FISCAL STUDIES, vol. 41, no. 1, pp. 221–250 (2020) 0143-5671

Health and Employment amongst Older Workers*

JACK BRITTON† and ERIC FRENCH‡

†*Institute for Fiscal Studies*
(jack_b@ifs.org.uk)
‡*University College London; Institute for Fiscal Studies*
(eric.french@ucl.ac.uk)

Abstract

Health and employment are strongly correlated. This paper reviews the existing evidence and brings in new evidence on the following issues: (a) the measurement of health; (b) the impact of health on employment rather than just the association between health and employment; (c) the mechanisms by which health impacts employment; and (d) the likely effect of recent retirement and disability policy changes in the UK. Although the magnitude of the estimated effect of health on employment varies greatly from study to study, some of this variation is driven by the health measure used. Given our preferred measure, the evidence suggests that 5–10 per cent of the employment decline between ages 50 and 70 is due to declining health in England, with the largest effects among low-educated men. Most of the effect comes through declining preferences for work and lower productivity when in bad health, although some of the effect

* Submitted April 2019.

The authors would like to thank the editor James Banks and a very helpful referee, the Social Security Administration (SSA) through the Michigan Retirement Research Center (MRRC grant UM17-15) and the Economic and Social Research Council (Centre for Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (grant ES/M010147/1) and grant ‘Inequality and the Insurance Value of Transfers across the Life Cycle’ (ES/P001831/1)). Financial support for this work is gratefully acknowledged. The opinions and conclusions expressed are solely those of the authors and do not represent the opinions or policy of SSA or any agency of the US federal government.

Keywords: health, retirement, labour supply.

JEL classification numbers: I10, I18, J26.

is from government-provided incentives to not work when in bad health, such as from disability benefits.

I. Introduction

Those in poor health are much less likely to be working than those in good health. In data from the English Longitudinal Study of Ageing (ELSA), those aged 50–60 who report that their health limits their ability to work have employment rates that are 44 percentage points lower than those who do not. This suggests that health is a key driver of employment. The risk of poor health, and the lost ability to work, exposes households to extreme financial risk, making the relationship between health and employment an important policy concern.

Meanwhile, numerous countries around the world are engaging in reforms to encourage later retirement. The UK government is no different, and has recently tightened eligibility thresholds for disability benefits and increased the state pension age. However, disability benefits and state pensions were implemented to insure individuals against the risk of health declines that accompany old age. Consequently, these reforms run the risk of reducing benefits for those who need them with little labour supply response if most of those receiving benefits are too ill to work. Understanding the relationship between health and labour supply is important for understanding the welfare consequences of these reforms.

In this paper, we consider the measurement of health, the causal impact of health on employment, and the channels through which that causal impact occurs. We draw on the recent academic literature and our own evidence to expand on recent reviews of the evidence,[1] focusing the discussion on the key policy issues currently in the UK. Although the discussion is heavily UK focused, the issues that we outline apply in almost all developed countries around the world.

The estimated effect of health on employment varies greatly from study to study. We show that many of these differences result from either how the study measures health or how the study attempts to address reverse causality and omitted variables bias problems.

We first discuss the measurement issues. There has been a long-running debate about whether self-reported 'subjective' measures, or 'objective' measures that identify specific conditions (such as cancer or stroke) or health outcomes (such as mortality), should be used to measure health when estimating the effect of health on employment. We confirm that these different measures result in very different estimated effects. For example, using a single

[1] O'Donnell, van Doorslaer and van Ourti, 2015; Blundell, French and Tetlow, 2016; French and Jones, 2017.

objective health measure leads to significantly smaller estimates than when using subjective measures. However, we also show that when a large number of different health measures are used, the difference in estimated effects between using subjective measures, objective measures or subjective measures instrumented by objective measures is modest. This is consistent with the view that using a single objective measure understates the full impact of health on employment.

We then discuss issues with identifying the causal effects, rather than just associations, between health and employment. The association between the two variables may arise from reverse causality (i.e. employment causing better or worse health) and/or biases caused by unobserved variables that are correlated with both health and employment. We highlight some approaches that have been taken in the literature to circumvent these issues. Although the association between health and employment overstates the causal effect of health on employment, the causal effect is still strong.

Nevertheless, in England, declines in health typically explain just 5–10 per cent of the declines in employment for different gender and education groups between ages 50 and 70 (with the largest effects for low-educated men). That is because although health has a large impact on employment, health only declines modestly within this age range. For example, the share of people who self-report work limitations only rises from 19 per cent to 33 per cent between ages 50 and 70.

These estimates do not tell us the channels through which health impacts employment, which might be important to understand in a changing policy environment. We therefore discuss in detail the mechanisms through which health affects labour supply in old age. We provide a simple structural framework for this and draw on the academic literature to discuss the relative importance of each of the channels. We estimate that around one-fifth of the employment gap between healthy and unhealthy people in England can be explained by the presence of disability benefits, which suggests that other factors such as preferences and reduced productivity are more important in explaining the differences in employment rates.

Drawing on the evidence on the importance of health for employment and the channels through which the effect occurs, we discuss the key policy issues in this area that are currently relevant in the UK – namely, the increase in the state pension age and the design of disability insurance. Given that health only declines modestly between ages 50 and 70, it appears that there is a good deal of work capacity at these ages, suggesting there is scope for reforms that boost employment of older people. Nevertheless, disability prevalence is non-trivial in this age range, suggesting that disability benefits provide valuable insurance to those in need. The continued difficulty in reducing the number of disability benefit claimants speaks to the fact that many are in poor health and find it difficult to work.

The paper is set out as follows: Section II outlines the descriptive relationship between health and employment near retirement; Section III considers the measurement of health; Section IV discusses the causal relationship between health and employment; Section V explores the different channels through which health impacts employment; and Section VI applies the lessons learned to the policy environment in the UK. Section VII concludes.

II. The association between health and employment

We start by outlining some stylised facts about the relationship between health and employment in England. For this, we use data from the English Longitudinal Study of Ageing (ELSA). This is a biennial survey of over-50s (and their partners) which began in 2002. It is based heavily on the Health and Retirement Study (HRS) in the US, which we also draw upon at various points throughout this paper.[2]

We start by presenting the relationship between health and employment in Figure 1 by looking at the employment rates of the healthy and unhealthy, and how they change with age. We measure health using the self-reported

FIGURE 1

Employment rate by age and health status

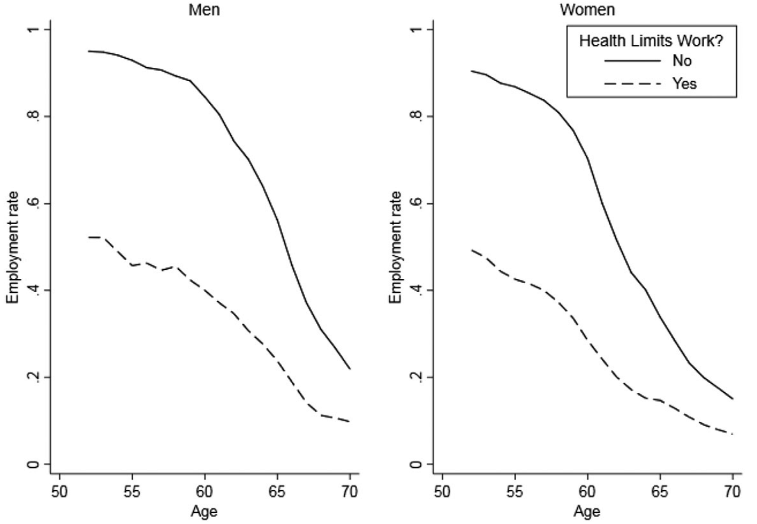

Note: Figures show a three-year moving average.
Source: ELSA, waves 1–6 pooled.

[2]We use waves 1–6 of ELSA, meaning we use it from 2002 to 2012. When we use the HRS, we use waves 3–11, which are from 1996 to 2012.

FIGURE 2

Work hours per week for those in work by age and health status

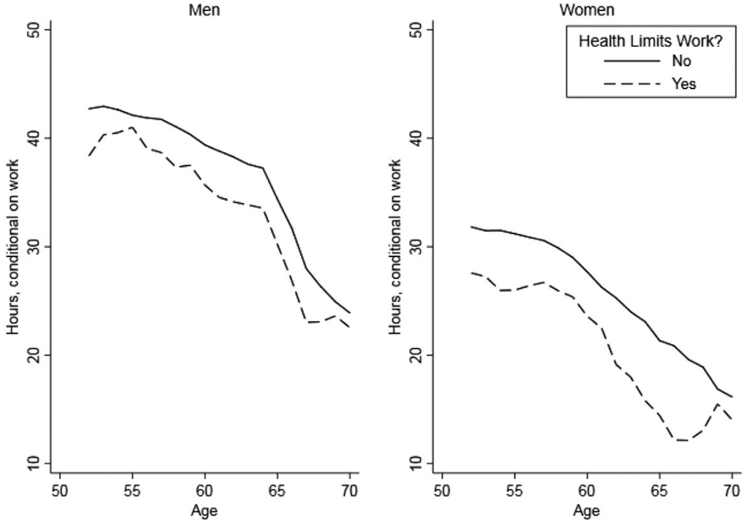

Note: Figures show a three-year moving average.
Source: ELSA, waves 1–6 pooled.

FIGURE 3

Proportion saying that health limits their ability to work, by age and education

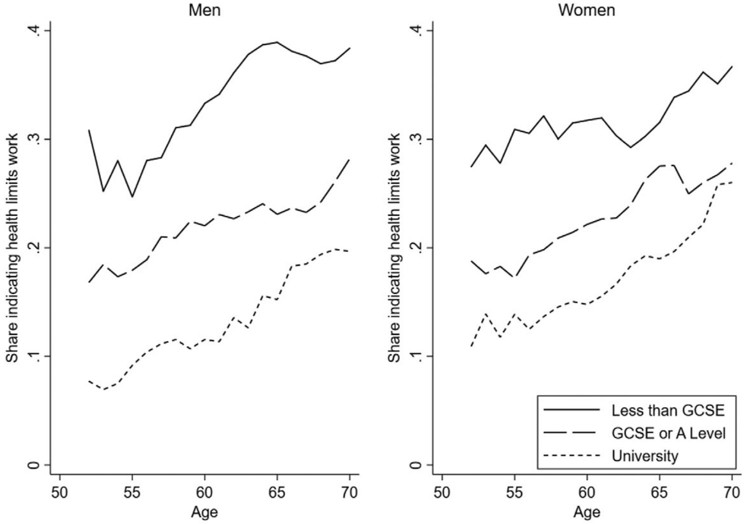

Note: Figures show a three-year moving average. GCSE and A-level qualifications include equivalent qualifications. 'University' indicates graduating rather than simply entering university.
Source: ELSA, waves 1–6 pooled.

'health limits my ability to work' indicator variable from ELSA. The figure shows extremely large differences in employment rates between the healthy and unhealthy, with the unhealthy only around half as likely to be employed as the healthy, for both men and women.

Figure 1 simply shows the employment rate based on a binary indicator of whether somebody is in employment, but Figure 2 looks at the intensive margin by considering differences in work hours. It shows that of those in employment, the unhealthy also tend to work fewer hours per week. Interestingly, the gap is larger for women (the unhealthy work on average 18 per cent fewer hours between ages 50 and 65) than it is for men (8 per cent fewer hours for the same age range).

Of course, these simple descriptive relationships may mask some important differences between the healthy and unhealthy. For example, Figure 3 shows how health varies by education, for which we divide people into 'Less than GCSE' (no GCSE or equivalent qualifications), 'GCSE or A level' (GCSE or A-level or equivalent qualifications) or 'University' (university degree or equivalent) groups.

There is a very clear ordering by education, with the lowest-educated consistently having the worst health and the highest-educated consistently having the best health at all ages, for both men and women. This highlights that the association between health and employment may be confounded by many other factors, which we explore further in Section IV. It is also notable that declines in health are actually relatively modest between ages 50 and 70, with the share reporting that their health limits their ability to work increasing by no more than around 10 percentage points for any gender–education group.

III. The measurement of health

As highlighted in a recent review of the literature by O'Donnell, van Doorslaer and van Ourti (2015), estimates of the effects of health on employment differ enormously from study to study, sometimes by as much as a factor of 10. O'Donnell et al. and French and Jones (2017) advance some potential explanations for the discrepancies between estimates, many of which relate to the measurement of health. The problem faced by the researcher is that health cannot be perfectly observed and often must be based on limited survey responses. These issues were revisited recently in Blundell et al. (2017) (henceforth BBCF).

Table 1 summarises a range of health measures taken from ELSA (England) and the HRS (US). It includes three 'subjective' measures of health that are self-reported, comprehensive measures of health. These are the binary 'health limits work' variable that we used in the previous section, a binary 'health limits activities' variable and a self-reported general health measure that is graded from 1 (excellent) to 5 (poor). We see that these subjective measures

TABLE 1

Subjective and objective health measures for 50- to 70-year-olds in England and the US

	England		US	
	Men	*Women*	*Men*	*Women*
Subjective measures				
Health limits activities	0.41	0.54	0.54	0.67
Self-reported health	2.61	2.57	2.75	2.78
Health limits work	0.24	0.25	0.25	0.27
Objective measures				
Blood pressure	0.30	0.26	0.50	0.50
Arthritis	0.23	0.34	0.44	0.57
Psychiatric	0.05	0.08	0.12	0.21
Lung disease	0.04	0.04	0.08	0.10
Cancer	0.03	0.03	0.08	0.11
Diabetes	0.09	0.06	0.19	0.17
Stroke	0.02	0.01	0.06	0.04
Heart attack	0.03	0.01	0.02	0.01
Sight	0.02	0.02	0.04	0.05
Hearing	0.05	0.02	0.06	0.02
No. of individuals	*6,267*	*7,204*	*11,407*	*13,955*
Total individual–wave obs.	*18,913*	*22,482*	*44,499*	*58,764*

Note: Individuals can be observed multiple times across different waves. The 'Self-reported health' measure uses a scale of 1–5, where 1 represents excellent health and 5 represents poor health. All other variables are binary. More detail on the objective conditions is provided in the online appendix.
Source: BBCF, tables 2 and 3. Data from ELSA (England) and the HRS (US).

give broadly similar averages for men and for women, although women are more likely to report that their health limits their activities.

Table 1 also summarises a set of binary 'objective' measures of health, which are typically based on specific conditions that the individual currently has.[3] Here we see that women are more likely to suffer from arthritis, which might explain why their health is more likely to limit their activities. Women are also more likely to report that they have had psychiatric problems.

Although US respondents look considerably worse than their English counterparts in terms of their *objective* measures, they are actually very similar in terms of their *subjective* measures. One possible reason for this is that US healthcare is better at detecting health conditions before they become severe. Banks et al. (2006), however, show that the differences are not just due to reporting differences: Americans are in worse health according to multiple

[3]See the online appendix for more detail on the definitions of these variables.

biological markers such as their high-density lipoprotein levels. Banks et al. also provide evidence on some of the behavioural risk factors that potentially lead to these health differences, such as the higher incidence of obesity in the US. A second potential reason is that Americans are objectively less healthy, but they set a higher threshold for considering themselves in bad health. Kapteyn, Smith and van Soest (2007) provide some evidence that Americans do indeed set high thresholds for reporting bad subjective health by using health vignettes to show that Americans set higher thresholds for being disabled than the Dutch.

Each of the explanations for the cross-country differences alludes to issues with the measurement of health, which will create issues when we are estimating the impact of health on employment, which is our focus here. Specifically, we are interested in identifying the parameter β in the following:

$$(1) \qquad E_{it} = \alpha + \beta h_{it} + X_{it}^{'}\gamma + \epsilon_{it},$$

where E_{it} is some measure of employment for individual i at time t, h_{it} is a measure of health, X_{it} is a set of important covariates that might drive both health and employment (such as age), and ϵ_{it} is random variation (such as preference for work) that affects employment.

The researcher would ideally like to observe a measure of h_{it} that fully captures work capacity, but must instead draw on the limited set of survey responses, such as those given in Table 1. Many researchers have made use of one of the subjective health measures, such as health limiting work. However, in practice, this is only a proxy for work capacity, and we expect the subjective health measure, h_{it}^{s}, to be related to true work capacity h_{it} as follows:

$$(2) \qquad h_{it}^{s} = h_{it} + u_{it},$$

where u_{it} represents measurement error.

It is likely that these subjective measures are going to be affected by two sources of measurement error. The first is classical measurement error caused by random variation in reporting – Crossley and Kennedy (2002) suggest this is significant by showing that 28 per cent of all respondents change their reported health status when asked the same self-assessed health question twice in an interview. This classical measurement error will, of course, bias β towards zero when h_{it}^{s} is used in place of h_{it}.

A second source of measurement error is justification bias. This is discussed in Bound (1991) and Dwyer and Mitchell (1999). Justification bias occurs if people justify their current situation through their responses to questions on their health status – for example, somebody saying they have a bad health condition to explain why they are currently unemployed. In this case, measured health will be spuriously correlated with employment, biasing β away from zero and resulting in an overstatement of the true effect of health on employment.

One approach to deal with the measurement issues is to draw on more 'objective' measures of health. These typically describe specific conditions, such as cancer or stroke. Since they describe specific conditions, they are likely less subject to justification bias, as – for example – it is unlikely that someone would claim to have cancer in order to justify being out of the labour force. Yet this typically does not fully resolve the measurement issues because measuring objective measures appropriately in a survey is highly challenging. Cancer, for example, comes in many different forms and severities, and even within one person will have highly heterogeneous effects on work capacity depending on the stage of the treatment and development of the illness. Measuring cancer with one binary variable is therefore also fraught with measurement issues. Furthermore, each specific condition only captures one aspect of health. For this reason, we would suspect any individual objective measure to lead to a downward-biased estimate of the overall impact of health on employment, whereas a subjective measure could lead to either an upward- or downward-biased estimate.

BBCF investigate this by estimating a model similar to that described in equation 1 using a range of health measures. They estimate everything separately by gender and education, and make the estimates comparable across health measures by estimating the share of the decline in employment that is explained by declines in health. This is given by

$$(3) \qquad \gamma = \frac{\beta \Delta h}{\Delta E},$$

where Δh is the change in health between ages 50 and 70, ΔE is the change in employment rate over the same period, and β is the estimated effect of health on employment as in equation 1.[4]

Table 2 shows a range of estimates from BBCF. The first panel shows estimates of the share of the employment declines between 50 and 70 explained by health, using the 'health limits work' variable. We see that the estimated share explained is small for all groups, averaging around 2 per cent.

The second panel uses just one of the objective measures – namely, an indicator for whether the individual has high blood pressure. Again the estimates are small. In the following panels, more objective conditions are added, with the estimated effects increasing. However, by the time sight and

[4]In some cases, BBCF use multiple health measures in one model, estimating $E_{it} = \alpha + \sum_j \beta_j h_{it}^j + X_{it}'\gamma + \epsilon_{it}$, where h^1, h^2 etc. are different health measures. In these cases, $\gamma = \sum_j (\frac{\beta_j \Delta h^j}{\Delta E})$.

TABLE 2

Share of employment decline explained by health using a range of health measures

	Men			Women		
	Less than GCSE	*GCSE / A level*	*University*	*Less than GCSE*	*GCSE / A level*	*University*
Subjective health: health limits work						
Health	0.036**	0.024***	–0.003	0.015*	0.022***	0.004
	(0.014)	(0.007)	(0.008)	(0.008)	(0.007)	(0.007)
Objective health: blood pressure only						
Health	0.036**	0.011	0.006	0.012	0.005	0.008
	(0.015)	(0.010)	(0.009)	(0.010)	(0.008)	(0.010)
Objective health: add arthritis, psychiatric, lung disease						
Health	0.066***	0.023*	0.027*	0.043***	0.039***	0.009
	(0.018)	(0.013)	(0.014)	(0.013)	(0.011)	(0.016)
Objective health: add cancer, diabetes, stroke, heart attack						
Health	0.087***	0.033**	0.040**	0.062***	0.037***	0.016
	(0.021)	(0.015)	(0.017)	(0.015)	(0.012)	(0.017)
Objective health: add sight, hearing (full set)						
Health	0.087***	0.033**	0.040**	0.061***	0.037***	0.022
	(0.021)	(0.015)	(0.017)	(0.015)	(0.012)	(0.019)
Combined subjective measures						
Health	0.087***	0.048***	0.045***	0.048***	0.052***	0.024***
	(0.015)	(0.009)	(0.013)	(0.008)	(0.008)	(0.008)
Combined subjective measures instrumented						
Health	0.092***	0.056***	0.057**	0.052***	0.057***	0.030*
	(0.021)	(0.017)	(0.024)	(0.013)	(0.014)	(0.016)
N	*4,692*	*6,327*	*3,362*	*6,957*	*7,911*	*2,759*

Note: Coefficients show the estimated share of employment declines explained by declines in health. Regressions are linear probability models and include controls for age and age squared, survey wave dummies, and initial conditions that include childhood health, prior work experience, and initial health and cognition (when the individual first entered the survey). ***, ** and * indicate statistical significance at the 1 per cent, 5 per cent and 10 per cent levels. Standard errors (shown in parentheses) are bootstrapped with 500 repetitions. Uses ELSA, waves 1–6.
Source: BBCF, tables 8, 10 and 13.

hearing are added, the estimates seem to be settled. The estimates are largest for low-educated men, for whom declines in health explain 9 per cent of the declines in employment between ages 50 and 70. They are typically larger for men than for women and decrease with education.

The fact that adding in more conditions increases the apparent importance of health for employment aligns with Blau and Gilleskie (2001), who draw similar conclusions. While the objective measures are noisy and incomplete, they should be uncorrelated with ϵ_i from equation 1. The assumption here (as discussed above) is that while the more subjective conditions are susceptible to justification bias, the more objective ones are not. Further, while subjective measures might be highly susceptible to – for example – how somebody is feeling on the day, this might be less true of objective measures: people are likely to be much more consistent in their reporting of whether they have had cancer over the last year from one day to the next, than they are in reporting about their general health. This theory is corroborated by Dwyer and Mitchell (1999), who argue that 'there is little evidence of measurement error in the more objective health measures'. They make use of early waves of the HRS to compare the impact of using objective and subjective health measures relative to their instrumented values, and show that although the impact of health varies from health measure to health measure, this is because the different health measures capture different dimensions of health. The authors argue that in fact neither measurement error nor justification bias is a major issue in either the objective or subjective health measures.

Table 2 also shows that there is not very much difference between the estimates using the full set of objective conditions and the estimates using the first principal component of three subjective measures. The final panel shows estimates from instrumenting the subjective measures with the full set of objective measures, following Stern (1989) and Bound (1991), and again the estimates do not change very much.

BBCF conclude that it does not matter very much whether subjective measures (so long as several subjective measures are included), objective measures (so long as several objective measures are included) or subjective measures instrumented with objective measures (so long as several of both measures are included) are used when investigating the effect health has on employment. However, using a small subset of these conditions can lead to very different estimates, which is potentially a key reason why the estimates given in O'Donnell, van Doorslaer and van Ourti (2015) vary so dramatically.[5]

Consistent with the view that using only one objective measure understates the effect of health on employment, several early papers[6] have used subsequent mortality as a measure of health. Bound (1991) shows that this approach yields much smaller estimates than using subjective health (or subjective health instrumented using objective measures such as subsequent mortality), which

[5]BBCF also conclude that once one of these composite measures of health that combines several measures is used, adding an additional index does not make much difference. While there is a small minority of papers (e.g. Gustman and Steinmeier, 2018; Capatina, Keane and Maruyama, 2018) using multiple health measures that all separately affect employment, making use of a single index of health is extremely valuable in the case of structural models, where having multiple indices can hugely increase the computational burden. Related to this, BBCF also investigate the effect of cognition on employment. They find that while cognition is correlated with employment, declines in cognition do not explain any additional declines in employment over and above those explained by declines in health.

[6]For example, Anderson and Burkhauser (1985).

is perhaps unsurprising as many of the conditions that people die from in older age are not well correlated with earlier health which is likely to matter for employment.

The estimates reported in Table 2 are within the range of estimates in the literature. Other studies have also found that declining health can explain only a modest share of the employment decline between ages 50 and 70. For example, Banks, Emmerson and Tetlow (2017) show, using a slightly different methodology, that declining health can explain a 6 percentage point drop in employment between ages 55–59 and ages 70–74 for both men and women. That paper is one of several chapters in Wise (2017) in which similar calculations are performed for multiple countries. It is shown to be consistently true across these countries that declining health can explain only a modest share of the employment decline near retirement.

IV. The causal effect of health on employment

In Section II, we presented some evidence on the association between health and employment. This section explores the extent to which these associations represent the causal effect of health on employment. It discusses some of the challenges faced when trying to estimate the causal effect of health on employment and describes how various papers have attempted to circumvent the methodological challenges.

Many of the papers that have estimated this relationship do so through a regression of some measure of labour supply (E) on health (h), as in equation 1, which we repeat here for convenience:

$$(4) \qquad E_{it} = \alpha + \beta h_{it} + X_{it}^{'}\gamma + \epsilon_{it}.$$

In Section III, we discussed the issues with the measurement of health, which may bias estimates of β. However, even if the researcher is able to observe and measure health perfectly, there are two clear endogeneity problems that may also result in biased estimates. Figure 4 highlights some of the pathways that health and employment are linked. It highlights that health and employment are likely to be correlated for multiple reasons, and thus the strong relationship between health and employment highlighted in the descriptive graphs in Section II should not be interpreted as causal.

First, there might be other variables that are correlated with both health and employment that the researcher might find difficult to control for directly. For example, household resources in childhood are correlated with both health outcomes and economic behaviour in adulthood.[7]

[7]Currie, 2009.

FIGURE 4

The linkages between health and employment

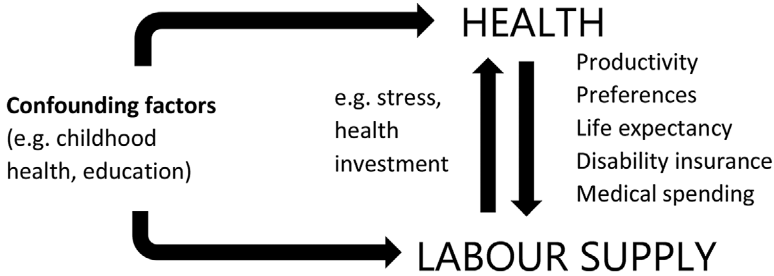

Second, employment may affect health through two channels. One possibility suggested by the Grossman (1972) model is that employment increases investment in health human capital – for example, through better access to medical care, higher-quality food and various other factors associated with better health.[8] Another possibility is that employment may also directly affect stress levels or social behaviours. Several papers have exploited exogenous employment shocks to investigate this relationship. Banks et al. (2019) exploit increases in the British state pension age for women from 60 to 63 (which boosted employment by 11 percentage points) to show that extending working life results in increased scores in cognitive tests and fewer signs of physical disability. Using data from multiple countries, Fonseca, Kapteyn and Zamarro (2017) present evidence that cognition falls after retirement. Black et al. (2018) show that labour market absence driven by receipt of disability benefits increases mortality rates for some groups of people. However, the evidence is not conclusive; for example, Gilleskie (1998) shows that a short period of time away from work can often lead to more rapid recovery from illness.

The reduced form literature has attempted to get around these main sources of bias in various ways. The dominant approach is to focus on how health changes lead to employment changes. The key assumption is that health shocks potentially affect employment, but employment shocks do not affect health. This is done by using first differences, by using fixed effects estimators or by controlling for initial health and employment. For example, Disney, Emmerson and Wakefield (2006) estimate an employment equation using British Household Panel Survey (BHPS) data to estimate a logit model with fixed effects predicting whether the individual is economically inactive or not, controlling for the presence of partners and children, the regional

[8]Recent evidence on the Grossman model includes Fonseca et al. (2009), Hugonnier, Pelgrin and St-Amour (2012), Scholz and Seshadri (2016) and Halliday et al. (2019).

TABLE 3

Regression of employment on health with and without initial conditions

	Men			Women		
	Less than GCSE	*GCSE / A level*	*University*	*Less than GCSE*	*GCSE / A level*	*University*
Without initial conditions						
Health	0.164***	0.123***	0.112***	0.128***	0.137***	0.095***
	(0.006)	(0.008)	(0.015)	(0.007)	(0.007)	(0.014)
σ_{health}	1.027	0.854	0.658	0.885	0.835	0.770
N	*6,555*	*8,007*	*4,103*	*9,417*	*9,614*	*3,226*
With initial conditions						
Health	0.082***	0.054***	0.066***	0.050***	0.075***	0.050***
	(0.008)	(0.009)	(0.017)	(0.007)	(0.008)	(0.014)
σ_{health}	1.153	0.973	0.745	0.994	0.949	0.877
N	*6,555*	*8,007*	*4,103*	*9,417*	*9,614*	*3,226*

Note: Coefficients show the impact of a one-unit change in health on the probability of employment. 'Health' is the first principal component of subjective health, instrumented with objective health. The standard deviation of this measure is given by σ_{health}. Regressions are linear probability models and include controls for age and age squared and ELSA wave dummies. Initial conditions include childhood health, prior work experience, and initial health and cognition (when the individual first entered the survey). ***, ** and * indicate statistical significance at the 1 per cent, 5 per cent and 10 per cent levels. Standard errors (shown in parentheses) are bootstrapped with 500 repetitions. Uses ELSA, waves 1–6.

unemployment rate, housing equity and age. Their explanatory variable of interest is subjective health instrumented with objective measures, following Bound et al. (1999). They find health to be strongly correlated with economic activity, and importantly that for those who transition between states, changes in health are associated with changes in economic activity. In a second approach, Disney, Emmerson and Wakefield (2006) again exploit variation in individuals' health by estimating a hazard function which incorporates both lagged and current health. They conclude from this that their finding that health matters for employment is robust. This is similar to Siddiqui (1997), who also uses a hazard model to show that chronic complaints and disability have a significant positive effect on the probability of early retirement.

BBCF use the 'initial conditions' approach and control for initial health and cognition, as well as childhood health and prior work experience. Table 3 repeats a similar set of estimates to those given in BBCF, estimating a regression model as in equation 4, where E_{it} is a binary employment indicator, h_{it} is subjective health instrumented with objective health (as in the final set of estimates in Table 2) and, in the top panel, X_{it} is simply age and age squared and a set of dummies for the different ELSA waves. In the bottom panel, a set of initial conditions is added, including health and cognition in the first wave the individual is observed in the survey, childhood health and employment history (specifically, the total number of years spent in employment prior to entering the survey).

The table shows that the large and statistically significant relationship between bad health and employment is roughly halved once controls for initial conditions are included. This is consistent with the view that recent health shocks that impact employment can explain part, but not all, of the relationship between health and employment. This aligns with arguments that there are other factors driving some of the cross-sectional relationship between health and employment, such as childhood factors that affect both health and employment, and suggests that the descriptive association overstates the causal relationship.

V. Pathways for health to affect retirement

A major drawback of the reduced form approaches discussed in the previous section is that they give the researcher an estimated effect size without revealing more on the mechanisms through which that effect is occurring. In practice, health can affect retirement through many channels and the overall effect can be ambiguous. A structural model, such as the one used in French (2005), can allow the unpicking of those channels. Here we present a simple structural model that allows health to affect employment through several channels and we discuss evidence from the literature about the importance of those different channels.

Our basic structural framework is given below. We consider a dynamic labour supply model where, given the constraints described below, individuals choose contingency plans of consumption C_{it} and work hours L_{it} to maximise lifetime utility:

$$(5) \qquad \max_{\{C_{it},\ L_{it}\}_{t=0}^{T}} E \sum_{t=0}^{T} \beta^t S_{it} u\left(C_{it}, L_{it},\ h_{it}\right),$$

where β is the discount factor and S_{it} is the probability of surviving from time 0 to time t, which potentially depends on health. The individual maximises utility subject to a budget constraint:

$$(6) \qquad A_{it+1} = (1+r)\, A_{it} - C_{it} + W_{it} L_{it} + B_{it} - \tau\left(W_{it} L_{it}, B_{it}\right) - m_{it},$$

where A is assets, r is the interest rate, W is wages, B is benefits, τ is taxes and m is medical spending. Individuals can accumulate assets which allows them to consume in retirement, although they cannot borrow against future income. We now explain the various channels through which health can affect labour supply in this simple model.

1. Productivity

Poor health can make individuals less productive. This loss of productivity is potentially reflected in lower wages. Most of the literature on the impact of health on wages allows contemporaneous realisations of health, and also age and a shock, to affect wages, such that

(7) $$W_{it} = W\left(t, h_{it},\ \nu_{it}\right).$$

Lower wages will affect the choices of consumption and labour supply. The effects are theoretically ambiguous: a reduction in wages will have a substitution effect that encourages a reduction in labour supply and an income effect that encourages an increase in labour supply. However, to the extent that health shocks are relatively short lived, the income effect from health shocks will be modest and the substitution effect likely dominates.

In our ELSA data, hourly wages of the unhealthy are around 10 per cent lower than those of the healthy (not conditioning on any observable differences in characteristics of the healthy and unhealthy). Using US data, French (2005) finds that when controlling for individual heterogeneity using a fixed effects estimator and for measurement error using different health measures to instrument for one another, those in bad health have approximately 5 per cent lower wages. A key issue is selection – i.e. many of the unhealthy will drop out of the labour force if the offered wage is low. French estimates that the wage penalty rises to 10 per cent when controlling for selection. While non-trivial, the impact of health on wages is not large enough to drive the observed labour supply effect of health.

However, there are multiple reasons that we might think it is not only current health that impacts wages, but past health as well. One potential reason is through work experience, or ‘learning by doing’. In many professions, skills are learned on the job, and time out of the labour force may lead to an atrophy of skills. In a learning-by-doing model, the wage function might allow current wages to be increasing in lagged employment:

(8) $$W_{it} = W\left(t, h_{it},\ L_{it-1}, W_{it-1},\ \nu_{it}\right).$$

This means that wages are not only affected by health directly, but also indirectly as past health affects past employment, which in turn impacts current wages. To see this more closely, note that in the model above, labour supply will depend on the following terms:

(9) $$L_{it} = L\left(t, h_{it}, W_{it}, A_{it}, \nu_{it}\right).$$

Inserting the *lagged* labour supply function into the wage equation and using recursion yields

$$(10) \qquad W_{it} = W^{*}\left(t, h_{it}, h_{it-1}, \ldots, h_{i0}, A_{it-1}, \ldots, A_{i0}, \nu_{it}, \ldots, \nu_{i0}\right),$$

where we would expect that wages would depend not just on current health, but on lagged health also. This suggests that conditioning on only current health may understate the full impact of health on wages.

In Table 4, we use ELSA data in a simple regression framework to explore this issue and to show that it might be important. Specifically, we regress employment on health and lagged health (where, as before, health is subjective health instrumented with objective measures), controlling for a set of possible covariates, such as age and the initial conditions included previously, in Table 3.

Table 4 shows that lagged health matters for employment over and above the effect of current health. This suggests there are indeed channels through which past health affects current labour supply, other than through current health. Calculating the true human capital impact of health shocks remains an open challenge for research in this area.

Capatina, Keane and Maruyama (2018) estimate a life-cycle labour supply model that includes health and learning-by-doing. Their analysis suggests that declining health can explain a much bigger share of employment declines than

TABLE 4

Regression of employment on health and lagged health

	Men			Women		
	Less than GCSE	*GCSE / A level*	*University*	*Less than GCSE*	*GCSE / A level*	*University*
Health	0.094***	0.037***	0.085***	0.056***	0.077***	0.026
	(0.014)	(0.012)	(0.026)	(0.010)	(0.012)	(0.023)
Lag health	0.027*	0.042***	0.003	0.024**	0.035***	0.041**
	(0.014)	(0.012)	(0.025)	(0.011)	(0.011)	(0.019)
σ_{health}	1.153	0.973	0.745	0.994	0.949	0.877
N	*3,766*	*5,079*	*2,694*	*5,689*	*6,277*	*2,088*

Note: Coefficients show the impact of a one-unit change in health on the probability of employment. ‘Health’ is the first principal component of subjective health, instrumented with objective health. The standard deviation of this measure is given by σ_{health}. Regressions are linear probability models and include controls for age, age squared, wave dummies, initial health, initial cognition, prior work experience and childhood health. ***, ** and * indicate statistical significance at the 1 per cent, 5 per cent and 10 per cent levels. Standard errors (shown in parentheses) are bootstrapped with 500 repetitions. Uses ELSA, waves 1–6.

previously thought.[9] However, more research is needed to fully quantitatively assess these mechanisms.

2. Preferences

Equation 5 accounts for preferences by allowing health to directly enter the utility function $u(C_{it}, L_{it}, h_{it})$. Bad health may raise the marginal utility of leisure relative to that of consumption or may reduce the time available for work, which encourages people to drop out of the labour market.

Estimating the importance of preferences is difficult because it is extremely difficult to observe them directly. As a result, structural estimation strategies are used to capture preferences indirectly. The usual approach in structural models is to specify other drivers of employment (such as wages, survival, and disability benefits) in a first stage, and then to estimate preferences by allowing them to vary to match patterns seen in the data.

Most structural papers allow for preferences to be affected by health, but few papers show the relative importance for health relative to other channels. Capatina (2015) and De Nardi, Paschenko and Porapakkarm (2017) both estimate the importance of preferences in structural models. However, neither paper accounts for all the channels through which health might affect employment, which might bias their estimates. For example, neither model includes out-of-work benefits, which provide insurance against bad health shocks and which may result in more people dropping out of the labour force as a result of a bad shock. Not having this mechanism in the structural model could mean the model overstates the importance of preferences in order to match patterns seen in the data. Nevertheless, the finding of both Capatina and De Nardi et al. that preferences are important is likely to hold.

3. Life expectancy

We have life expectancy appearing as S_{it} in equation 5, and we allow it to be directly related to health. Several papers have investigated how health affects longevity. De Nardi, French and Jones (2016) show that lifespan is 3.3 years shorter for those with bad health than for those with good health, while Pijoan-Mas and Ríos-Rull (2014) show the equivalent numbers are 5.6 for men and 4.7 for women at age 50.

A key mechanism for health to affect labour supply is through life expectancy: poor health implies a shorter life span, increasing the incentive to retire early. However, Haan and Prowse (2014) suggest the effects are small: a 6.4-year improvement in life expectancy is associated with just a 0.4-year increase in employment. They add job-loss and job-finding rates to the model

[9]See also Hokayem and Ziliak (2014), Hai and Heckman (2015) and Gilleskie, Han and Norton (2017).

above and find that job prospects in old age are poor. Thus, the scope for added labour in old age is limited, which explains the small effects of this channel in their paper. In models without these employment frictions, however, the employment effects are likely to be much larger.

It is potentially the case that it is not only actual mortality risk which is important for employment, but perceived mortality risk as well. Evidence from the UK in O'Dea and Sturrock (2018) indicates that people are aware that various risk factors, such as smoking and early death of parents, shorten life expectancy, suggesting that it is likely people will adjust their life expectancy in response to shocks to their health. However, this paper highlights a crucial issue in the UK that should concern policymakers. Specifically, people in their 50s and 60s very commonly *underestimate* their survival chances – for example, men (women) born in the 1940s who were interviewed at age 65 reported a 65 per cent (65 per cent) chance of making it to age 75, whereas the official estimate was 83 per cent (89 per cent). This could potentially mean people under-save for retirement, increasing the likelihood of them becoming a burden on the state.

4. Out-of-work benefits

Poor health can allow people to qualify for extra state benefits. This is incorporated in equation 6, where we allow out-of-work benefits B to appear in the assets equation. This creates an income effect for individuals, which may result in a reduction in labour supply. Furthermore, in most cases – such as for disability insurance (DI) in the US and employment & support allowance (ESA) in the UK – the individual will lose those benefits if they are earning above a certain level or working more than a certain number of hours. This creates a substitution effect, which will also reduce labour supply.

Estimating the causal impact of disability benefits on labour supply and wages is a major empirical challenge. Very few people receiving such benefits work, but it is not clear whether this is due to the work incentives of disability benefits or the poor health of those receiving the benefits.[10] Wise (2012) considers natural experiments from DI reforms across multiple countries, including Belgium, Canada, Denmark, France, Germany and Sweden. He concludes that these reforms can have very large effects on the employment of older workers. Another approach is to compare those receiving disability benefits with those who applied for benefits and were denied them. Bound (1989) does this and finds that the employment rate of those allowed benefits versus those denied them was 34 percentage points lower, suggesting that

[10]The structural literature (Kitao, 2014; Low and Pistaferri, 2015; French, von Gaudecker and Jones, 2016) does investigate the importance of disability benefits for labour supply, but this research is only in its infancy.

benefit receipt could reduce the employment of those receiving benefits by up to 34 percentage points. Von Wachter, Song and Manchester (2011) find that these labour supply responses have, if anything, grown over time because applicants are now younger and have potentially less severe health impairments. Of course, as acknowledged by Bound (1989), those allowed and denied benefits are different from one another, creating a selection problem: in particular, those who were allowed benefits are more likely to have serious health problems. Nevertheless, studies that attempt to address the selection problem in the US[11] and other countries[12] find estimates that are only modestly smaller than estimates that do not account for selection.

5. Medical spending and health insurance

The final channel through which health can affect labour supply that we will discuss here is medical spending. Healthcare spending can be significant, especially in the US. De Nardi et al. (2016) show that medical spending is responsible for 18 per cent of US GDP and that much of it is paid privately: for example, the government pays for 65 per cent of healthcare spending by the elderly aged 65+, while 13 per cent is financed out-of-pocket and 13 per cent by private insurance. Among the non-elderly, private insurance becomes a larger payer: 14 per cent is paid out-of-pocket, 45 per cent is paid by private insurance and 25 per cent is paid by the government.

Medical spending generates a negative income effect, which should increase the need to work. However, in practice, many Americans do not pay directly for their medical costs and instead use their health insurance. Prior to the Affordable Care Act (ACA, 'Obamacare'), many US workers could only receive actuarially fair health insurance while they continued to work. In such cases, expensive medical conditions could induce a form of 'job lock', whereby workers delayed retirement to maintain their health insurance coverage.[13] For example, French and Jones (2011) show that those who have access to employer-provided health insurance plans that provide insurance both when working and when retired tend to retire approximately 6 months earlier than those whose plans only provide health insurance when working. On the other hand, disabled individuals are eligible for government-provided health insurance (Medicaid or Medicare, depending on whether they had low or high income prior to disability), but only if they do not work. Thus these benefits provide an incentive not to work among those who are disabled. Therefore, while it is likely the case that the affluent who are in bad health have an

[11]For example, French and Song (2018) use the random assignment of judges to DI cases and exploit variation in their leniency to estimate the causal effects of being assigned DI relative to just missing out. See also Maestas, Mullen and Strand (2013).

[12]Dahl, Kostøl and Mogstad, 2014.

[13]Rust and Phelan, 1997; Blau and Gilleskie, 2006; French and Jones, 2011.

incentive to work in order to receive private healthcare, their less affluent but unhealthy counterparts have an incentive not to work in order to receive government-provided healthcare or free healthcare through default on medical bills.

The Affordable Care Act has given new evidence about the importance of health insurance for employment. Most studies so far have found only modest employment responses to the reforms.[14] It is still not clear why these estimated responses are so small.[15]

In the UK, universal healthcare through the National Health Service means that individuals are insured against medical spending shocks, which means they do not need to work so much to insure themselves against these shocks, and also do not need to work in order to gain access to coverage. On the other hand, poorer individuals in bad health do not need to worry about losing their access to government-provided healthcare if they decide to work.

In short, the impact of medical spending risk on labour supply is ambiguous. For those who rely on employer-provided insurance, medical spending risk provides an incentive for those with high medical spending to stay at their jobs. For those who can qualify for means-tested government insurance, it provides an incentive to stay out of the labour force.

VI. Policy

In this section, we turn our attention to two of the key current policy issues in the UK – namely, choice of the state pension age and the design of disability benefits. We draw on both the empirical evidence and the theory from the previous sections to consider the potentially wide-ranging implications of changes to these policies.

1. Increasing the state pension age

In March 2019, the state pension age was 65 years and 3 months for both men and women, having risen from 60 for women and 65 for men in 2010. This was the first time the age had been the same for men and women since the 1930s.

[14]See, for example, Levy, Buchmueller and Nikpay (2015), French, von Gaudecker and Jones (2016) and Aizawa and Fu (2018).

[15]French, von Gaudecker and Jones (2016) provide some evidence that, for those with low levels of wealth, default on medical bills is a substitute for formal health insurance. This provides households an incentive to have low levels of wealth and not work, since default is only possible for those with low levels of resources. The ACA provided low-cost health insurance to low- and middle-income families, causing default on medical bills to fall. Thus default on medical bills works in much the same way as Medicaid: it provides an incentive to stop working in order to receive either free or heavily subsidised medical care. The ACA reduced this work disincentive, thus meaning that the overall impact of the ACA on labour supply is ambiguous.

TABLE 5

Estimated percentage declines in employment explained by decline in health

	Men			Women		
	Less than GCSE	*GCSE / A level*	*University*	*Less than GCSE*	*GCSE / A level*	*University*
England						
% explained	9.2***	5.6***	5.7**	5.2***	5.7***	3.0*
	(2.1)	(1.7)	(2.4)	(1.3)	(1.4)	(1.6)
US						
% explained	14.8***	12.4***	14.9***	14.2***	13.9***	10.3***
	(2.5)	(1.3)	(2.2)	(2.5)	(1.0)	(1.9)

Note: Coefficients show the estimated share of employment declines explained by declines in health. Health is measured using subjective measures instrumented by objective measures. ***, ** and * indicate statistical significance at the 1 per cent, 5 per cent and 10 per cent levels. Standard errors (shown in parentheses) are bootstrapped with 500 repetitions. English results are based on ELSA, waves 1–6. US results are based on the HRS, waves 3–11. 'Less than GCSE' equates to high school dropout, 'GCSE / A level' to high school and 'University' to college.
Source: BBCF, table 10.

The government has set out its target of increasing the state pension age to 68 for both men and women over the next 20 years. This policy has been driven by concerns about affordability, with an ageing population and increased life expectancy of older people.

As discussed earlier, the empirical evidence suggests that work capacity amongst older people is high. Table 5 shows estimates from BBCF of the share of declines in employment between 50 and 70 that can be explained by declines in health (using subjective health instrumented with objective measures). In England, this share is estimated to be around 5 per cent, although for low-educated men it is closer to 10 per cent. These shares are lower than those in the US, which are consistently between 10 and 15 per cent, similar to the estimates from French (2005).

These estimates suggest that most individuals leave the labour force for reasons other than declining health and are thus still capable of continued work. In fact, the evidence suggests that as the state pension age is raised, people will indeed work for longer: Banks et al. (2019) estimate that each lost year of state pension benefits increases employment by 11 percentage points amongst the women affected by the reforms. However, for the sizeable minority of individuals who are in bad health and are unable to work, the loss of benefits will be a significant financial burden.

An important consideration, however, is that as people increasingly retire later, health will become a more important driver of employment decisions. To provide evidence on this, Figure 5 shows employment declines in the US and

FIGURE 5

Employment changes through people's 60s: England and the US

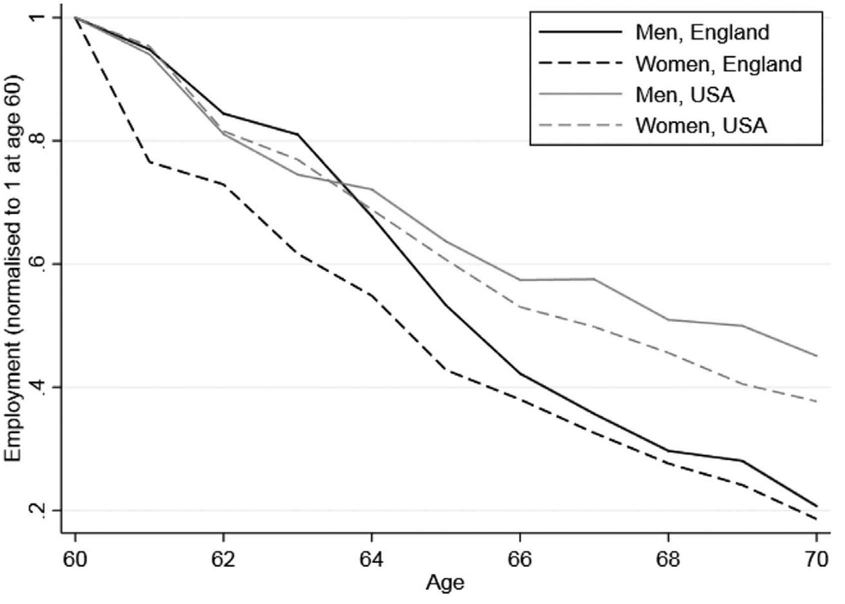

Source: England – ELSA, waves 1–6; US – HRS, waves 3–11.

the UK for men and for women. We see that the rate of decline of employment after age 60 is much faster in the UK than it is in the US. One possible reason for the smaller effects of health on employment in England that we observe in Table 5 is that older people in the US are more likely to be in work and thus more likely to be in work if they experience health deteriorations. It is possible that in England, many people have already left the labour market when they experience health declines.

If increasing the state pension age does indeed mean that more people experience health declines while still in work, this could have important implications for disability benefits. As discussed below, there were wide-scale reforms to disability benefits in England in 1995 that resulted in large reductions in spending on such benefits. However, as discussed in Banks, Blundell and Emmerson (2015), a large part of this reduction was due to people shifting from disability benefits to the state pension. Increases in the state pension age will likely lead to increases in disability benefit claims. This discussion highlights the fact that the two policies cannot be thought of in isolation; we therefore now look at the policy environment for disability benefits in the UK.

2. Disability benefits

There are two main types of disability benefit in the UK: incapacity benefit, which is for those who have experienced short- or long-term health issues that preclude them from working, and disability benefits, such as disability living allowance and the personal independence payment, which are paid to help people with the additional costs of being disabled and (unlike incapacity benefit) is paid independent of work status.

A longstanding policy concern for the UK government has been the number of people on a disability benefit. In the early 1990s, an extremely high number of people were on some form of these benefits, leading to a big reform in 1995, which rebranded the 'invalidity benefit' as 'incapacity benefit' and changed several details to make the assessment criteria stricter.[16] As shown in Banks, Blundell and Emmerson (2015), this resulted in very large reductions in the number of people on disability benefits.

In 2008, there was another reform, which rebranded incapacity benefit as employment & support allowance (ESA). This did not reduce public spending on incapacity by as much as the government had hoped. Emmerson, Joyce and Sturrock (2017) show that government spending on incapacity benefits in 2015–16 was £15 billion, around 50 per cent more than the Office for Budget Responsibility had forecast it would be at that stage in 2012. The lack of saving was partly due to people being assessed as having worse health than the government expected and partly due to moving goalposts. There have been large regional shifts, as well as demographic shifts: in fact, receipt of incapacity benefits has declined significantly amongst high-educated 60- to 65-year-olds, while it has increased significantly for the low-educated young.

As discussed in the previous section, it is a concern for policymakers that benefits such as incapacity benefit generate disincentives to work. In fact, it is potentially true that these disincentives explain a non-trivial share of the employment gap between healthy and unhealthy people. Figure 1 showed that the employment difference between the healthy and the unhealthy is about 40 percentage points amongst those aged 50–65. Based on our own calculations from ELSA, approximately 25 per cent of the unhealthy receive incapacity benefits. As discussed in Section V.4, Bound (1989) showed that the employment rate of those who qualified for disability insurance in the US (which has similarly strict assessment criteria to incapacity benefit) was 34 percentage points lower than the rate for those who did not qualify. More recent estimates that have adjusted for selection into disability insurance have suggested the causal estimates are lower than this but in a similar range.[17] While

[16]The main difference was a switch to assessing individuals' ability to do *any* kind of paid work rather than the kind of work appropriate to them given their skills. Furthermore, there was a change to the rules for assessment, so that it had to be done by regional-level medical staff rather than a personal doctor.

[17]For example, Von Wachter, Song and Manchester (2011) and French and Song (2018).

no directly comparable result is available from the UK, Dal Bianco (2019) uses a structural model and finds that a continuous eligibility reassessment would force about half of disability benefit recipients to exit the programme. Of these, 31 per cent would return to work, which is similar to the US evidence cited above. Based on this, if we assume that employment among incapacity benefit recipients is 30 percentage points lower due to benefit receipt, this would imply that incapacity benefit reduces employment rates by $25\% \times 30 = 7.5$ percentage points, or $100 \times 7.5/40 = 19$ per cent of the difference in employment rates between the healthy and the unhealthy.[18]

This is a crude prediction and the actual impact could be different in either direction. However, it is a useful benchmark for policymakers in the absence of more robust causal evidence from the UK. The government has been explicit about its desire to increase the labour market attachment of those in poor health: it previously pledged to reduce the gap in employment by one-half, and while this was recently removed as an official target, the issue remains a key challenge. Our calculation suggests that the majority of the employment gap is driven by factors other than disability benefits creating a disincentive to work, which might help to explain why employment rates of people with disabilities have remained stubbornly low.

As mentioned above, an important concern for government is that increasing the state pension age will put more pressure on the disability benefits system. But perhaps a more important concern is that vulnerable people will slip through the net. In Table 6, we investigate the role of transitory shocks to the health of older workers by education type. Specifically, the table shows the instrumented 'with initial conditions' results from Table 3 and also the results from a similar estimation with fixed effects. In almost all cases, the coefficients are considerably smaller in the fixed effects specification. As a possible explanation of this, imagine health follows the following process:

$$h_{it} = \alpha_i + \pi_{it} + \epsilon_{it}, \tag{11}$$

where α_i is an individual fixed effect, π_{it} is an AR process such that $\pi_{it} = \rho\pi_{it-1} + \omega_{it}$, and ϵ_{it} and ω_{it} are white noise shocks. In this case, it is trivial to show that a fixed effects regression will put more weight on the transitory components than on the permanent components when the variance of the transitory shocks, ϵ_{it}, is larger than the variance of the permanent shocks, ω_{it}, and ρ is close to 1, which Blundell et al. (2016) suggest is likely. Consequently, an explanation for the much smaller fixed effects estimates in Table 6 is that

[18]Another issue in interpreting the reduced-form estimates above is that they do not take into account dynamics. Many people not yet receiving benefits may be out of the labour force in order to qualify for benefits (French and Song, 2018). For this reason, we may be understating the importance of disability benefits in explaining the employment differences between the healthy and the unhealthy.

TABLE 6

Regression of employment on health with and without individual fixed effects

	Men			Women		
	Less than GCSE	*GCSE / A level*	*University*	*Less than GCSE*	*GCSE / A level*	*University*
Without fixed effects (with initial conditions)						
Health	0.082***	0.054***	0.066***	0.050***	0.075***	0.050***
	(0.008)	(0.009)	(0.017)	(0.007)	(0.008)	(0.014)
σ_{health}	1.153	0.973	0.745	0.994	0.949	0.877
N	*6,555*	*8,007*	*4,103*	*9,417*	*9,614*	*3,226*
With fixed effects						
Health	0.078***	0.020*	0.035	0.041***	0.044***	–0.006
	(0.014)	(0.012)	(0.028)	(0.011)	(0.014)	(0.025)
σ_{health}	1.153	0.973	0.745	0.994	0.949	0.877
N	*6,555*	*8,007*	*4,103*	*9,417*	*9,614*	*3,226*

Note: Coefficients show the impact of a one-unit change in health on the probability of employment. ‘Health’ is the first principal component of subjective health, instrumented with objective health. The standard deviation of this measure is given by σ_{health}. Regressions are linear probability models and include controls for age, age squared, wave, childhood health, prior work experience, and initial health and cognition (when the individual first entered the survey). ***, ** and * indicate statistical significance at the 1 per cent, 5 per cent and 10 per cent levels. Standard errors (shown in parentheses) are bootstrapped with 500 repetitions. Uses ELSA, waves 1–6.

health follows a process consisting of both persistent and transitory shocks and that the transitory shocks matter much less than the permanent shocks for employment.

The obvious exceptions in Table 6 are the low-educated, where the coefficients do not decline very much when we switch to the fixed effects specification, and particularly for low-educated men. One explanation for this is that transitory health shocks do matter for this group. Perhaps they are more likely to be doing manual work and are less able to continue working in the case of a minor injury or illness. Or perhaps transitory shocks mean that they learn about their own mortality and decide to stop working. Or perhaps lots of absences from work put off potential employers who think about the risk of absenteeism.

The concern for policy is that this group might be adversely affected by the increase in state pension age, as their health is too bad for them to carry on working in the industry that they know and in which they have built up contacts, but not bad enough for them to qualify for incapacity benefits. This is particularly relevant given the policy changes that have made eligibility stricter by assessing the applicant’s ability to do *any* work rather than work appropriate to their previous work experience. Banks, Blundell and Emmerson (2015) give the example of an economics professor unable to continue doing that job but not considered eligible for incapacity benefits if she could do some other job. But the same could apply to a manual labourer.

This evidence is only suggestive, and the flip side is that both the state, through lower welfare payments, and individuals would potentially benefit from increased engagement with work, which has been argued to have wider benefits for the individual.[19] This is therefore a highly important topic for future research.

VII. Conclusions

This paper surveys the literature on the relationship between health and employment, discusses in detail the pathways through which health affects employment, and considers the policy environment in the UK, drawing on the empirical evidence and economic theory. We show that health and employment are very closely related, with those reporting that their health limits their ability to work having roughly half the employment rate between ages 50 and 70 of those who report that they are healthy.

We show that much, but not all, of this relationship can be interpreted as causal. Through some back-of-the-envelope calculations, we suggest that most of the effect comes through declining preferences for work and lower productivity when in bad health, although some of the effect is from government-provided incentives to not work when in bad health, such as from disability benefits.

Despite this large effect of health on employment, health is only one of many determinants of retirement. For example, in England, declines in health explain only 5–10 per cent of the declines in employment between ages 50 and 70, with the largest effects among low-educated men. This suggests that the work capacity of retired individuals remains fairly high, which is promising for policymakers considering increases in the state pension age.

Supporting information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

- Appendix

References

Aizawa, N. and Fu, C. (2018), 'Local market equilibrium and the design of the public health insurance system', mimeo.

[19]For example, in Banks et al. (2019).

Anderson, K. H. and Burkhauser, R. V. (1985), 'The retirement-health nexus: a new measure of an old puzzle', *Journal of Human Resources*, vol. 20, pp. 315–30.

Banks, J., Blundell, R. and Emmerson, C. (2015), 'Disability benefit receipt and reform: reconciling trends in the United Kingdom', *Journal of Economic Perspectives*, vol. 29, pp. 173–90.

—, Cribb, J., Emmerson, C. and Sturrock, D. (2019), 'The impact of work on cognition and physical disability: evidence from English women', Institute for Fiscal Studies (IFS), Working Paper no. W19/13.

—, Emmerson, C. and Tetlow, G. (2017), 'Health capacity to work at older ages: evidence from the United Kingdom', in D. Wise (ed.), *Social Security Programs and Retirement around the World: The Capacity to Work at Older Ages*, Chicago, IL: University of Chicago Press.

—, Marmot, M., Oldfield, Z. and Smith, J. P. (2006), 'Disease and disadvantage in the United States and in England', *Journal of the American Medical Association*, vol. 295, pp. 2037–45.

Black, B., French, E., McCauley, J. and Song, J. (2018), 'The effect of disability insurance receipt on mortality', mimeo.

Blau, D. M. and Gilleskie, D. B. (2001), 'The effect of health on employment transitions of older men', in S. Polachek (ed.), *Worker Wellbeing in a Changing Labor Market*, Bingley: Emerald Group Publishing Limited.

— and — (2006), 'Health insurance and retirement of married couples', *Journal of Applied Econometrics*, vol. 21, pp. 935–53.

Blundell, R. W., Britton, J., Costa Dias, M. and French, E. (2016), 'The dynamic effects of health on the employment of older workers', Michigan Retirement Research Center, Research Paper no. 2016-348.

—, —, — and — (2017), 'The impact of health on labor supply near retirement', Institute for Fiscal Studies (IFS), Working Paper no. W17/18.

—, French, E. and Tetlow, G. (2016), 'Retirement incentives and labor supply', in J. Piggott and A. Woodland (eds), *Handbook of the Economics of Population Aging*, Amsterdam: Elsevier.

Bound, J. (1989), 'The health and earnings of rejected disability insurance applicants', *American Economic Review*, vol. 79, pp. 482–503.

— (1991), 'Self-reported versus objective measures of health in retirement models', *Journal of Human Resources*, vol. 26, pp. 106–38.

—, Schoenbaum, M., Stinebrickner, T. R. and Waidmann, T. (1999), 'The dynamic effects of health on the labor force transitions of older workers', *Labour Economics*, vol. 6, pp. 179–202.

Capatina, E. (2015), 'Life-cycle effects of health risk', *Journal of Monetary Economics*, vol. 74, pp. 67–88.

—, Keane, M. P. and Maruyama, S. (2018), 'Health shocks and the evolution of consumption and income over the life-cycle', University of New South Wales Business School, Research Paper no. 2018-14a.

Crossley, T. F. and Kennedy, S. (2002), 'The reliability of self-assessed health status', *Journal of Health Economics*, vol. 21, pp. 643–58.

Currie, J. (2009), 'Healthy, wealthy, and wise: socioeconomic status, poor health in childhood, and human capital development', *Journal of Economic Literature*, vol. 47, pp. 87–122.

Dahl, G. B., Kostøl, A. R. and Mogstad, M. (2014), 'Family welfare cultures', *Quarterly Journal of Economics*, vol. 129, pp. 1711–52.

Dal Bianco, C. (2019), 'Labour supply and welfare effects of disability insurance: a survey', *Italian Economic Journal*, vol. 5, pp. 161–89.

De Nardi, M., French, E. and Jones, J. B. (2016), 'Medicaid insurance in old age', *American Economic Review*, vol. 106, pp. 3480–520.

—, —, — and McCauley, J. (2016), 'Medical spending of the US elderly', *Fiscal Studies*, vol. 37, pp. 717–47.

—, Pashchenko, S. and Porapakkarm, P. (2017), 'The lifetime costs of bad health', National Bureau of Economic Research (NBER), Working Paper no. 23963.

Disney, R., Emmerson, C. and Wakefield, M. (2006), 'Ill health and retirement in Britain: a panel data-based analysis', *Journal of Health Economics*, vol. 25, pp. 621–49.

Dwyer, D. S. and Mitchell, O. S. (1999), 'Health problems as determinants of retirement: are self-rated measures endogenous?', *Journal of Health Economics*, vol. 18, pp. 173–93.

Emmerson, C., Joyce, R. and Sturrock, D. (2017), 'Working-age incapacity and disability benefits', in C. Emmerson, P. Johnson and R. Joyce (eds), *The IFS Green Budget: February 2017*, London: Institute for Fiscal Studies.

Fonseca, R., Kapteyn, A. and Zamarro, G. (2017), 'Retirement and cognitive functioning: international evidence', in O. S. Mitchell, P. B. Hammond and S. P. Atkus (eds), *Financial Decision Making and Retirement Security in an Aging World*, Oxford: Oxford University Press.

—, Michaud, P-C., Galama, T. and Kapteyn, A. (2009), 'On the rise of health spending and longevity', RAND, Working Paper no. WR-722.

French, E. (2005), 'The effects of health, wealth, and wages on labour supply and retirement behaviour', *Review of Economic Studies*, vol. 72, pp. 395–427.

— and Jones, J. B. (2011), 'The effects of health insurance and self-insurance on retirement behavior', *Econometrica*, vol. 79, pp. 693–732.

— and — (2017), 'Health, health insurance, and retirement: a survey', *Annual Review of Economics*, vol. 9, pp. 383–409.

— and Song, J. (2018), 'The effect of disability insurance on labor supply: a dynamic analysis', mimeo.

—, von Gaudecker, H-M. and Jones, J. B. (2016), 'The effect of the Affordable Care Act on the labor supply, savings, and social security of older Americans', Michigan Retirement Research Center, Research Paper no. 2016-354.

Gilleskie, D. B. (1998), 'A dynamic stochastic model of medical care use and work absence', *Econometrica*, vol. 66, pp. 1–45.

—, Han, E. and Norton, E. C. (2017), 'Disentangling the contemporaneous and dynamic effects of human and health capital on wages over the life cycle', *Review of Economic Dynamics*, vol. 25, pp. 350–83.

Grossman, M. (1972), 'On the concept of health capital and the demand for health', *Journal of Political Economy*, vol. 80, pp. 223–55.

Gustman, A. L. and Steinmeier, T. L. (2018), 'The role of health in retirement', *Research in Labor Economics*, vol. 46, pp. 229–97.

Haan, P. and Prowse, V. (2014), 'Longevity, life-cycle behavior and pension reform', *Journal of Econometrics*, vol. 178, pp. 582–601.

Hai, R. and Heckman, J. J. (2015), 'A dynamic model of health, education, and wealth with credit constraints and rational addiction', mimeo.

Halliday, T. J., He, H., Ning, L. and Zhang, H. (2019), 'Health investment over the life-cycle', *Macroeconomic Dynamics*, vol. 23, pp. 178–215.

Hokayem, C. and Ziliak, J. P. (2014), 'Health, human capital, and life cycle labor supply', *American Economic Review*, vol. 104, pp. 127–31.

Hugonnier, J., Pelgrin, F. and St-Amour, P. (2012), 'Health and (other) asset holdings', *Review of Economic Studies*, vol. 57, pp. 779–813.

Kapteyn, A., Smith, J. P. and van Soest, A. (2007), 'Vignettes and self-reports of work disability in the United States and the Netherlands', *American Economic Review*, vol. 97, pp. 461–73.

Kitao, S. (2014), 'A life-cycle model of unemployment and disability insurance', *Journal of Monetary Economics*, vol. 68, pp. 1–18.

Levy, H., Buchmueller, T. C. and Nikpay, S. (2015), 'The effect of health reform on retirement', Michigan Retirement Research Center, Research Paper no. 2015-329.
Low, H. and Pistaferri, L. (2015), 'Disability insurance and the dynamics of the incentive insurance trade-off', *American Economic Review*, vol. 105, pp. 2986–3029.
Maestas, N., Mullen, K. J. and Strand, A. (2013), 'Does disability insurance receipt discourage work? Using examiner assignment to estimate causal effects of SSDI receipt', *American Economic Review*, vol. 103, pp. 1797–829.
O'Dea, C. and Sturrock, D. (2018), 'Subjective expectations of survival and economic behaviour', Institute for Fiscal Studies (IFS), Working Paper no. W18/14.
O'Donnell, O. A., van Doorslaer, E. K. A. and van Ourti, T. (2015), 'Health and inequality', in A. B. Atkinson and F. J. Bourguignon (eds), *Handbook of Income Distribution*, vol. 2B, Amsterdam: Elsevier.
Pijoan-Mas, J. and Ríos-Rull, J-V. (2014), 'Heterogeneity in expected longevities', *Demography*, vol. 51, pp. 2075–102.
Rust, J. and Phelan, C. (1997), 'How social security and Medicare affect retirement behavior in a world of incomplete markets', *Econometrica*, vol. 65, pp. 781–831.
Scholz, J. K. and Seshadri, A. (2016), 'Health and wealth in a lifecycle model', mimeo.
Siddiqui, S. (1997), 'The impact of health on retirement behaviour: empirical evidence from West Germany', *Health Economics*, vol. 6, pp. 425–38.
Stern, S. (1989), 'Measuring the effect of disability on labor force participation', *Journal of Human Resources*, vol. 24, pp. 361–95.
Von Wachter, T., Song, J. and Manchester, J. (2011), 'Trends in employment and earnings of allowed and rejected applicants to the social security disability insurance program', *American Economic Review*, vol. 101, pp. 3308–29.
Wise, D. A. (ed.) (2012), *Social Security Programs and Retirement around the World: Historical Trends in Mortality and Health, Employment, and Disability Insurance Participation and Reforms*, Chicago, IL: University of Chicago Press.
— (ed.) (2017), *Social Security Programs and Retirement around the World: The Capacity to Work at Older Ages*, Chicago, IL: University of Chicago Press.