarkers." *Economics & Human Biology*, 40: 100949.

Gorry, Devon, and Sita Nataraj Slavov. 2023. "The Effect of Retirement on Health Behaviors." *Health Economics*, 32(10): 2234–2259.

Gupta, Nabanita Datta, and Mona Larsen. 2010. "The Impact of Health on Individual Retirement Plans: Self-Reported versus Diagnostic Measures." *Health Economics*, 19(7): 792–813.

Gustman, Alan L., and Thomas L. Steinmeier. 2018. "The Role of Health in Retirement." In *Transitions through the Labor Market: Work, Occupation, Earnings and Retirement.* 229–297. Emerald Publishing Limited.

Haider, Steven J., and Melvin Stephens Jr. 2007. "Is There a Retirement-Consumption Puzzle? Evidence Using Subjective Retirement Expectations." *Review of Economics and Statistics*, 89(2): 247–264.

Hosseini, Roozbeh, Karen A Kopecky, and Kai Zhao. 2022. "The Evolution of Health Over the Life Cycle." *Review of Economic Dynamics*, 45: 237–263.

Hosseini, Roozbeh, Karen A. Kopecky, and Kai Zhao. 2024. "How Important is Health Inequality for Lifetime Earnings Inequality?"

Hsu, De Fen, Melinda Morrill, and Aditi Pathak. 2024. "Health and retirement: Heterogeneity in the responsiveness to pension incentives." *Economics Letters*, 238: 111677.

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

McClellan, Mark. 1998. "Health Events, Health Insurance, and Labor Supply: Evidence from the Health and Retirement Survey." *Frontiers in the Economics of Aging*, 301–346.

McGarry, Kathleen. 2004. "Health and Retirement: Do Changes in Health Affect Retirement Expectations?" *Journal of Human Resources*, 39(3): 624–648.

Munnell, Alicia H., Geoffrey T. Sanzenbacher, and Matthew S. Rutledge. 2018. "What Causes Workers to Retire Before they Plan?" *Journal of Retirement*, 6(2): 35–52.

Nielsen, Nick Fabrin. 2019. "Sick of Retirement?" *Journal of Health Economics*, 65: 133–152.

O'Donnell, Owen, Eddy Van Doorslaer, and Tom Van Ourti. 2015. "Health and Inequality." In *Handbook of Income Distribution*. Vol. 2, 1419–1533. Elsevier.

Searle, Samuel D, Arnold Mitnitski, Evelyne A Gahbauer, Thomas M Gill, and Kenneth Rockwood. 2008. "A standard procedure for creating a frailty index." *BMC Geriatrics*, 8: 1–10.

Shoven, John B, and Daniel B Walton. 2021. "An Analysis of the Performance of Target Date Funds." *The Journal of Retirement*.

Figure 1: Effects of Main Health Shocks on Retirement Expectations

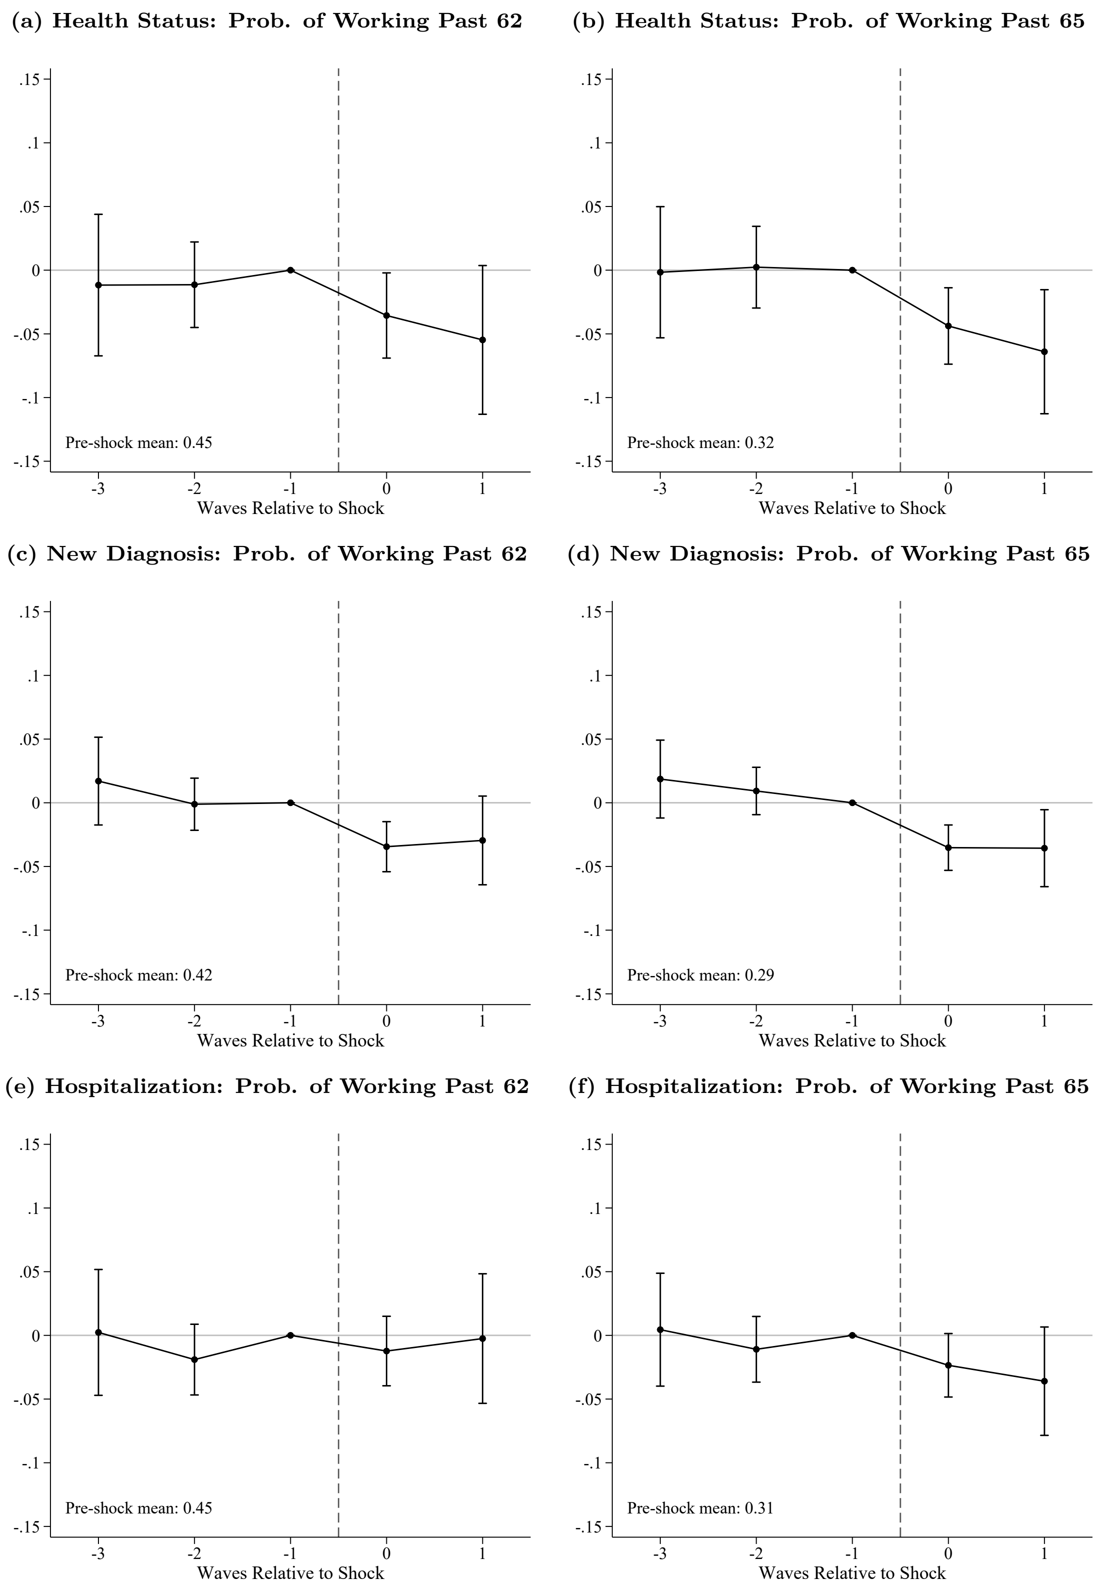

Notes: This figure presents the nonparametric event study estimates for our three main health shocks (declines in health status, new diagnoses, and hospitalizations) and our two main outcome variables (the probabilities of working past 62 and 65). The graphs plot point estimates and confidence intervals for the δ_τ coefficients in equation (1).

Figure 2: Effects of Specific Health Conditions on the Probability of Working Past 65

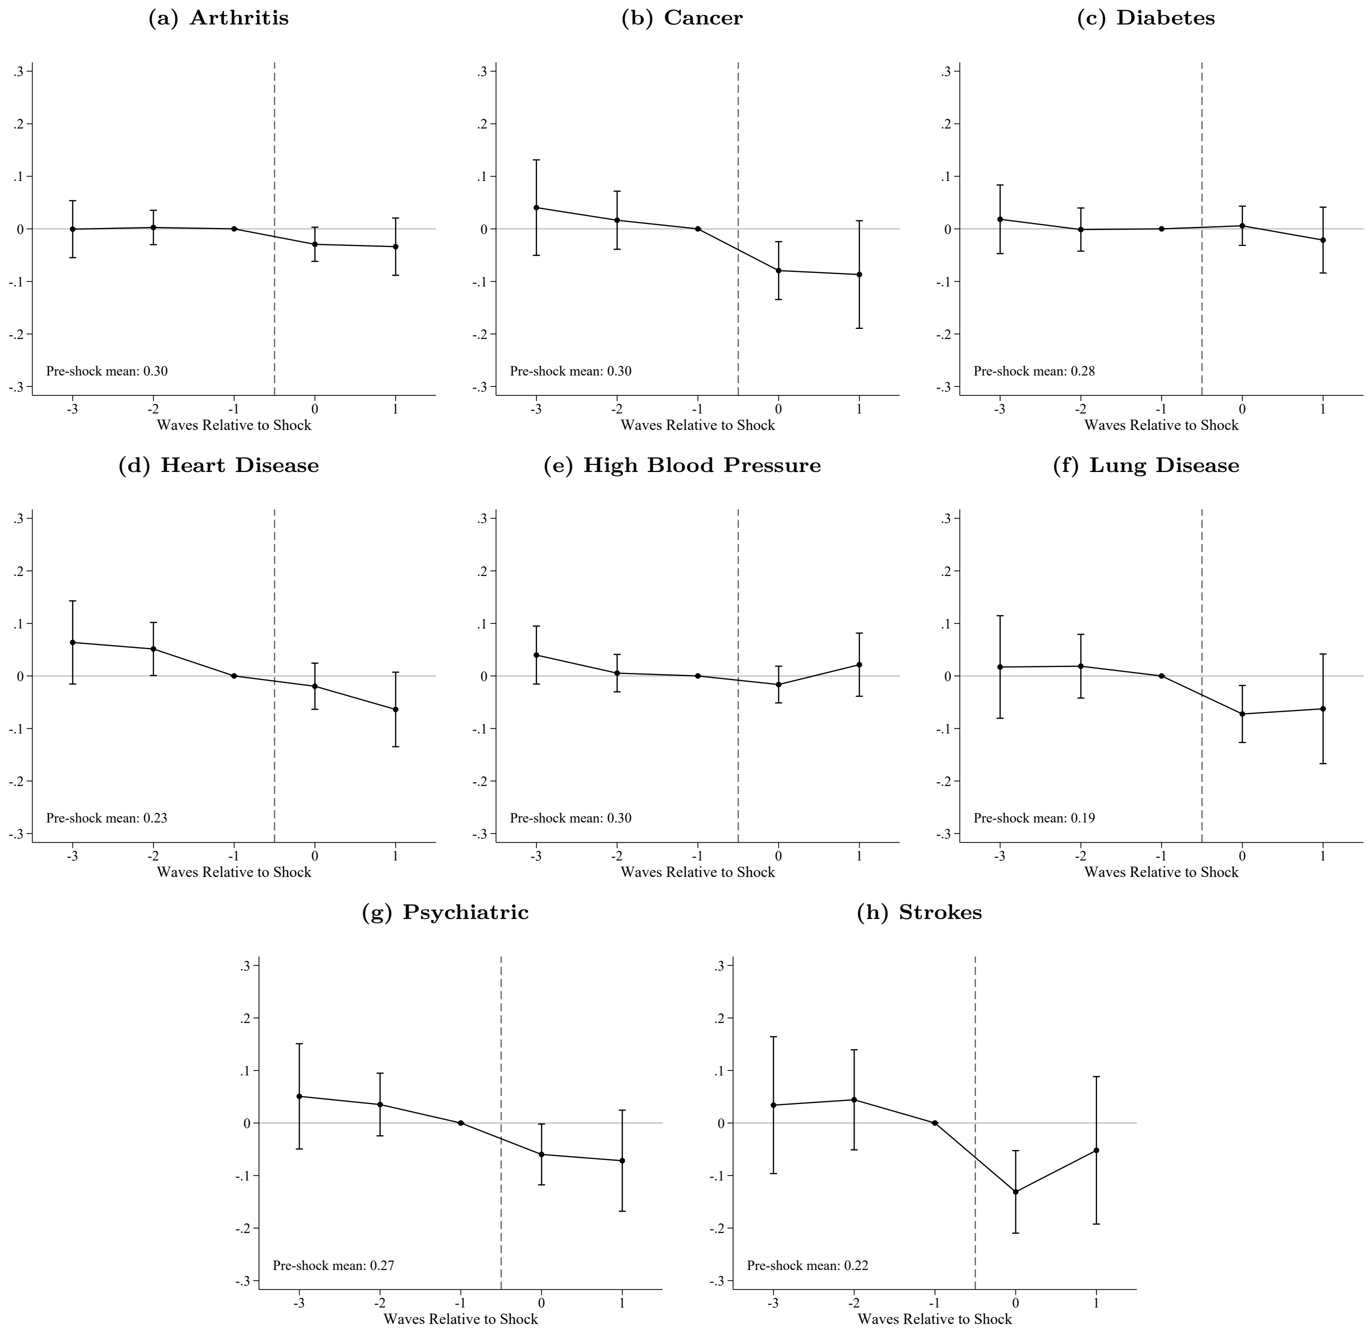

Notes: This figure presents the nonparametric event study estimates for each of the eight specific health conditions that we study and for the probability of working past 65. The graphs plot point estimates and confidence intervals for the δ_τ coefficients in equation (1).

Table 1: The Effect of Health Shocks on Retirement Expectations

	Probability of Working Past 62				Probability of Working Past 65			
	Estimate (1)	Mean (2)	Clusters (3)	Obs. (4)	Estimate (5)	Mean (6)	Clusters (7)	Obs. (8)
A: Main Shocks								
Health Status Decline	-0.040**	0.45	855	3,153	-0.045**	0.32	854	3,116
	(0.019)				(0.018)			
New Diagnosis	-0.025**	0.42	2,597	9,416	-0.026**	0.29	2,598	9,328
	(0.012)				(0.011)			
Hospitalization	-0.008	0.45	1,210	4,463	-0.019	0.31	1,210	4,431
	(0.016)				(0.015)			
B: Specific Diagnoses								
Arthritis	-0.047**	0.46	815	2,988	-0.030	0.30	816	2,959
	(0.021)				(0.019)			
Cancer	-0.072**	0.44	233	854	-0.058*	0.30	233	849
	(0.034)				(0.032)			
Diabetes	0.020	0.41	521	1,932	0.017	0.28	521	1,905
	(0.023)				(0.021)			
Heart Disease	0.046	0.35	351	1,275	0.009	0.23	351	1,266
	(0.030)				(0.029)			
High Blood Pressure	-0.022	0.45	692	2,496	0.006	0.30	692	2,470
	(0.022)				(0.020)			
Lung Disease	-0.074**	0.28	208	751	-0.065**	0.19	208	744
	(0.037)				(0.030)			
Psychiatric	-0.041	0.40	316	1,144	-0.036	0.27	316	1,139
	(0.038)				(0.034)			
Stroke	-0.021	0.29	127	436	-0.119***	0.22	127	431
	(0.051)				(0.045)			

Notes: This table reports the parametric event study estimates for the probabilities of working past 62 and 65. Panel A presents estimates for our main health shocks of interest (declines in health status, new diagnoses, and hospitalizations). Panel B presents estimates for specific health conditions. Column (1) present estimates of δ_0 from estimating equation (2). Standard errors clustered at the individual level are in parentheses. Column (2) presents the dependent variable mean in the survey wave before the health shock. Column (3) presents the number of clusters. Column (4) presents the number of observations.
$^{***}\, p < 0.01$, $^{**}\, p < 0.05$, $^{*}\, p < 0.1$

Table 2: Robustness of Main Estimates

		Health Status		New Diagnosis		Hospitalization	
		Age 62 (1)	Age 65 (2)	Age 62 (3)	Age 65 (4)	Age 62 (5)	Age 65 (6)
A.	Baseline Specification	-0.040**	-0.045**	-0.025**	-0.026**	-0.008	-0.019
		(0.019)	(0.018)	(0.012)	(0.011)	(0.016)	(0.016)
B.	Include Control Variables	-0.034*	-0.039**	-0.024**	-0.026**	-0.011	-0.022
		(0.019)	(0.018)	(0.011)	(0.011)	(0.016)	(0.016)
C.	Use Survey Weights	-0.047	-0.045*	-0.022	-0.021	0.006	-0.006
		(0.029)	(0.026)	(0.015)	(0.014)	(0.021)	(0.021)
D.	Include Wave Fixed Effects	-0.052***	-0.050***	-0.025**	-0.021**	-0.022	-0.030**
		(0.019)	(0.018)	(0.012)	(0.011)	(0.016)	(0.016)
E.	Include Individual Fixed Effects	-0.040**	-0.034**	-0.033***	-0.022**	-0.038***	-0.038***
		(0.017)	(0.015)	(0.009)	(0.009)	(0.013)	(0.013)
F.	Subsample: No Missing Values	-0.035*	-0.041**	-0.021*	-0.021*	-0.015	-0.017
		(0.021)	(0.020)	(0.012)	(0.011)	(0.016)	(0.016)
G.	Subsample: Balanced Panel	-0.046*	-0.039	-0.029**	-0.020	-0.023	-0.017
		(0.025)	(0.024)	(0.015)	(0.014)	(0.019)	(0.019)

Notes: This table reports results from assessing the robustness of our main estimates, which correspond to our main health shocks (declining health status, new diagnoses, and hospitalizations) and our main outcome variables (the probabilities of working past 62 and 65). The estimates presented are the δ_0s from estimating equation (2). Each column corresponds to a different health shock and outcome variable. Each row corresponds to a different robustness check. Row A reproduces the baseline estimates for ease of comparison. Row B adds control variables to the regressions. Row C uses survey weights when estimating the regressions. Row D uses wave fixed effects instead of HRS-cohort-by-wave fixed effects. Row E uses individual fixed effects and two-wave fixed effects. Row F focuses on a subsample of people who never have missing values for any of the outcome variables. Row G focuses on the balanced panel of observations of individuals without missing values in survey waves corresponding to event times between $\tau = -2$ and $\tau = 0$.
$^{***}\, p < 0.01$, $^{**}\, p < 0.05$, $^{*}\, p < 0.1$

Appendix A Additional Figures and Tables

Table A.1: Summary Statistics in the Survey Wave Before Each Main Health Shock

	Health Status		New Diagnosis		Hospitalization	
	Mean (1)	Obs. (2)	Mean (3)	Obs. (4)	Mean (5)	Obs. (6)
Age	55.69	862	55.56	2,623	55.79	1,222
Male	0.43	862	0.39	2,623	0.42	1,222
White	0.56	862	0.59	2,623	0.64	1,222
Married	0.64	862	0.66	2,623	0.67	1,222
Any College	0.47	862	0.51	2,623	0.53	1,222
Probability of Working Past 62	0.45	819	0.42	2,520	0.45	1,183
Probability of Working Past 65	0.32	816	0.29	2,508	0.31	1,179
Probability of Working Past 70	0.12	635	0.11	1,923	0.12	843

Notes: This table reports summary statistics for the samples of people who experience our main health shocks of interest. For each analysis sample, the underlying data consists of observations of people between ages 50 and 60 in survey waves 8 through 15 who experience the corresponding health shock. The sample means and observations presented in the table are for the survey wave before the shock.

Figure A.1: Missing Data by Work Status Across Survey Waves

(a) Prob. of Working Past 62 **(b) Prob. of Working Past 65** **(c) Prob. of Working Past 70**

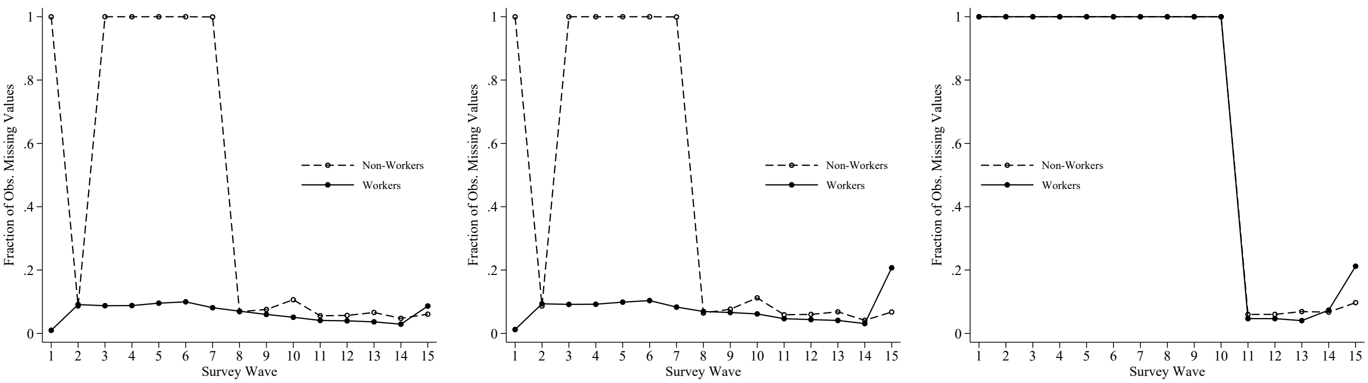

Notes: This figure illustrates the missing data issues with the earlier sample waves. The underlying sample consists of people in the HRS data between ages 50 and 60. Each graph corresponds to a different outcome variable of interest and plots the fraction of observations that have missing values for that outcome variable across survey waves for workers and non-workers separately. Graphs (a) and (b) show how the probabilities of working past 62 and 65 are missing for all non-workers in the earlier survey waves (except for wave 2); the rates of missing data for non-workers then decline to rates similar for workers starting with wave 8, which is when the underlying survey questions started being consistently asked to workers and non-workers. Graph (c) shows how the probability of working past 70 is missing for workers and non-workers before wave 11, before the outcome variable existed.

Figure A.2: Effects of Main Shocks on the Probability of Working Past 70

(a) Health Status Decline

(b) New Diagnosis

(c) Hospitalization

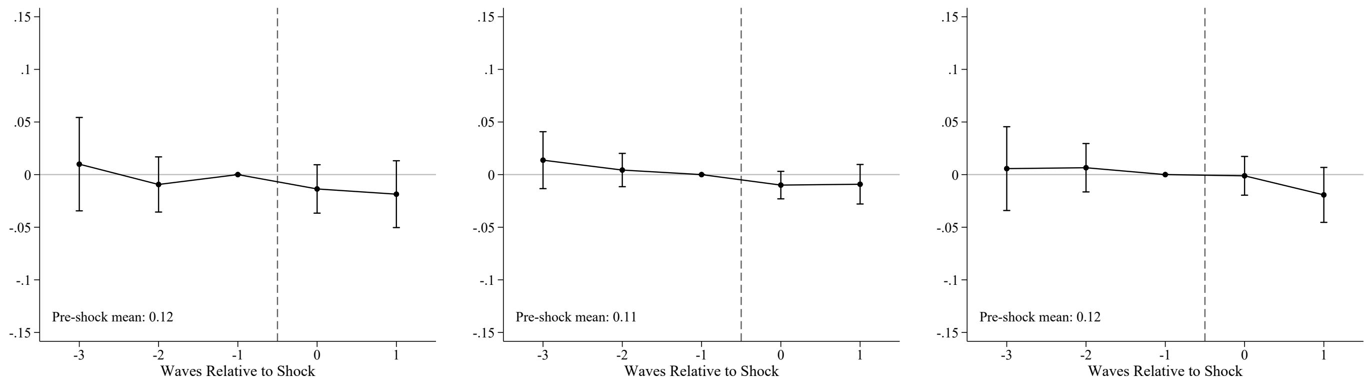

Notes: This figure presents the nonparametric event study estimates for our three main health shocks (declines in health status, new diagnoses, and hospitalizations) and for the probability of working past 70. The graphs plot point estimates and confidence intervals for the δ_τ coefficients in equation (1).

Figure A.3: Effects of Specific Health Conditions on the Probability of Working Past 62

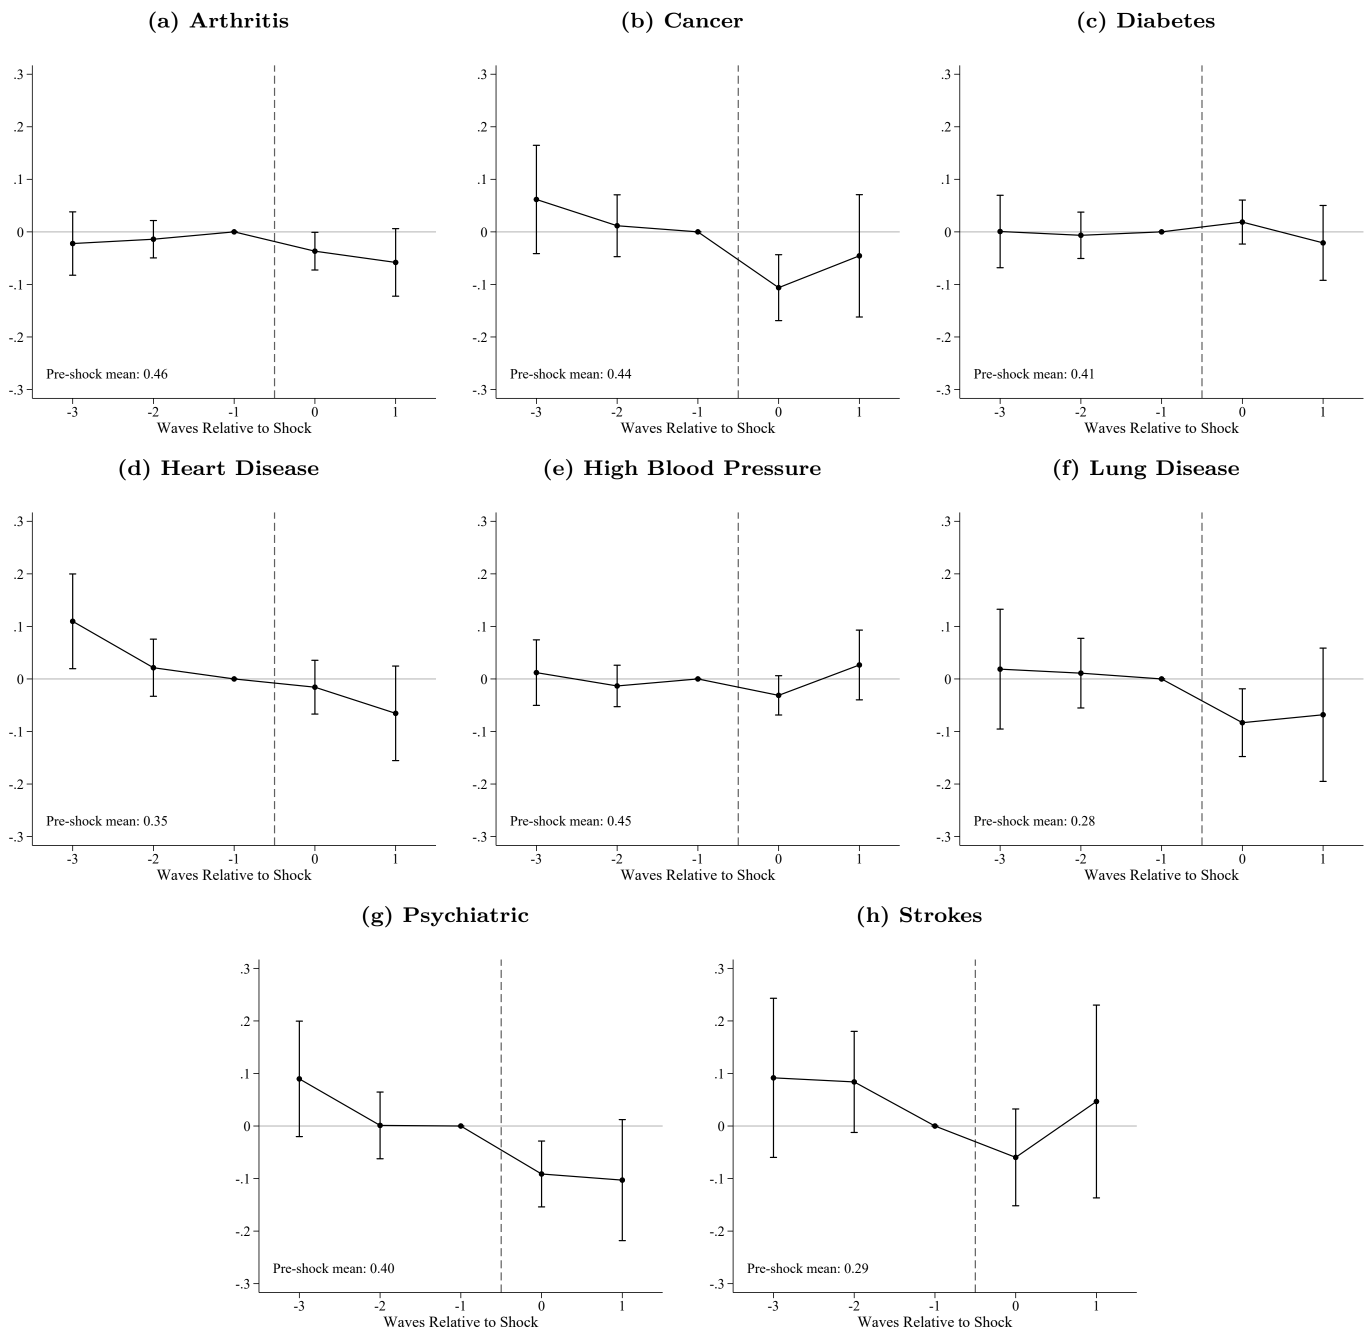

Notes: This figure presents the nonparametric event study estimates for each of the eight specific health conditions that we study and for the probability of working past 62. The graphs plot point estimates and confidence intervals for the δ_τ coefficients in equation (1).

Figure A.4: Effects of Specific Health Conditions on the Probability of Working Past 70

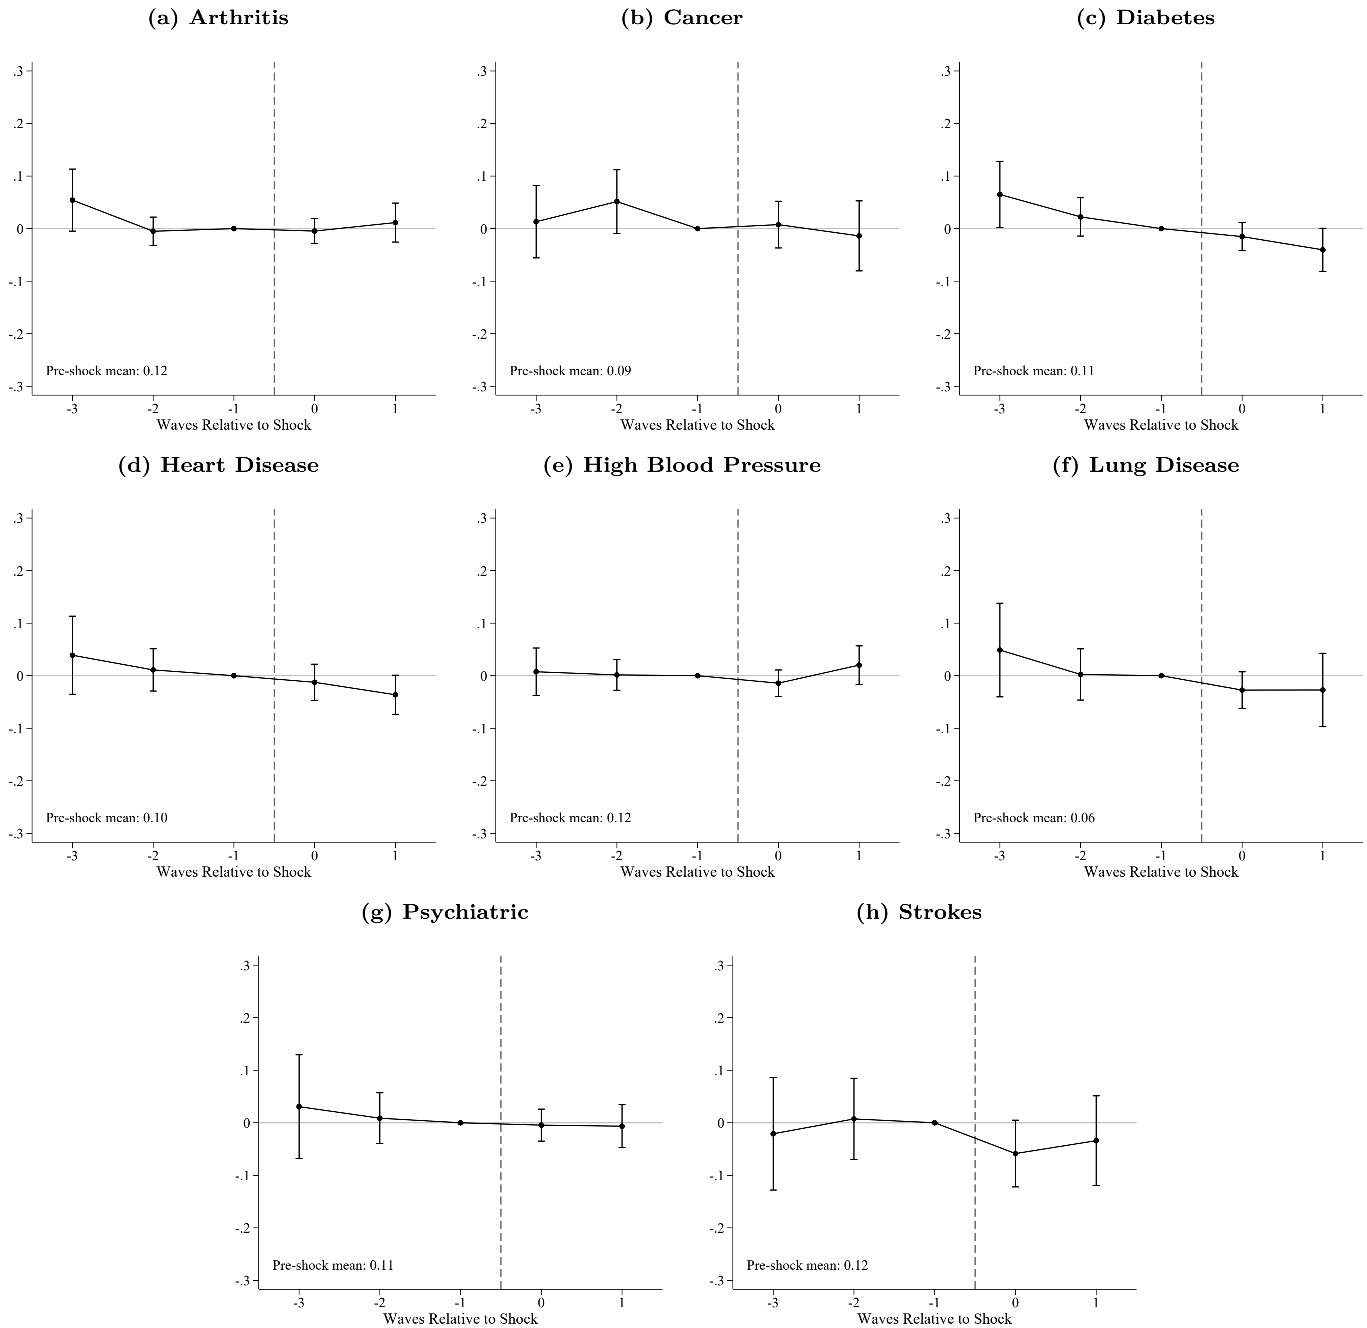

Notes: This figure presents the nonparametric event study estimates for each of the eight specific health conditions that we study and for the probability of working past 70. The graphs plot point estimates and confidence intervals for the δ_τ coefficients in equation (1).

Figure A.5: Assessing the Importance of Avoiding Sample Selection Issues: Results Using Earlier Survey Waves

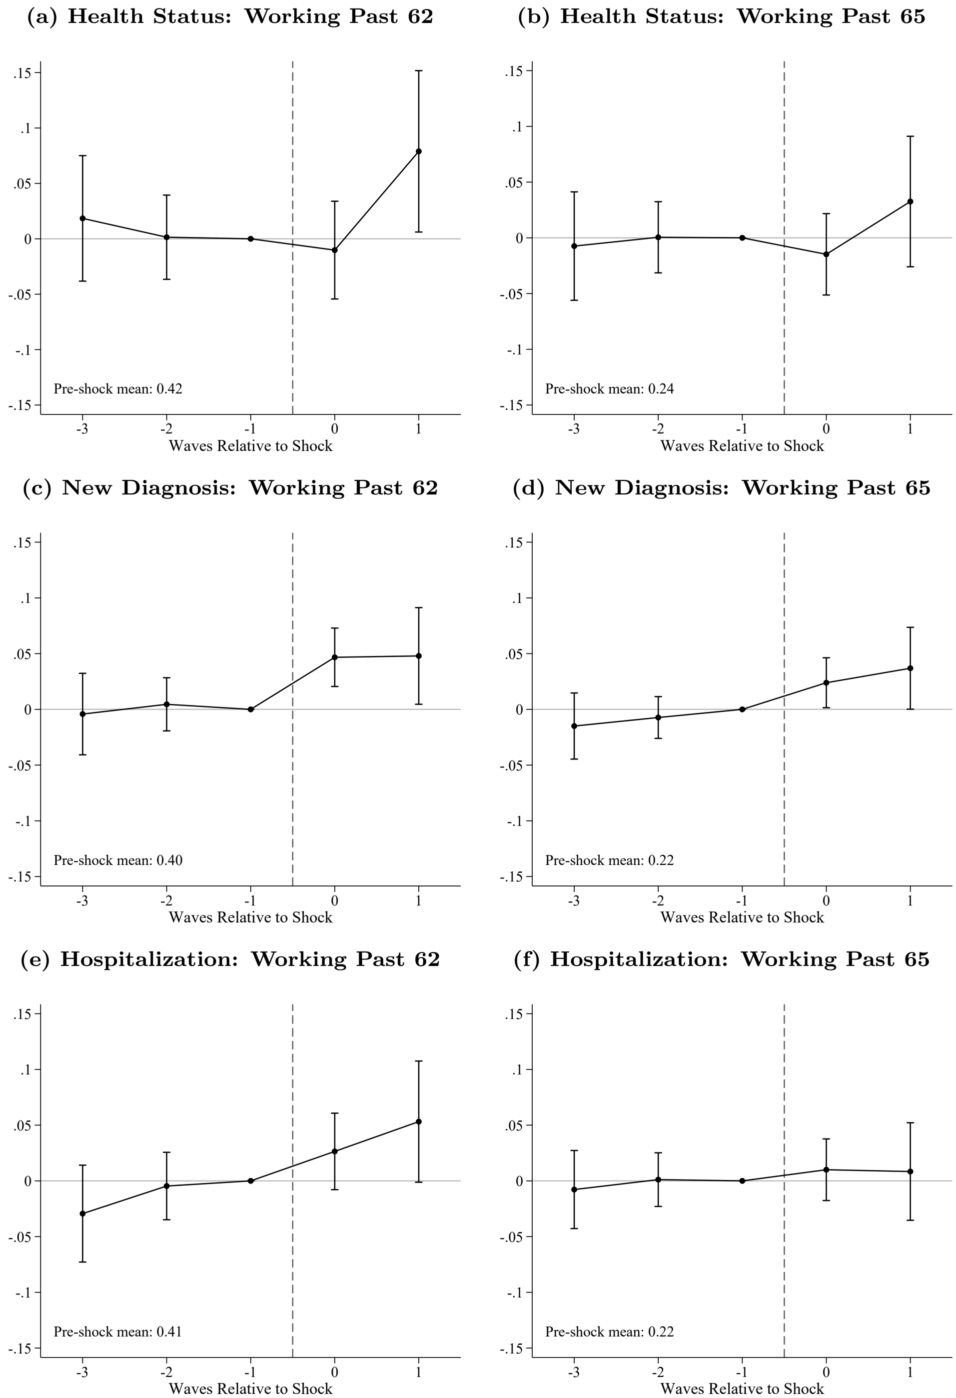

Notes: This figure presents the nonparametric event study estimates for our three main health shocks (declines in health status, new diagnoses, and hospitalizations) and our two main outcome variables (the probabilities of working past 62 and 65) when we use earlier sample waves, which are subject to sample selection concerns, instead of later sample waves. The graphs plot point estimates and confidence intervals for the δ_τ coefficients in equation (1).

Figure A.6: Assessing the Importance of Avoiding Sample Selection Issues: Results Using Workers in Later Survey Waves

(a) Health Status: Working Past 62

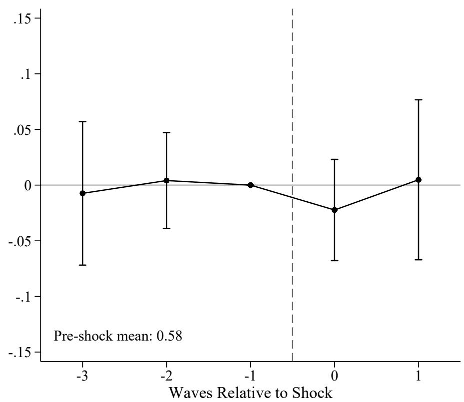

(b) Health Status: Working Past 65

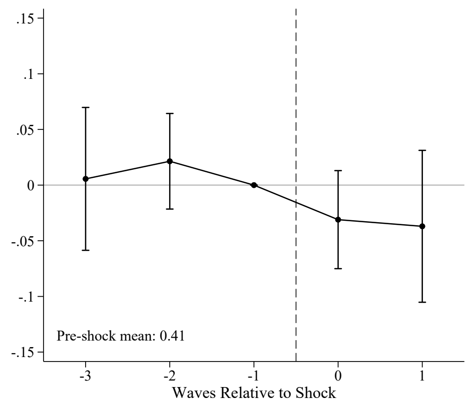

(c) New Diagnosis: Working Past 62

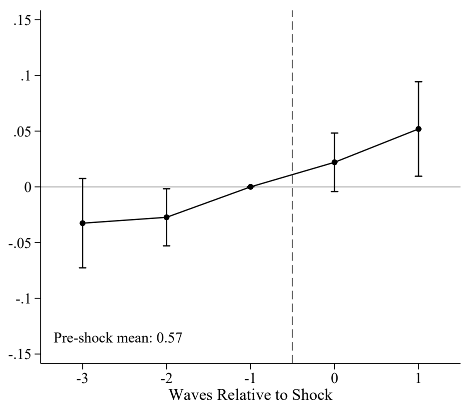

(d) New Diagnosis: Working Past 65

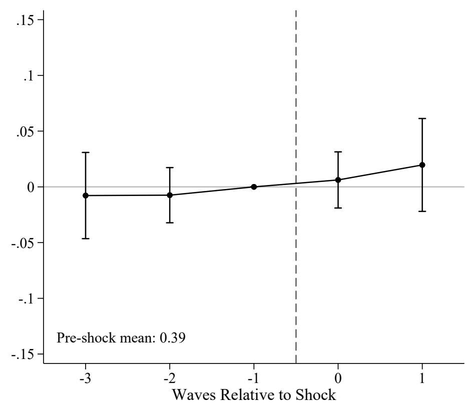

(e) Hospitalization: Working Past 62

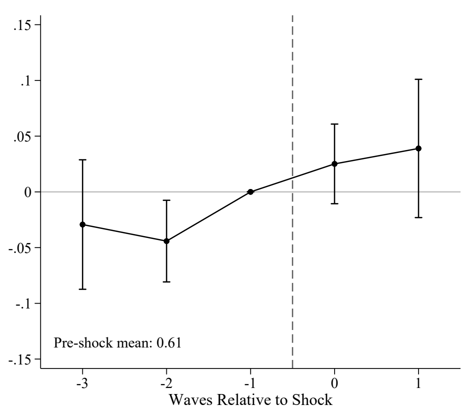

(f) Hospitalization: Working Past 65

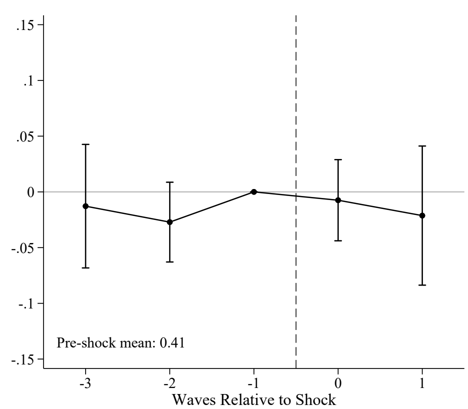

Notes: This figure presents the nonparametric event study estimates for our three main health shocks (declines in health status, new diagnoses, and hospitalizations) and our two main outcome variables (the probabilities of working past 62 and 65) when we restrict our analysis samples to workers, which creates sample selection concerns, instead of studying workers and non-workers. The graphs plot point estimates and confidence intervals for the δ_τ coefficients in equation (1).

Table A.2: The Effect of Health Shocks on the Probability of Working Past 70

	Probability of Working Past 70			
	Estimate (1)	Mean (2)	Clusters (3)	Obs. (4)
A: Main Shocks				
Health Status Decline	-0.010	0.12	792	2,170
	(0.015)			
New Diagnosis	-0.004	0.11	2,386	6,444
	(0.009)			
Hospitalization	0.002	0.12	1,087	2,914
	(0.013)			
B: Specific Diagnoses				
Arthritis	0.018	0.12	754	2,082
	(0.016)			
Cancer	0.018	0.09	214	580
	(0.029)			
Diabetes	0.016	0.11	493	1,406
	(0.020)			
Heart Disease	0.006	0.10	315	890
	(0.027)			
High Blood Pressure	-0.011	0.12	620	1,656
	(0.017)			
Lung Disease	-0.008	0.06	193	519
	(0.025)			
Psychiatric	0.009	0.11	294	786
	(0.025)			
Stroke	-0.068*	0.12	118	305
	(0.037)			

Notes: This table reports the parametric event study estimates for the probability of working past 70. Panel A presents estimates for our main health shocks of interest (declines in health status, new diagnoses, and hospitalizations). Panel B presents estimates for specific health conditions. Column (1) present estimates of δ_0 from estimating equation (2). Standard errors clustered at the individual level are in parentheses. Column (2) presents the dependent variable mean in the survey wave before the health shock. Column (3) presents the number of clusters. Column (4) presents the number of observations.
$^{***}\,p < 0.01$, $^{**}\,p < 0.05$, $^{*}\,p < 0.1$

Table A.3: Robustness of Age 70 Estimates for the Main Shocks

		Health Status	New Diagnosis	Hospitalization
		Age 70	Age 70	Age 70
		(1)	(2)	(3)
A.	Baseline Specification	-0.010	-0.004	0.002
		(0.015)	(0.009)	(0.013)
B.	Include Control Variables	-0.007	-0.004	0.000
		(0.015)	(0.009)	(0.013)
C.	Use Survey Weights	-0.009	0.013	0.012
		(0.020)	(0.013)	(0.018)
D.	Include Wave Fixed Effects	-0.013	-0.009	-0.012
		(0.015)	(0.009)	(0.013)
E.	Include Individual Fixed Effects	0.006	-0.006	-0.001
		(0.014)	(0.008)	(0.012)
F.	Subsample: No Missing Values	-0.014	-0.005	-0.003
		(0.016)	(0.009)	(0.014)
G.	Subsample: Balanced Panel	-0.023	-0.004	0.001
		(0.019)	(0.011)	(0.017)

Notes: This table reports results from assessing the robustness of our estimates for the main health shocks (declining health status, new diagnoses, and hospitalizations) and for the probability of working past 70. The estimates presented are the δ_0s from estimating equation (2). Each column corresponds to a different health shock. Each row corresponds to a different robustness check. Row A reproduces the baseline estimates for ease of comparison. Row B adds control variables to the regressions. Row C uses survey weights when estimating the regressions. Row D uses wave fixed effects instead of HRS-cohort-by-wave fixed effects. Row E uses individual fixed effects and two-wave fixed effects. Row F focuses on a subsample of people who never have missing values for any of the outcome variables. Row G focuses on the balanced panel of observations of individuals without missing values in survey waves corresponding to event times between $\tau = -2$ and $\tau = 0$.
$^{***} p < 0.01$, $^{**} p < 0.05$, $^{*} p < 0.1$

Table A.4: Robustness of Age 62 Estimates for Specific Health Conditions

		Arthritis (1)	Cancer (2)	Diabetes (3)	Heart Disease (4)	Blood Pressure (5)	Lung Disease (6)	Psych. (7)	Stroke (8)
A.	Baseline Specification	-0.047**	-0.072**	0.020	0.046	-0.022	-0.074**	-0.041	-0.021
		(0.021)	(0.034)	(0.023)	(0.030)	(0.022)	(0.037)	(0.038)	(0.051)
B.	Include Controls	-0.045**	-0.078**	0.016	0.041	-0.023	-0.072*	-0.043	-0.022
		(0.021)	(0.035)	(0.023)	(0.030)	(0.021)	(0.037)	(0.038)	(0.054)
C.	Use Survey Weights	-0.007	-0.102**	0.008	0.031	-0.019	-0.126*	0.002	0.015
		(0.026)	(0.045)	(0.034)	(0.043)	(0.028)	(0.067)	(0.047)	(0.072)
D.	Include Wave Fixed Effects	-0.054**	-0.041	-0.008	0.058*	-0.021	-0.068**	-0.041	-0.034
		(0.021)	(0.034)	(0.024)	(0.030)	(0.022)	(0.035)	(0.037)	(0.049)
E.	Include Individual Fixed Effects	-0.050***	-0.080**	0.023	0.034	-0.027	-0.052*	-0.057**	-0.068*
		(0.017)	(0.031)	(0.019)	(0.024)	(0.019)	(0.029)	(0.028)	(0.039)
F.	Subsample: No Missing Values	-0.043**	-0.064*	0.023	0.060*	-0.014	-0.075*	-0.022	-0.082
		(0.022)	(0.036)	(0.024)	(0.033)	(0.023)	(0.041)	(0.041)	(0.057)
G.	Subsample: Balanced Panel	-0.064**	-0.037	0.023	0.000	-0.024	-0.075	-0.077	-0.039
		(0.026)	(0.045)	(0.031)	(0.040)	(0.028)	(0.048)	(0.049)	(0.074)

Notes: This table reports results from assessing the robustness of our estimates for the specific health conditions and for the probability of working past 62. The estimates presented are the δ_0s from estimating equation (2). Each column corresponds to a different health shock. Each row corresponds to a different robustness check. Row A reproduces the baseline estimates for ease of comparison. Row B adds control variables to the regressions. Row C uses survey weights when estimating the regressions. Row D uses wave fixed effects instead of HRS-cohort-by-wave fixed effects. Row E uses individual fixed effects and two-wave fixed effects. Row F focuses on a subsample of people who never have missing values for any of the outcome variables. Row G focuses on the balanced panel of observations of individuals without missing values in survey waves corresponding to event times between $\tau = -2$ and $\tau = 0$.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table A.5: Robustness of Age 65 Estimates for Specific Health Conditions

		Arthritis (1)	Cancer (2)	Diabetes (3)	Heart Disease (4)	Blood Pressure (5)	Lung Disease (6)	Psych. (7)	Stroke (8)
A.	Baseline Specification	-0.030	-0.058*	0.017	0.009	0.006	-0.065**	-0.036	-0.119***
		(0.019)	(0.032)	(0.021)	(0.029)	(0.020)	(0.030)	(0.034)	(0.045)
B.	Include Controls	-0.029	-0.063*	0.012	0.005	0.007	-0.064**	-0.038	-0.127***
		(0.019)	(0.033)	(0.021)	(0.029)	(0.020)	(0.030)	(0.034)	(0.047)
C.	Use Survey Weights	-0.014	-0.061	0.025	0.001	-0.012	-0.104**	0.029	-0.072
		(0.023)	(0.043)	(0.031)	(0.040)	(0.027)	(0.049)	(0.045)	(0.054)
D.	Include Wave Fixed Effects	-0.028	-0.045	0.006	0.018	0.007	-0.071**	-0.049	-0.102**
		(0.019)	(0.032)	(0.021)	(0.029)	(0.020)	(0.028)	(0.033)	(0.043)
E.	Include Individual Fixed Effects	-0.018	-0.063**	0.033*	0.006	-0.001	-0.038	-0.036	-0.095***
		(0.016)	(0.030)	(0.018)	(0.024)	(0.017)	(0.026)	(0.026)	(0.032)
F.	Subsample: No Missing Values	-0.024	-0.044	0.030	0.014	0.013	-0.074**	-0.032	-0.160***
		(0.019)	(0.033)	(0.022)	(0.031)	(0.022)	(0.033)	(0.036)	(0.054)
G.	Subsample: Balanced Panel	-0.030	-0.020	0.017	0.032	0.013	-0.069	-0.029	-0.118*
		(0.024)	(0.040)	(0.028)	(0.036)	(0.025)	(0.045)	(0.045)	(0.071)

Notes: This table reports results from assessing the robustness of our estimates for the specific health conditions and for the probability of working past 65. The estimates presented are the δ_0s from estimating equation (2). Each column corresponds to a different health shock. Each row corresponds to a different robustness check. Row A reproduces the baseline estimates for ease of comparison. Row B adds control variables to the regressions. Row C uses survey weights when estimating the regressions. Row D uses wave fixed effects instead of HRS-cohort-by-wave fixed effects. Row E uses individual fixed effects and two-wave fixed effects. Row F focuses on a subsample of people who never have missing values for any of the outcome variables. Row G focuses on the balanced panel of observations of individuals without missing values in survey waves corresponding to event times between $\tau = -2$ and $\tau = 0$.
$^{***} p < 0.01$, $^{**} p < 0.05$, $^{*} p < 0.1$

Table A.6: Robustness of Age 70 Estimates for Specific Health Conditions

		Arthritis (1)	Cancer (2)	Diabetes (3)	Heart Disease (4)	Blood Pressure (5)	Lung Disease (6)	Psych. (7)	Stroke (8)
A.	Baseline Specification	0.018	0.018	0.016	0.006	-0.011	-0.008	0.009	-0.068*
		(0.016)	(0.029)	(0.020)	(0.027)	(0.017)	(0.025)	(0.025)	(0.037)
B.	Include Controls	0.016	0.016	0.013	0.005	-0.008	-0.007	0.010	-0.069*
		(0.015)	(0.029)	(0.020)	(0.026)	(0.017)	(0.025)	(0.025)	(0.038)
C.	Use Survey Weights	0.054**	0.026	0.055*	0.039	-0.031	-0.032	0.042	-0.009
		(0.022)	(0.037)	(0.032)	(0.034)	(0.024)	(0.033)	(0.030)	(0.046)
D.	Include Wave Fixed Effects	0.007	0.014	0.011	0.002	-0.010	-0.007	-0.017	-0.060
		(0.016)	(0.028)	(0.019)	(0.024)	(0.016)	(0.023)	(0.026)	(0.037)
E.	Include Individual Fixed Effects	0.002	0.004	0.017	0.015	-0.003	0.003	0.014	-0.050
		(0.014)	(0.028)	(0.018)	(0.023)	(0.015)	(0.021)	(0.023)	(0.033)
F.	Subsample: No Missing Values	0.027*	0.013	0.018	0.003	-0.016	-0.006	-0.001	-0.092**
		(0.016)	(0.030)	(0.022)	(0.029)	(0.018)	(0.029)	(0.028)	(0.045)
G.	Subsample: Balanced Panel	0.003	0.074	-0.000	-0.012	-0.012	-0.027	0.014	-0.050
		(0.018)	(0.045)	(0.025)	(0.031)	(0.022)	(0.035)	(0.032)	(0.070)

Notes: This table reports results from assessing the robustness of our estimates for the specific health conditions and for the probability of working past 70. The estimates presented are the δ_0s from estimating equation (2). Each column corresponds to a different health shock. Each row corresponds to a different robustness check. Row A reproduces the baseline estimates for ease of comparison. Row B adds control variables to the regressions. Row C uses survey weights when estimating the regressions. Row D uses wave fixed effects instead of HRS-cohort-by-wave fixed effects. Row E uses individual fixed effects and two-wave fixed effects. Row F focuses on a subsample of people who never have missing values for any of the outcome variables. Row G focuses on the balanced panel of observations of individuals without missing values in survey waves corresponding to event times between $\tau = -2$ and $\tau = 0$.
$^{***} p < 0.01$, $^{**} p < 0.05$, $^{*} p < 0.1$

Table A.7: Assessing the Importance of Avoiding Sample Selection Issues

	Probability of Working Past 62				Probability of Working Past 65			
	Estimate (1)	Mean (2)	Clusters (3)	Obs. (4)	Estimate (5)	Mean (6)	Clusters (7)	Obs. (8)
A: Earlier Survey Waves								
Health Status Decline	0.000	0.58	902	2,591	-0.019	0.24	900	2,582
	(0.026)				(0.021)			
New Diagnosis	0.043***	0.57	2,219	6,526	0.016	0.22	2,216	6,501
	(0.015)				(0.013)			
Hospitalization	0.010	0.61	1,467	4,290	0.005	0.22	1,467	4,280
	(0.019)				(0.015)			
B: Workers in Later Survey Waves								
Health Status Decline	-0.028	0.58	661	1,940	-0.032	0.41	660	1,910
	(0.026)				(0.026)			
New Diagnosis	0.008	0.57	1,879	5,614	0.003	0.39	1,877	5,541
	(0.016)				(0.015)			
Hospitalization	0.016	0.61	890	2,674	-0.010	0.41	888	2,644
	(0.023)				(0.022)			

Notes: This table reports the parametric event study estimates for the probabilities of working past 62 and 65 when we study our main health shocks using alternative samples that are subject to sample selection concerns. Panel A presents estimates for our main health shocks of interest when we use earlier survey waves instead of later survey waves. Panel B presents estimates for our main health shocks of interest when we restrict our main analysis samples to workers instead of studying workers and non-workers. Column (1) present estimates of δ_0 from estimating equation (2). Standard errors clustered at the individual level are in parentheses. Column (2) presents the dependent variable mean in the survey wave before the health shock. Column (3) presents the number of clusters. Column (4) presents the number of observations. $^{***}\,p < 0.01$, $^{**}\,p < 0.05$, $^{*}\,p < 0.1$