The Effect of Cancer on the Labor Supply of Employed Men over the Age of 65

David Candon *†‡

Abstract

This paper investigates the relationship between cancer diagnosis and the labor supply of employed men over the age of 65. While almost 60% of male cancers are diagnosed in men over the age of 65, no previous research has examined the effect that cancer has on this age group, which is surprising given the relevance of this group to public policy. With data from the Health and Retirement Study, I show that cancer has a significant negative effect on the labor supply of these workers. Using a combination of linear regression models and propensity score matching, I find that respondents who are diagnosed with cancer work 3 fewer hours per week than their non-cancer counterparts. They are also 10 percentage points more likely to stop working. This reduction seems to be driven by a deterioration in physical and mental health.

JEL classification: I10; J10; J22

Keywords: Cancer; Health shock, Employment; Labor supply

*A previous version of this paper was circulated under the title "The effect of cancer on the employment of older males: Attenuating selection bias using a high risk sample".

†School of Economics, University of Edinburgh, 30 Buccleuch Place, Edinburgh, EH8 9JT, UK. Email: david.candon@ed.ac.uk

‡I would like to thank Erlend Berg, Cathy Bradley, Paul Devereux, Steven Dieterle, Frank Dudbridge, Kanika Kapur, and Alessandro Tarozzi for helpful comments. I am also grateful to seminar participants at University College Dublin, the 29th Annual Irish Economic Association conference, and the 6th Biennial American Society of Health Economists conference for their suggestions. All remaining errors are my own. I have no conflicts of interest to declare.

Highlights

- This paper studies the effect of cancer on the labor supply of employed men over the age of 65.

- Respondents diagnosed with cancer work 3 fewer hours per week than non-cancer ones.

- They are also 10 percentage points more likely to stop working.

- This reduction seems to be driven by a deterioration in physical and mental health.

- The effect doesn't differ by the employer-sponsored insurance status of the respondents.

1 Introduction

One of the most common health shocks that a worker can face is cancer. As of 2014, the American Cancer Society (ACS) estimates that there are 14.5 million Americans alive who have been diagnosed with cancer at some point in their lives (American Cancer Society, 2016a). Concerning yearly diagnoses, they expect that almost 1.7 million people will be diagnosed with cancer in 2016 (American Cancer Society, 2016a). For comparison with the other severe health shocks, the Centers for Disease Control and Prevention report that 800,000 people have a stroke and 735,000 have a heart attack every year (Centers for Disease Control and Prevention, 2015, 2016). Due to the sheer number of diagnoses, it is important for policy makers to be aware of the impact that cancer has on employment.

Since cancer is the most common severe health shock, a large literature has developed over the past 15 years which focuses exclusively on the effect that cancer diagnosis has on labor supply. However, despite this previous research, one area has received little to no attention: the effect of cancer on the labor supply of workers who are employed past the age of 65 at the time of diagnosis. This is surprising for a number of reasons. The first is that the majority of cancer diagnoses occur in people over the age of 65. The National Cancer Institute estimates that the median age of cancer diagnosis is 66 (National Cancer Institute, 2015). Focusing on males specifically, the ACS estimates that the majority of new cancers, 56%, will be diagnosed in men who are at least 65 years of age (American Cancer Society, 2016b). The second reason is that the labor force participation rate for older workers has been steadily increasing for a number of years. According to the Bureau of Labor Statistics, men between the ages of 65 to 74 have a labor force participation rate of 31% (Bureau of Labor Statistics, 2015). Given that the population of males between the ages of 65 to 74 is over 12 million, this is not a negligible portion of the labor market (Bureau of Labor Statistics, 2015). Finally, due to the persistent threats of cuts to Social Security and Medicare, this is a group whose labor market response to health shocks is extremely policy relevant.

In spite of the lack of research on this topic, a robust examination of the effect of cancer on the labor supply of workers over the age of 65 is required for the following

reasons. First, the percentage of the labor force made up of workers over the age of 65 is growing. In 2024, workers between the ages of 65-74 will make up 6.5% of the labor force, compared to 4.4% in 2014 (Bureau of Labor Statistics, 2015)[1]. Given that people who work past the age of 65 are likely to be healthier, better educated, and more motivated than the workers who are already retired, the results from previous studies of younger cancer survivors cannot simply be extrapolated to this population. While workers who work past the age of 65 in the future may differ from workers who work past 65 today (since the regular retirement age is being raised from 65 to 67) knowing whether their labor supply response is similar to contiguous age groups can help us understand whether the mechanisms that lead to reductions in labor supply at younger ages are exacerbated or improved at older ages.

Second, even though the regular retirement age is being raised, 65 is still the age at which people become eligible for Medicare. This is an important consideration given that it has been shown that having access to multiple health insurance programs leads to different labor supply responses (Bradley et al., 2013). The results from this study may differ from previous ones due to the fact that the entire population is eligible for Medicare, in addition to any employer provided insurance they have. Finally, the results from this study could be particularly informative for public policy in the future. The increased economic pressures that stem from an ever aging population can be alleviated by knowing how to stimulate the labor market activity of the elderly. According to the Congressional Budget Office, the aging of the population is expected to account for 55% of the projected growth in federal spending on Social Security and the major health care programs as a share of GDP through 2039 (Congressional Budget Office, 2014). If cancer is causing labor market exits, it will help policymakers to know whether these exits are being driven by cancer related morbidity, changes in labor-leisure preferences, or by health insurance status. From the point of view of cancer survivors, returning to work may relieve any financial strain associated with illness and improve both their physical and mental health.

[1]While these numbers are for the labor force as a whole they are very similar when broken down by gender.

This paper contributes to the literature in numerous ways. Firstly, it exploits the fact that age is the biggest cancer risk factor. Within this high risk group all other risk factors appear to have a negligible relationship with future cancer diagnosis. With data from the Health and Retirement Study I demonstrate this by showing that the cancer and non-cancer respondents look identical across multiple dimensions. This provides appropriate justification for the assumptions behind the econometric analysis. Secondly, I find that males, who are employed at the time of cancer diagnosis, work 3 hours fewer per week in the period after diagnosis when compared to their non-cancer counterparts. This reduction is driven by changes at the extensive margin. Respondents diagnosed with cancer are 10 percentage points more likely to leave the labor market than non-cancer respondents. I find no evidence of any reduction at the intensive margin. While these estimates are of the same magnitude as estimates for younger survivors it is important to remember that those who work past the age of 65 are a highly selected sample. If the general population had been made to work past the age of 65 due to increases in the retirement age then the estimates would likely be much larger. This is an important consideration for policy makers due to the increasing retirement age. Thirdly, I find that the presence of employer provided health insurance (EPHI) has no effect on the probability of leaving the labor market. Since the previous literature on EPHI and job mobility offers no clear pattern for male cancer survivors or men who are eligible for Medicare, this provides more empirical evidence that male cancer survivors with EPHI are not job-locked once they are eligible for Medicare. However, it does appear as if the results are driven in part by the health of the respondents as I find that cancer has a negative effect on both mental and physical health.

The rest of the paper is organized as follows: A brief review of the related literature is contained in Section 2; Section 3 contains information on the methods and the data used in the analysis; The validity of the assumptions for the econometric models are scrutinized in Section 4; In Section 5, I present the main results of the analysis; Section 6 examines the mechanisms by which the respondents reduce their labor supply; Section 7 discusses the implications of the results and Section 8 concludes.

2 Related literature

While I am not aware of any other studies which explicitly examine the effect of cancer on male labor supply for this age group, there are papers which have addressed this question at other ages. In the US, Coile (2004) shows that married men (between 50 and 69 years of age) diagnosed with cancer are 8 percentage points more likely to leave the labor market than workers without cancer. Coile (2004) also uses HRS data to examine the effect of health shocks on labor supply but my paper differs in a number of substantive ways. Firstly, I exclusively focus my analysis on workers over the age of 65 while workers over the age of 65 in Coile (2004) are most likely a small group given that the mean age of males is 60. Secondly, Coile uses data from waves 1 - 6 of the HRS while I use data from waves 1 - 11. Coile includes respondents who have reported previous conditions of severe illnesses (such as cancer) while these people are excluded from my analysis. Coile also uses linear probability models while I augment my analysis with propensity score matching and rigorous checks of the conditional independence and overlap assumptions. Finally, Coile groups together all health shocks while I focus specifically on cancer, for reasons justified in the Introduction. In fact, the result quoted above is the only cancer specific result reported in the entire paper as it is from a specification check of a main result.

Datta Gupta et al. (2015) shows that older men (between 55 and 64 years of age) diagnosed with cancer are 9.8 percentage points more likely to not work than non-cancer survivors two years after diagnosis. Bradley et al. (2005) and Bradley et al. (2007) show that older men (with a mean age of 56) who have been diagnosed with prostate cancer are 10 percentage points less likely to work in the 6 months following prostate cancer diagnosis than healthy controls. Looking at longer term outcomes, Short et al. (2008) show that, between two to six years after diagnosis, there is no significant difference in employment for older male (between 55 and 65 years of age) cancer survivors versus healthy workers, though they are less likely to work full-time and they work fewer hours compared to control workers. The longer term impact for younger workers (with a mean

age of 45) is similar, with Moran et al. (2011) showing that male survivors are less likely to work, less likely to work full-time, and work fewer hours than controls.

There have also been several studies which have examined the effect of cancer on labor supply around the world[2]. Evidence from Canada (Jeon, 2017), Denmark (Datta Gupta et al., 2015; Heinesen and Kolodziejczyk, 2013; Kolodziejczyk and Heinesen, 2016), Finland (Taskila-Brandt et al., 2004), Norway (Torp et al., 2013), the UK (Candon, 2015), and a collection of 16 European countries (Trevisan and Zantomio, 2016), all show that cancer has a negative effect on labor supply over various time horizons.

In this group, the most similar paper to this one is Candon (2015), though there are still a number of sizable differences between the papers. With regards to the populations of interest, the major difference between these two studies is that all of the respondents in this paper are at least 65 years of age whereas only 4% of the sample in Candon (2015) is at least 65. As mentioned earlier, there is no reason to believe that the labor supply response of those below the traditional age of retirement would be the same as those who are above it. With regards to the location, Candon (2015) is set in the UK which has both a markedly different labor market and a markedly different healthcare market compared to the US. This means that, even if the age groups were comparable, the results from one setting cannot be generalized to the other. This also allows me to examine the role that health insurance status plays in the labor supply decision, something not done in Candon (2015) but something which is extremely important in US labor markets, as the large literature on this topic demonstrates (Madrian (1994), Buchmueller and Valletta (1996), and Kapur (1998)). Also, the aim of Candon (2015) was to determine how long the negative labor supply shock from cancer diagnosis lasted for, rather than a comprehensive understanding of the mechanisms related to the labor supply reduction, which is done here. Finally, Candon (2015) only examines the effect of cancer diagnosis on total hours of work and probability of working, rather than whether the reductions

[2]There are other studies which have examined the effect of health shocks on labor supply in the US such as McClellan (1998) and Smith (2005) and in Europe such as Garcia-Gomez and Lopez-Nicolas (2006), Garcia-Gomez (2011), and Garcia-Gomez et al. (2013). These studies may include cancer shocks but no specific cancer effects are reported.

were taking place at the intensive or extensive margin. Again, this is something which is done in this study.

With regards to the econometric methods, all of the papers mentioned here deal with the endogeneity of cancer diagnosis by using some combination of controlling for observable differences between groups, differencing away unobservables which are constant over time, and using matching to construct comparable control groups. This is the standard for the literature since there are no obvious discontinuities or instruments to exploit with the data at hand. The methods used here are discussed in more detail in the next section.

3 Methods and Data

3.1 Methods

In order to estimate the effect of cancer on labor supply I use a variety of simple two period models in which the periods are referred to as period $t-1$ and period t. The first method involves trying to eliminate any unobserved heterogeneity which is constant across both periods. While it is possible to control for many known cancer risk factors such as smoking and alcohol consumption, the worry when estimating the effect of cancer on labor supply is the presence of unobserved confounding variables which are correlated with cancer status but also determine labor market outcomes. The estimator of the effect of cancer on labor supply will be biased if such a variable is omitted from the model. However, if such variables are constant between the two time periods then they can be differenced out, even if they are unobserved. To do this, I create a first-difference (FD) model. Consider the following model,

$$Y_{i,t} = \beta_0 + \beta_1 D_{i,t} + \delta_0 + \alpha_i + u_{i,t}, \tag{1}$$

$$Y_{i,t-1} = \beta_0 + \beta_1 D_{i,t-1} + \alpha_i + u_{i,t-1}, \tag{2}$$

where Y represents employment outcomes, D is a binary variable representing cancer status, α is time-invariant unobserved heterogeneity, δ is a period t indicator, and u is the traditional idiosyncratic error. If we subtract Equation (2) from Equation (1) we get,

$$\Delta Y_i = \delta_0 + \beta_1 \Delta D_i + \Delta u_i. \tag{3}$$

This allows the time-invariant unobservable variables that affect both cancer status and employment outcomes to be differenced out. Because β_1 represents the difference in outcomes, where the outcomes are already in differences, an estimator of β_1 is a difference-in-differences estimator (Wooldridge, 2016). This equation can then be estimated via OLS.

Other control variables are omitted from (3) in order to prevent any contamination in their differences that would stem from cancer diagnosis in period t. A different method that would allow for the inclusion of control variables to account for observable differences between the cancer and non-cancer respondents is a lagged dependent variable (LDV) model. Such a model could be represented as

$$Y_{i,t} = \beta_0 + \beta_1 D_{i,t} + \beta_2 Y_{i,t-1} + \gamma' X_{i,t-1} + u_{i,t}, \tag{4}$$

where X is a (k x 1) vector of additional control variables measured in period $t-1$. Again, this model can then be estimated with OLS. In short, the identification of the effect of cancer on labor market outcomes relies on either differencing out any unobservable factors affecting both cancer and employment or controlling for pre-cancer observables. As noted by Angrist and Pischke (2009), because the FD and LDV models rely on different identifying assumptions, they have a useful bracketing property that provides upper and lower bounds to the causal effect. Obtaining similar results using both models provides a simple demonstration of the robustness of the results.

The final method focuses on finding a suitable counterfactual for the respondents who are diagnosed with cancer. Many previous papers in this field have used matching or propensity score methods to examine the causal effect of cancer on labor market outcomes.

While these methods rely on the same assumptions as linear regression models to identify causal effects they are less reliant on extrapolation across differing covariate distributions and do not impose strong functional form restrictions (Imbens, 2015). Because it is a common belief that the treated group and the control group are drawn from different populations in nonexperimental studies, the average treatment effect on the treated (ATT) is the parameter which has received the most attention in this literature, and it is the parameter that the analysis will focus on.

Using notation from Heckman et al. (1997), the observed outcome for an individual, Y, is a combination of their potential outcomes $\{Y(1), Y(0)\}$ where the parentheses denote the outcome when an individual receives the treatment (1) and does not receive the treatment (0). If D is an indicator of whether the individual receives the treatment (where $D = 1$ if the individual receives the treatment; $D = 0$ otherwise) this can be expressed more succinctly as

$$Y = DY(1) + (1 - D)Y(0). \tag{5}$$

The conventional method of matching rests on two assumptions[3]. The first is that, conditional on the control variables, X, the outcome for those who do not receive the treatment and treatment assignment are independent. This is known as the conditional independence assumption (CIA) and can be denoted as

$$Y(0) \perp D|X. \tag{6}$$

The second assumption is that there is overlap between the treatment and control units for all values of X. This prevents the case where receiving the treatment can be perfectly predicted and can be denoted as

$$\Pr(D = 1|X) < 1. \tag{7}$$

[3]Because I am interested in the average treatment effect on the treated, and not the average treatment effect, the weak versions of these assumptions can be used. The strong versions of these assumptions would require having to justify $Y(1), Y(0) \perp D|X$ and $0 < \Pr(D = 1|X) < 1$. For more on the difference between the strong and weak versions please consult Caliendo and Kopeinig (2008).

If these two conditions are met then there is no need to condition on treatment status. This is known by many names such as ignorability, selection-on-observables, or exogeneity. Rosenbaum and Rubin (1983) show that if treatment assignment is ignorable for the finest level of balance, the covariates themselves, then it must also be ignorable for coarser measures of balance such as the propensity score. The propensity score is defined by

$$p(x) \equiv \Pr(D = 1|X). \tag{8}$$

Conditioning on the propensity score, rather than the covariates themselves, reduces the dimensionality problem. If these assumptions are satisfied, then the ATT can be estimated as the difference in observed responses to the treatment indicator, conditional on the propensity score, for those in the treated group,

$$\begin{aligned} E[E\{Y(1)|p(X), D = 1\} - E\{Y(0)|p(X), D = 0\}|D = 1] = \\ E[E\{Y(1) - Y(0)|p(X)\}|D = 1] = \text{ATT} \end{aligned} \tag{9}$$

where the outer expectation is over the distribution of $p(X)$ in the sub-population of treated respondents.

If two periods of data are available, then the previous formula can be augmented to include this information. The ATT then becomes,

$$E[E\{Y_t(1) - Y_t(0)|p(X_{t-1})\}|D_t = 1] = \text{ATT} \tag{10}$$

where $Y_t(1)$ and $Y_t(0)$ refer to the outcomes for individuals who receive and do not receive the treatment in period t, D_t denotes cancer status in period t, and X_{t-1} denotes control variables measured in period $t-1$. This means that I am only using pre-treatment variables as part of the conditioning set, in order to ensure that they cannot have been affected by the treatment. It is also possible to use lagged values of the outcome variable as part of the conditioning set. In this case, all the variables used in X, whether lagged outcome variables or other control variables, have been measured in period $t-1$, before

anyone is diagnosed with cancer. In Section 4, I present empirical evidence to support the assumptions which are required to use these models.

3.2 Data

In order to obtain the requisite data on older workers with which to do this analysis, I use data from the Health and Retirement Study (HRS)[4]. The HRS is a large, longitudinal data set which contains information on the respondents' health, wealth, employment, and other demographic information. The first wave was collected in 1992 using a representative sample of 51 to 61 year olds and, since then, it has been collected every two years. For this analysis, I use the RAND version of the data set which has already been cleaned and compiled for ease of use[5]. This data set contains the first 11 survey waves, spanning the years 1992 to 2012. With regards to the cancer variable, a respondent is considered as having cancer if they respond in the affirmative to the question indicating whether or not a doctor "has ever told them that they have cancer or a malignant tumor of any kind except skin cancer". They do not have cancer if they respond in the negative. As with any study which examines the impact that cancer has on labor supply, it is important to observe the respondents' pre-cancer behavior while they are in employment. To do this, information on the respondents before and after they are diagnosed with cancer is required. This means that information in two different time periods is required for each observation. Regarding the HRS, having 11 waves of data allows me to observe 10 potential non-cancer to cancer transitions: wave 1 to wave 2, wave 2 to wave 3, . . . , wave 10 to wave 11.

[4]The Health and Retirement Study (2015) is sponsored by the National Institute on Aging (grant number NIA U01AG009740) and is conducted by the University of Michigan.

[5]Due to it's survey design the HRS oversamples Black and Hispanic respondents and weights are provided for use in data analysis to account for this. The main reason that the weights are often used is to obtain the correct descriptive statistics of the target population. However, the issue of whether to use them when estimating causal effects is more nuanced. Solon et al. (2015) outline some cases where weighting may do more harm than good, particularly if the oversampling is done exogenously rather than endogenously. They recommend in cases like this to present both the unweighted and weighted results as the contrast can serve as a useful diagnostic against a misspecified model or misunderstanding of the sampling process. In this case, the unweighted results and weighted results were so similar that I decided to just use the unweighted results. The weighted results are available in Table A7 of the Appendix.

Four main restrictions are placed on the overall data set. The respondents must be working in period $t-1$ so we can observe the effect of cancer on employment; they must be at least 65 years of age in period $t-1$ as this is the age group that the analysis will focus on; the respondents must all be men as fewer women work past the age of 65; and they cannot be diagnosed with cancer in period $t-1$, or at any time in the past, so we can observe their pre-cancer behavior. This means that respondents are removed if they answer "Yes" to the cancer question listed above in period $t-1$[6]. This means that any non-cancer to cancer transition is unique, though respondents may appear multiple times with regards to their non-cancer work behavior. Respondents are categorized as working if they are working full-time, part-time, or are part-retired. They are categorized as not working if they are unemployed, retired, disabled, or not in the labor force. Some other minor restrictions are also placed on the data[7]. Imposing these restrictions leaves a final sample of 5,602 person-wave observations, of which, 238 will be diagnosed with cancer in period t. Of the 5,602 observations, 2,436 are unique respondents[8]. As noted by Bertrand et al. (2004), the serial correlation from using multiple waves of information on the same respondents can severely understate the standard errors of the estimators. However, they recommend aggregating the data into two simple before and after time periods, which is done here, in order to help mitigate the problem. Another way to deal with the serial correlation is to follow a similar two period analysis by Finkelstein (2004) and cluster the standard errors at the individual level. This is done for the linear regression models.

In addition to using the employment and cancer information, I also use information on pre-diagnosis variables that are measured in period $t-1$. This includes demographic

[6]There are two potential drawbacks to the way that the HRS asks the cancer question which should be highlighted here. The first is that respondents with a past diagnosis who are now cancer free are eliminated because of this past diagnosis. So respondents who are cancer free in period $t-1$ will be removed if they had cancer at some point earlier in their lives. This is why the sample I focus on is respondents who have never had cancer, rather than simply respondents who do not have cancer in period $t-1$. The second is that I can't distinguish between a single cancer diagnosis occurring between period $t-1$ and t or multiple diagnoses. So a respondent may be diagnosed with two cancers but it will be treated the same as a respondent diagnosed with one form of cancer. However, including the respondents with previous diagnoses has no effect on the overall results (see Table A6).

[7]These restrictions, and the number of observation lost by imposing them, are presented in Table A3.

[8]While it is possible for a respondent to appear more than once with regards to their non-cancer work status, each of the 238 cancer cases represents a unique individual. This is because respondents are removed from the analysis if they have been previously diagnosed with cancer.

information such as age, whether the respondent is non-white, has some or a full college education, or is married. The behavioral cancer risk factors, as documented by Danaei et al. (2005), that are taken into account are whether the respondent is in poor health, has ever smoked, currently smokes, ever drinks any alcohol, and if they are obese[9]. Regarding their employment situation, I include variables indicating whether they have a working spouse, are self-employed, and their hours of work per week. I also control for their earnings, household income, income from a pension or annuity, and income from Social Security. These variables are all transformed using the inverse hyperbolic sine (IHS) transformation[10]. This transformation is preferred to the logarithmic transformation since it behaves like the logarithmic function for positive values but does not exclude zero or negative values. Finally, dummy variables for HRS wave and census division are also included. All the variables are binary except for age, the income variables, and hours of work. Information on whether or not the respondents have employer provided health insurance is also available. However it is not included in the main analysis due to the fact that there were a large number of missing observations. Nevertheless, analyses which include this variable are presented in Table 7 as part of subgroup analyses. These results are largely similar to the main results.

4 Ignorability

4.1 Overlap

Using different econometric models to identify the effect of cancer on labor supply means that different assumptions are required to be justified. The FD model relies on the change in cancer status being uncorrelated with changes in the error term while the LDV and PSM models rely on the unconfoundedness (conditional independence) assumption in Equation (6). The PSM model also relies on the overlap assumption in Equation (7).

[9]Poor health is defined as being in poor or fair health as opposed to good, very good or excellent in a self-reported health measure.

[10]The transformation of wealth, W, is $w = \ln\left(W + \sqrt{(W^2 + 1)}\right)$. For more on this transformation see Pence (2006). The results do not change when the natural log is used as the wealth transformation, even with the reduced sample size (see Table A8).

Therefore, I spend this section examining the overlap assumption, the unconfoundedness assumption, and the trends of the respondents' over time. The first assumption, whether there is sufficient overlap in the covariate distributions, is relatively straightforward to test. One way to do this is to compare the means of these covariates between groups. The descriptive statistics for both the cancer and non-cancer groups are presented in Table 1, along with t tests for the equality of their means. Regarding the largest differences, the cancer group is about 5 percentage points less likely to drink alcohol and about 4 percentage points less likely to have ever smoked. The only differences which are statistically significant at conventional levels are the alcohol difference and the difference in pension income. In terms of demographic characteristics, cancer risk factors, employment related characteristics, and earnings, the groups appear balanced. In the final column, I present the normalized differences between the groups which provides a scale and sample size free way of assessing overlap (Imbens, 2015). Again, a similar pattern emerges, the variables with the largest differences are the alcohol and pension variables. Looking at the period t outcomes, there is a 10 percentage point difference between the groups with regards to whether or not the respondents are working, which is significant at the 1% level. The cancer group also works almost 3 hours a week fewer than the non-cancer group. This difference is significant at the 5% level.

While it may be surprising to see no significant differences between the cancer and non-cancer groups with regards to specific cancer risk factors, this is actually a common epidemiological observation. Rothman and Poole (1988) noted that the association between a disease and a single specific risk factor could be made stronger by examining the relationship within a group where the effect of all other possible risk factors has been eliminated. The relationship is made stronger because the other mechanisms of disease development have been shut off and disease development can only take place through this specific risk factor. Similarly, they also noted that this finding held in reverse. Strong associations between risk factors and diseases became weak when analyzed in high risk groups with many potential mechanisms of development. In this case, because age is the single biggest risk factor in developing cancer, both smokers and non-smokers (and

obese respondents versus non-obese respondents etc.) are very likely to get cancer anyway, simply because of their age (American Society of Clinical Oncology, 2012; National Cancer Institute, 2015). This then makes the relationship between cancer and its risk factors appear weak. There are numerous explanations as to why old age is the most important cancer risk factor. Historically, it has been posited that longevity allows cells to build up enough mutations over time to form tumors (Cancer Research UK, 2016). More recently, new research has also shown tissue changes which occur in old age can make cells more susceptible to cancer and that DNA repair, which combats mutations, weakens in old age (DeGregori, 2013; Li et al., 2017).

As a final inspection of the overlap assumption, I follow Trevisan and Zantomio (2016) and graph the density of the propensity score for the cancer and non-cancer groups (more information on the generation of the propensity score is given in Section 5.2). This allows for a visual inspection of the overlap between the two groups and to see if matching can help make the groups more comparable. In Figure 1 we can see that a simple comparison of the two densities shows that the groups are already very similar, even before matching.

4.2 Unconfoundedness

The identification of the effect also relies on whether or not the CIA assumption can be satisfied. While it is fundamentally impossible to test the CIA, steps can be taken to assess whether or not it is plausible. One way to assess it would be to check what factors help select respondents into the treatment by regressing the treatment variable on the control variables. In this case, the coefficients reported give the effect of a particular variable on future cancer diagnosis after the effect of all of the other control variables have been partialled out. The results of this logistic regression are reported in column (1) of Table 2. The results reported are average marginal effects. Again, the ability of these variables to predict future diagnosis is weak. Even for the cancer risk factors, the magnitude of the coefficients is rarely greater than 0.01. Due to the limited association between the other variables and cancer diagnosis, the null hypothesis of a Wald test for the joint insignificance of the variables cannot be rejected, meaning that there is no

evidence that they jointly help explain future cancer diagnosis. I report the p-value for the Wald test, rather than the Wald statistic, as it is easier to interpret.

With the lack of a relationship between cancer diagnosis and the control variables, it may be tempting to say that the variables used in the analysis are simply poor measures of the underlying variables. For example, it may be the case that the cancer group are unhealthier than the non-cancer group but, because of the way the variable is measured or the way the survey questions are asked, it does not show up in this analysis. I attempt to address this issue in the remaining columns in Table 2. Here, I regress period t diagnosis of high blood pressure, heart problems, lung disease, arthritis, and diabetes on the same period $t-1$ variables. The questions in the HRS that ask the respondents about these health conditions are worded the same as the question that asks the respondents about their cancer status. It is apparent that, even after controlling for the fact that the respondents may have been diagnosed with this condition in period $t-1$ or earlier, all of the known risk factors (except for whether the respondents drink alcohol) are statistically significant in determining these diseases. Since these variables help explain variation in future diagnoses of other diseases, it supports the idea that the covariates are in fact balanced between the cancer and non-cancer groups, rather than the idea that these variables are simply coarse measures of underlying variables. While the comparison across health conditions isn't perfect, since people with previous cancer diagnosis are removed in the Table 1 analysis, similar results can be found for these other health conditions when restricting the sample to those respondents with no previous record of the condition.

4.3 Trends

Finally, in order to justify the use of the FD model, I exploit the panel nature of HRS and compare the past labor market outcomes and health outcomes of both sets of respondents. If the outcomes are similar in the lead up to cancer diagnosis then we can be confident that any difference found after diagnosis is due to cancer. The two outcomes I follow are the weekly hours of work of the respondents and the proportion of respondents reporting

themselves in poor health. A t test for the equality of their means is tested in each period. It is apparent in both Figures 2 and 3 that in the three periods before cancer diagnosis there is no difference in either outcome across groups. The only statistically significant difference between the two groups arises in the period after diagnosis. Taken all together, this evidence supports the plausibility of the ignorability assumptions and the use of the different models.

5 Results

5.1 Linear regression models

I begin by estimating a simple regression of the hours of work in period t on cancer status in period t. The results are presented in Panel A of Table 3. Cancer diagnosis leads to a 3.2 hour reduction in the number of hours worked per week. When the FD model in (3) is estimated the results show that cancer diagnosis leads to a 3.3 hour reduction in hours of work. No observations are lost when using the FD model since one of the restrictions placed on the respondents is that they must be working in period $t-1$. HRS wave dummy variables are also included in the FD model to capture any secular changes in hours of work. The LDV model from (4) is estimated with and without the other pre-diagnosis covariates. In both cases, cancer reduces hours of work by approximately 3.2 hours per week. Given the lack of correlation between cancer status and the other covariates it is not surprising to see similar estimates across different specifications. All results are statistically significant at the 1% level.

In order to get a more comprehensive understanding of the effect that cancer has on labor market outcomes, I check whether the previous results are due to reductions at the intensive margin or the extensive margin. This will tell us whether the respondents are reducing their labor supply within employment or whether they simply stop working. I begin with the extensive margin. Since the respondents may stop working in different ways, I use two different variables to capture the negative effect that cancer may have. The first variable is whether or not the respondent is working in period t. This variable

is equal to 1 if they are working full-time, part-time, or are part-retired and equal to 0 if the respondent is retired, unemployed, disabled, or not in the labor force. The second variable is whether or not the respondent is in the labor force. This variable is equal to 1 if the respondent is working full-time, part-time, part-retired, or unemployed and 0 if they are retired, disabled, or not in the labor force. Because all of the respondents are working in period $t-1$ the FD and LDV models are not appropriate for these outcomes. Instead, I simply run pooled OLS with and without the control variables (which include period $t-1$ hours of work). In this case the identifying assumption is the same as the one which will be used in the PSM section: conditional on observable characteristics in period $t-1$, cancer diagnosis in period t is assumed to be exogenous. The results of this analysis are presented in Panel B. Cancer leads to an approximate 10 percentage point decrease in the probability of working and an 11 percentage point decrease in the probability of being in the labor force. Results are significant at the 1% level. Because the number of respondents who are unemployed in period t is so small, respondents who are not working will be referred to as out of the labor force when discussing the remaining results.

I now check the intensive margin and estimate the same four models as in Panel A, but with the added restriction that all the respondents must be employed in period t. The results are presented in Panel C. In no model is the effect of cancer on the hours of work of the respondents who are employed statistically significant. The respondents who do return appear to work the same number of hours as those respondents who have not been diagnosed with cancer. Together with the results from Panel B, this suggests that the results presented in Panel A are being driven by reductions at the extensive margin, rather than the intensive margin[11].

5.2 Propensity score matching

I now use PSM to examine the effect of cancer diagnosis on employment outcomes. The specific matching algorithm I use is kernel matching. This involves using a weighted

[11] I also estimate similar results with other panel data methods such as random effects and fixed effects. For more information see Tables A1 and A2 of the Appendix.

average of the entire control group to generate a counterfactual for each treatment observation. While many different matching algorithms exist, kernel matching has been shown to work well in situations where there are many control observations per treatment observation and the propensity score distributions are well behaved (Frolich, 2004)[12]. For a more detailed review of this, and other matching algorithms, see Caliendo and Kopeinig (2008). I use a Gaussian kernel and, to select the bandwidth, I follow Heckman et al. (1998) and use a rule of thumb. In this case, the rule of thumb is defined as

$$\text{bandwidth} \approx 0.9An^{-\frac{1}{5}}, \tag{11}$$

where $A = \min\{\text{sample standard deviation, sample interquartile range}/1.34\}$ and n is the number of observations in the control group (Sheather, 2004). The rule of thumb above suggests a bandwidth of 0.003. I also estimate the effect using half and double this bandwidth to test how sensitive the results are to this specification. To generate the propensity score, I follow Dehejia and Wahba (1999, 2002) and use a logit model where the covariates from Table 2 are entered linearly along with dummy variables for HRS waves and census divisions. A balancing test confirms that the covariates are evenly distributed across both groups within different strata of the estimated propensity score[13]. As was shown in Figure 1 previously, the distribution of the propensity score looks similar for both the cancer and non-cancer groups so the results from the matching process should give similar results to those in Table 3.

The results of the PSM are presented in Table 4. Panel A shows that cancer diagnosis leads to an approximate 3.3 hour reduction in the number of hours worked per week for each of the three different bandwidths. Results are significant at the 1% level. In Panel B, I check the extensive margin and find that cancer diagnosis leads to an approximate 10 percentage point decrease in the probability of working. Results are again significant at the 1% level. Finally, I turn to the intensive margin in Panel C. In no model is the effect

[12]I also estimate the results with methods which use the propensity score in different ways. These methods provide almost identical estimates. For more information see Tables A4 and A5 of the Appendix.

[13]To implement these methods, I use the ATT commands developed by (Becker and Ichino, 2002) for Stata.

of cancer on the hours of work of those who are employed statistically significant. In addition, the coefficients are very close to zero. The respondents who do return work the same number of hours as those respondents who have not been diagnosed with cancer. As predicted, the results from PSM perfectly mirror the results from the other two period models presented in Table 3.

5.3 Sensitivity

In this scenario, I hypothesize the presence of an unobserved binary confounding variable, U, whose exclusion results in selection bias. This means that assignment to treatment is not independent given X, but it is independent once U is taken into account,

$$\Pr(D|Y(1), Y(0), X) \neq \Pr(D|X), \tag{12}$$

$$\Pr(D|Y(1), Y(0), X, U) = \Pr(D|X, U). \tag{13}$$

If the omission of the variable U does result in selection bias then it must be the case that U fulfills the following two conditions. The first condition is that U selects respondents into the treatment. This means that a respondent with the confounding variable is more likely to be in the treatment group. The second is that U leads to worse labor market outcomes for those respondents who are not in the treatment group. This means that, in the non-treatment group, the respondents with the confounding variable are more likely to have worse labor market outcomes.

If we let

$$p_{ij} = \Pr(U = 1|D = i, Y = j, X) = \Pr(U = 1|D = i, Y = j), \tag{14}$$

$$p_{i.} = \sum_{j=0}^{1} p_{ij} \cdot \Pr(Y = j|D = i), \tag{15}$$

with $i, j \in \{0, 1\}$, then these two conditions can be expressed more precisely as the following two equations:

$$s = p_{1.} - p_{0.} > 0, \tag{16}$$

$$d = p_{01} - p_{00} < 0. \tag{17}$$

Equation (16) states that the probability of having the confounder given you are in the treated group is greater than the probability of having the confounder given you are in the untreated group. Equation (17) states that the probability of having the confounder given you have a negative work outcome is greater than the probability of having the confounder given you have a positive work outcome, conditional on being in the untreated group. The idea behind this model is that, if this confounder is the true cause of the treatment effect, it needs to be the case that we are mistaking the confounder for the treatment variable ($s > 0$) and that the confounder has it's own negative effect on outcomes that is separate from the treatment ($d < 0$). For a more mathematically rigorous explanation of this model, and the Stata commands to implement it, readers should consult both Ichino et al. (2008) and Nannicini (2007)[14]. The work outcome that I choose to focus on in this section is the probability of working given that cancer has been shown to have a large negative effect on the extensive margin rather than the intensive margin. Since the results in the previous section did not vary with different bandwidths I restrict my analysis to the 0.003 bandwidth.

A potential candidate for this unobserved confounder could be a hazardous working environment which simultaneously makes workers more likely to get cancer ($s > 0$) and more likely to stop working even if they don't get cancer ($d < 0$). If we imagine that this hazardous working environment is the confounder, we can generate this selection bias with just the four p_{ij} parameters that it would produce from Equation (14) if it was U. A new variable, which is characterized by these four p_{ij} parameters, is then added to the matching process and a new ATT is calculated. This is done 200 times and an average of the ATTs is taken. This test is similar to that of Altonji et al. (2005) but with the benefit of not having to rely on a specific model to characterize the bias.

In the top panel of Table 5, I hypothesize that the unobserved confounder takes on a distribution similar to some of the binary variables already in the model. For

[14]The SENSATT command builds on the existing ATT commands developed by Becker and Ichino (2002).

example, if the omitted variable were to take on a distribution similar to the binary variable indicating whether the respondent is a drinker, then its inclusion in the model would result in the kernel matching estimate of the negative effect of cancer on working decreasing in magnitude from 10.2 percentage points to 10 percentage points, a negligible change. In fact, this is the largest change that happens when the omitted variable takes on a distribution like one of the binary variables already in the model. Again, this is to be expected given the weak relationship between these variables and cancer diagnosis.

I also vary the values of both d and s manually and examine how the matching estimate changes in response. This allows me to demonstrate how strong both of the effects would have to be in order to drive the estimate to zero. In order to reduce the dimensionality of the problem in the search for the "killer" confounders, I need a system of four equations to identify the four p_{ij} parameters that characterize the confounder's distribution. To do this, I fix at pre-determined values the parameters $\Pr(U = 1)$ (the prevalence of the confounder in the whole sample) and $p_{11} - p_{10}$ (the difference in the probability of having the confounder given that you have a positive labor market outcome in the presence of the treatment and the probability of having the confounder given that you have a negative labor market outcome in the presence of the treatment). The values they are fixed at are

$$\begin{aligned} \Pr(U = 1) = p_{11} \cdot \Pr(Y = 1|D = 1) \cdot \Pr(D = 1) + \\ p_{10} \cdot \Pr(Y = 0|D = 1) \cdot \Pr(D = 1) + \\ p_{01} \cdot \Pr(Y = 1|D = 0) \cdot \Pr(D = 0) + \\ p_{00} \cdot \Pr(Y = 0|D = 0) \cdot \Pr(D = 0) = 0.2. \end{aligned} \tag{18}$$

$$p_{11} - p_{10} = 0. \tag{19}$$

These values indicate that the prevalence of the confounder in the sample is 20%, while the effect of the confounder on the treated outcome is normalized to zero. Ichino et al. (2008) notes that these parameters can be held constant at fixed values as they are not expected to threaten the validity of the original estimates. In the case of Equation (18),

it is not the prevalence of the confounding variable in the sample that is the issue but the fact that it causes selection into the treatment variable. We have already represented this selection problem as s. The value of 20% is chosen simply because the prevalence of the other health conditions in Table 1 is between 10 - 60%. Similarly, with regards to Equation (19), selection bias does not stem from the fact that the confounder affects the work outcomes of those in the treatment group so we have normalized this effect to be zero. In order for the confounder to be the true cause of the effect it needs to have it's own negative effect on the outcome, separate from the treatment. We have already represented this outcome problem as d. So while the values chosen for the parameters in Equation (18) and Equation (19) are somewhat arbitrary, the important thing is that they are held constant at fixed, known values, in order to simulate values of d and s. Equation (18) and Equation (19) are combined with Equation (16) and Equation (17) to solve for the four p_{ij} parameters when certain numerical values are entered for d and s.

In the second panel of Table 5, the values of both d and s that are required to produce substantial changes in the estimate are far greater than any variable already in the model. This means that there would have to be an unobserved variable which strongly makes the respondents more likely to get cancer and less likely to work if they did not have cancer. Given the major cancer risk factors and other important determinants of work are already included in the model, this seems unlikely. Taken altogether, the analyses presented in Sections 4 and 5 robustly demonstrate that cancer has a negative effect on the labor supply of men over the age of 65.

6 Mechanisms

6.1 Health mechanism

The next step in the analysis is to find out why the respondents who are diagnosed with cancer are reducing their labor supply. The most obvious mechanism would be that cancer affects the respondents' health and that this decrease in health causes the reduction. In order to examine whether it is the change in the health of the respondents

which is driving the reduction in labor supply I replace the outcome variable with a variety of health variables and estimate the effect using the same four models that were used in Table 3.

In particular, I focus on binary outcomes indicating whether the respondents are in poor health and whether health is limiting their work. The results in Table 6 show that cancer increases the probability of reporting these negative health outcomes by between 14 - 16 percentage points. I also examine whether it is mental health or physical health which is affected. These results should be interpreted with care as the original outcome variables are not binary. For mental health, I use a Center for Epidemiologic Studies Depression scale (CES-D) which is a 9-point scale with each point corresponding to a negative response to one of these eight CES-D questions: whether the respondent feels depressed; feels everything is an effort; has restless sleep; feels sad; feels lonely; cannot get going; doesn't feel happy; and doesn't enjoy life. Cancer diagnosis leads to an increase of between 0.2 to 0.3 of one of these points. For physical health, I use a summary of the Activities of Daily Living (ADL) variables. This variable uses a 4-point scale which is the sum of three binary variables indicating whether the respondent has trouble bathing, eating, and dressing. Cancer diagnosis leads to an increase of 0.08 of one of these points. Overall, the results show that cancer has a negative effect on general, mental, and physical health with each result in columns (2), (3), and (4) significant at the 5% level.

Previous research has cautioned against interpreting self-reported health problems in this scenario due to the fear that people may exaggerate illness to explain labor force absence (Anderson and Burkhauser, 1985; Stern, 1989; Bound, 1991). However, this is unlikely to be the case here since a) the workers do not need to exaggerate illness to gain access to Social Security or Medicare at this age and b) they do not face societal pressures to justify leaving the labor force in the way a younger worker would.

6.2 Insurance mechanism

An important factor which may affect the labor supply of workers is health insurance. There already exists a large literature which examines the effect of health on worker

mobility in the presence of employer provided health insurance (EPHI) (see Madrian (1994), Buchmueller and Valletta (1996), and Kapur (1998) for some of the original papers in this area). With regards to cancer, more recent studies have found evidence of job-lock for women (Bradley et al., 2006; Tunceli et al., 2009; Bradley et al., 2013) though the evidence for men is mixed (Tunceli et al., 2009; Bradley et al., 2012). There is the potential for this to be mitigated in this sample since the respondents are all eligible for Medicare and are not relying on continuous employment to maintain their health insurance. However, previous studies which have examined the effect of Medicare eligibility on worker mobility have also found mixed results (Fairlie et al., 2011, 2016). Since the previous literature offers no clear pattern for male cancer survivors and men who are eligible for Medicare, the question of whether we expect the effect to differ based on EPHI status will have to be answered empirically.

There is a variable in the HRS data set which indicates EPHI status but this was not used in the main analysis due to a large number of missing observations. In Table 7, I present the results of extra analyses using this EPHI variable. Column (1) gives the results when EPHI status in period $t-1$ is included as an extra control in the LDV model. The new effects calculated are almost identical to the ones in Table 3. The effects are slightly smaller when the sample is restricted to those respondents who did not have EPHI in period $t-1$ and slightly larger when the sample is restricted to those with EPHI in period $t-1$. While the difference between these coefficients may seem large they are not statistically different from zero. For this sample of workers, it appears as if their mobility is not affected by EPHI.

6.3 Hours mechanism

A final way in which the labor supply decision of the respondents may be affected is by the amount of hours they are already working before they are diagnosed. Respondents who work relatively few hours per week may feel a weaker attachment to labor market and may decide to reduce their labor supply at a greater rate in the event of illness. In Table 8, I run the LDV model separately for respondents who work fewer than 30 hours

per week and respondents who work at least 30 hours per week. While the results in Panel A show no difference between the models, the difference in the coefficients in Panel B is quite large. It shows that respondents diagnosed with cancer who were working fewer than 30 hours per week are almost 16 percentage points less likely to be working compared to their non-cancer counterparts. However, the coefficient for those who work at least 30 hours is only 5 percentage points. The difference between these coefficients is also large enough that it is statistically different from zero. For those who do return to work no difference is found, as has been the case throughout the analysis. Similar results are found if the workers are classified as full-time versus not full-time instead of at least 30 hours per week versus fewer than 30 hours per week. It may be tempting to say that the respondents working fewer hours to begin with may have been sicker or more susceptible to disease but this does not appear to be the case. In the Appendix, Table A9 gives some descriptive statistics for both of the groups. The two groups are very well balanced across the demographic and health variables that we saw in Table 1. The only major difference appears to be that those who were working fewer than 30 hours are slightly older than the other group. While this analysis does not say anything about cancer diagnoses changing the respondents' labor-leisure preferences, it does suggest that those who had a greater preference for leisure over labor to begin with were the ones who had the greatest labor supply response.

6.4 Propensity score matching results

Table 9 presents PSM estimates for all of the LDV and Pooled OLS estimates calculated in Section 6. As was the case with the results in Section 5, the PSM estimates generally give the same results as the linear regressions.

7 Discussion

How do these results compare with the results from other studies in the field? With regards to studies on slightly younger survivors (Bradley et al., 2005; Short et al., 2008;

Datta Gupta et al., 2015) the results are very similar. However, care should be taken to not generalize these results to other groups, particularly future cohorts of workers who will work past the age of 65. Those who are working past the age of 65 in this data set are a highly selected sample: they are likely to be healthier than the people who have already left the labor market at this age, or have jobs that require less physical exertion. Given these differences it is also likely that their labor supply response to this shock will be different to the people who have already left the market. For example, we may expect that a worker with an average level of health may retire after a cancer diagnosis, whereas a worker who is in excellent health may still return to work after a cancer diagnosis. However, in the future, the people who would have left the market at 65 will need to remain in order to receive full Social Security benefits. Estimating this model for workers who need to work past the age of 65 for their full retirement benefits would likely give results which are larger in magnitude since the sample would now include workers who are not as healthy, who are in more physical jobs, or who are not as attached to their jobs as the workers in this sample. Nevertheless, policy makers can use this result as a lower bound of the true effect for future cohorts.

If the labor supply reductions are being driven by changes in the health of the respondents, as the results suggest, then it has important implications for future policy decisions. For example, flexible working conditions such as shorter work days or the ability to work from home are likely to benefit both workers who are physically impaired, as well as workers who have had their labor-leisure preferences changed. However, more targeted approaches, such as helping the workers find rehabilitative services from an external provider, would only be of benefit to workers whose health had suffered. Neumark et al. (2015) recently showed that working women who had been diagnosed with breast cancer were more likely to stay working if their employer provided such services. Given that the results in this study are similar in magnitude to those from other studies on younger cohorts, such as Bradley et al. (2005) and Datta Gupta et al. (2015), policies which have improved the labor market activity of these survivors may be able to help older cohorts too.

This paper also contains some limitations. The first is that, as with any study of health shocks on employment, I can only observe the effect of cancer on labor supply for those who are healthy enough to stay in the survey. Therefore, attrition is likely to be an issue. It is easy to imagine that the respondents who have been severely affected by cancer would be more likely to leave the survey than the respondents who have only been mildly affected. This again raises the issue of how representative the remaining sample is of employed men who work past the age of 65. What are the implications of these attriters for the results and future policies? With regards to the results we can again think of these estimates as lower bounds of the true effects. If the attriters had been included in the analysis then the effect of cancer on labor supply would likely be greater in magnitude, not closer to zero. This is because we are not including the people who would have had the biggest reduction in labor supply. So while attrition is a limitation of this paper, solving the attrition problem would reinforce the results rather than undermine them. We have a similar situation regarding policy: any remedial programs arising from these results are likely only to be able to help respondents who are healthy enough to take advantage of them and could be of less help to people who are so sick that they drop out of the survey.

The second is that I am unable to explicitly test whether the reduction in labor supply is driven by physical impairment, a change in labor-leisure preferences, or both. However, given that I do find a negative effect on physical and mental health it does suggest at least some of the effect is driven by health impairment. This is in line with recent research from Trevisan and Zantomio (2016) who find that physical impairment is a major driver of labor market exits for men. Also, it should be stated that the results found here are short term results, where workers must be within two years of diagnosis. If the respondents were followed over a longer period of time, as in studies such as Moran et al. (2011) or Short et al. (2008), then the effect of cancer on the labor supply of this age group may be different. Finally, I do not have access to the type of cancer that the respondents are diagnosed with so the results cannot be broken down by cancer subgroup. However, this is not an uncommon limitation as studies such as Datta Gupta et al. (2015), Moran et al.

(2011), Short et al. (2008), and Trevisan and Zantomio (2016) also report the effect for all cases pooled together. Despite these limitations, this paper contributes to the literature by providing a comprehensive analysis of the effect of cancer on the labor supply of men over the age of 65, including demonstrating that EPHI has no effect on the labor supply reduction of these workers.

8 Conclusion

While previous research has documented the negative effect that cancer has on male labor supply, no paper has examined this effect on those who work past the age of 65. Due to the number of new cancer diagnoses within this age group every year, and the relevance of this age group to public policy, addressing this question can provide valuable information as both age and labor demographics shift dramatically. I contribute to this literature by providing a robust analysis of the effect of cancer on male workers over the age of 65. I find that respondents who are diagnosed with cancer work 3 hours fewer per week when compared with workers who were not diagnosed with cancer. This reduction stems from workers leaving the labor market. For the workers who continue to work, no difference is observed between the cancer and non-cancer respondents. Cancer has a negative effect on both the physical and mental health of the respondents. Finally, while the results differ based on whether or not the respondents were working fewer than 30 hours per week as opposed to at least 30 hours per week, the effect of cancer on labor supply does not appear to differ by EPHI status. In the future, further research can help shed light on the exact mechanism by which these reductions occur.

Figures

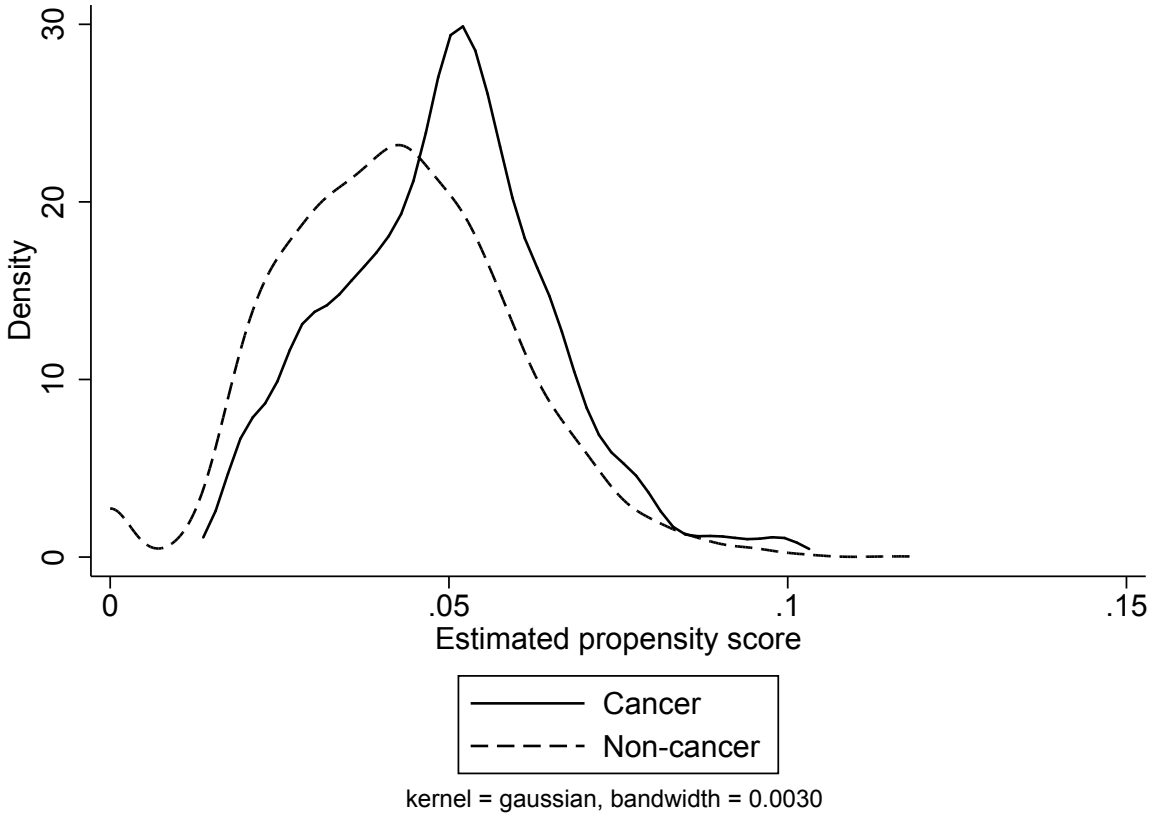

Figure 1: Kernel density estimates

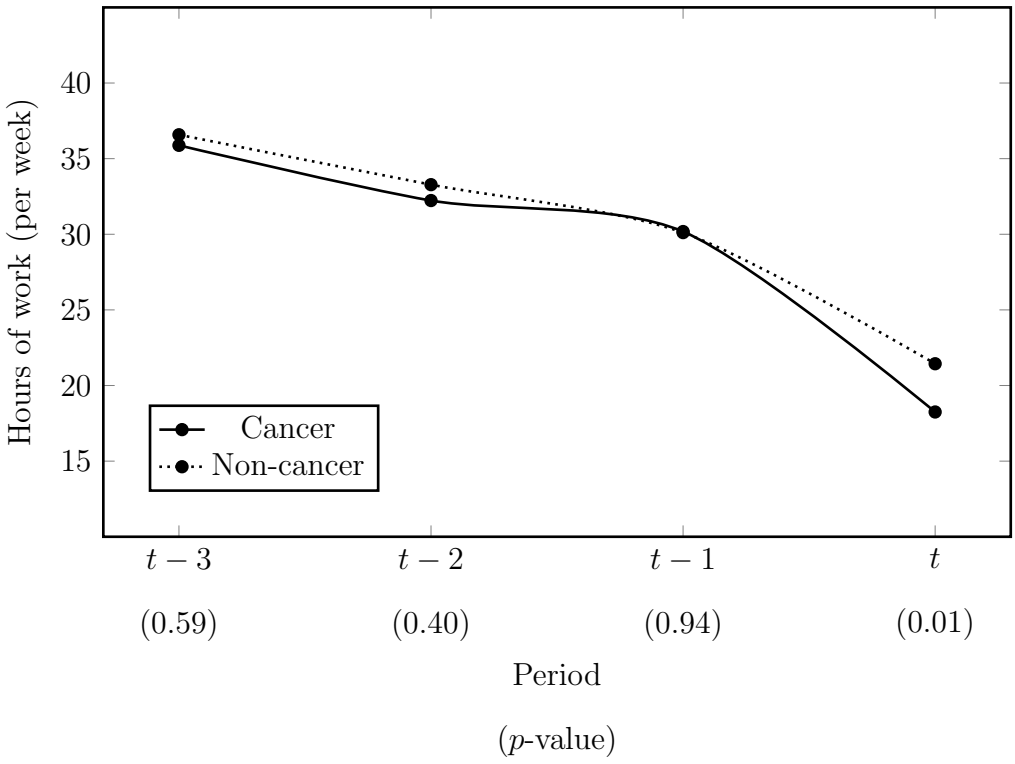

Figure 2: Employment comparison

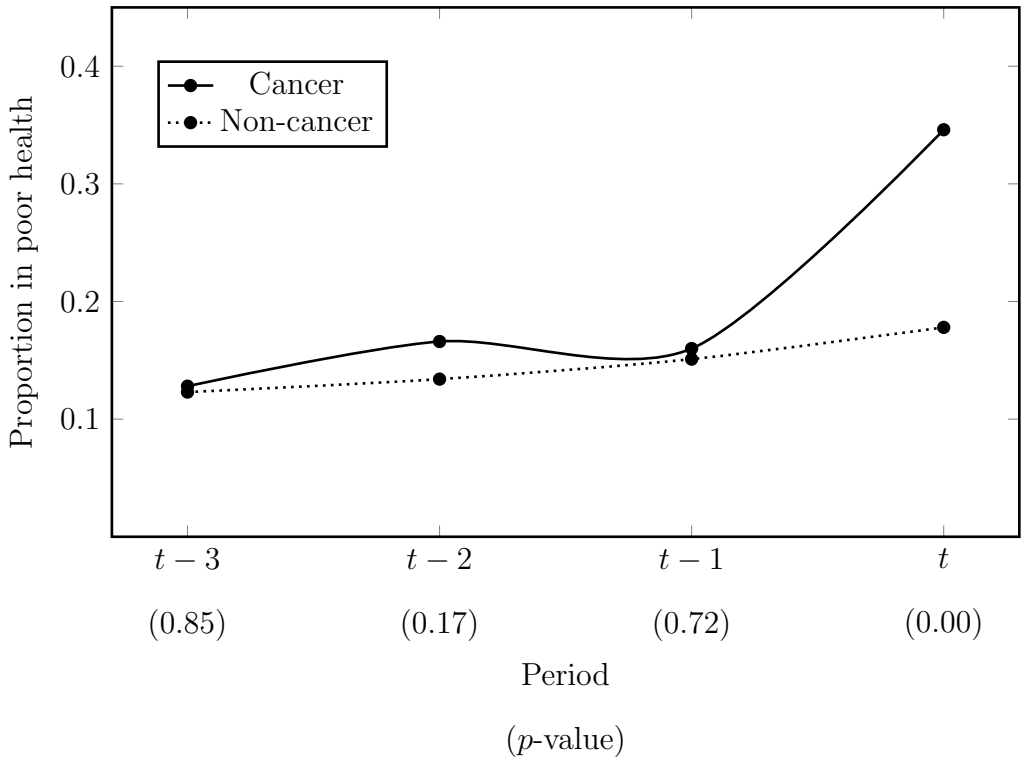

Figure 3: Health comparison

Tables

Table 1: Summary statistics

	(1)	(2)	(3)	(4)	(5)	(6)
	Cancer		Non-cancer			
	Mean	Standard deviation	Mean	Standard deviation	t stat	Nor-diff
Period $t-1$						
Age	69.85	3.66	69.54	3.68	1.30	0.09
Non-white	0.13	0.34	0.14	0.34	-0.15	-0.01
College	0.48	0.50	0.47	0.50	0.20	0.01
Married	0.88	0.32	0.87	0.34	0.77	0.05
Poor health	0.16	0.37	0.15	0.36	0.36	0.02
Smoker (ever)	0.64	0.48	0.68	0.47	-1.19	-0.08
Smoker (now)	0.13	0.33	0.11	0.32	0.65	0.04
Drink alcohol	0.55	0.50	0.60	0.49	-1.68	-0.11
Obese	0.25	0.43	0.25	0.43	-0.02	-0.00
Spouse working	0.37	0.48	0.36	0.48	0.47	0.03
Self-employed	0.39	0.49	0.39	0.49	-0.05	-0.00
Hours of work	30.18	15.73	30.11	16.23	0.07	0.00
IHS (earnings)	7.17	5.02	7.09	5.05	0.25	0.02
IHS (HH income)	11.87	0.81	11.79	0.87	1.33	0.09
IHS (pension)	4.08	5.02	3.53	4.84	1.72	0.11
IHS (SS)	8.80	3.45	8.78	3.48	0.06	0.00
Period t						
Working	0.63	0.48	0.73	0.45	-3.39	-0.22
Hours of week	18.25	18.60	21.44	18.95	-2.55	-0.17
Observations	238		5,364			

Note: Nor-diff - Normalized differences; IHS - Inverse hyperbolic sine; HH - Household.

Table 2: Selection into diseases

	(1)	(2)	(3)	(4)	(5)	(6)
	Cancer	High blood pressure	Heart problems	Lung disease	Arthritis	Diabetes
Period $t-1$						
Past diagnoses	-	0.371***	0.351***	0.164***	0.399***	0.278***
	-	(0.016)	(0.012)	(0.014)	(0.011)	(0.021)
Age	0.001	0.000	0.003***	0.001	0.002**	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Non-white	0.000	0.005	-0.033***	-0.007	-0.003	0.010
	(0.008)	(0.009)	(0.011)	(0.006)	(0.010)	(0.007)
College	-0.000	-0.009	0.004	-0.001	-0.004	0.001
	(0.006)	(0.007)	(0.007)	(0.004)	(0.008)	(0.006)
Married	0.005	-0.012	0.003	0.001	-0.007	0.004
	(0.009)	(0.009)	(0.010)	(0.006)	(0.011)	(0.007)
Poor health	0.003	0.007	0.037***	0.014***	0.031***	0.002
	(0.007)	(0.010)	(0.010)	(0.005)	(0.011)	(0.007)
Smoker (ever)	-0.007	0.021***	0.003	0.018***	0.005	0.003
	(0.006)	(0.007)	(0.007)	(0.006)	(0.008)	(0.005)
Smoker (now)	0.011	-0.024**	-0.009	0.026***	-0.009	-0.007
	(0.009)	(0.011)	(0.011)	(0.005)	(0.012)	(0.009)
Drink alcohol	-0.009*	-0.005	-0.012*	-0.000	0.003	-0.003
	(0.006)	(0.006)	(0.007)	(0.004)	(0.008)	(0.005)
Obese	0.001	0.021***	0.008	0.007	0.022**	0.027***
	(0.006)	(0.008)	(0.008)	(0.004)	(0.009)	(0.005)
Spouse working	0.002	-0.013*	0.002	0.006	0.009	-0.008
	(0.006)	(0.007)	(0.007)	(0.004)	(0.008)	(0.006)
Self-employed	-0.002	0.003	-0.002	-0.008	-0.005	-0.001
	(0.007)	(0.008)	(0.009)	(0.005)	(0.009)	(0.007)
Hours of work	0.000	0.000	-0.000	-0.000	-0.000*	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
IHS (earnings)	-0.000	-0.000	0.000	-0.001*	0.001	-0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
IHS (HH income)	0.005	-0.000	-0.001	0.002	-0.007	-0.004
	(0.004)	(0.003)	(0.004)	(0.003)	(0.005)	(0.003)
IHS (pension)	0.001	-0.001	0.001	0.000	0.002***	0.000
	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)	(0.001)
IHS (SS)	-0.000	0.000	0.001	0.000	0.001	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
p-value	0.493	0.000	0.000	0.000	0.000	0.000
Observations	5,602	5,373	5,477	5,506	5,328	5,498

Note: All models estimated by maximum likelihood. Clustered standard errors (by individual) in parentheses. * Result significant at the 10% level. ** Result significant at the 5% level. *** Result significant at the 1% level. The p-value is for the Wald test for the joint insignificance of all the variables in the logistic regression. IHS - Inverse hyperbolic sine; HH - Household.

Table 3: Linear models

	(1)	(2)	(3)	(4)
	Panel A: Hours of work			
	Pooled	FD	LDV	LDV
Cancer	-3.196***	-3.298***	-3.240***	-3.176***
	(1.233)	(1.172)	(1.084)	(1.074)
Wave dummies	No	Yes	No	Yes
Region dummies	No	No	No	Yes
Period $t-1$ controls	No	No	No	Yes
Observations	5,602	5,602	5,602	5,602
	(5)	(6)	(7)	(8)
	Panel B: Working		Panel C: In labor force	
	Pooled	Pooled	Pooled	Pooled
Cancer	-0.100***	-0.100***	-0.111***	-0.110***
	(0.032)	(0.031)	(0.032)	(0.031)
Wave dummies	No	Yes	No	Yes
Region dummies	No	Yes	No	Yes
Period $t-1$ controls	No	Yes	No	Yes
Observations	5,602	5,602	5,602	5,602
	(9)	(10)	(11)	(12)
	Panel D: Hours of work (if working)			
	Pooled	FD	LDV	LDV
Cancer	-0.368	-1.461	-1.094	-1.025
	(1.277)	(1.008)	(0.931)	(0.907)
Wave dummies	No	Yes	No	Yes
Region dummies	No	No	No	Yes
Period $t-1$ controls	No	No	No	Yes
Observations	4,046	4,046	4,046	4,046

Note: All coefficients are from separate models. All models estimated by OLS. Clustered standard errors (by individual) in parentheses. * Result significant at the 10% level. ** Result significant at the 5% level. *** Result significant at the 1% level. FD - First-differences; LDV - Lagged dependent variable.

Table 4: Propensity score matching

	(1)	(2)	(3)
	Panel A: Hours of work		
	Bandwidth: 0.0015	Bandwidth: 0.003	Bandwidth: 0.006
Cancer	-3.271***	-3.257***	-3.227***
	(1.154)	(1.146)	(1.134)
Observations	5,602	5,602	5,602
	Panel B: Working		
	Bandwidth: 0.0015	Bandwidth: 0.003	Bandwidth: 0.006
Cancer	-0.102***	-0.102***	-0.102***
	(0.032)	(0.032)	(0.031)
Observations	5,602	5,602	5,602
	Panel C: Hours of work (if working)		
	Bandwidth: 0.0015	Bandwidth: 0.003	Bandwidth: 0.006
Cancer	-0.299	-0.313	-0.328
	(1.222)	(1.202)	(1.192)
Observations	4,046	4,046	4,046

Note: All coefficients are from separate models. Bootstrapped standard errors in parentheses. * Result significant at the 10% level. ** Result significant at the 5% level. *** Result significant at the 1% level.

Table 5: Sensitivity of matching estimates to unobserved confounder

Original 'Working' estimate = -0.102

Panel A: 'Calibrated' confounder

Confounder like	New estimate	d	s
Non-white	-0.102	-0.03	-0.01
College	-0.103	0.10	0.01
Married	-0.102	0.02	0.01
Poor health	-0.102	-0.05	0.01
Smoker (ever)	-0.103	-0.05	-0.04
Smoker (now)	-0.101	-0.04	0.02
Alcohol drinker	-0.100	0.03	-0.05
Obese	-0.102	-0.02	0.00

Panel B: 'Killer' confounder

Confounder like	New estimate	d	s
Manually constructed	-0.089	-0.10	0.10
Manually constructed	-0.077	-0.10	0.20
Manually constructed	-0.064	-0.10	0.30
Manually constructed	-0.077	-0.20	0.10
Manually constructed	-0.056	-0.20	0.20
Manually constructed	-0.025	-0.20	0.30
Manually constructed	-0.063	-0.30	0.10
Manually constructed	-0.031	-0.30	0.20
Manually constructed	0.008	-0.30	0.30

Note: $s = p_{1.} - p_{0.} > 0$ - the probability of having the confounder given you are in the treated group is greater than the probability of having the confounder given you are in the untreated group; $d = p_{01} - p_{00} < 0$ - the probability of having the confounder given you have a negative work outcome is greater than the probability of having the confounder given you have a positive work outcome, conditional on being in the untreated group.

Table 6: Health mechanism

	(1)	(2)	(3)	(4)
	Panel A: Poor health			
	Pooled	FD	LDV	LDV
Cancer	0.168***	0.155***	0.164***	0.164***
	(0.031)	(0.031)	(0.029)	(0.029)
Observations	5,601	5,601	5,601	5,601
	Panel B: Health limits work			
	Pooled	FD	LDV	LDV
Cancer	0.120***	0.158***	0.143***	0.137***
	(0.031)	(0.032)	(0.031)	(0.030)
Observations	5,237	4,978	4,978	4,978
	Panel C: CES-D			
	Pooled	FD	LDV	LDV
Cancer	0.159	0.306***	0.224**	0.219**
	(0.110)	(0.100)	(0.097)	(0.097)
Observations	5,076	4,842	4,842	4,842
	Panel D: ADL			
	Pooled	FD	LDV	LDV
Cancer	0.082**	0.086**	0.083***	0.081**
	(0.033)	(0.034)	(0.032)	(0.032)
Observations	5,600	5,490	5,490	5,490
Wave dummies	No	Yes	No	Yes
Region dummies	No	No	No	Yes
Period $t-1$ controls	No	No	No	Yes

Note: All coefficients are from separate models. All models estimated by OLS. Clustered standard errors (by individual) in parentheses. * Result significant at the 10% level. ** Result significant at the 5% level. *** Result significant at the 1% level. FD - First-differences; LDV - Lagged dependent variable; CES-D - Center for Epidemiologic Studies Depression Scale; ADL - Activities of Daily Living.

Table 7: Insurance mechanism

	(1)	(2)	(3)	(4)
	Panel A: Hours of work			
	LDV	LDV	LDV	p-value for (2) - (3) = 0
Cancer	-3.066***	-2.854**	-4.065**	0.5937
	(1.132)	(1.443)	(1.771)	
Observations	5,274	3,299	1,975	5,274
	Panel B: Working			
	Pooled	Pooled	Pooled	p-value for (2) - (3) = 0
Cancer	-0.100***	-0.093**	-0.127**	0.6016
	(0.032)	(0.040)	(0.054)	
Observations	5,274	3,299	1,975	5,274
	Panel C: Hours of work (if working)			
	LDV	LDV	LDV	p-value for (2) - (3) = 0
Cancer	-0.705	-0.320	-1.758	0.4086
	(0.947)	(1.297)	(1.180)	
Observations	3,819	2,342	1,477	3,819
EPHI_{t-1} included	Yes	No	No	-
$\text{EPHI}_{t-1} = 0$	No	Yes	No	-
$\text{EPHI}_{t-1} = 1$	No	No	Yes	-
Wave dummies	Yes	Yes	Yes	-
Region dummies	Yes	Yes	Yes	-
Period $t-1$ controls	Yes	Yes	Yes	-

Note: All coefficients are from separate models. All models estimated by OLS. Clustered standard errors (by individual) in parentheses. * Result significant at the 10% level. ** Result significant at the 5% level. *** Result significant at the 1% level. LDV - Lagged dependent variable; EPHI - Employer Provided Health Insurance.

Table 8: Hours mechanism

	(1)	(2)	(3)
	Panel A: Hours of work		
	LDV	LDV	p-value for (1) - (2) = 0
Cancer	-3.046**	-3.379*	0.8739
	(1.212)	(1.732)	
Observations	2,556	3,046	5,602
	Panel B: Working		
	Pooled	Pooled	p-value for (1) - (2) = 0
Cancer	-0.157***	-0.054	0.0943*
	(0.047)	(0.041)	
Observations	2,556	3,046	5,602
	Panel C: Hours of work (if working)		
	LDV	LDV	p-value for (1) - (2) = 0
Cancer	-0.405	-1.565	0.5461
	(1.587)	(1.118)	
Observations	1,729	2,317	4,046
$\text{Hours}_{t-1} < 30$	Yes	No	-
$\text{Hours}_{t-1} \geq 30$	No	Yes	-
Wave dummies	Yes	Yes	-
Region dummies	Yes	Yes	-
Period $t-1$ controls	Yes	Yes	-

Note: All coefficients are from separate models. All models estimated by OLS. Clustered standard errors (by individual) in parentheses. * Result significant at the 10% level. ** Result significant at the 5% level. *** Result significant at the 1% level. LDV - Lagged dependent variable.

Table 9: Propensity score matching estimates for subsections

	(1)	(2)	(3)	(4)
			Panel A: Table 6	
	Poor health	Health limits work	CES-D	ADL
Cancer	0.165***	0.136***	0.220**	0.081***
	(0.029)	(0.032)	(0.097)	(0.031)
			Panel B: Table 7	
		Hours of work	Working	Hours of work (if working)
Cancer ($EPHI_{t-1}$ included)		-3.020***	-0.100***	-0.718
		(1.079)	(0.032)	(1.106)
Cancer ($EPHI_{t-1}$=0)		-3.003**	-0.095**	0.592
		(1.348)	(0.038)	(1.691)
Cancer ($EPHI_{t-1}$=1)		-3.408	-0.116*	-2.161
		(2.106)	(0.061)	(1.984)
			Panel C: Table 8	
		Hours of work	Working	Hours of work (if working)
Cancer ($Hours_{t-1} < 30$)		-3.240**	-0.162***	-0.448
		(1.375)	(0.052)	(1.778)
Cancer ($Hours_{t-1} \geq 30$)		-3.075	-0.046***	-1.464
		(1.936)	(0.045)	(1.372)

Note: All coefficients are from separate models. Bootstrapped standard errors in parentheses. * Result significant at the 10% level. ** Result significant at the 5% level. *** Result significant at the 1% level.

References

Altonji, J. G., Elder, T. E. and Taber, C. R. (2005), 'Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools.', *Journal of Political Economy* **113**(1), 151 – 184.

American Cancer Society (2016a), 'Cancer facts and figures', `https://www.cancer.org/content/dam/cancer-org/research/cancer-facts-and-statistics/annual-cancer-facts-and-figures/2016/cancer-facts-and-figures-2016.pdf`. Accessed: 2017-10-17.

American Cancer Society (2016b), 'Estimated new cancer cases by sex and age', `https://www.cancer.org/content/dam/cancer-org/research/cancer-facts-and-statistics/annual-cancer-facts-and-figures/2017/estimated-new-cases-for-the-four-major-cancers-by-sex-and-age-group-2017.pdf`. Accessed: 2017-10-17.

American Society of Clinical Oncology (2012), 'Aging and cancer', `http://www.cancer.net/navigating-cancer-care/older-adults/aging-and-cancer`. Accessed: 2017-10-17.

Anderson, K. H. and Burkhauser, R. V. (1985), 'The retirement-health nexus: A new measure of an old puzzle.', *Journal of Human Resources* **20**(3), 315 – 330.

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

Becker, S. O. and Ichino, A. (2002), 'Estimation of average treatment effects based on propensity scores.', *The Stata Journal* **2**(4), 358 – 377.

Bertrand, M., Duflo, E. and Mullainathan, S. (2004), 'How much should we trust differences-in-differences estimates?', *Quarterly Journal of Economics* **119**(1), 249 – 275.

Bound, J. (1991), 'Self-reported versus objective measures of health in retirement models.', *Journal of Human Resources* **26**(1), 106 – 138.

Bradley, C. J., Neumark, D. and Barkowski, S. (2013), 'Does employer-provided health insurance constrain labor supply adjustments to health shocks: New evidence on women diagnosed with breast cancer.', *Journal of Health Economics* **32**(5), 833 – 849.

Bradley, C. J., Neumark, D., Luo, Z. and Bednarek, H. (2006), 'Employment-based health insurance, illness, and labor supply of women: evidence from married women with breast cancer.', *Health Economics* **16**(7), 719–34.

Bradley, C. J., Neumark, D., Luo, Z., Bednarek, H. and Schenk, M. (2005), 'Employment outcomes of men treated for prostate cancer.', *Journal of the National Cancer Institute* **97**(13), 958–65.

Bradley, C. J., Neumark, D., Luo, Z. and Schenk, M. (2007), 'Employment and cancer: findings from a longitudinal study of breast and prostate cancer survivors.', *Cancer Investigations* **25**(1), 47 – 54.

Bradley, C. J., Neumark, D. and Motika, M. (2012), 'The effects of health shocks on employment and health insurance: The role of employer-provided health insurance.', *International Journal of Health Care Finance and Economics* **12**(4), 253 – 267.

Buchmueller, T. C. and Valletta, R. G. (1996), 'The effects of employer-provided health insurance on worker mobility', *Industrial and Labor Relations Review* **49**(3), 439–455.

Bureau of Labor Statistics (2015), 'Labor force projections to 2024: the labor force is growing, but slowly.', `https://www.bls.gov/opub/mlr/2015/article/pdf/labor-force-projections-to-2024.pdf`. Accessed: 2017-10-17.

Caliendo, M. and Kopeinig, S. (2008), 'Some practical guidance for the implementation of propensity score matching.', *Journal of Economic Surveys* **22**(1), 31 – 72.

Cancer Research UK (2016), 'Age and cancer', `http://www.cancerresearchuk.org/about-cancer/causes-of-cancer/age-and-cancer`. Accessed: 2018-02-27.

Candon, D. (2015), 'The effects of cancer on older workers in the english labour market.', *Economics and Human Biology* **18**, 74 – 84.

Centers for Disease Control and Prevention (2015), 'Heart disease facts.', `http://www.cdc.gov/heartdisease/facts.htm`. Accessed: 2017-10-17.

Centers for Disease Control and Prevention (2016), 'Stroke.', `http://www.cdc.gov/stroke/`. Accessed: 2017-10-17.

Coile, C. C. (2004), Health shocks and couples' labor supply decisions. National Bureau of Economic Research, working paper no. 10810.

Congressional Budget Office (2014), 'What are the causes of projected growth in spending for social security and major health care programs?', `https://www.cbo.gov/publication/45543`. Accessed: 2017-10-17.

Danaei, G., Hoorn, S. V., Lopez, A. D., Murray, C. J. and Ezzati, M. (2005), 'Causes of cancer in the world: comparative risk assessment of nine behavioural and environmental risk factors', *The Lancet* **366**(9499), 1784 – 1793.

Datta Gupta, N., Kleinjans, K. J. and Larsen, M. (2015), 'The effect of a severe health shock on work behavior: Evidence from different health care regimes.', *Social Science and Medicine* **136-137**, 44–51.

DeGregori, J. (2013), 'Challenging the axiom: does the occurrence of oncogenic mutations truly limit cancer development with age?', *Oncogene* **32**, 18691875.

Dehejia, R. H. and Wahba, S. (1999), 'Causal effects in nonexperimental studies: Reevaluating the evaluation of training programs.', *Journal of the American Statistical Association* **94**(448), 1053 – 1062.

Dehejia, R. H. and Wahba, S. (2002), 'Propensity score-matching methods for nonexperimental causal studies.', *Review of Economics and Statistics* **84**(1), 151 – 161.

Fairlie, R. W., Kapur, K. and Gates, S. (2011), 'Is employer-based health insurance a barrier to entrepreneurship?', *Journal of Health Economics* **30**(1), 146 – 62.

Fairlie, R. W., Kapur, K. and Gates, S. (2016), 'Job lock: Evidence from a regression discontinuity design.', *Industrial Relations* **55**(1), 92 – 121.

Finkelstein, A. (2004), 'The interaction of partial public insurance programs and residual private insurance markets: evidence from the us medicare program', *Journal of Health Economics* **23**(1), 1 – 24.

Frolich, M. (2004), 'Finite-sample properties of propensity-score matching and weighting estimators.', *Review of Economics and Statistics* **86**(1), 77 – 90.

Garcia-Gomez, P. (2011), 'Institutions, health shocks and labour market outcomes across europe.', *Journal of Health Economics* **30**(1), 200 – 213.

Garcia-Gomez, P. and Lopez-Nicolas, A. (2006), 'Health shocks, employment and income in the spanish labour market.', *Health Economics* **15**(9), 997 – 1009.

Garcia-Gomez, 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.

Health and Retirement Study (2015). RAND version N, public use dataset. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740). Ann Arbor, MI.

Heckman, J., Ichimura, H., Smith, J. and Todd, P. (1998), 'Characterizing selection bias using experimental data.', *Econometrica* **66**(5), 1017 – 1098.

Heckman, J. J., Ichimura, H. and Todd, P. E. (1997), 'Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme.', *Review of Economic Studies* **64**(4), 605 – 654.

Heinesen, E. and Kolodziejczyk, C. (2013), 'Effects of breast and colorectal cancer on labour market outcomes–average effects and educational gradients.', *Journal of Health Economics* **32**(6), 1028 – 1042.

Ichino, A., Mealli, F. and Nannicini, T. (2008), 'From temporary help jobs to permanent employment: What can we learn from matching estimators and their sensitivity?.', *Journal of Applied Econometrics* **23**(3), 305 – 327.

Imbens, G. W. (2015), 'Matching methods in practice: Three examples', *Journal of Human Resources* **50**(2), 373 – 419.

Jeon, S.-H. (2017), 'The long-term effects of cancer on employment and earnings.', *Health Economics* **26**(5), 671 – 684.

Kapur, K. (1998), 'The impact of health on job mobility: A measure of job lock.', *Industrial and Labor Relations Review* **51**(2), 282 – 298.

Kolodziejczyk, C. and Heinesen, E. (2016), 'Labour market participation after breast cancer for employees from the private and public sectors: Educational and sector gradients in the effect of cancer.', *Economics and Human Biology* **21**, 33 – 55.

Li, J., Bonkowski, M. S., Moniot, S., Zhang, D., Hubbard, B. P., Ling, A. J. Y., Rajman, L. A., Qin, B., Lou, Z., Gorbunova, V., Aravind, L., Steegborn, C. and Sinclair, D. A. (2017), 'A conserved nad+ binding pocket that regulates protein-protein interactions during aging', *Science* **355**(6331), 1312–1317.

Madrian, B. C. (1994), 'Employment-based health insurance and job mobility: Is there evidence of job-lock?', *Quarterly Journal of Economics* **109**(1), 27 – 54.

McClellan, M. (1998), Health events, health insurance, and labor supply: Evidence from the health and retirement survey., *in* D. A. Wise, ed., 'Frontiers in the Economics of Aging', Chicago University Press, Chicago, pp. 301 – 346.

Moran, J. R., Short, P. F. and Hollenbeak, C. S. (2011), 'Long-term employment effects of surviving cancer.', *Journal of Health Economics* **30**(3), 505 – 514.

Nannicini, T. (2007), 'Simulation-based sensitivity analysis for matching estimators.', *The Stata Journal* **7**(3), 334 – 350.

National Cancer Institute (2015), 'Age.', `http://www.cancer.gov/about-cancer/causes-prevention/risk/age`. Accessed: 2017-10-17.

Neumark, D., Bradley, C. J., Henry, M. and Dahman, B. (2015), 'Work continuation while treated for breast cancer.', *Industrial and Labor Relations Review* **68**(4), 916–954.

Pence, K. M. (2006), 'The role of wealth transformations: An application to estimating the effect of tax incentives on saving.', *The BE Journal of Economic Analysis & Policy* **5**(1), –.

Rosenbaum, P. and Rubin, D. (1983), 'The central role of the propensity score in observational studies for causal effects.', *Biometrika* **70**(1), 41–55.

Rothman, K. and Poole, C. (1988), 'A strengthening programme for weak associations', *International journal of epidemiology* **17**(4), 955959.

Sheather, S. J. (2004), 'Density estimation', *Statist. Sci.* **19**(4), 588–597.

Short, P. F., Vasey, J. J. and Moran, J. R. (2008), 'Long-term effects of cancer survivorship on the employment of older workers.', *Health Services Research* **43**(1), 193–210.

Smith, J. P. (2005), Consequences and predictors of new health events., *in* D. A. Wise, ed., 'Analyses in the Economics of Aging', Chicago University Press, Chicago, pp. 213 – 237.

Solon, G., Haider, S. J. and Wooldridge, J. M. (2015), 'What are we weighting for?.', *Journal of Human Resources* **50**(2), 301 – 316.

Stern, S. (1989), 'Measuring the effect of disability on labor force participation.', *Journal of Human Resources* **24**(3), 361 – 395.

Taskila-Brandt, T., Martikainen, R., Virtanen, S. V., Pukkala, E., Hietanen, P. and Lindbohm, M.-L. (2004), 'The impact of education and occupation on the employment status of cancer survivors.', *European Journal of Cancer* **40**(16), 2488–93.

Torp, S., Nielsen, R. A., Gudbergsson, S. B., Fossa, S. D. and Dahl, A. A. (2013), 'Change in employment status of 5-year cancer survivors.', *European Journal of Public Health* **23**(1), 116–22.

Trevisan, E. and Zantomio, F. (2016), 'The impact of acute health shocks on the labour supply of older workers: Evidence from sixteen european countries.', *Labour Economics* **43**, 171–85.

Tunceli, K., Short, P. F., Moran, J. R. and Tunceli, O. (2009), 'Cancer survivorship, health insurance, and employment transitions among older workers.', *Inquiry* **46**(1), 17 – 32.

Wooldridge, J. (2016), *Introductory Econometrics: A Modern Approach*, Cengage Learning, Boston, MA.

Table A1: LDV with random effects

	(1) Hours	(2) Working	(3) Hours (if working)
Cancer	-3.338***	-0.109***	-1.074
	(1.079)	(0.031)	(0.949)
Observations	5,602	5,602	4,046

Note: All models are estimated by OLS. Clustered standard errors in parentheses. All models also include the control variables from Table 1, dummy variables for HRS waves, and census regions. * Result significant at the 10% level. ** Result significant at the 5% level. *** Result significant at the 1% level.

Table A2: Fixed effects model

	(1) Hours	(2) Working	(3) Hours (if working)
Cancer	-4.342***	-0.138***	-1.062
	(1.186)	(0.033)	(1.026)
Observations	8,038	8,038	6,482

Note: Before estimating the fixed effects model it is important to remember that each observation in the original sample requires two periods of information to construct it. Therefore, each observation represents a two-wave transition which involves a respondent transitioning from a non-cancer state to a cancer state, or remaining without cancer. This means that for each unique respondent in the main sample, the first observation is lost by constructing these two-wave transitions. In order to run the correct fixed effects model, I need to add this first observation for each unique person back into the sample. Because there are 2,436 unique respondents in the main sample that means adding 2,436 observations to the original sample of 5,602 observations for a total of 8,038. All models are estimated by OLS. Clustered standard errors in parentheses. All models also include dummy variables for HRS waves. * Result significant at the 10% level. ** Result significant at the 5% level. *** Result significant at the 1% level.

Table A3: Sample information

Exclusion criteria	Observations
Unrestricted sample	207,816
Female	120,956
Younger than 65 or older than 78 in period $t-1$	60,853
Previous cancer diagnosis in period $t-1$ (or before)	4,238
Not working in period $t-1$	15,319
Reporting no hours when working (and vice versa) in either period	458
Working more than 80 hours a week in either period	63
Missing information for control variables in $t-1$	327
Restricted sample	5,602

Note: Observations refers to person-wave observations (i.e one individual in the survey for 10 waves appears as 10 observations). While the final sample used is only a fraction of the unrestricted sample, the overwhelming majority of observations excluded results from restricting the data to the population of interest: working men who are over the age of 65 and who have never been diagnosed with cancer before.

Table A4: Stratification matching

	(1) Hours	(2) Working	(3) Hours (if working)
Cancer	-3.229***	-0.102***	-0.319
	(1.133)	(0.031)	(1.198)
Observations	5,602	5,602	4,046

Note: All coefficients are from separate models. Three blocks are used for stratification. Bootstrapped standard errors in parentheses. * Result significant at the 10% level. ** Result significant at the 5% level. *** Result significant at the 1% level.

Table A5: Inverse probability weighting

	(1) Hours	(2) Working	(3) Hours (if working)
Cancer	-3.122***	-0.093***	-0.606
	(1.175)	(0.033)	(1.298)
Observations	5,602	5,602	4,046

Note: All coefficients are from separate models. Bootstrapped standard errors in parentheses. * Result significant at the 10% level. ** Result significant at the 5% level. *** Result significant at the 1% level.

Table A6: Including respondents with previous cancer diagnoses

	(1)	(2)	(3)	(4)
	Panel A: Hours of work			
	Pooled	FD	LDV	LDV
Cancer	-2.932**	-2.887**	-2.897***	-2.810***
	(1.194)	(1.166)	(1.074)	(1.060)
Wave dummies	No	Yes	No	Yes
Region dummies	No	No	No	Yes
Period $t-1$ controls	No	No	No	Yes
Observations	6,487	6,487	6,487	6,487
	(5)	(6)	(7)	(8)
	Panel B: Working		Panel C: In labor force	
	Pooled	Pooled	Pooled	Pooled
Cancer	-0.086***	-0.085***	-0.095***	-0.094***
	(0.032)	(0.031)	(0.032)	(0.031)
Wave dummies	No	Yes	No	Yes
Region dummies	No	Yes	No	Yes
Period $t-1$ controls	No	Yes	No	Yes
Observations	6,487	6,487	6,487	6,487
	(9)	(10)	(11)	(12)
	Panel D: Hours of work (if working)			
	Pooled	FD	LDV	LDV
Cancer	-0.602	-1.031	-0.882	-0.790
	(1.226)	(1.018)	(0.936)	(0.912)
Wave dummies	No	Yes	No	Yes
Region dummies	No	No	No	Yes
Period $t-1$ controls	No	No	No	Yes
Observations	4,689	4,689	4,689	4,689

Note: All coefficients are from separate models. All models estimated by OLS. Clustered standard errors (by individual) in parentheses. * Result significant at the 10% level. ** Result significant at the 5% level. *** Result significant at the 1% level. FD - First-differences; LDV - Lagged dependent variable.

Table A7: Using HRS weights to account for over-represented populations

	(1)	(2)	(3)	(4)
	Panel A: Hours of work			
	Pooled	FD	LDV	LDV
Cancer	-3.486**	-4.197**	-3.886**	-3.872***
	(1.490)	(1.754)	(1.523)	(1.498)
Wave dummies	No	Yes	No	Yes
Region dummies	No	No	No	Yes
Period $t-1$ controls	No	No	No	Yes
Observations	5,306	5,306	5,306	5,306
	(5)	(6)	(7)	(8)
	Panel B: Working		Panel C: In labor force	
	Pooled	Pooled	Pooled	Pooled
Cancer	-0.114***	-0.119***	-0.124***	-0.128***
	(0.037)	(0.037)	(0.037)	(0.037)
Wave dummies	No	Yes	No	Yes
Region dummies	No	Yes	No	Yes
Period $t-1$ controls	No	Yes	No	Yes
Observations	5,306	5,306	5,306	5,306
	(9)	(10)	(11)	(12)
	Panel D: Hours of work (if working)			
	Pooled	FD	LDV	LDV
Cancer	-0.168	-0.611	-0.477	-0.401
	(1.503)	(1.247)	(1.172)	(1.114)
Wave dummies	No	Yes	No	Yes
Region dummies	No	No	No	Yes
Period $t-1$ controls	No	No	No	Yes
Observations	3,833	3,833	3,833	3,833

Note: All coefficients are from separate models. All models estimated by OLS. Clustered standard errors (by individual) in parentheses. * Result significant at the 10% level. ** Result significant at the 5% level. *** Result significant at the 1% level. FD - First-differences; LDV - Lagged dependent variable.

Table A8: Using natural log transformation of wealth instead of IHS

	(1)	(2)	(3)	(4)
	Panel A: Hours of work			
	Pooled	FD	LDV	LDV
Cancer	-3.888**	-3.246**	-3.452**	-3.473**
	(1.527)	(1.509)	(1.390)	(1.374)
Wave dummies	No	Yes	No	Yes
Region dummies	No	No	No	Yes
Period $t-1$ controls	No	No	No	Yes
Observations	3,179	3,179	3,179	3,179
	(5)	(6)	(7)	(8)
	Panel B: Working		Panel C: In labor force	
	Pooled	Pooled	Pooled	Pooled
Cancer	-0.124***	-0.120***	-0.139***	-0.133***
	(0.042)	(0.041)	(0.042)	(0.041)
Wave dummies	No	Yes	No	Yes
Region dummies	No	Yes	No	Yes
Period $t-1$ controls	No	Yes	No	Yes
Observations	3,179	3,179	3,179	3,179
	(9)	(10)	(11)	(12)
	Panel D: Hours of work (if working)			
	Pooled	FD	LDV	LDV
Cancer	-0.678	-0.926	-0.764	-0.868
	(1.614)	(1.351)	(1.254)	(1.241)
Wave dummies	No	Yes	No	Yes
Region dummies	No	No	No	Yes
Period $t-1$ controls	No	No	No	Yes
Observations	2,312	2,312	2,312	2,312

Note: All coefficients are from separate models. All models estimated by OLS. Clustered standard errors (by individual) in parentheses. * Result significant at the 10% level. ** Result significant at the 5% level. *** Result significant at the 1% level. FD - First-differences; LDV - Lagged dependent variable.

Table A9: Demographic and health variables for hours groups

	(1)	(2)	(3)	(4)	(5)	(6)
	< 30 hours per week		≥ 30 hours per week			
	Mean	Standard deviation	Mean	Standard deviation	t stat	Nor-diff
Period $t-1$						
Age	70.29	3.76	68.93	3.49	14.05	0.38
Non-white	0.12	0.32	0.16	0.36	-4.51	-0.12
College	0.49	0.50	0.46	0.50	2.86	0.08
Married	0.88	0.33	0.86	0.35	1.90	0.05
Poor health	0.15	0.36	0.15	0.36	0.38	0.01
Smoker (ever)	0.68	0.47	0.67	0.47	1.21	0.03
Smoker (now)	0.10	0.30	0.12	0.33	-2.36	-0.06
Drink alcohol	0.62	0.49	0.59	0.49	1.79	0.05
Obese	0.24	0.43	0.26	0.44	-1.81	-0.05
Spouse working	0.33	0.47	0.39	0.49	-4.84	-0.13
Self-employed	0.40	0.49	0.38	0.49	1.19	0.03
Observations	2,556		3,046			

Note: Nor-diff - Normalized differences.