Review of Economic Studies (2024) **00**, 1–35 doi:10.1093/restud/rdae071
© The Author(s) 2024. Published by Oxford University Press on behalf of The Review of Economic Studies Limited.
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.
Advance access publication 25 June 2024

A Structural Analysis of Mental Health and Labour Market Trajectories

Grégory Jolivet
Department of Economics, University of Bristol, UK

and

Fabien Postel-Vinay
Department of Economics, UCL and IFS, UK

First version received August 2020*; Editorial decision December* 2023*; Accepted June* 2024 *(Eds.)*

We analyse the joint life-cycle dynamics of labour market and mental health outcomes while allowing for two-way interactions between work and mental health. We model selection into jobs on a labour market with search frictions, accounting for the level of exposure to stress in each job using data on occupational health contents. Taking our model to British data from *Understanding Society* combined with information from O*NET, we estimate the impact of job characteristics on health dynamics and the effects of health and job stress contents on career choices. We use our model to quantify the effects of job loss, health shocks, or job stress shocks that propagate over the life cycle through both health and work channels. We also estimate the (large) values workers attach to health, employment, or nonstressful jobs.

Key words: Job search, Mental health, Life cycle, Structural estimation

JEL codes: I12, I14, J62, J64

1. INTRODUCTION

Tackling the personal and economic costs of mental ill health is making its way up the list of priorities of social scientists and policy makers alike.[1] As a preliminary step toward the design of efficient policies addressing mental ill health, this paper aims to contribute to a better understanding of the link between individual labour market trajectories (employment, wages, occupations) and mental health outcomes.

1. The World Health Organization has launched a 5-year special initiative for mental health in 2019 (WHO, 2019). The OECD calls for a stronger policy response to mental ill-health, outlining its economic costs (OECD, 2015). See also the recent article by Layard (2017).

The editor in charge of this paper was Kurt Mitman.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

The interaction between work and health is a two-way one: ill health affects labour supply, and conversely working in a stressful job likely affects future mental health. This two-way interaction has been acknowledged in the health economics literature (see Currie and Madrian, 1999 for a survey) and has recently been factored into structural models of health and labour supply (discussed below).

We add to that literature by modelling that two-way interaction jointly with key drivers of labour market careers such as the cost of labour supply, search frictions, and mobility on the job ladder, and by fully estimating our model using longitudinal worker data linked to new measures of the health contents of occupations. Our focus on mental health will allow us to inform recent debates on the mental health consequences of recessions or the increase in work-related stress. Using our model for treatment evaluation, we speak to important questions such as the impact of job loss on mental health, the effect of a mental health shock on careers, or the cost of working in a stressful job.

The main ingredients of our model are as follows. Jobs are characterised by a wage, working hours, and a "health content"—*i.e.* a measure of how much of a strain the job puts on workers' mental health. Individuals' mental health evolves stochastically over time depending on their current job characteristics. Individuals self-select in and out of employment as well as across jobs subject to search frictions. Job characteristics affect individual utility (through the wage or through the disutility of working) but also future health. Conversely, health affects the decision to work through the disutility of labour as working in a stressful job is more difficult when in poor health. Hence, the feedback loop between health and work is built into our model.

We conduct our empirical analysis on a sample of men aged 25–55 on the UK labour market (we later discuss why we focus on this country, gender, and age group). We collate information from two different data sets: individual-level health and labour market data from the UK Household Longitudinal Survey (UKHLS) and data on the health contents of occupations from the O*NET project. UKHLS is a large panel data set with detailed information on individual labour market status, earnings, and labour market transitions. Crucially, UKHLS also conveys information on individual mental health. That information is based on answers to the widely used 12-item Short-Form (SF-12) health questionnaire, which are conveniently aggregated into summary measures of mental (or physical) health. As for occupation data, the O*NET project contains a large set of descriptors conveying information about the mental health contents of occupations. We run a principal component analysis on specific descriptors to construct an indicator of the stress contents of occupations in O*NET. That indicator can then be linked to UKHLS data using standardised occupation codes.

We develop a multi-stage procedure to estimate our structural model. First, the discrete distribution of individual heterogeneity driving labour and health outcomes is estimated building on the recent clustering approach of Bonhomme *et al.* (2019, 2022). The dynamic process of health is estimated in the second stage by maximising the likelihood of health transitions conditionally on individual and job characteristics. In the third stage, we estimate the remaining parameters relating to wage offers, search frictions, and preferences, by matching simulated moments to their empirical counterparts. Importantly, we prove the identification of the model parameters, conditional on the distribution of individual heterogeneity.[2] Even though the model allows us to follow individual trajectories along two dimensions (work and health), the computational cost of

2. In the absence of a formal proof of identification for the distribution of individual heterogeneity, we perform a (tough) numerical test of internal consistency by comparing the clusters obtained on simulated trajectories from the estimated model with those obtained from the data. The very high overlap between these two distributions leads us to conclude that our approach to estimating individual heterogeneity is, at least, internally consistent.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

our estimation procedure remains very manageable. Our approach would thus easily lend itself to useful extensions for further policy analysis. These will be discussed in the conclusion.

Our first set of results follows directly from the estimation of the model. We find that being older, earning a higher wage or working in a less stressful job have a positive causal impact on future mental health. We further estimate individual preference parameters showing that the job stress content, poor mental health, and age all increase the cost of labour supply.

We use the estimated structure of our model to study three types of shocks: a job loss shock, a health shock, and an increase in job stress content. All treatments are allowed to vary with individual heterogeneity, age, health, or labour market status. We find that, while the health effect of a one-time mental health shock dissipates after 3–4 years on average, the effects of such a shock on employment and income can last longer. Moreover, a job loss shock has lasting negative effects not only on labour market outcomes, consistently with the recent labour economics literature (see Davis and Von Wachter, 2011, or Jarosch, 2015), but also on health.

The two-way interaction between health and work is key to the propagation of both health and labour market shocks. For instance, we show that, while an adverse health shock has little impact on the labour supply decisions of workers in less stressful jobs, it can induce workers employed in more difficult jobs to quit and enter the health and career dynamics of an unemployed worker, making them more likely to stay in poor health and finding certain jobs too costly to accept. Our treatment analysis allows for, and quantifies, this type of interaction.

Alongside these treatment evaluations, we can also estimate the (expected) utility cost of being in poor health, of becoming unemployed, or of an increase in job stress content. Consistently with the recent macro-labour literature, we find that the cost of job loss is very high. The value loss due to a severe health shock, while lower than that of job loss, is also substantial, in particular for older workers or for workers employed in stressful jobs. Moreover, we find that workers in stressful jobs are willing to give up between 5% and 14% of their wage (depending on their age) to avoid an adverse health shock. Job health content is also important to workers: 50-year-old workers are willing to forego 6–12% of their wage (depending on their health) to go from a high-stress to a medium-stress job.

While our approach provides a detailed description of health and labour market trajectories, it has certain limitations. First, health affects an individual's flow utility only through the disutility of labour: we cannot separately identify the effects of health on the flow utility of employed and unemployed workers. We can thus quantify the extent to which poor health makes it more costly to work or affects a worker's labour market decisions but we do not capture the utility cost of being in poor health while not working. Our predictions on the welfare costs of poor health should thus be understood as lower bounds. Second, we perform our analysis on UK data where the National Health Service offers universal coverage, and abstract from any health insurance choice issues (see Currie and Madrian, 1999, or Gruber and Madrian, 2004). Third, the effect of job characteristics on health are modelled in a causal reduced-form fashion: we do not delve into individual decisions to invest in medical treatment or healthier living (Gilleskie, 1998 has a structural analysis of medical treatment decisions). Lastly, our analysis focuses on men aged 25–55. We discuss important extensions of our approach to older workers and to women in the conclusion as they are next on our research agenda.

We end the paper with an extension of our model where we allow workers to direct their search to specific job types (stress levels). This serves two purposes: to check the robustness of our results to endogenising state-dependence in job sampling and to tackle the estimation of a rich directed on-the-job search structure which as far as we know, is new. Estimation results suggest that the effects of adverse health or high job stress shocks are still substantial when allowing workers to direct their search, but less so than in the benchmark random search model.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

Related literature. This paper is an attempt to bridge the gap between the structural labour and health literatures by explicitly modelling the two-way interaction between work and health in a quantitative dynamic model with labour supply, labour market frictions, and on-the-job search. Structural models have been used in many ambitious empirical evaluations of the effects of labour market shocks (job loss, recessions, changes in task contents) on inequality, individual trajectories, and welfare. Recent contributions include Lise and Robin (2017) who analyse aggregate shocks, Jarosch (2015) and Burdett *et al*. (2020) who study the effect of job loss or Low *et al*. (2010) who look at productivity and job destruction shocks. The causal effects of those shocks on health have been documented by both the economic (Sullivan and Wachter, 2009; Davis and Von Wachter, 2011) and medical (Fryers *et al*., 2003) literatures, but attempts at fully incorporating them into structural models of individual careers are still relatively scarce.

A branch of the structural literature investigates the effect of health on retirement decisions (French, 2005; French and Jones, 2011) or on long-term labour earnings and welfare (de Nardi *et al*., 2017). Apart from very few exceptions, that literature tends to stay silent on the feedback effect of labour market outcomes on health. Exceptions include Jacobs and Piyapromdee (2016), who seek to explain the reverse retirement behaviour, and Salvati (2020) who studies the (physical) health, labour supply, and savings of women at older ages. Aside from the difference in focus, both these papers differ from ours in that neither features labour market frictions or heterogeneity across occupations in job health contents.

The two-way interaction between work and health has been considered in structural models of health and medical treatment decisions (*e.g.* Papageorge, 2016). These models are geared to the description of health investments rather than career dynamics: due to their design and purpose, they do not include features such as search frictions or job ladder transitions that not only drive labour market turnover and inequality but are also likely to play a role in the propagation of health shocks.

Finally, a few recent structural contributions on health shocks and labour supply allow for a reciprocal link between health and work. Capatina *et al*. (2020) calibrate a sophisticated life-cycle model with human capital accumulation to analyse the contribution of health shocks to earnings inequality and labour market outcomes. In a similar vein, Harris (2019) and Tran (2017) study, respectively, the effect of body weight and mental health on labour market choices while reciprocally allowing occupations to impact weight and mental health. Lastly, Michaud and Wiczer (2018a) calibrate a dynamic labour supply model to study the effect of macroeconomic shocks on workers who can choose between occupations with different disability risks.

Our paper differs from these important contributions in several key dimensions.[3] First, we allow for heterogeneity in the health contents of jobs and introduce an original occupation-level measure of exposure to stress. Our model thus explicitly features a disamenity (stress on the job) that affects both workers' labour supply and their mental health. Second, our model includes labour market frictions and a job ladder mechanism driving worker turnover. The combination of frictions and heterogeneity in job stress contents is important to evaluate the consequences

3. While this section focuses on what our paper has that existing papers do not, we obviously acknowledge that those papers have features that are absent from ours, due to their different focus. For instance, Michaud and Wiczer (2018a) model disability insurance and retirement decisions (see also Michaud and Wiczer, 2018b, for a related paper). Capatina *et al*. (2020) also model retirement and, due to their application in the US context, account for employer-provided health insurance, which is not a first-order issue in our analysis on UK data.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

of health and labour market shocks.[4] We believe that our model is the first to capture the propagation mechanism of health shocks that operate through labour market frictions.[5]

Outline. The paper is organised as follows. Section 2 presents the data, including our new measure of the mental health content of jobs. Section 3 presents our model and Section 4 details our multi-step identification and estimation procedure. Estimation results and model fit are analysed in Section 5. Section 6 presents various counterfactual analyses. In Section 7, we briefly explore an alternative assumption about workers' ability to direct their search to specific job types. Finally, Section 8 concludes and discusses potential extensions. Details on the data, estimation procedure, and identification proofs are gathered in the Appendix.

2. DATA AND DESCRIPTIVE STATISTICS

2.1. *Labour force and health data: UKHLS*

Our main data source is the UK Household Longitudinal Survey (UKHLS), a.k.a. *Understanding Society*. UKHLS is a yearly household panel started in 2009 as the follow up to the British Household Panel Survey (BHPS).[6] It contains yearly observations of respondents' labour market status, wage, and occupation (if any), age, as well as high-frequency observations of any labour market transitions respondents may have experienced between interviews.

We use the ten waves of data available at the time of writing and focus on men aged 25 to 55. We leave women out of this analysis as we do not model fertility decisions, which are likely to affect both labour market and health trajectories. We stop following individuals after 55 because we do not model retirement decisions. We drop individuals who are never employed during our observation period and we stop following individuals after they become self employed. We exclude observations based on proxy interviews as health and some employment questions are not asked in proxy interviews (we keep all of an individual's nonproxy observations). We further delete a small number of observations with inconsistent information on labour market transitions. We will refer to the data set consisting of only the first observation for each individual as the "initial cross section".[7]

Importantly, UKHLS provides information on individuals' health every year along two dimensions, mental and physical. We focus on mental health (our approach could easily be applied to physical health). Mental health is measured by a continuous summary variable, from 0 to 100, based upon the 12-item Short-Form survey (SF-12), on which we provide more details in Appendix A. For tractability and interpretation, we derive four discrete health states from the continuous health indicator, which we label as "Good," "Average," "Poor," and "Severe". As far

4. For example, workers who suffer an adverse health shock causing them to quit their job then have to resume climbing the job ladder from the bottom and, because of search frictions, may have to accept more stressful or less well-paid jobs before going back to a job that is less damaging for their health.

5. On a more technical note, we should also mention that while many studies cited above rely on the calibration of an elaborate model, the results in this paper follow from the full structural estimation of our model, supported by a formal proof of identification.

6. Interviews for each wave of UKHLS take place over a period of 24 months, but these 24 month periods overlap to ensure that each individual is interviewed once a year. So, for example, wave 1 interviews took place in 2009 and 2010, wave 2 interviews took place in 2010 and 2011, and wave 10 (the most recent at the time of writing) was conducted in 2018 and 2019.

7. The number of observations (individuals) for 25- to 55-year-old men is 83,027 (18,000). These numbers go down to 65,786 (15,335) when cutting trajectories after the first self-employment spell and to 59,923 (13,398) after further removing individuals who are never observed in employment. The last substantial sample selection step consists in removing proxy interviews as key variables are unavailable for these, as well as individuals with just one observation at this stage as we cannot construct their trajectory. Our final sample has 33,886 observations for 8,228 individuals.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

as we know, there is no consensus on where to set in SF-12 indicator cutoffs to define specific health states. Several studies from the medical literature (for instance Ware *et al*., 1996; Salyers *et al*., 2000) have assessed the validity of the SF-12 mental health indicator against other indicators of depression, anxiety or more severe mental disorders and showed that it performs well in screening different diagnoses. In particular, the study by Gill *et al*. (2007) shows that thresholds of 36 and 50 are adequate to identify severe psychological conditions and common mental disorders (depression or anxiety), respectively. We use these two threshold to define the "Severe" (score below 36) and "Poor" (between 36 and 50) health states.[8] We could not find references for a cutoff between an "Average" and a "Good" state so we set it at the 80% quantile (57.2) of the score distribution in the initial cross section. The distribution of the continuous mental health indicator and the three cutoffs are shown in Appendix A.

Our model is written at a yearly frequency because health is only measured once a year. Identification of the health transition process at a higher frequency would require unwarranted formal restrictions: we prefer to leave the health process relatively unconstrained and thus work with a yearly frequency. This level of time aggregation prevents our model from replicating observations with more than two job transitions in a year. We remove those observations which represent 0.8% of our overall observations and 6.8% of observations with at least one transition.

Our final work sample has 33,886 observations (relating to 8,228 individuals). Each observation consists of the individual's labour market status (nonemployed, part-time, or full-time employed), their earnings, their occupation, their age, their mental health, their labour market transitions (between full time and part time and/or between employers and/or in and out of employment) and, if they are not employed, the reason why their previous job ended (thus allowing us to separate quits from job losses). Importantly, we will account for time-invariant individual heterogeneity in the estimation using clustering methods, presented in Section 4.1.

2.2. *Occupation data: O*NET*

The health content of jobs is measured based on data from the O*NET program. O*NET, a.k.a. the *Occupational Information Network*, is a database describing occupations.[9] It comes as a list of around 300 descriptors, with ratings of importance, level, relevance, or extent, for almost 1,000 different occupations. O*NET data are compiled in the US. Since, to our knowledge, no comparable data set exists for the UK, we have to work on the assumption that the contents of occupations, as described by O*NET, are identical in the UK and US.

O*NET descriptors are organised into ten broad categories: *Skills*, *Abilities*, *Knowledge*, *Experience/Education Levels Required*, *Job Interests*, *Work Activities*, *Work Context*, *Work Values*, *Work Styles*, and *Tasks*. We retain the 16 descriptors from the category *Work Styles* and the 13 descriptors from the *Structural Job Characteristics* section of the *Work Context* category. These provide the descriptors that are most germane to the stress content of the job; descriptors

8. Scores of 36 and 50 are respectively at the 7th and 36th quantile of the mental health score distribution in the initial cross section.

9. O*NET is developed by the North Carolina Department of Commerce and sponsored by the US Department of Labor. More information is available on https://www.onetcenter.org, or on the related Department of Labor site https://www.doleta.gov/programs/onet/eta_default.cfm.

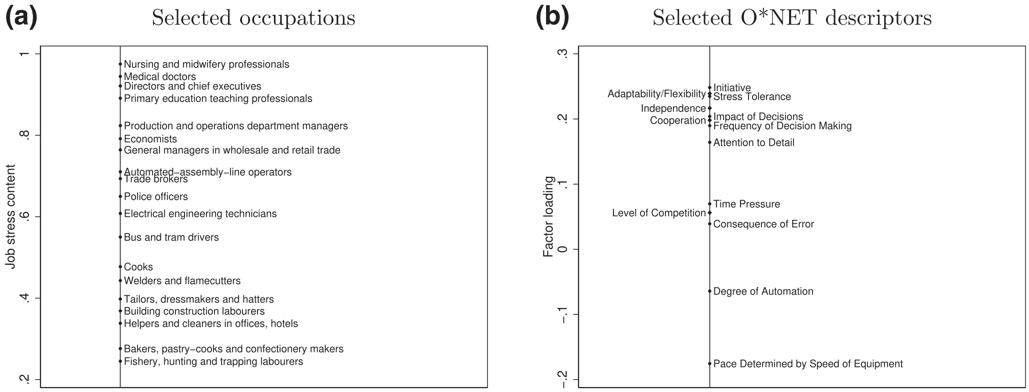

FIGURE 1

Job health content – Illustrations.

contained in the other files are less directly interpretable in terms of mental health contents. This leaves us with 29 descriptors.[10, 11]

O*NET data have been updated at least once a year since 2003. However, not all occupations or descriptors are updated every year: this is done on a rotating basis and it takes about 5 years for all occupation/descriptor pairs to be updated. We process the O*NET data to construct a panel of occupations covering the same time period as the UKHLS and then take averages over time for each occupation/descriptor pair. This leaves us with a cross section of occupations with 29 averaged descriptors. We run a principal component analysis of these 29 descriptors (in the population of occupations) and keep the first component, to which we apply an affine transform to confine it to [0, 1]. We then merge the resulting table into the UKHLS sample described in the previous subsection.

To illustrate our job health content indicator, we show in Figure 1(a), the value of this variable for a selected sample of occupations. Medical, executive or teaching jobs are classified as high-stress while tailors, bakers, or cleaners tend to be low-stress occupations. To further illustrate the job stress indicator, Figure 1(b) shows the factor loadings on a few selected O*NET descriptors (recall that the job stress variable is the first principal component of our set of descriptors). "Stress tolerance," "adaptability/flexibility," or "impact of decisions" have relatively large positive loadings while the "degree of automation" or "pace determined by speed of equipment" both have negative loadings. Figure 1 thus suggests that our job stress measure is more related to the cognitive and interpersonal dimensions of jobs than to their manual skill requirements.

10. From the *Work styles* category, we have: *Achievement/Effort*, *Adaptability/Flexibility*, *Analytical Thinking*, *Attention to Detail*, *Concern for Others*, *Cooperation*, *Dependability*, *Independence*, *Initiative*, *Innovation*, *Integrity*, *Leadership*, *Persistence*, *Self Control*, *Social Orientation* and *Stress Tolerance*. From the *Work context (Structural Job Characteristics)* category we have: *Consequence of Error*, *Degree of Automation*, *Duration of Typical Work Week*, *Freedom to Make Decisions*, *Frequency of Decision Making*, *Impact of Decisions on Co-workers or Company Results*, *Importance of Being Exact or Accurate*, *Importance of Repeating Same Tasks*, *Level of Competition*, *Pace Determined by Speed of Equipment*, *Structured versus Unstructured Work*, *Time Pressure* and *Work Schedules*

11. While O*NET uses many different rating systems, only two are relevant to the *Work Styles* and *Work Context* files that we use. Descriptors from the *Work Styles* files are rated in terms of "importance". The "importance" rating measures how important an activity or ability is in exercising the occupation being described. Descriptors in the *Work Context* file give a "context" rating, which, according to the O*NET documentation, "includes a variety of scales with some unique and specific work context variables". More details are available on www.onetonline.org/help/online/scales. We rescale all descriptors linearly to take values inside [0, 1].

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

TABLE 1
Distributions of wages, job health contents, and health

	Monthly wage quantiles					Job health content distribution					
	Q_{10}	Q_{25}	Q_{50}	Q_{75}	Q_{90}	Value	0.39	0.52	0.64	0.77	0.89
Full Time	1,083	1,328	1,750	2,306	3,039	Share	0.14	0.22	0.24	0.34	0.07
Part Time	347	456	630	850	1,365	Mean wage	1,405	1,486	1,908	2,244	2,437

	Health distribution				Health elasticities				
	Severe	Poor	Average	Good		Severe	Poor	Average	Good
All	0.07	0.29	0.45	0.19	Age	−0.40	−0.25	0.06	0.34
Employed	0.06	0.29	0.45	0.19	Wage	−0.20	−0.12	0.03	0.17
Unemployed	0.16	0.34	0.34	0.17	Job content	0.37	0.23	−0.06	−0.32

Notes: Distributions and elasticities in the initial cross section.
The bottom right panel shows the elasticities of the probability of each health state with respect to age, wage, and job content.
For example: $\partial \ln [\Pr(Severe)] / \partial \ln(wage) = -20\%$.
All elasticities estimated by an ordered logit and significantly different from 0 at 1%.

Finally, for computational reasons, we discretise the support of the job stress content variable into five equally sized bins (with mid-points 0.39, 0.52, 0.64 and 0.77 and 0.89).

2.3. *Descriptive statistics*

The work sample has 33,886 observations, relating to 8,228 men aged 25–55 during our observation period. In the initial cross section, 90% of workers are employed full time, 3% are employed part time, and 7% are not employed.

Table 1 shows distributions of the main outcomes of interest (health, wage, and job content) in the initial cross section. The top left panel shows the extent of wage dispersion in the data, both within and between full-time and part-time jobs, with $Q_{90/10}$ ratios of 2.8 and 3.9 respectively. The top right panel of Table 1 shows the distribution of the job mental health content variable. The higher the value, the more stressful the job. The five values of job health contents are in the top row, their shares in the initial cross section of employed workers is on the second row and the last row shows the average wage among workers in jobs with that level of stress, pointing at a positive relationship in the data between wages and job health contents. We should not yet take this as evidence of compensating differentials as no individual heterogeneity has been controlled for. Our structural estimation will allow us to pin down the wage compensation for health content in job offers.

The bottom left panel of Table 1 shows the proportion of individuals in each of the four health states in the initial cross section. The most noticeable fact is that unemployed workers are more often in the worse health states, for example, the share of individuals in severe health almost triples (from 6% to 16%) when going from employed to unemployed workers. As a first-pass investigation of the link between health and job outcomes over the career, we run an ordered logit regression of health on age, the (monthly) wage and the job health content. The resulting elasticities are in the bottom right panel and show a clear negative relationship between poor health and both age and wage while job stress content is positively linked to the probabilities of severe or poor health. These estimates are only descriptive: our structural analysis will aim to identify the two-way link between health and labour market outcomes.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

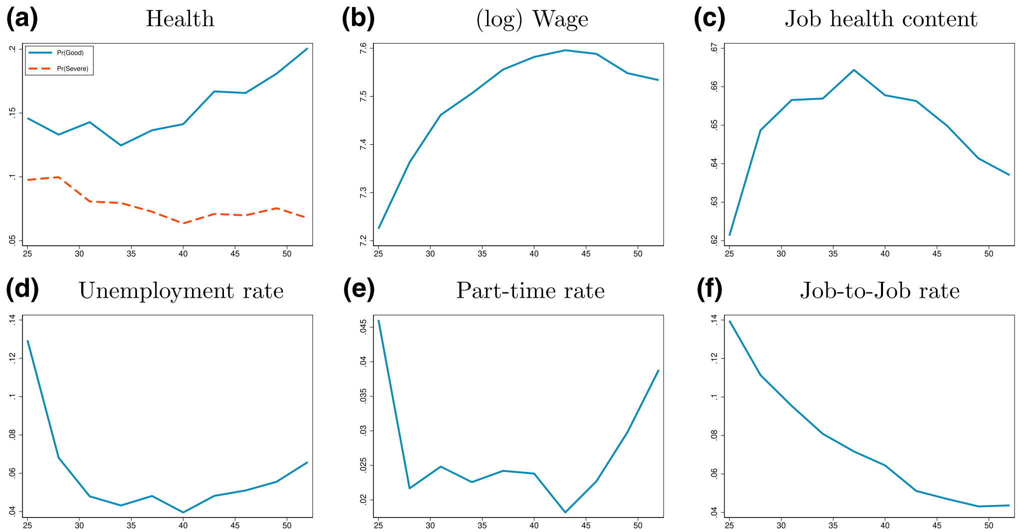

FIGURE 2
Health and labour market outcomes by age

Note: All graphs show 3-year averages (at age 25, 28, 31, etc.)

Since our focus is on workers' health and labour market trajectories, we first describe the relationship of several variables of interest with age. In Figure 2, we plot the average of health and labour market outcomes by age from the data, using a 3-year average for smoothing. These moments will be targeted in the estimation of our structural model.

Figure 2(a) shows that in our sample, older individuals are more likely to be in the "Good" health state and less likely to be in "Severe" health.[12] (Log) wages display an increasing and concave age profile (Figure 2(b)) while the job health content increases up to age 37–39 and decreases afterwards (Figure 2(c)). Both the unemployment and part-time rates are U-shaped with the former being much higher for younger than for older workers. The job-to-job transition rate declines with age, consistently with a standard job ladder model.

In order to interpret these stylised facts further, we need some structure on worker selection in and out but also between jobs (Figure 2(d,f)), to allow job characteristics (wage, health content, full/part-time) to influence both workers' decisions and their health (Figure 2(b,c,e)), and to let mental health affect job decisions (Figure 2(a)). These concerns guide the model structure that we present in the next section and which will include both on-the-job search and a labour supply margin (for selection in/out/across jobs), jobs as characteristics goods (wage, health content and part/full time) affecting worker utility and health dynamics, and health as a dynamic state variable affecting worker utility.

12. A legitimate concern would be that the mental health correlation with age shown in Figure 2(a) depends on sample selection or on the use of the SF-12 indicator. We show in Appendix A that the descriptive patterns are robust to using either the raw data sample or alternative measures of mental health or well being.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

3. THE MODEL

3.1. *Environment*

We write a discrete-time partial-equilibrium model of the labour market with (random) search on and off the job, an extensive employment margin, and health dynamics. Workers are characterised by a triple (x, t, h) where $x \in \mathbb{R}$ is a constant attribute (individual heterogeneity), t is the worker's age, and h is their time- (or age-)varying health status. Jobs are characterised by an attribute $\mathbf{y} = (w, \ell, s) \in \mathbb{R}^3$, where w is the monthly wage attached to that job, ℓ is the number of working hours in the job, and s is a measure of the (mental) health content of the job. We use $\mathbf{y} = \emptyset$ as a placeholder for job characteristics for nonemployed workers. Moreover, we restrict working hours to $\ell \in \mathcal{L} = \{0, 1/2, 1\}$ where $\ell = 0$ means that the worker is not employed, $\ell = 1/2$ means part-time work, and $\ell = 1$ means full-time work.

The cost to the worker of supplying ℓ is $c(\ell, s, t, h) + \varepsilon$, independent of x or w, where ε is a labour supply shock, i.i.d. across workers and time with cdf $H(\cdot)$, taking values in $\mathbb{R}_+$. Nonemployed workers receive a flow income of b and incur no explicit labour supply cost.

Each period unemployed and employed workers receive job offers with probability λ_0 and λ_1, respectively. An offer consists of a triple $\mathbf{y}^o = (w^o, \ell^o, s^o)$ drawn from a state-dependent sampling distribution. Specifically, we assume that type-x unemployed workers draw their offers from a baseline sampling distribution $F(\cdot|x)$ conditional on their type, and that type-x employed workers draw from a mixture of $F(\cdot|x)$ and a distribution depending on their current job health content. With probability ρ, the offer has the same job health content s as the worker's current job and the wage w^0 and working hours ℓ^o are drawn from $F(w^o, \ell^o|x, s^o = s)$. With probability $1 - \rho$, the job offer is drawn from the same baseline $F(\cdot|x)$ as unemployed workers. The terms of the offer (wage, working hours, health contents) are nonnegotiable, and the recipient's only choice is whether to accept or reject it.

Our modelling of job offers is motivated by the empirical observation that the majority of job-to-job transitions occur with no change in job health content s. A fully random search model, with a completely state-independent sampling distribution of job offers, would struggle to capture this specific feature of the data. An alternative would be a (partially) directed search model, where workers are allowed to choose the type s of jobs on which to focus their search. One of the main theoretically appealing features of such a model would be to partly endogenise state-dependence of the sampling distribution. Unfortunately, empirical implementation of such a model is problematic for several reasons, the chief one being that it may not be identified.[13] Those issues notwithstanding, we implement a partially directed search version of our model in Section 7 as a robustness check. For the main part of the paper, though, we use the reduced-form way of capturing state-dependence in job offer sampling described above. This modelling device, inspired by Bonhomme *et al*. (2016), allows the model to capture the inertia in job health content while maintaining identification.

A job's health content s stays constant over time. Working hours are subject to occasional random changes following a first-order Markov process with transition probabilities $\Pr\{\ell'|\ell\}$ which is independent of worker and job attributes. The wage offer is drawn conditionally on the worker's characteristics and on the job's health content and working time. The wage attached to a job then remains constant *conditional on working hours*, *i.e.* it only changes when working hours change. Thus, the characteristics of an ongoing job change from $\mathbf{y} = (w, \ell, s)$ to $\mathbf{y}' = (w', \ell', s)$

13. Intuitively, since the determinants of a worker's utility will enter both the job acceptance decision and, because of directed search, the offer distribution, there is no exclusion restriction to separate these two components of the selection model.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

with probability $\Pr\{\ell'|\ell\}$: the health content of the job always stays constant, and the change in wage is entirely driven by the change in working hours. When an employer seeks to impose a change of working hours (from full time to part time or vice versa), and a corresponding change in wage, the worker can either accept the change or quit.

Employed workers further face a per-period job destruction probability of δ, and always have the option of quitting into unemployment.

As we discussed in Section 2.1, health status is only observed once a year in our data, which constrains us to set our model period to 1 year. While the vast majority of workers experience no more than one job transition within a typical year, some move either from job to unemployment back into a job, or from unemployment to job back into unemployment between two consecutive annual interviews. Section 2.1 showed that multiple transitions within a year are rare, but nonnegligible as a fraction of all labour market transitions. To accommodate those observations, we introduce two additional reallocation shocks. With probability $\widetilde{\lambda}$, type-x workers who were just hit by a job destruction shock δ draw an offer $\mathbf{y}^o$ from $F(\cdot|x)$ before the end of the current period and start the following period as employed if they accept that job. Symmetrically, initially unemployed workers who draw a job offer (probability λ_0) face a probability $\widetilde{\delta}$ of that job being terminated before the end of the period. Those features of the model have no particular theoretical content and are meant to handle time aggregation.

Health evolves over a finite set of values following a first-order Markov process with transition probabilities $\Pr\{h'|\mathbf{y}, x, t, h\}$ which depend on the worker's age and type, as well as the wage, health content, and hours of the job the worker is currently employed at. These probabilities are $\Pr\{h'|\emptyset, x, t, h\}$ for nonemployed workers. In the application, we restrict health to take on one of four values: $h \in \mathcal{H} = \{\text{G, A, P, S}\}$ for Good, Average, Poor, and Severe.

Finally, workers have a working life of T periods, at the end of which they retire and enjoy an exogenous terminal value.

3.2. *Dynamics*

Workers start each period with knowledge of their type (x), current age (t), health status (h), labour supply shock (ε), job status and job attributes ($\mathbf{y} = (w, \ell, s)$, or b if unemployed). Then, the period unfolds in a way that depends on the worker's employment status.

Unemployed workers. For initially unemployed workers, the sequence of events is as follows:

1. The worker receives his unemployment income b.
2. At the end of the period, he draws his next-period health status $h' \sim \Pr\{h'|\emptyset, x, t, h\}$, labour supply shock $\varepsilon' \sim H$, and his age is incremented.
3. With probability λ_0, he receives a job offer $\mathbf{y}^o \sim F(\cdot|x)$. With probability $\widetilde{\delta}$, this job offer gets hit by a destruction shock straightaway.[14]
4. The worker chooses the better option between unemployment (always available) and the job offer (if he received a job offer which was not immediately destroyed).

14. An implicit assumption in our within-period timing is that, whenever two transitions occur within a period, both occur at the end of the period, the intermediate spell being of negligible length. As a consequence, the attributes of that intermediate spell have no impact on current-period values or next-period health. We take this shortcut because our data lack information on short spells that start and end between two consecutive interviews. In particular, when a worker moves from unemployment to job to unemployment, the data convey no information about the characteristics (wage, hours, occupation) of that intervening job spell.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

Employed workers. If the worker is employed at a job of type $\mathbf{y} = (w, \ell, s)$ at the start of the period, the sequence of events is the following:

1. The worker incurs a labour supply cost $c(\ell, s, t, h) + \varepsilon$, and receives the wage w.
2. At the end of the period, he draws his next-period health status $h' \sim \Pr\{h'|\mathbf{y}, x, t, h\}$, labour supply shock $\varepsilon' \sim H$, and working hours $\ell' \sim \Pr\{\ell'|\ell\}$. Age is incremented.
3. With probability δ, his current job is destroyed. In this case, the worker draws a new job offer $\mathbf{y}^o \sim F(\cdot|x)$ with probability $\widetilde{\lambda}$ before the start of the new period.
4. If his current job has not been destroyed, the worker receives an outside job offer $\mathbf{y}^o$ with probability λ_1. With probability $1-\rho$ this offer is drawn from $F(\cdot|x)$ and with probability ρ this offer has the same s as the worker's current job and the wage and working hours are drawn from $F(w^o, \ell^o|x, s^o = s)$.
5. The worker chooses the best option between unemployment (always available), his old job (unless that job was destroyed by a δ-shock), and the job offer (if he has received one).

Value functions. In what follows, we denote the current characteristics of a worker's job by $\mathbf{y} = (w, \ell, s)$, and the future characteristics of the same job by $\mathbf{y}' = (w', \ell', s)$. We denote the characteristics of a job offer by $\mathbf{y}^o$, and the value of unemployment to a worker with characteristics (x, t, h) by $U(x, t, h)$. The value to a worker with characteristics (x, t, h) and current labour supply shock ε of being employed in a job with attributes $\mathbf{y}$ is $V(x, t, h, \mathbf{y}) - \varepsilon$, where V is the worker's value gross of the labour supply shock. The dynamics of the model described above imply the following value functions for unemployed workers:

$$U(x,t,h) = b + (1+r)^{-1} \sum_{h' \in \mathcal{H}} \Pr\left\{h'|\emptyset, x, t, h\right\} \Big[\left(1 - \lambda_0\left(1-\widetilde{\delta}\right)\right) U(x, t+1, h') \\ + \lambda_0\left(1-\widetilde{\delta}\right) \iint \max\left\{U(x,t+1,h'); V(x,t+1,h',\mathbf{y}^o) - \varepsilon'\right\} dH(\varepsilon') dF(\mathbf{y}^o|x)\Big] \tag{1}$$

On the right-hand side of (1), the first term is unemployment flow income. From the second term, we reach the end of the period so discounting is applied and the new health and labour supply shocks are drawn. With probability $1-\lambda_0(1-\widetilde{\delta})$, the worker either fails to receive a job offer or receives one which is immediately destroyed, and in both cases remains unemployed at the start of $t+1$ with value $U(x, t+1, h')$. With probability $\lambda_0(1-\widetilde{\delta})$, the worker receives an offer that is not destroyed during the period, he may choose to accept it or stay unemployed, thus receiving the greater of the value of unemployment and that of his job offer.

For employed workers, the value function is:

$$\begin{aligned} &V(x,t,h,\mathbf{y}) - \varepsilon \\ &\quad = w - c(\ell, s, t, h) - \varepsilon \\ &\qquad + (1+r)^{-1} \sum_{h' \in \mathcal{H}} \sum_{\ell' \in \mathcal{L}} \Pr\left\{h'|\mathbf{y}, x, t, h\right\} \Pr\left\{\ell'|\ell\right\} \Big[\delta\left(1-\widetilde{\lambda}\right) U\left(x, t+1, h'\right) \\ &\qquad + \delta\widetilde{\lambda} \iint \max\left\{U(x,t+1,h'); V(x,t+1,h',\mathbf{y}^o) - \varepsilon'\right\} dH(\varepsilon') dF(\mathbf{y}^o|x) \\ &\qquad + (1-\delta)\lambda_1 \iint \max\left\{U(x,t+1,h'); V(x,t+1,h',\mathbf{y}') - \varepsilon'; V(x,t+1,h',\mathbf{y}^o) - \varepsilon'\right\} dH(\varepsilon') dF(\mathbf{y}^o|x,\mathbf{y}) \\ &\qquad + (1-\delta)(1-\lambda_1) \int \max\left\{U(x,t+1,h'); V(x,t+1,h',\mathbf{y}') - \varepsilon'\right\} dH(\varepsilon')\Big] \end{aligned} \tag{2}$$

The right-hand side of (2) has the instantaneous utility of being employed (wage minus disutility of labour) on the first line. The discounted, end-of-period continuation value starts on the second

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

line, where a new health status, new labour supply shock and new working hours in the current job are drawn. With probability $\delta(1-\widetilde{\lambda})$, the current job is destroyed and no offer is drawn before the new period starts so the worker becomes unemployed (second line). On the third line, the current job is destroyed and the worker received an outside offer so he can take this new job or become unemployed. If the current job is not destroyed (probability $1-\delta$), the worker may receive an outside offer with probability λ_1 and he can then choose between his current job, the outside offer or unemployment (fourth line). Note that we use the shortcut $dF(\mathbf{y}^o|x, \mathbf{y})$ to denote that the offer is drawn from a mixture that depends on (the health content of) the current job. On the last line, the worker did not draw an outside offer so he can choose to stay in his current job or become unemployed.

4. ESTIMATION AND IDENTIFICATION

4.1. *First estimation step: individual heterogeneity*

Worker types x take values on a discrete set $\{1, \ldots, K\}$. We build on the approach of Bonhomme *et al*. (2022), (BLM thereafter, see also Bonhomme *et al*., 2019) who estimate the classes of individual heterogeneity by k-means clustering based on relevant moments of outcome variables, where the moments are computed at the individual level. Let individuals be denoted by $i \in \{1, \ldots, N\}$ and let $\mathbf{m}_i$ be a given vector of M moments of individual i's outcomes. We estimate a discrete distribution of individual heterogeneity, with a fixed, finite number K of classes. A partition assigns a group $x_i \in \{1, \ldots, K\}$ to each individual i. The classification step of BLM consists of finding the partition $\{x_i\} \in \{1, \ldots, K\}^N$ that minimises $\sum_{i=1}^N \|\mathbf{m}_i - \widetilde{\mathbf{m}}(x_i)\|^2$, where $\|\cdot\|$ is the Euclidian norm and $\widetilde{\mathbf{m}}(k)$ is the mean of vector $\mathbf{m}$ in group k. Once the classification step is completed, the estimation step (presented in the next subsections) can be carried out, treating individual heterogeneity as observed group dummy variables.

The BLM approach hinges on the existence of an injective map between individual unobserved heterogeneity and the vector of moments $\mathbf{m}_i$, essentially meaning that the moments used are informative about individual heterogeneity. Consistency of the BLM approach assumes that individual heterogeneity is identified, something we cannot prove formally. We will, however, use simulations to assess the ability of our approach to recover the estimated distribution of individual heterogeneity (see Section 5.4). In our application, the two main sources of individual heterogeneity are productivity and health. Accordingly, we use observed wages, job health contents, and health status to compute the individual moments $\mathbf{m}$. More precisely, the moments we use are the average wage, the average job health content and three health dummy averages ("Poor," "Average," and "Good").

Our particular application of the BLM clustering technique raises a specific data issue which requires a modification of the classification step. We observe an individual's job characteristics and health only for up to 10 years and our panel data have many different cohorts. Hence, we observe some individuals' wage, job stress, and health in their early thirties while for other individuals these outcomes will be observed when they are in their fifties. Using moments of the wage, job stress, and/or health over the observation period to form the individual moment vector $\mathbf{m_i}$ will thus lead to clusters that capture a mix of individual heterogeneity and age effects. To overcome this, we consider the following minimization problem:

$$\min_{\{x_i\}\in\{1,\ldots,K\}^N,\ \mathbf{g}:\{1,\ldots,K\}\times\mathbb{R}\to\mathbb{R}^M} \sum_{i=1}^N \left\|\mathbf{m}_i - \mathbf{g}\left(x_i, T_i^0\right)\right\|^2, \tag{3}$$

where T_i^0 is the age of individual i at the beginning of the observation period (thus identifying his cohort) and $\mathbf{g}(1,\cdot), \ldots, \mathbf{g}(K,\cdot)$ are parametric functions of the individual's cohort, accounting for within-class differences in moment outcomes due to age. In practice, we specify each $\mathbf{g}(x,\cdot)$ as a (2nd-order) polynomial of T.[15] Estimating the $\mathbf{g}$ functions in the classification step prevents us from using standard k-means algorithms as in BLM. Details are provided in Appendix B.

4.2. *Second estimation step: health dynamics*

We specify health transition probabilities as follows:

$$\begin{aligned}
\Pr\left\{h'=\mathrm{G}|\mathbf{y},x,t,h\right\} &= 1-\Lambda\left(\tau_{\mathrm{G}}+\tau\left(\mathbf{y},x,t,h\right)\right)\\
\Pr\left\{h'=\mathrm{A}|\mathbf{y},x,t,h\right\} &= \Lambda\left(\tau_{\mathrm{G}}+\tau\left(\mathbf{y},x,t,h\right)\right)-\Lambda\left(\tau_{\mathrm{A}}+\tau\left(\mathbf{y},x,t,h\right)\right)\\
\Pr\left\{h'=\mathrm{P}|\mathbf{y},x,t,h\right\} &= \Lambda\left(\tau_{\mathrm{A}}+\tau\left(\mathbf{y},x,t,h\right)\right)-\Lambda\left(\tau_{\mathrm{P}}+\tau\left(\mathbf{y},x,t,h\right)\right)\\
\Pr\left\{h'=\mathrm{S}|\mathbf{y},x,t,h\right\} &= \Lambda\left(\tau_{\mathrm{P}}+\tau\left(\mathbf{y},x,t,h\right)\right)
\end{aligned}\tag{4}$$

where Λ is the logistic cdf, τ_G, τ_A, τ_P are parameters and the threshold function τ is given by:

$$\tau\left(\mathbf{y},x,t,h\right)=\eta^{(t)}\cdot t+\eta^{(t,\ell)}\cdot t\cdot\mathbf{1}\{\ell>0\}+\eta^{(w)}\cdot w+\eta^{(s)}\cdot s+\mathbf{I}_{\ell}\cdot\eta^{(\ell)}+\mathbf{I}_{h}\cdot\eta^{(h)}+\mathbf{I}_{x}\cdot\eta^{(x)}.\tag{5}$$

Throughout the paper, we use $\mathbf{I}_h$ to denote a vector of dummy variables for all health values, and use similar notation for heterogeneity classes ($\mathbf{I}_x$) or working times ($\mathbf{I}_\ell$). The first two terms on the right-hand side of (5) reflect the effect of age t on health dynamics, allowing for this effect to depend on employment ($\ell>0$). The third and fourth terms on the right-hand side of (5) pertain to the wage and job health content. The next two terms, respectively, capture the effect of the employment status and the effect of past health. The last term is the effect of individual heterogeneity x. Whenever a worker is unemployed ($\mathbf{y}=\emptyset$), we set $\ell=w=s=0$.

As (4) shows, health transitions are governed by a first-order Markov process which only involves observable variables $(\mathbf{y},x,t)$ and an i.i.d. shock. Thus, (4) can be estimated by maximising the likelihood of observed health transitions conditionally on $(\mathbf{y},x,t)$. Given the logistic assumption made on the i.i.d. shock, this boils down to an ordered logit regression.[16]

4.3. *Third estimation step: job characteristics and utility*

Specification. We start with the specification of the distribution of job offer characteristics. The offer distribution of $\mathbf{y}$ is modelled as follows: s and ℓ are drawn independently from their respective marginal distributions. The sampling distribution of s has a finite five-point support $s_1<\cdots<s_5$, with $\Pr\{s_j\}=s_j^{\alpha_1-1}\cdot(1-s_j)^{\alpha_2-1}/\sum_{k=1}^{5}[s_k^{\alpha_1-1}\cdot(1-s_k)^{\alpha_2-1}]$. For employed workers receiving an offer (probability λ_1), the offered job health content could either be equal to the current one (probability ρ) or (probability $1-\rho$) be drawn randomly from the sampling distribution (as is the case for unemployed workers). Working time ℓ takes on two possible values: $\ell=1$ for full-time work and $\ell=1/2$ for part-time work. The Markov process governing working hours in ongoing matches is characterised by two transition probabilities, intuitively denoted as $\Pr\{\mathrm{FT}|\mathrm{PT}\}$ and $\Pr\{\mathrm{FT}|\mathrm{FT}\}$. The working time associated with any new job offer is part-time

15. Formally: $\mathbf{g}(x,T)=\mathbf{g}_{x,0}+\mathbf{g}_{x,1}\cdot T+\mathbf{g}_{x,2}\cdot T^2$, $\quad x\in\{1,\ldots,K\}$, where, for any $x\in\{1,\ldots,K\}$, $\mathbf{g}_{x,0}$, $\mathbf{g}_{x,1}$ and $\mathbf{g}_{x,2}$ are vectors of M parameters estimated by solving (3). The first-order conditions of this problem imply that if $\mathbf{g}_{x,1}=\mathbf{g}_{x,2}=\mathbf{0}_M$ for all x then the optimal $\mathbf{g}_{x,0}$ is the average of individual moments in class x, and the solution coincides with the standard BLM classification.

16. We also have to assign a value of $\mathbf{y}$ to workers experiencing a labour market transition (*i.e.* a change of labour market spell) within the year. We use the $\mathbf{y}$ of the job where the worker spent more time during the year.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

with probability $\Pr_0\{\text{PT}\}$. Then, conditional on x, s and ℓ, the wage offer is:

$$\log w = \beta^{(0)} + \mathbf{I}_x \cdot \beta^{(x)} + (\mathbf{I}_x * s) \cdot \beta^{(s)} + \beta^{(\text{pt})} \cdot \mathbf{1}\{\ell = \text{PT}\} + \sigma^{(\omega)} \cdot \omega \tag{6}$$

where $\omega \sim \mathcal{N}(0,1)$ independently of worker characteristics, s or ℓ. The term $\mathbf{I}_x * s$ means that we interact individual heterogeneity with job stress. The parameters $\beta^{(s)}$ (one for each heterogeneity class) thus capture compensating wage differentials for stressful work in job offers.

We next present the specification of preferences. The distribution of the labour supply shock is specified as log-normal: $\ln \varepsilon \sim \mathcal{N}(-\frac{\sigma^{(\varepsilon)2}}{2}, \sigma^{(\varepsilon)})$. This specification implies that the mean of ε is normalised at 1 as said mean is not separately identified from the scale of unemployment income b. Recall that the labour supply shock is assumed i.i.d. across periods (persistence in observed labour supply is produced in the model by search frictions, health, and infrequent "forced" changes in working hours within ongoing jobs). The labour supply shock process is therefore characterised by a single parameter, $\sigma^{(\varepsilon)}$. The cost of labour supply is otherwise specified as:

$$c(\ell, s, t, h) = \mathbf{I}_h \cdot \kappa^{(h)} \times \mathbf{I}_\ell \cdot \kappa^{(\ell)} \times \left(1 + \kappa^{(s)}(s - s_1)\right) \times \left(1 + \kappa^{(t)}(t - 25)\right). \tag{7}$$

This specification allows for interactions between each of the four inputs so that, for example, the cost of working in a stressful job varies with age as well as health or working hours.[17]

Lastly, we fix the discount rate r at 10% per annum.

Parameters estimated directly from the data. Several parameters pertaining to job characteristics, job dynamics, and preferences have direct data counterparts, and can therefore be estimated directly. We list those parameters here and give some intuition about their identification. Technical details are in Appendix C.

The coefficients $\beta^{(x)}$ and $\beta^{(s)}$ in (6) can be obtained from regressions across (x, s) cells of the maximum wage observed in these cells on worker type dummies and their interaction with job stress. This is because focusing on highest wage ensures that the unobserved job ladder rung ω is equal across cells.

The within-job full-time/part-time transition probabilities, $\Pr\{\text{FT}|\text{FT}\}$ and $\Pr\{\text{FT}|\text{PT}\}$, can be estimated from observed within-job changes in hours.

The exogenous job destruction rate δ is estimated directly as the probability that a worker employed in a given year either experiences a job-to-unemployment-to-job transition or is non employed the following year, and says that the reason that their previous job ended was one of "made redundant," "dismissed/sacked," "temporary job ended," or "other reason".[18]

Lastly, the first reallocation shock $\widetilde{\delta}$ can be directly estimated from the probabilities of making an unemployment-to-job transition or of making an unemployment-to-job-to-unemployment transition. The other reallocation shock, $\widetilde{\lambda}$ cannot be estimated directly; however, the ratio $\widetilde{\lambda}/\lambda_0$ can, from the probabilities of making an unemployment-to-job transition or a job-to-unemployment-to-job transition. For these and all other parameters discussed above, the actual moments used for identification are formally derived in Appendix C.

Indirect inference. At this point, we are left with the following vector of parameters to estimate: the job offer arrival rates λ_0 and λ_1, the standard deviation $\sigma^{(\omega)}$ of the log-wage offer residual, the parameters $(\alpha_1, \alpha_2, \rho)$ of the sampling distribution of health contents in job

17. The factor specific to health (resp. to job health content, resp. to age) is equal to 1 when the health state is "Good" (resp. when in the least stressful type of job s_1, resp. at the youngest age in our sample, 25 years old.)

18. This leaves out the following other possible reasons: "left for better job," "took retirement," "health reasons," "left to have baby," "look after family," "look after other person," and "moved area".

offers, the standard deviation of the labour supply shock $\sigma^{(\varepsilon)}$, the share of part-time job offers $\Pr_0(\ell' = \mathrm{PT})$, and the effect of age ($\kappa^{(t)}$), job content ($\kappa^{(s)}$), full/part-time ($\kappa^{(\ell)}_{\mathrm{FT}}$ and $\kappa^{(\ell)}_{\mathrm{PT}}$) and health ($\kappa^{(h)}_G, \kappa^{(h)}_A, \kappa^{(h)}_P, \kappa^{(h)}_S$) on the disutility of labour (where $\kappa^{(h)}_G$ is normalised to 1).

We estimate those parameters using indirect inference. We simulate a panel of just over 160,000 workers (20 times the size of our initial cross section) over 9 years (eight full years, plus the initial observation) to replicate our data sample. The initial state of each worker (age, health, employment state, job type, wage) is taken directly from the initial cross section in the data. We then arrange observations, both in simulated and actual data, into ten 3-year worker age classes —age classes thus contain workers that are 25–27, 28–30, etc. up to 55 years old at the time of observation— and compute averages by age group as moments to match. This indirect inference step is the computationally costly part of our estimation and motivates many of the simplifying assumptions of the model (for instance, the Markovian dynamic process of health) and discretisation of some state variables (health and job health content).

The set of targeted moments includes averages by age class of log-wage and job health content (2×10 moments); unemployment and part-time rates (2×10 moments); mean health of unemployed, part-time, and full-time workers (3×10 moments, where health encoded as $1 =$ "Severe," $2 =$ "Poor," $3 =$ "Average," $4 =$ "Good"); quit, unemployment-to-job and job-to-job transition rates (3×10 moments).

In addition, we match the coefficients obtained from the following pooled OLS regressions, to capture covariances between variables of interest: log wage on worker type, health, and part-time dummies, age, job content, and a constant (10 moments); squared residual of the previous wage regression on worker type, health and unemployment-to-job transition dummies, and a constant (8 moments); yearly log-wage change on job content both at the start and at the end of the period and a constant for job changers (three moments), employment dummy on worker type and health dummies, age, and a constant (eight moments); current job type s on worker type and health dummies, age, and a constant (eight moments). We also target the proportion of job-to-job transitions with a strict increase (resp. strict decrease) in health content (two moments). Finally, to inform the job-type sampling distribution, we target the distribution of job types s accepted by workers upon exiting unemployment (four additional moments).

Standard errors. We estimate standard errors by bootstrap. Our multi-step estimation procedure is applied to 200 bootstrap samples and we take the standard errors from the resulting distribution of parameter estimates. All our estimation steps, from the clustering procedure described in Section 4.1 to the indirect inference step presented above are run on each bootstrap samples. Hence, the standard errors shown in the results section account for all the variation arising at each stage of our estimation procedure.

4.4. *Identification*

As discussed in Section 4.1, we do not formally prove identification of the unobserved heterogeneity classes, but we will conduct a numerical test in Section 5.4 to gauge the internal consistency of our estimation approach. The second step, presented in Section 4.2, is a straightforward ordered logit regression using observed regressors only and allows us to identify and estimate all the parameters in the health dynamic process. In the last step, as we explain in Section 4.3, several parameters have direct data counterparts and are thus identified by construction. However identification is less straightforward for the final set of parameters which are estimated by indirect inference. These parameters pertain to the arrival rates and distribution of job offers and to individual preferences (namely, the disutility of work).

An important feature of our analysis is that we are able to formally prove identification of those parameters from our data under weak assumptions, and conditional on knowledge of the

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

TABLE 2
Individual heterogeneity classes

Class	Share (%)	Health (%) Severe	Poor	Average	Good	Employed (%)	Average... Wage	Content	Age
1	17.2	8.1	32.7	39.7	19.5	89.8	1,478	68.6	44
2	33.5	7.3	28.9	40.5	23.2	89.8	1,444	47.3	39
3	35.3	6.9	30.0	46.0	17.0	95.2	1,991	72.6	37
4	14.0	3.4	25.2	56.3	15.1	97.6	3,143	75.2	35

distribution of individual heterogeneity (and the variance of the labour supply shock). The technical proof is in Appendix C and consists of three steps. Essentially, we face a selection model where workers sort themselves between unemployment and different types of jobs. The first step identifies the outcome equation (and consequently the selection equation). The second step exploits the Bellman equations (1)–(2) to show that the instantaneous utility functions follow from the value functions in the selection equation by using a mapping identified from the data. The last step then combines the first two results to identify the preference parameters.

5. ESTIMATION RESULTS

5.1. *Model estimates: individual heterogeneity*

The first step of our estimation procedure produces a partition of the sample of workers into four heterogeneity classes (see Section 4.1). In Table 2, we describe these classes by showing class-level averages of various relevant variables.

Class 4 is a group of high-skill/high-wage workers who tend to be in average health and work in more stressful jobs relative to other groups. At the other end of the spectrum, Classes 1 and 2 earn similarly low wages but differ markedly with respect to their health and job stress content as workers in Class 1 are more likely to be in severe or poor health and have more stressful jobs. Workers in Class 2 have by far the least stressful jobs. Finally, Class 3 is in between Class 4 and Classes 1 and 2 in terms of wages, health, and job stress content (with the latter being almost as high as that of Class 4).

In short, our first estimation step has produced four groups of workers along essentially three directions: high-productivity/stressful-job/healthy workers (Class 4), average-productivity/stressful-job/average-health workers (Class 3), low-productivity/low-job-stress/average-health workers (Class 2) and low-productivity/average-job-stress/poor-health workers (Class 1).

5.2. *Model estimates: health dynamics*

The second estimation step maximises the likelihood of health transitions given the model structure (4) and the logistic assumption on health shocks. Parameter estimates of our reduced-form specification (4) of the health dynamic process are reported in Table 3. While the magnitudes of those estimates are not straightforward to interpret, their signs mostly are.

The estimated coefficients of the wage, $\eta^{(w)}$, and of the job content, $\eta^{(s)}$, in the threshold (5) are, respectively, significantly positive and negative, indicating that better pay increases the probability of health improving across years while the converse is true for job stress levels. Also, since $\eta^{(t)}$ is significantly positive, ageing improves mental health—a result paralleling the descriptive fact shown on Figure 2(a). While negative, the coefficients on full or part-time

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

TABLE 3

Health dynamics - Ordered logit estimates

$\eta^{(t)}$	0.018	(0.006)	$\eta^{(P)}$	1.352	(0.060)	$E_x(\tau_P)$	0.190	(0.167)
$\eta^{(t,\ell)}$	−0.005	(0.006)	$\eta^{(A)}$	2.714	(0.060)	$E_x(\tau_A)$	2.700	(0.169)
$\eta^{(w)}$	0.116	(0.023)	$\eta^{(G)}$	3.624	(0.066)	$E_x(\tau_G)$	5.119	(0.170)
$\eta^{(s)}$	−0.338	(0.160)	$\eta^{(PT)}$	−0.129	(0.115)	$\eta^{(FT)}$	−0.144	(0.113)

Notes: The τ threshold parameters are averaged over the distribution of individual heterogeneity x.
The baseline thresholds estimates are $\tau_P = 0.117$, $\tau_A = 2.627$ and $\tau_G = 5.046$.
The heterogeneity threshold estimates are $\eta^{(2)} = 0.129$, $\eta^{(3)} = 0.035$ and $\eta^{(4)} = 0.125$.

TABLE 4

Estimated probabilities of going to the severe or poor health state

Employment status						Age					
40 years old						Median wage, medium stress					
Med. wage/stress			Unemployed			30 years old			50 years old		
	S	P		S	P		S	P		S	P
S	0.35	0.52	S	0.44	0.46	S	0.38	0.50	S	0.32	0.53
P	0.12	0.51	P	0.17	0.55	P	0.14	0.53	P	0.11	0.49
A	0.03	0.27	A	0.05	0.34	A	0.04	0.30	A	0.03	0.25
G	0.01	0.14	G	0.02	0.19	G	0.02	0.15	G	0.01	0.12
Wage						Job health content					
40 years old, medium stress						40 years old, median wage					
1st wage decile			9th wage decile			low-stress job			high-stress job		
	S	P		S	P		S	P		S	P
S	0.36	0.51	S	0.34	0.52	S	0.33	0.53	S	0.37	0.51
P	0.13	0.52	P	0.12	0.50	P	0.11	0.50	P	0.13	0.52
A	0.04	0.28	A	0.03	0.26	A	0.03	0.26	A	0.04	0.29
G	0.02	0.14	G	0.01	0.13	G	0.01	0.13	G	0.02	0.15

Notes: Origin health states in rows, destination states in columns.
S: Severe; P: Poor; A: Average; G: Good. Employed means full-time employment.
Transition probabilities averaged over individual heterogeneity.

work, $\eta^{(FT)}$ and $\eta^{(PT)}$, are more difficult to interpret. Indeed, they are relative to the nonemployment state where workers earn a low income b which is itself detrimental to health. The overall impact of not being employed (as opposed to being employed, say, full time at a certain wage) is therefore ambiguous, which motivates the next set of results.

To better appreciate the magnitude of the impact of the various determinants of health dynamics, we show in Table 4 a selection of transition probabilities based upon the estimated model. We compute the yearly probabilities of going to the severe or poor health state from any of the four health states for a given employment status, age, wage, and job content, where each probability is averaged over the distribution of individual heterogeneity x in the initial cross section (estimated in the first step).

The top left panel shows how health transition probabilities vary for 40-year-old men between working full time at a median-wage/medium-stress job and being unemployed. We note that unemployment lowers the probability to leave the bad (severe or poor) health state and increases the rate at which an individual in good or average health enters the poor or severe

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

TABLE 5

Job offer characteristics and arrival rates

Job offer characteristics								
$\beta^{(0)}$	6.150	(0.374)	$\beta_1^{(s)}$	0.433	(0.541)	$\beta^{(\mathrm{pt})}$	−1.197	(0.041)
$\beta_2^{(x)}$	0.531	(0.470)	$\beta_2^{(s)}$	−0.080	(0.309)	$\sigma^{(\omega)}$	0.568	(0.026)
$\beta_3^{(x)}$	0.672	(0.365)	$\beta_3^{(s)}$	0.053	(0.195)	α_1	1.765	(0.159)
$\beta_4^{(x)}$	0.907	(0.370)	$\beta_4^{(s)}$	0.372	(0.184)	α_2	1.406	(0.111)
$\mathbf{E}_x(\beta^{(0)}+\beta^{(x)})$	6.692	(0.084)	$\mathbf{E}_x(\beta^{(s)})$	0.119	(0.122)	ρ	0.591	(0.018)
Arrival rates								
λ_0	0.760	(0.022)	$\Pr\{\ell^o\text{=PT}\}$	0.115	(0.003)	$\widetilde{\lambda}$	0.262	(0.036)
λ_1	0.607	(0.022)	$\Pr\{\ell'\text{=PT}\vert\ell\text{=FT}\}$	0.002	(0.000)	$\widetilde{\delta}$	0.037	(0.010)
δ	0.017	(0.001)	$\Pr\{\ell'\text{=PT}\vert\ell\text{=PT}\}$	0.880	(0.025)			

Notes: $\mathbf{E}_x(\beta^{(0)}+\beta^{(x)})$, resp. $\mathbf{E}_x(\beta^{(s)})$, is the average of $\beta^{(0)}+\beta^{(x)}$, resp. $\beta^{(s)}$, over the distribution of x.
$\Pr\{\ell^o\text{=PT}\}$ is the probability to draw a part-time job offer.
$\Pr\{\ell'\text{=PT}|\ell\text{=FT}\}$ is the probability that a full-time job goes part-time.
$\Pr\{\ell'\text{=PT}|\ell\text{=FT}\}$ is the probability that a part-time job stays part-time.

health state. For example, the probability that a worker goes from the severe to the average or good state equals 13% ($= 1 - 0.35 - 0.52$) if they work full-time at a median wage/stress job and equals 10% if they are unemployed.

The top right panel shows how health dynamics changes with age (for someone employed at the median wage and middling job content). Older workers are more likely to recover from ill mental health and are less likely to experience health transitions in the other direction. For example, a 30-year-old worker in severe health (at a median wage/stress job) has a 12% ($= 1 - 0.38 - 0.50$) probability of leaving the severe or poor state within a year whereas this probability goes to 15% for a 50-year-old worker.

The bottom left panel shows health transition probabilities for 40-year-old men as a function of the wage. It reveals that the wage has a small impact on the probability of going to the severe or poor health state: going from the first to the ninth wage decile while keeping everything else constant only decreases the risk of being in severe or poor health the following year by at most three percentage points.

The bottom right panel shows that job stress contents have a more substantial impact on health transitions: working in the most stressful job type makes it less likely to leave the severe or poor health states and more likely to enter those states for workers currently in average or good health. Therefore, the reverse causality channel from labour market outcomes to mental health is active through the effect of employment, wages and, more noticeably, job stress contents.

5.3. *Model estimates: utility and job characteristics*

The remaining parameter estimates pertain to the sampling distribution and arrival rates of job offers (Table 5), and to worker preferences, including the labour supply cost function (Table 6).

First, recall that several parameters in the wage offer equation (6) are allowed to vary with individual heterogeneity x. We report in the first two columns of Table 5 the wage intercepts $\beta^{(0)}$, class effects $\beta^{(x)}$ (normalised to 0 for class 1) and wage offer compensating differentials for job stress $\beta^{(s)}$, together with their averages over the distribution of unobserved heterogeneity.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

TABLE 6
Utility estimates

Disutility of labour						Supply shock		
$\kappa^{(t)}$	0.281	(0.025)	$\kappa_S^{(h)}$	61.02	(2.18)	$\sigma^{(\varepsilon)}$	3.36	(0.130)
$\kappa^{(s)}$	3.673	(0.374)	$\kappa_P^{(h)}$	17.68	(0.90)			
$\kappa_{\mathrm{FT}}^{(\ell)}$	1.227	(0.066)	$\kappa_A^{(h)}$	9.78	(0.63)		Unemp. income	
$\kappa_{\mathrm{PT}}^{(\ell)}$	0.592	(0.037)	$\kappa_G^{(h)}$	1.00	(norm.)	b	429.0	(30.3)

Notes: Estimated standard errors in parenthesis.

TABLE 7
Labour supply cost estimates (GBP/month)

	age 30					age 50				
	s_1	s_2	s_3	s_4	s_5	s_1	s_2	s_3	s_4	s_5
Severe health	159	233	306	380	453	580	849	1, 117	1, 385	1, 654
Poor health	46	67	89	110	131	168	246	324	401	479
Average health	25	37	49	61	73	93	136	179	222	265
Good health	3	4	5	6	7	10	14	18	23	27

Notes: Units are GBP/month. Job stress grows from s_1 (lowest) to s_5 (highest).

Our results show that the effect $\beta^{(s)}$ of job stress content on wage offers is significantly positive for one group (Class 4, which consisted of high-productivity/good health workers). On average, there is no significantly positive compensating differentials for job stress in wage offers. As expected, wages for part-time job offers are substantially lower than for full-time jobs.

The parameter estimates for the sampling distribution of job health contents $(\alpha_1, \alpha_2) = (1, 765, 1.406)$ imply that this distribution is inverse U-shaped, as was is distribution of job stress among employed workers (see Table 1). However, the two distributions are different as, for example, the estimated probability to receive a high-stress job (highest value of s) is 17% while only 7% of employed workers are in such jobs (in the initial cross section). This may indicate some degree of self-selection of workers away from higher-stress jobs.

We also note that the probability to draw a job offer with the same health content as one's current job is large, $\rho = 0.591$. This feature is important to have the model fit job outcomes (wage and stress content) following job-to-job transitions, especially the large share of job changes where the health content stays constant.

The utility parameter estimates in Table 6 show that, as expected, worse health, higher job stress levels, age, and full-time work increase the disutility of labour. For ease of interpretation, we can replace the estimated values from Table 6 into the labour cost formula (7) for different values of health, age and job content. The resulting cost estimates are shown in Table 7.

The cost of labour supply can more than triple from age 30 to 50. Controlling for age and health, the labour supply cost increases substantially when going up one level of job stress. Lastly, we note that being in severe or even poor health increases the labour supply cost dramatically relative to being in good health, regardless of age. These large effects of health and job stress on labour supply costs make health an important determinant of individual labour market decisions. We further investigate this point in counterfactual analyses in the next section.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

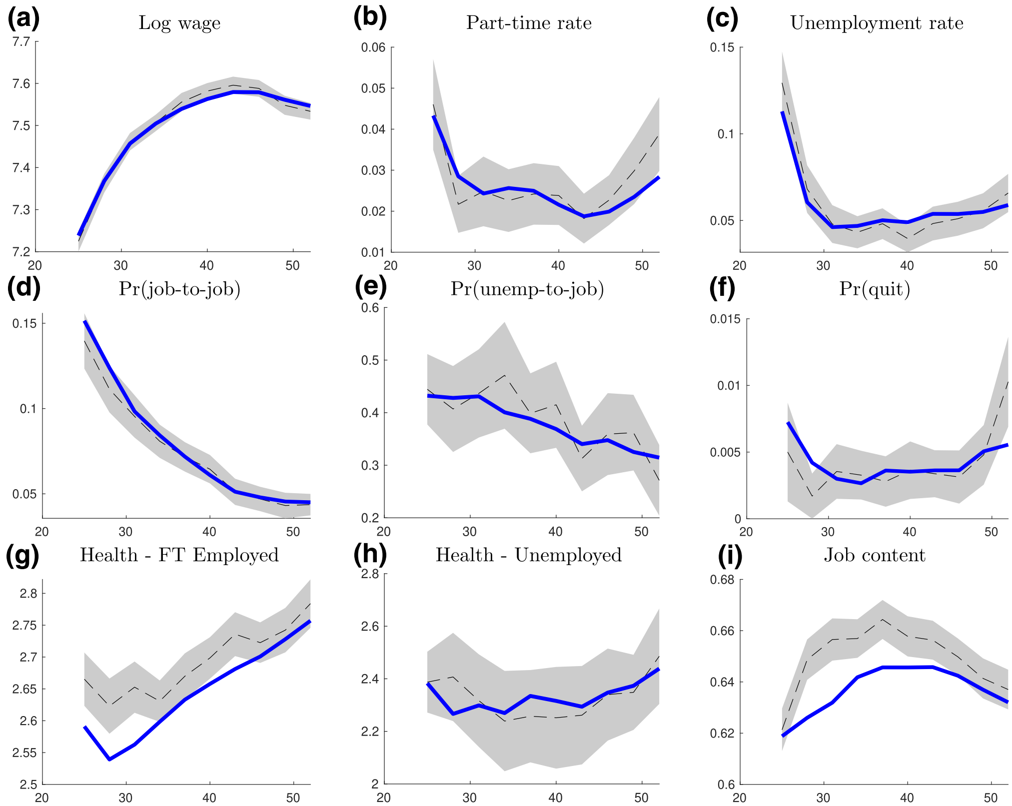

FIGURE 3
Model fit - Work and health outcomes by age

Notes: Age on horizontal axis. Solid line: Model. Dashed line: Data. Shaded area: 95% confidence bands around empirical moments. The health outcome is an average of a variable equal to 1, 2, 3, or 4 if health is S, P, A, or G respectively.

5.4. *Model fit*

Targeted moments. As explained in Section 4.3, we match two types of moments in the last estimation step. The first set of moments is a series of age profiles for health, job characteristics and labour turnover. The second set of moments are estimates from auxiliary regressions. Figure 3 presents the fit of the model for the first type of moments: model-predicted profiles are shown together with their empirical counterparts and associated 95% confidence bands.

Overall our model offers a very good fit to the profiles of the wage, unemployment and part-time rates, and health conditional on employment status (despite a slight tendency to under-predict the health of full-time workers under the age of 35). We do not show health for part-time workers to save space but the fit is as good as for the other two employment statuses. The middle row of Figure 3 shows that the model also captures the profiles of labour turnover, both between employment and unemployment and across jobs. The U-shaped pattern of quits is captured by our model. The last graph shows that while the inverse U-shaped profile of job stress is predicted by our model, we tend to under-predict the average job stress content. This may be due to an overly restrictive specification of the job offer distribution (only two parameter govern job health content offers) and could be fixed by allowing for a more flexible parametric distribution. However, looking at the magnitude of the difference between the data and model prediction, the

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

TABLE 8

Fit of the re-estimated heterogeneity classes

	% in simulated class			
Empirical class	1	2	3	4
1	75.7	12.9	11.4	0.0
2	8.6	90.9	0.4	0.1
3	36.4	2.8	60.2	0.7
4	1.7	0.7	5.6	92.0

model only understates s by less than 0.02 (less than 3% of the mean stress content). We thus deemed the computational cost of introducing more parameters to be unjustified as the model offers an accurate enough prediction of job stress content by age.

The fit to auxiliary regression parameters is shown in Appendix D, Figure A3, to save space in the main text. All regression coefficients are reasonably well captured by the model, which is thus able to fit a large set of moments relating to health, job outcomes or labour turnover with a relatively small number of parameters (only 17 in the indirect inference step of the estimation).

Unobserved heterogeneity. To conclude this section on model fit, we investigate the internal consistency of our estimated distribution of unobserved heterogeneity. As discussed in Section 4.1, we cannot formally prove that the set of individual outcomes used for clustering identifies heterogeneity classes. What we can show, however, is that our clustering is internally consistent in that it can recover classes used as inputs into the model. To that end, we simulate individual health and labour market trajectories using the estimated model with the classes produced in the first estimation step from the real data (see Section 5.1), then run our clustering algorithm on simulated data. Each "simulated" class should overlap strongly with only one of the "empirical" classes that were used as inputs.

Table 8 shows the mapping between the two partitions. Each number is the probability that an individual from the "empirical" class (rows) be assigned to a "simulated" class (columns). The results show that the minimal "matching rate" is 60% for one class but then the rate goes up to 75% for the next class and above 90% for the last two. Hence, while the mapping is not perfect, there is a very substantial overlap between the two clusters and we view these results as encouraging about the validity of our estimation of unobserved individual heterogeneity.

6. THE EFFECTS OF HEALTH AND EMPLOYMENT SHOCKS

6.1. *Job loss*

The value of employment. We begin our counterfactual analysis by quantifying the utility cost of job loss at ages 30 and 50. We do so by computing the difference in expected utility between working and being unemployed. Formally, the value of being employed in a type-$\mathbf{y}$ job relative to being unemployed for a type-x worker of age t and with health h can be written as:

$$\Delta^{(\ell)} V(t, \mathbf{y}, h) = \mathbf{E}_{x,\varepsilon}\left[V(x, t, h, \mathbf{y}) - \varepsilon - U(x, t, h)\right], \tag{8}$$

where V is defined by (2), U is defined by (1), and we average the difference over the distribution of individual heterogeneity x and labour supply shock ε. We report the value of employment $\Delta^{(\ell)} V(t, \mathbf{y}, h)$ for various health states, wage deciles, and job stress contents in Table 9, together with the relative welfare loss due to unemployment (in parentheses).

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

TABLE 9

Value of a full-time job (versus unemployment) at age 30 and 50

	Medium-stress job (s_3)			High-stress job (s_5)		
	Wage decile			Wage decile		
Initial health	2nd	5th	8th	2nd	5th	8th
Age 30						
Severe	15,864 (9.3%)	40,952 (20.2%)	80,695 (32.6%)	9,790 (5.9%)	33,079 (17.0%)	71,721 (30.2%)
Poor	19,066 (10.8%)	44,146 (21.3%)	83,877 (33.3%)	14,853 (8.5%)	38,172 (19.0%)	76,822 (31.5%)
Average	19,915 (11.2%)	44,986 (21.6%)	84,705 (33.5%)	16,387 (9.2%)	39,727 (19.5%)	78,380 (31.8%)
Good	20,611 (11.5%)	45,677 (21.8%)	85,390 (33.6%)	17,568 (9.8%)	40,918 (19.9%)	79,572 (32.1%)
Age 50						
Severe	5,114 (5.3%)	42,343 (27.9%)	111,330 (49.3%)	−7,728 (−7.0%)	21,744 (17.1%)	89,242 (44.1%)
Poor	16,800 (13.7%)	54,901 (32.5%)	123,023 (51.4%)	8,410 (7.4%)	39,947 (26.5%)	107,448 (48.2%)
Average	19,927 (15.5%)	57,232 (33.5%)	126,119 (51.7%)	12,710 (10.5%)	45,282 (28.7%)	112,775 (49.1%)
Good	22,501 (17.0%)	59,804 (34.3%)	128,670 (52.1%)	16,534 (13.1%)	49,461 (30.3%)	116,950 (49.9%)

Note: Value differences $\Delta^{(\ell)}V(t,\mathbf{y},h)$ as defined in (8), measured in GBP.
Relative differences in parentheses: $\Delta^{(\ell)}V(t,\mathbf{y},h)/\mathbf{E}_{x,\varepsilon}\left[V(x,t,h,\mathbf{y})-\varepsilon\right]$.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

We estimate the value of employment at age 30 to be substantial, ranging from £15,864 for a low-wage, medium-stress job (when in severe health) to £80,695 for a high-wage and medium-stress job (when in good health). The relative welfare loss due to unemployment is between 9 and 33% and we note that it decreases when health deteriorates *i.e.* jobs are less valuable to workers in poorer health (because of the interaction between health and disutility of labour).

The bottom panel of Table 9 shows that the range of employment values widens when going from age 30 to 50. For workers in good health and working in high-wage and high-stress jobs, age amplifies the negative effect of job loss (the value of employment goes up to £116,950). This may be a consequence of the job ladder: losing a job at a high rung of the ladder is more damaging for older workers as they will have less time to climb the ladder again and also their disutility of working in high-stress (and potentially better paid) jobs increases as they get older. Interestingly, we note that low-paid jobs with a high-stress content are worse than unemployment for 50-year-old workers in severe mental health (the relative difference in values is 7% that is £7,728 in level).

The effect of job loss on labour market and health trajectories. Next, we take a more dynamic perspective and evaluate the effect of job loss at age 30 and 50 on future paths of individual careers and health. First, we simulate the health and career trajectories of a representative[19] worker who at age 30 (resp. 50) is working full time at the median wage and a middling stress level (s_3); these workers are our “control group”. Second, we simulate the same trajectories for a representative worker who is unemployed at age 30 (resp. 50), the “treatment” group. We define the effect of job loss on future labour market and health paths as the difference between the age profiles of the treatment and control groups.

Figure 4 shows the impact of job loss on future health outcomes. Specifically, it shows the difference between treatment and control groups in the probability of being either in the severe or poor health state (which we collectively refer to as “bad health”), conditional on initial health. We first note that the initial health state affects the impact of unemployment on health.[20]

From Figure 4, we see that losing a median-wage/medium-stress job at age 30 causes a very persistent increase in the probability of being in poor or severe health. This effect peaks 1 year after job loss and decreases slowly afterwards. The moderate but noticeable increase in the risk of bad health is consistent with our estimates of the health dynamic process (see the top left panel of Table 4). The effect is slightly lower for 50-year-old workers but still persistent.

The persistence of the effect of job loss on health may arise from differences in labour market trajectories between the treatment and control groups, which cause different subsequent health outcomes. To assess that, we turn to labour market outcomes in Figure 5. In the first row, we focus on the average treatment effect of going from a median wage/medium-stress full-time job to unemployment at age 30 for a worker in average health. Figure 5(a) shows that the unemployment rate remains higher for the treated group after 5 years. The income loss, in Figure 5(b), shows even greater persistence, consistently with the recent macro-labour literature. Indeed, 5 years after the shock, income (wage or unemployment income) is still 30 log points lower for workers who lost their job at age 30 than for those who kept it. This is partly explained by a slow return to employment and by the job ladder mechanism, whereby workers gradually select into better jobs through on-the-job search. Interestingly, we note in Figure 5(c) that workers who lost their job at a medium-stress level return to work in slightly less stressful jobs than those who

19. We simulate trajectories for each worker type x, then take the average over the estimated distribution of x.

20. Importantly, recall that this effect is measured as a first difference (along the employment margin). Hence, Figure 4 shows that the increase in the risk of bad health due to unemployment (*versus* employment) is affected by the worker’s health at the time of job loss and, for example, that this increase is relatively smaller for workers who were in severe health than for those in average health.

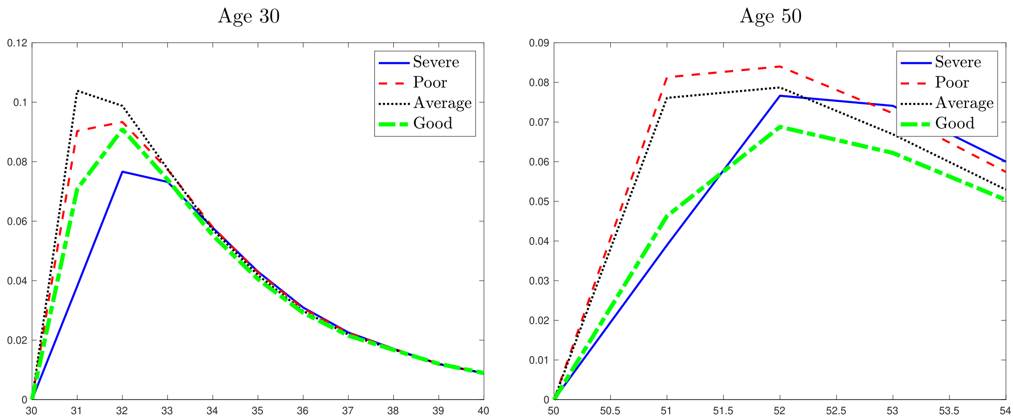

FIGURE 4

Health effects of job loss at age 30 and 50

Note: The vertical axis shows the difference in the probability of being in severe or poor health, given initial health at age 30 (or 50), between an unemployed worker and a full-time worker at the median wage/medium job stress content. The horizontal axis shows age (in years).

initially stayed employed. The combined effect of income loss and persistent unemployment contribute to the persistent health effect shown in Figure 4.[21]

The third row of Figure 5 shows similar effects of job loss for 50-year-old workers with even more persistence for income (almost 50 log points after 4 years) and unemployment. The increased persistence of averse income shocks with age is consistent with recent findings from the macro-labour literature (see Karahan and Ozkan, 2013).

Figure 5(a–c) (resp. Figure 5(g–i)) showed the average treatment effect of job loss on labour market outcomes conditionally on being in average health at age 30 (resp. age 50). To further explore the two-way relationship between health and work, we show in Figure 5(d–f) how these effects vary with the worker's initial health, taking average health as the reference. We can see that the short-term effect of job loss is worse when the worker's health is poorer. For example, we see in Figure 5(d) that being in severe rather than average health adds up to one percentage point to the increase in unemployment risk of treated workers during the 5 years following job loss. Consistently, the drop in income is larger for workers initially in severe health relative to average health. Interestingly, the effect on job stress content is lower for individuals in bad health. This can be explained by the fact that stressful jobs are more costly to workers in bad health so unemployed workers will be less likely to accept stressful job offers when their health is worse. The results are qualitatively identical for 50-year-old workers (Figure 5(j–l)).

6.2. *Health shocks*

The value of health for employed workers. We now quantify the utility cost of health shocks, proceeding as we did for job loss by computing the values of being in different health states while controlling for individual and job attributes. We define the relative value for an individual of age

21. The short-term income loss is found to be larger than those estimated by studies on US which, following Davis and Von Wachter (2011), focus on displacement shocks. This can be explained by the fact that we do not consider job reallocation (job loss instantaneously followed by a new offer) in this counterfactual exercise, or by the US having a more fluid labour market than the UK. In fact, our result is closer to those of studies on UK data such as Hijzen *et al.* (2010).

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

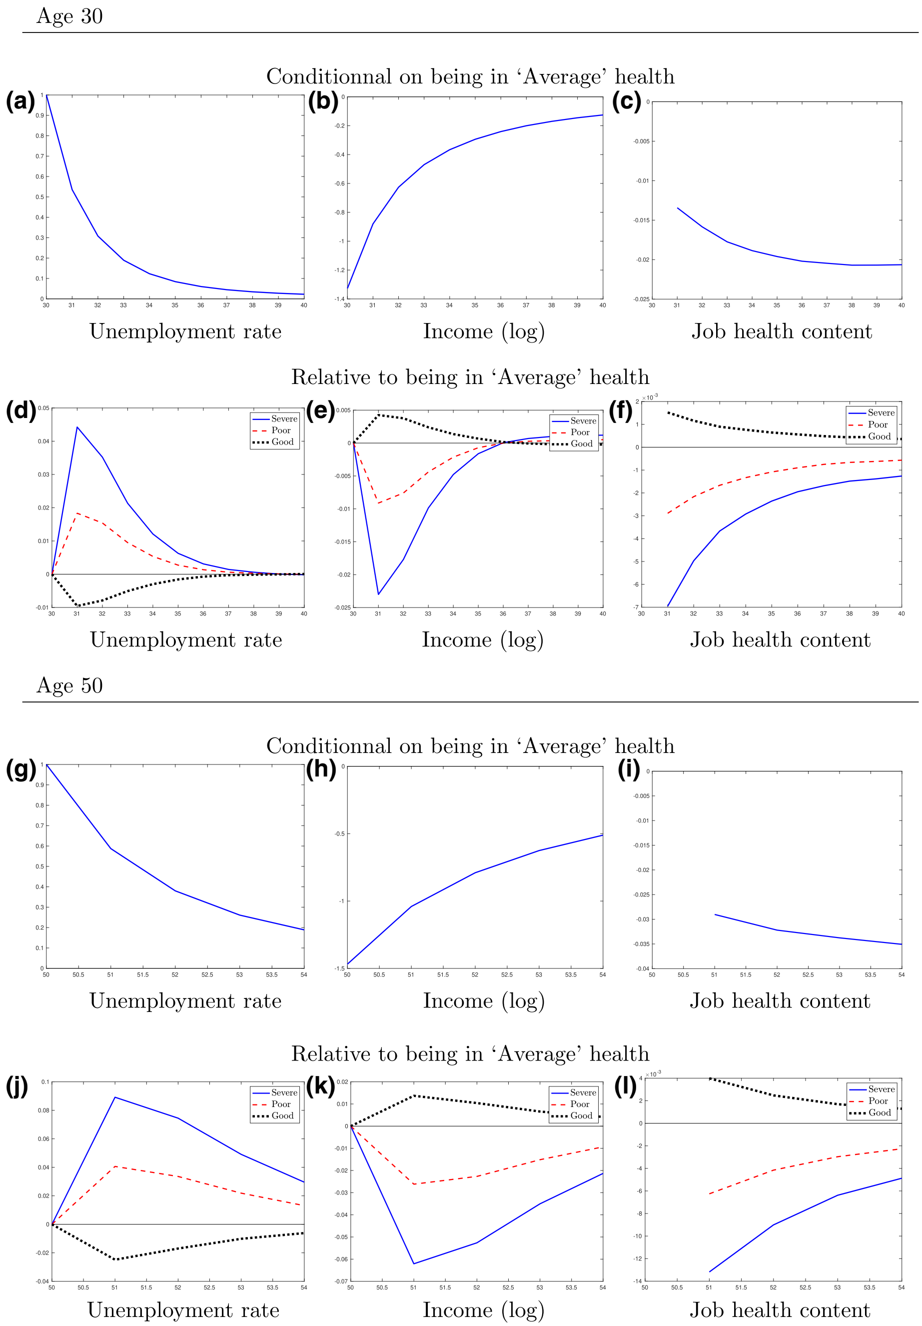

FIGURE 5

Labour market effects of job loss at age 30 and 50

Note: Age on horizontal axis. The vertical axis shows the labour market outcome difference between an unemployed worker and a full-time worker at the median wage/medium job stress content.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

TABLE 10
Value of being in the average versus severe health state at age 30 and 50

	Age 30		Age 50	
	Employed full-time at the median wage in a...		Employed full-time at the median wage in a...	
	medium-stress job	high-stress job	medium-stress job	high-stress job
$\Delta^{(h)}c$	3,086	4,570	11,257	16,666
$\Delta^{(h)}V$	5,428	7,925	18,176	26,270
$MWP^{(h)}$	0.049	0.073	0.124	0.179

Notes: Differences in costs $\Delta^{(h)}c$ and values $\Delta^{(h)}V$ are in GBP.
$\Delta^{(h)}c$ is the labour cost in average health minus the labour cost in severe health.
$\Delta^{(h)}V$ is the job value in average health minus the job value in severe health.
Medium, resp. high, stress refers to the 3rd (out of 5), resp. highest, stress level.

t and working in job $\mathbf{y} \neq \emptyset$, of being in health state h compared to state h' as $\Delta^{(h)}V(t, \mathbf{y}, h, h') = \mathbf{E}_x[V(x, t, h, \mathbf{y}) - V(x, t, h', \mathbf{y})]$. The value difference $\Delta^{(h)}V$ combines differences in instantaneous utility and differences in continuation values, as health affects future career dynamics. The former can be written as the labour supply cost difference $\Delta^{(h)}c(\ell, s, t, h, h') = c(\ell, s, t, h) - c(\ell, s, t, h')$. Note that since the wage enters the utility function additively and health is a discrete variable in our model, one can interpret $\Delta^{(h)}c$ as a structural marginal willingness to pay *i.e.* the wage increase required to equalise the instantaneous utility of two workers with different health statuses and otherwise similar characteristics.

Although $\Delta^{(h)}c$ is a model primitive (and thus invariant to policy), we may be more interested in the wage compensation needed to equalise value functions across different health states, since welfare would be measured by those values in our dynamic model. We then define $MWP^{(h)} = MWP^{(h)}(t, \mathbf{y}, h, h') = \mathbf{E}_x[MWP_x^{(h)}(x, t, \mathbf{y}, h, h')]$, where $MWP_x^{(h)}$ is implicitly defined by $V(x, t, h, \mathbf{y}) = V(x, t, h', (w \cdot e^{MWP_x^{(h)}(x,t,\mathbf{y},h,h')}, \ell, s))$. In words, $MWP^{(h)}$ is the average (in the population) log-wage increase required to set the value function of a worker in severe health equal to that of a worker in average health (and otherwise similar characteristics).

Results are shown in Table 10 for full-time workers in a median-wage job with either a middling or a high-stress content. We consider two ages when the health shock can occur: 30 and 50. The values shown are the differences between the average and the severe health states. The values of average (*versus* severe) health are large, in excess of £18,000 for older workers. In Table 9, we evaluated the value of a median wage, high-stress job (*versus* unemployment) at around £45,000 for a 50-year old in average health. The cost of severe health shocks is, therefore, substantial for these workers, in the order of 40% of the cost of job loss.

Comparing $\Delta^{(h)}V$ and $\Delta^{(h)}c$ in Table 10, we note that a large share of the value loss is incurred at the time of the health shock. The remaining loss is in the continuation value and is therefore incurred in subsequent years through two channels. The first is persistence in bad health and its related utility (labour supply) cost. As we will see in Figure 6, this effect is short lived. The second channel is triggered when the health shock induces workers to quit into unemployment. Quitting their job harms their future income and employment prospects, thus leading to a large drop in continuation value relative to workers who were not hit by the severe health shock.

Hence, even though our model does not explicitly account for the (flow) utility cost of being in poor health while unemployed, our estimates of the value of health for workers highlight the importance of taking health shocks into account when studying labour market trajectories. This is further supported by our estimates of workers' marginal willingness to pay for health

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

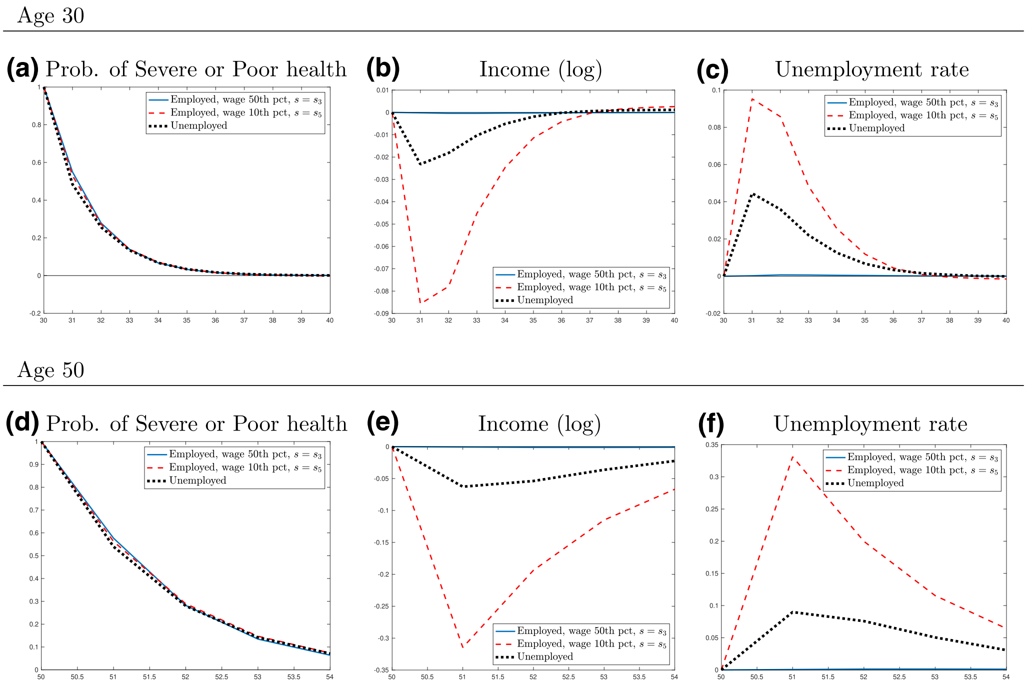

FIGURE 6

Effects of being in severe versus average health at age 30 and 50

Notes: Difference between the outcomes of individuals in severe versus average health at age 30 or 50 on horizontal axis.

(see the third row in Table 10): 30-year-old workers in average health and employed full time in a median-wage medium-stress (resp. high-stress) job would give up 5% (resp. 7%) of their wage to avoid a severe mental health shock. The marginal willingness to pay of older workers for average (*versus* severe) health goes up to 18% for a 50-year-old worker in a high-stress job.

The effect of a severe health shock on labour market and health trajectories. We next evaluate the dynamic effects of a severe health shock at ages 30 and 50 on future health and labour market paths. The “control group” now consists of a representative individual in labour market state **y** (employment, wage and stress content) at age 30 (or 50) who is in the average health state. The “treatment group” consists of the same individual starting out in the severe health state. We use our estimated model to simulate the health, income and unemployment trajectories for these two individuals and show the difference between treated and control groups in Figure 6. Three cases are considered: one where the individual is initially working full time at the median wage in a medium-stress (s_3) job (solid line), one where the individual is initially working full time at a low wage (first decile) in the most stressful (s_5) job type (dashed line) and one where the individual is initially unemployed (dotted line).

First, we note that it takes 4–5 years for differences in the probability to be in bad (*i.e.* severe or poor) health to disappear, regardless of the initial employment status. Then, the interaction between health and work which is at the core of our model is further illustrated in the second and third columns in Figure 6. We see that a severe health shock creates income losses that can persist for up to 5 years for workers who are unemployed or working in a stressful and poorly paid job. This is largely explained by the extensive margin of labour supply: Figure 6(f) shows that workers hit by the health shock are less likely to be employed over the following 5

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

years. The severe health shock increases the disutility of working, especially in a stressful job, leading some employed workers to quit, especially those in low-paying jobs, and preventing unemployed workers from accepting job offers they would take if they were in better health. Workers who quit their job following the health shock then see their health and labour market outcomes start down the unemployment dynamics shown in the previous subsection, with very persistent income losses and potentially long unemployment duration, especially for workers in bad health (as is the case for treated workers in our exercise).

We note that these affects are present for 30- and 50-year-old workers, with a lower magnitude for younger workers. The employment and income effects of adverse health shocks for younger workers are still substantial, with an unemployment probability increasing by almost 10% and an income loss of 8 log points (see Figure 6(b,c)).

Summing up, an important result of this analysis is that even though the effect of a severe health shock on individual health has mostly disappeared after 4–5 years, such a shock can generate persistent income losses and higher unemployment risks for workers who are either unemployed or employed in stressful or poorly paid jobs at the time of the shock. This result complements the ones shown in the previous subsection, providing further evidence of feedback effects between health and labour market outcomes.[22]

6.3. *Stressful jobs*

The willingness to pay for a nonstressful job. We now investigate workers' valuations of job health contents. Here again, we define the value for an individual of age t, health h, of working in job $\mathbf{y}=(w,s,\ell)$ relative to job $\mathbf{y}'=(w,s',\ell)$ as $\Delta^{(s)}V(t,h,\mathbf{y},\mathbf{y}')=\mathbf{E}_x[V(x,t,h,\mathbf{y})-V(x,t,h,\mathbf{y}')]$. The difference in instantaneous utility will be captured by $\Delta^{(s)}c(\ell,s,s',t,h)=c(\ell,s,t,h)-c(\ell,s',t,h)$.

We can further define $MWP^{(s)}=MWP^{(s)}(t,h,\mathbf{y},\mathbf{y}')=\mathbf{E}_x[MWP_x^{(s)}(x,t,h,\mathbf{y},\mathbf{y}')]$, where $MWP_x^{(s)}$ is implicitly defined by $V(x,t,h,\mathbf{y})=V(x,t,h,(w\cdot e^{MWP_x^{(s)}(x,t,h,\mathbf{y},\mathbf{y}')},s',\ell))$. In words, $MWP^{(s)}$ is the average log-wage change required to set the value of worker in a high-stress job s' equal to that of a worker in a medium-stress job s (and with otherwise similar characteristics).

We report the value of a high- *versus* medium-stress job in the first row Table 11 for full-time workers at the median wage and in each of the four health states. In both panels, the first row shows the difference in instantaneous utility. Working in a more stressful job substantially increases the disutility of working, especially for workers in severe health and for older workers. The second row shows that workers attach a substantial value to the health content of their job, with differences in values between a medium and a high-stress job ranging from £4,700 for 30-year workers in good health to more than £20,000 for older workers in severe health. As expected, the relative value of less stressful jobs decreases with better health and increases with age. The order of magnitude is large, as is reflected on the third row of each panel by the marginal willingnesses to pay. We find that 30-year-old workers are willing to give up between 3% (when in good health) and 6% of their wage (when in severe health) to work in a less stressful job. Note

22. We also investigated the potential role of search frictions in the propagation of health shocks by running counterfactuals where $\lambda_0=\lambda_1=1$ *i.e.* workers received job offers at every period. Results (not shown here, available on demand) show that in this environment with fewer search frictions, the effect of an adverse health shock is much weaker for 30-year workers in low-paid stressful jobs as, even if they quit their current job, they will find a new job fairly quickly. For 50-year-old workers we also observe a weaker effect, but only slightly so.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

TABLE 11

Value of a medium- versus high-stress job at age 30 and 50

	Health			
	Severe	Poor	Average	Good
		Age 30		
$\Delta^{(s)}c$	1,766	512	283	29
$\Delta^{(s)}V$	7,816	5,918	5,203	4,703
$MWP^{(s)}$	0.059	0.042	0.035	0.030
		Age 50		
$\Delta^{(s)}c$	6,442	1,866	1,032	106
$\Delta^{(s)}V$	20,580	14,125	11,932	10,325
$MWP^{(s)}$	0.139	0.097	0.081	0.070

Note: Differences $\Delta^{(s)}c$ and $\Delta^{(s)}V$ are in GBP.
$\Delta^{(s)}c$ is the labour cost in a stressful job minus the cost in a medium-stress job.
$\Delta^{(s)}V$ is the value in a median-stress job minus the value in a stressful job.

FIGURE 7

Effects of going from a medium to most stressful job type

Notes: In all graphs, workers are in "average" health at the starting age. Age on the horizontal axis. At the starting age, workers are employed full time at the 2nd, 5th, or 8th wage percentile (resp. w2, w5, and w8).

that these quantities are evaluated at the median wage. For 50-year-old workers, these marginal willingnesses to pay go up to 7% and 14% respectively.

The effect of working in a stressful job. We finally turn to the dynamic effects of an increase in job stress content on future health and career trajectories. We consider a representative full-time worker at a specific wage level and in the average health state, at two different ages: 30 and 50. In the control group, the worker is in a medium-stress job. In the treated group, the worker is in the most stressful type of job. We simulate the health and job health content trajectories for these two groups and show the differences in age profiles in Figure 7.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

The first column in Figure 7 shows that working in a more stressful job permanently raises the probability to be in bad (severe or poor) health by up to 3–4%. The effect is also persistent so that over time the proportion of workers whose health has deteriorated can become substantial. Interestingly, the persistence of the job stress shock depends positively on the initial wage for younger workers.

The heterogeneity in persistence is due to worker selection across jobs, as illustrated by the second column in Figure 7. Compared to the control group (workers in a medium-level stress job), workers in a high-stress job at age 30 or 50 will on average spend a long time in more stressful jobs. The effect decreases over time but is quite persistent. As mentioned above, this may come from search frictions. The persistence varies positively with the initial wage. Workers who earn a low wage when they are hit by a high job stress shock tend to go back to less stressful jobs more quickly than workers who initially earned higher wages. This explains the heterogenous responses of health dynamics. Indeed, a high wage may compensate for the disutility of working in a stressful job so workers feel less of an urge to change for a less stressful (and potentially less well-paid) job. This may have a negative effect on their future health but this effect is small, discounted as part of the continuation value and the higher wage will partly alleviate the disutility of working while in poor health.[23]

More stressful jobs also affect the extensive margin of labour supply, as illustrated in the last column in Figure 7. When the wage is not too low (at the median or higher quantiles), working in a medium or high-stress job makes little difference to quits or unemployment. However, a combination of low-wage and high-job stress does increase the risk of unemployment after a few years for workers aged 30 while the increase occurs right away for 50-year-old workers. This is because the wage no longer compensates for the higher disutility of working in a high-stress job, making workers likely to quit following an adverse health or labour supply shock. The increase in quits will set workers on a different income path: the average treatment effect on income, not shown here, is negative (an income loss of up to 30 log points during the subsequent 4 years) for 50-year-old workers at a low wage when the high job stress shock occurs.

7. A PARTIALLY DIRECTED SEARCH MODEL

In this final section, we explore a variant of our model where workers can partially direct their search to a specific job type (hence to a specific job stress level). From a modelling perspective, this gives foundations to the reduced-form feature we introduced in our benchmark model whereby employed workers have a higher probability to draw offers with a stress level similar to that of their current job. From a policy perspective, this extension assesses whether our counterfactual results on the effect of, say, health shocks, depend on our reduced-form way of capturing state-dependence in job offer sampling. For the reasons discussed in Section 3, we are unable to prove identification of this partially directed search model so we are treating it as an extension and, for brevity, we merely describe the model and a small selection of results in this section, confining the details to Appendix E.

23. Similarly to what we did with health shocks, we ran counterfactuals for job stress while setting the job offer arrival rates to 1, to simulate an environment with fewer search frictions. The results (not shown here, available on demand) show that up to 0.5% of the 3% increase in the probability of being in bad health following a stressful job shock in Figure 7(a) may be due to workers not always receiving job offers. The probability of being unemployed or the average job stress content also responds slightly less to job stress when offers arrive every period. This illustrates the interaction between search frictions and the health and employment effects of job stress shocks.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

7.1. *The model*

Workers can now direct their search to a specific job stress level s_k, but conditional on their targeted occupation, they draw wages and hours at random. We further assume that, before deciding where to search, workers draw a vector $\boldsymbol{\zeta} = (\zeta_1, \ldots, \zeta_K)$ of i.i.d job-type-specific shocks, which we interpret as a random "prominence shock" of each job type in the worker's information set. The labour supply shock ε is revealed after the worker has chosen where to search, but the prominence shock ($\boldsymbol{\zeta}$), the health shock (h'), and hours in the old job (ℓ') are revealed before search. The value of employment (gross of the labour supply shock) is:[24]

$$V(x,t,h,\mathbf{y}) = w - c(\ell,s,t,h) + (1+r)^{-1} \sum_{h'\in\mathcal{H}} \sum_{\ell'\in\mathcal{L}} \Pr\left\{h'|\mathbf{y},x,t,h\right\} \Pr\left\{\ell'|\ell\right\}$$
$$\times \left[\delta\left(1-\widetilde{\lambda}\right) U(x,t+1,h') + \delta\widetilde{\lambda}\mathbf{E}_\zeta \max_k \left\langle S\left(x,t+1,h',\emptyset,s_k\right) + \zeta_k\right\rangle\right.$$
$$\left. + (1-\delta)\mathbf{E}_\zeta \max_k \left\langle S\left(x,t+1,h',\mathbf{y}',s_k\right) + \zeta_k\right\rangle\right]$$

where $S(\cdot)$ is the expected value of (random) search in the market for jobs with stress level s_k. The latter object is defined as follows for employed workers ($\mathbf{y} \neq \emptyset$):

$$S(x,t,h,\mathbf{y},s_k) = (1-\lambda_1(s_k|s)) \int \max\left\{V(x,t,h,\mathbf{y}) - \varepsilon; U(x,t,h)\right\} dH(\varepsilon)$$
$$+ \lambda_1(s_k|s) \iint \max\left\{V\left(x,t,h,\mathbf{y}^o\right) - \varepsilon; V(x,t,h,\mathbf{y}) - \varepsilon; U(x,t,h)\right\}$$
$$\times dF(\mathbf{y}^o|x, s = s_k) dH(\varepsilon)$$

In the expression above, $\lambda_1(s_k)$ is the probability of a job seeker who is currently employed in a job with stress level s being offered a job with stress level s_k, conditional on applying for it, and $\widetilde{\delta}$ is the probability of a new job being destroyed within its first period of existence.

Next, if prominence shocks are i.i.d. extreme-value distributed, then the conditional probabilities of "target" job types are given by $\Pr\{s_k \mid x,t,h,\mathbf{y}\} = e^{S(x,t,h,\mathbf{y},s_k)} / \sum_{j=1}^K e^{S(x,t,h,\mathbf{y},s_j)}$ and the maximised value of search is $\mathbf{E}_\zeta \max_k \langle S(x,t+1,h',\mathbf{y}',s_k) + \zeta_k\rangle = \ln(\sum_{k=1}^K e^{S(x,t,h,\mathbf{y},s_k)}) + \gamma$, where γ is Euler's constant. This partially directed search model is therefore formally almost isomorphic to our main random search model, albeit with one important conceptual difference: the (nonnormalised) sampling density of job types is now endogenous and equal to $\lambda_1(s_k|s) \cdot \Pr\{s_k \mid x,t,h,\mathbf{y}\}$.

7.2. *Results*

Estimation results and results from the same counterfactual experiments as we conducted in Section 6 are shown in Appendix E. The fit to work and health outcomes age profiles and to auxiliary regressions is similar to our benchmark model. The criterion minimised in indirect inference (which is comparable across models as we are matching the exact same set of moments in both cases), is slightly lower (by 2%) in the directed search model. While it is not easy to compare point estimates across models (as parameters do not all have the same interpretation in

24. For brevity, we focus on employed workers in this section. Unemployed workers ($\mathbf{y} = \emptyset$) are covered in Appendix E.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

the two models), we can compare labour supply cost estimates. We find these values to be fairly close across models (see Table 7 and Table A4 in Appendix E), except in low-stress jobs where the directed search model implies larger labour supply costs (for example £238 per month for 30-year-old workers in severe health, compared to £159).

The counterfactual results on the effects of job loss are quantitatively similar between the benchmark and the directed search model, except for the health gradient of these effects which is more pronounced in the benchmark model (see Figure 5 and Figure A7 in Appendix E).

The impacts of health shocks remain qualitatively similar, although their scale is slightly lower in the directed search model for workers in high-stress job (for example the MWP for average versus severe health in stressful jobs is 1.5 percentage point larger in the benchmark model). Moreover, 30-year workers in low-wage stressful jobs are less likely to quit to unemployment when hit by a severe health shocks in the directed search model (see Figure 6 and Figure A8 in Appendix E). This is intuitive and motivates the analysis based on a (partially) directed search model. Indeed, if workers can direct their search towards specific job health contents then those who are in stressful jobs and suffer an adverse health shock will tend to look for less stressful jobs to reduce their disutility of working. This is not an option in the random search model as workers in high-stress jobs cannot direct their search so they are more likely to quit to unemployment when a severe health shock increases their labour supply cost.

A similar comment applies to the differences between models in the effects of job stress. In a directed search model, workers value job health content less as they can direct their search toward less stressful jobs if labour supply in their current job becomes too costly. The estimated MWPs for medium (*versus* high) job stress are still substantial in the directed search model (see Table A7 in Appendix E) but are lower than those from the benchmark model (see Table 11). For example, the MWP for a 50-year-old worker goes increases 7% (when in good health) to 14% (when in severe health) in the benchmark model and from 5% to 9% in the directed search model.

This extension can be seen as a robustness check for our counterfactual results. It shows that our qualitative conclusions about the importance of health and job content in workers' utility, and about work and health outcomes interacting to propagate of shocks, do not depend on whether search is random or directed. However, quantitative differences between the two models also shed light on the extent to which workers may mitigate health or job content shocks by searching directly for less stressful jobs.

8. CONCLUSION

We have constructed a structural model of the joint dynamics of individual careers and mental health trajectories. The model accounts for two-way interactions between work and health and for key features of the labour market, such as search frictions, on-the-job search, and selection of workers into jobs based on wages and the health contents of jobs. Taking the model to British data, we estimate the dynamic response of health to various job attributes, as well as worker selection into jobs as a function of health and job characteristics. We then use our estimated model to quantify the impact of job loss, health shocks, or an increase in job stress on individual health and labour market outcomes, highlighting the way in which the health (resp. employment) channel amplifies the adverse effects of job loss (resp. of a health shock).

Our model can be extended to address specific issues. For example, we could shift the focus to older workers and add more structure on the extensive employment margin (modelling pensions and savings) to study how an individual's employment, wage but also occupation history affects their late-career health and retirement decisions. This would extend the work of French (2005), French and Jones (2011), and Salvati (2020), and could prove useful for the design of pension systems that include occupational health contents as a factor impacting eligibility.[25]

Another important extension is to study the interaction between health and work for women. We focus on men in the present paper because a serious analysis of the health and careers of women would have to model fertility decisions and their impact on both health and career outcomes. One of our counterfactual results showed that men in low-wage, stressful jobs are especially vulnerable to health shocks as they can trigger an unemployment spell, with lasting effects on income and employment, which then slow down health recovery. We think these mechanisms may be even more salient for women as fertility may act as an additional propagation channel for health or employment shocks. The approach used here could thus inform the joint design of unemployment, parental leave and health insurance policies.

Acknowledgments. We thank the Editor and two anonymous referees for helping us improve our work. We would also like to thank Stéphane Bonhomme, Donna Gilleskie, Eric French, Amanda Michaud, Joe Mullins, Renata Narita, Andrew Oswald, Helen Simpson, Anthony Terriau, Marcos Vera-Hernandez, and Frank Windmeijer for very useful comments at various stages in the production of this paper. We are grateful to seminar and workshop participants at Aarhus (Sandbjerg Manor), Aix-Marseille (SaM), CREST, Edinburgh, Essex, Georgetown, Louvain-La-Neuve, Lyon, Nottingham, Royal Holloway, UCL, Warwick, and Yale for helpful feedback. The usual disclaimer applies.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability Statement

The data (except UKLHS files) and codes underlying this research are in a replication package available on Zenodo at https://doi.org/10.5281/zenodo.10462481. The UKLHS data files are available from the UK Data Archive (https://ukdataservice.ac.uk) for registered users subject to an end user licence agreement. Detailed instructions on how to access these data are given in the replication package.

REFERENCES

BONHOMME, S., JOLIVET, G. and LEUVEN, E. (2016), "School Characteristics and Teacher Turnover: Assessing the Role of Preferences and Opportunities", *The Economic Journal*, **126**, 1342–1371.
BONHOMME, S., LAMADON, T. and MANRESA, E. (2019), "A Distributional Framework for Matched Employer Employee Data", *Econometrica*, **87**, 699–739.
BONHOMME, S., LAMADON, T. and MANRESA, E. (2022), "Discretizing Unobserved Heterogeneity", *Econometrica*, **90**, 625–643.
BURDETT, K., CARILLO-TUDELA, C. and COLES, M. (2020), "The Cost of Job Loss", *The Review of Economic Studies*, **87**, 1757–1798.
CAPATINA, E., KEANE, M. and MARUYAMA, S. (2020), "Health Shocks and the Evolution of Earnings over the Life-Cycle" (mimeo, University of New South Wales).
CURRIE, J. and MADRIAN, B. (1999), "Health, Health Insurance and the Labor Market", in Ashenfelter, O. and Card, D. (eds) *Handbook of Labor Economics* (North-Holland, Amsterdam: Elsevier Science, chap. 50) 3309–416.
DAVIS, S. and VON WACHTER, T. (2011), "Recessions and the Cost of Job Loss", *Brookings Papers on Economic Activity*, **43**, 1–72.
DE NARDI, M., PASHCHENKO, S. and PORAPAKKARM, P. (2017), "The Lifetime Costs of Bad Health" (NBER working paper no. 23963, National Bureau of Economic Research).
FRENCH, E. (2005), "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behavior", *The Review of Economic Studies*, **72**, 395–427.
FRENCH, E. and JONES, J. (2011), "The Effects of Health Insurance and Self-Insurance on Retirement Behavior", *Econometrica*, **79**, 693–732.
FRYERS, T., MELZER, D. and JENKINS, R. (2003), "Social Inequalities and the Common Mental Disorders", *Social Psychiatry and Psychiatric Epidemiology*, **38**, 229–237.

25. For instance, such a reform was implemented in France in 2015 ("*compte pénibilité*"). The number of years spent in "strenuous" jobs can bring forward the age of full pension eligibility.

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024

GILL, S., BUTTERWORTH, P., RODGERS, B., *et al.* (2007), "Validity of the Mental Health Component Scale of the 12-Item Short-Form Health Survey (MCS-12) as Measure of Common Mental Disorders in the General Population", *Psychiatry Research*, **152**, 63–71.

GILLESKIE, D. (1998), "A Dynamic Stochastic Model of Medical Care Use and Work Absence", *Econometrica*, **66**, 1–45.

GRUBER, J. and MADRIAN, B. (2004), "Health Insurance, Labor Supply, and Job Mobility: A Critical Review of the Literature", in McLaughlin, C. (ed) *Health Insurance, Labor Supply, and Job Mobility: A Critical Review of the Literature* (Washington, DC: Urban Institute Press) 97–178.

HARRIS, M. (2019), "The Impact of Body Weight on Occupation Mobility and Career Development", *International Economic Review*, **60**, 631–660.

HIJZEN, A., UPWARD, R. and WRIGHT, P. (2010), "The Income Losses of Displaced Workers", *Journal of Human Resources*, **45**, 243–269.

JACOBS, L. and PIYAPROMDEE, S. (2016), "Labor Force Transitions at Older Ages: Burnout, Recovery, and Reverse Retirement" (Finance and Economics Discussion Series 2016-053, Washington: Board of Governors of the Federal Reserve System).

JAROSCH, G. (2015), "Searching for Job Security and the Consequences of Job Loss" (mimeo, University of Chicago).

KARAHAN, F. and OZKAN, S. (2013), "On the Persistence of Income Shocks over the Life Cycle: Evidence, Theory, and Implications", *Review of Economic Dynamics*, **16**, 452–476.

LAYARD, R. (2017), "The Economics of Mental Health" (IZA World of Labor).

LISE, J. and ROBIN, J.-M. (2017), "The Macro-Dynamics of Sorting Between Workers and Firms", *American Economic Review*, **107**, 1104–1135.

LOW, H., MEGHIR, C. and PISTAFERRI, L. (2010), "Wage Risk and Employment Risk over the Life Cycle", *American Economic Review*, **100**, 1432–1467.

MICHAUD, A. and WICZER, D. (2018a), "The Disability Option: Labor Market Dynamics with Economic and Health Risks" (mimeo, University of Western Ontario).

MICHAUD, A. and WICZER, D. (2018b), "Occupational Hazards and Social Disability Insurance", *Journal of Monetary Economics*, **96**, 77–92.

OECD (2015), *Fit Mind, Fit Job: From Evidence to Practice in Mental Health and Work, Mental Health and Work* (Paris: OECD Publishing).

PAPAGEORGE, N. (2016), "Why Medical Innovation is Valuable: Health, Human Capital, and the Labor Market", *Quantitative Economics*, **7**, 671–725.

SALVATI, F. (2020), "Female Health and Labor Supply at Older Ages" (mimeo, University College London).

SALYERS, M., BOSWORTH, H., SWANSON, J., *et al.* (2000), "Reliability and Validity of the SF-12 Health Survey among People with Severe Mental Illness", *Medical Care*, **38**, 1141–1150.

SULLIVAN, D. and WACHTER, T. VON (2009), "Job Displacement and Mortality: An Analysis Using Administrative Data", *Quarterly Journal of Economics*, **124**, 1265–1306.

TRAN, T. (2017), "Too Stressed to Work: The Effects of Job Stressors on Health and Employment Outcomes" (mimeo, University of North Carolina at Chapel Hill).

WARE, J., KOSINSKI, M. and KELLER, S. (1996), "A 12-Item Short-Form Health Survey: Construction of Scales and Preliminary Tests of Reliability and Validity", *Medical Care*, **34**, 220–233.

WHO (2019), "Special Initiative for Mental Health (2019–2023)" (World Health Organization, WHO/MSD/19.1).

Downloaded from https://academic.oup.com/restud/advance-article/doi/10.1093/restud/rdae071/7698384 by University College London user on 08 October 2024