



The health burden of job strain: evidence from Europe

Petru Crudu^{1,2} · Giacomo Pasini^{2,3}

Received: 12 November 2025 / Accepted: 2 September 2026

© The Author(s) 2026

Abstract

This paper studies how occupational stressors and job tasks over the life course are associated with health later in life. Using SHARE data and retrospective job histories for individuals across Europe, we construct career-long exposure to occupational stressors by linking 4-digit ISCO occupations to external measures of job demands. We find that physical job intensity, measured by the metabolic equivalent of task, is associated with worse later-life health, and this result remains stable across sensitivity analyses. Higher exposure is also associated with a greater incidence of some cardiovascular and respiratory diseases and with an earlier onset of several cardiovascular, metabolic, and respiratory conditions. We use detailed task information and a LASSO-based approach to identify the job tasks most strongly associated with poor health and to build a Job Strain Intensity index. This task-based framework offers a granular and scalable way to detect harmful dimensions of work and inform policy discussions.

Keywords Health · Working conditions · Job strain · Job tasks

Responsible editor: Klaus F. Zimmermann

This study was funded by the European Union - Next Generation EU, Mission 4 Component 2, within the NRRP (National Recovery and Resilience Plan) project “Age-It - Ageing well in an ageing society” (PE0000015 - CUP H73C22000900006). The views and opinions expressed are only those of the authors and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them. Petru Crudu gratefully acknowledges financial support from the Austrian Science Fund (FWF), project grant P36845-G (DOI: 10.55776/P36845). An older version of this paper based on a PhD thesis chapter of Petru Crudu is available in the Working Papers series of the Economics Department, Ca’ Foscari University [here](#). We thank the Editor, Klaus F. Zimmermann, and the two anonymous referees for their constructive comments and guidance, which helped us improve the paper substantially. We also thank Beatrix Eugster, Lorenzo Cappellari, and Agar Brugiavini for fruitful comments and discussions.

✉ Petru Crudu
Petru.Crudu@uibk.ac.at

Giacomo Pasini
giacomo.pasini@unive.it

¹ University of Innsbruck, Innsbruck, Austria

² Ca’ Foscari University of Venice, Venice, Italy

³ NETSPAR, Network for Studies on Pensions, Ageing and Retirement, Tilburg, Netherlands

JEL Classification I12 · I14 · J81**1 Introduction**

European citizens spend 26% of their adulthood at work.¹ Long hours of desk work, constant repetitive motions, physically demanding tasks, or high-stress decision-making do not just make the workday harder, but can also affect health. These effects can accumulate over time and become more apparent later in life. Understanding how working conditions over the life course affect health is therefore important.

Despite a growing literature in health economics and occupational epidemiology on the relationship between working conditions and health (Case and Deaton 2005; Holtermann et al. 2010; Fletcher et al. 2011; Samitz et al. 2011; Krause et al. 2015; Ravesteijn et al. 2018), several important questions remain open. In particular, the long-term health consequences of sustained exposure to workplace stressors over the full working life are still not well understood, and it remains unclear which specific job tasks are most harmful to health. Broad proxy measures may obscure substantial differences across activities, making it difficult to determine the aspects of work that matter most for health later in life.

In this paper, we use retrospective occupational histories to examine how occupational stressors and job tasks experienced over the working life are associated with health at older ages. This allows us to study exposure over the full career and to provide a more detailed account of the dimensions of work most strongly associated with later-life health.

We use microdata from the Survey of Health, Ageing and Retirement in Europe (SHARE), a multidisciplinary, cross-national survey of individuals aged 50 and over covering 27 European countries and Israel (Börsch-Supan et al. 2013). SHARE provides detailed information on health, socioeconomic characteristics, and, crucially for our purposes, retrospective life histories that allow us to reconstruct individuals' full occupational trajectories over the working life (Brugiavini et al. 2019). Occupations are reported at the 4-digit level of the International Standard Classification of Occupations (ISCO), which allows us to assign occupational stressors and tasks precisely to each job. Our main measure of physical job intensity is the Metabolic Equivalent of Task (MET) (Deyaert et al. 2017), which captures the energy expenditure associated with the tasks performed in each occupation. In addition, we consider exposure to psychosocial stress and carcinogenic agents (Kroll 2015), as well as the routine content of job tasks (Mihaylov and Tijdens 2019).

To estimate the relationship between occupational stressors and health later in life, we rely on a rich set of controls that is more detailed than is typical in this literature. In particular, beyond standard socio-demographic characteristics, our data allow us to account for detailed childhood circumstances, a broad set of working-life characteristics, and health-related behaviours.

¹ This calculation is an average taking into account several factors. The expected duration of working life in the European Union surpasses 36.5 years (Eurostat 2022). Average weekly working hours amount to approximately 36.2 h (Eurostat 2021). We define adulthood as the working-age period from age 20 to 65. Finally, we assume an average of 8 h of sleep per night.

Our results show that occupational physical intensity is the only occupational stressor that is consistently and significantly associated with worse health at older ages. By contrast, we find no consistent evidence of significant effects for the other occupational stressors. In the baseline specification, a one-unit increase in MET, which corresponds roughly to an interquartile-range increase in occupational physical intensity, is associated with a deterioration in self-perceived health equivalent to about 2.6 additional years of ageing. In practical terms, this corresponds roughly to moving from a relatively less physically demanding occupation, such as *driving instructors*, to a more physically intensive one, such as *shelf fillers*. We also document significant gender heterogeneity, with larger associations for women than for men.

Although these estimates are economically meaningful, identification remains challenging in this setting, and the results should therefore be interpreted cautiously. We therefore evaluate the sensitivity of our estimates to potential selection and reverse causality concerns through a range of diagnostics and robustness checks. A central advantage of the data is that we can control for an unusually rich set of predetermined characteristics, including childhood socioeconomic and health conditions. Within this rich-control framework, Oster bounds (Oster 2019) provide an informative, though not definitive, sensitivity check: selection on unobservables would need to come close to the magnitude of selection on observed characteristics to eliminate the estimated association between occupational physical intensity and later-life health.

We further assess reverse causality using evidence on job mobility and early-career exposure. These checks yield estimates around 80–90% of the baseline coefficient, suggesting that reverse causality is unlikely to account for the full association. We also implement the heteroskedasticity-based instrumental variable approach of Lewbel (2012) as an additional robustness check, obtaining estimates that remain close in magnitude and significance to the baseline results. Additionally, we consider a wider range of health outcomes, showing that the main findings are robust across different dimensions of later-life health. We further assess the robustness of the results using Causal Forest methods (Athey et al. 2019), which allow for a more flexible relationship between occupational stressors and health, and find that the main findings remain robust.

We also examine how prolonged exposure to occupational physical intensity is associated with the incidence and timing of specific diseases. The results indicate that higher MET exposure is linked to an increased probability of some conditions, particularly cardiovascular and respiratory diseases. For example, a one-unit increase in MET is associated with about a 1.3 percentage point higher probability of heart attack and a 0.8 percentage point higher probability of stroke, and is also positively associated with lung disease. Evidence from the duration models further suggests that higher exposure is associated with an earlier onset of certain conditions. For example, a one-unit increase in MET is associated with a 13% higher hazard of heart attack, an 11% higher hazard of stroke, and a 10% higher hazard of diabetes. Some gender differences also emerge. In particular, the point estimates are larger for heart attack, hypertension, and diabetes among women, and for lung disease among men. Overall, the results point to a meaningful relationship between sustained physical job intensity and later-life health deterioration, while remaining cautious about causal interpretation.

In the second part of the paper, we develop a novel approach based on the Least Absolute Shrinkage and Selection Operator (LASSO; Tibshirani 1996) and detailed task information for each 4-digit ISCO occupation. This approach allows us to detect the tasks most strongly associated with worse health outcomes from a sparse, high-dimensional task matrix. We then use this information to construct a new measure, which we call Job Strain Intensity, capturing the share of health-adverse task predictors within each 4-digit ISCO code. This data-driven index is designed to capture broader dimensions of job strain, including components that standard occupational stressor measures may fail to observe directly.

This second part of the paper shows that the association between MET and health is robust to controlling for Job Strain Intensity. It also offers a more granular view of the task content most strongly associated with adverse health outcomes. For instance, the five tasks that emerge as most predictive of worse later-life health are *cleaning kitchen, bathroom, and toilet; identifying, reporting, and documenting hazards; operating heavy-duty equipment; farming; and repair and maintenance*.

Previous research on the effects of occupational strain on health spans different approaches. In the medical and occupational epidemiology literature, Samitz et al. (2011) provide a meta-analysis of cohort studies and find an overall protective association between higher physical activity, including occupational activity, and all-cause mortality; Cillekens et al. (2022) review cohort studies of occupational physical activity and cardiovascular mortality and find no clear protective effect of physically demanding work; Holtermann et al. (2010) show in a cohort study that high physical work demands are associated with higher ischaemic heart disease and all-cause mortality, especially among less fit workers; and Krause et al. (2015) show in a prospective cohort study that higher relative occupational workload increases the long-run risk of acute myocardial infarction. These studies provide mixed evidence, likely due to substantial differences in data, exposure measurement, outcomes, and empirical design.

The literature closest to our paper lies in health economics and studies the relationship between working conditions and health using individual-level data and regression-based approaches. Case and Deaton (2005) use repeated cross-sections from the US National Health Interview Survey over 1986–2001 and show that workers in manual occupations report worse health and experience steeper health decline with age than non-manual workers. Fletcher et al. (2011) use US panel data from the PSID, merged with occupational characteristics from the Dictionary of Occupational Titles, to construct 5-year cumulative measures of physical demands and adverse environmental conditions. They find that exposure to these job characteristics worsens self-reported health, with stronger effects for some groups, especially women. Ravesteijn et al. (2018) use data from the German Socio-Economic Panel, linked to occupation-level measures of physical and psychosocial job stressors, and estimate a dynamic panel model to study the relationship between occupation and health. Their results indicate that physical strain and low job control are detrimental to health.

While this literature has provided important insights, several aspects of how occupation shapes health remain only partially examined. We address these gaps in several ways. First, existing studies typically observe occupational histories over only a limited portion of the working life, often while individuals are still in the early or middle stages of their careers. As a result, occupational exposure is measured over only part

of the career, and the long-run health consequences of occupational strain, especially those that become more apparent at older ages, may not yet have fully emerged. We address this limitation by exploiting retrospective information to reconstruct individuals' full occupational trajectories and measure exposure over the entire working life. This is important because the effects of occupational stressors are likely to depend on prolonged exposure and to accumulate over time through mechanisms such as allostatic load (McEwen 2000), so that their health consequences may become more evident at older ages. We therefore focus on health later in life, when individuals are close to or at the end of their careers, and most occupational exposure has already taken place. This allows us to study the long-run health consequences of working conditions once they are more likely to have materialised.

Second, the measurement of occupational stressors in the existing literature remains relatively coarse, ranging from broad manual versus non-manual categories to occupation-level indicators of physical demands or environmental conditions. We improve on this by using more granular measures of occupational stressors, in particular physical intensity based on Metabolic Equivalents of Task (MET), which capture energy expenditure rather than relying on coarse or binary proxies. In addition, we examine a broad range of later-life health outcomes, including specific diseases, which allows us to provide a more detailed picture of how occupational exposure is associated with health.

Third, the existing literature provides limited evidence on which specific task dimensions of occupations are most predictive of worse later-life health. We address this issue through a task-based approach. Using detailed task information at the 4-digit ISCO level, we detect the tasks most strongly associated with worse health and construct a new index, Job Strain Intensity, defined as the share of harmful tasks within each occupation.

Finally, whereas much of the existing evidence comes from single-country settings, especially the United States, we provide evidence from multiple European countries. This broader comparative perspective improves the generalisability of the findings relative to the existing literature and allows us to examine whether the relationship between occupational stressors and later-life health emerges across diverse institutional and labour-market settings.

This paper is organised as follows. Section 2 describes the data, and Section 3 outlines the empirical strategy. Section 4 presents the main results together with a range of sensitivity analyses. Section 5 then introduces the novel LASSO-based approach. Section 6 concludes.

2 Data

2.1 SHARE data

The Survey of Health, Ageing and Retirement in Europe (SHARE) (Börsch-Supan et al. 2013; Brugiavini et al. 2019) is a harmonised cross-national longitudinal survey designed to be representative of individuals aged 50 and over in participating European countries, as well as Israel. It is a multidisciplinary panel survey that col-

Table 1 Sample composition by country

Country	Freq.	Percent	Country	Freq.	Percent
Austria	3043	4.70	Hungary	684	1.06
Belgium	4967	7.67	Israel	1416	2.19
Czech Republic	4551	7.02	Latvia	660	1.02
Denmark	3321	5.13	Lithuania	1068	1.65
Estonia	3975	6.13	Luxembourg	938	1.45
Finland	1159	1.79	Malta	552	0.85
France	3534	5.45	Netherlands	1940	2.99
Germany	4204	6.49	Poland	4312	6.65
Greece	547	0.84	Portugal	775	1.20
Italy	4005	6.18	Romania	910	1.40
Slovakia	1598	2.47	Slovenia	2748	4.24
Spain	3873	5.98	Sweden	3622	5.59
Switzerland	2438	3.76	Bulgaria	1585	2.45
Croatia	1849	2.85	Cyprus	521	0.80
Total observations: 64,795 (100%)					

Notes: This reports the number and share of unique individuals by country in the analysis sample. Shares are computed relative to the full sample of 64,795 individuals. Country refers to the respondent's country of interview

lects harmonised microdata on health, socioeconomic status, and social and family networks among older individuals.

SHARE is particularly well suited to our analysis for three reasons. First, the retrospective life-history interviews collected in SHARELIFE Waves 3 and 7 provide detailed information on individuals' trajectories over the life course. Among other aspects, they record complete occupational histories, with each job episode coded at the 4-digit ISCO level, allowing us to reconstruct employment trajectories over time. Second, the regular SHARE interviews provide several measures of health observed at older ages and diagnosed diseases, which serve as our main outcomes. Third, it contains retrospective information on early-life circumstances, which allows us to account for childhood conditions. Combining these features, we use the retrospective life-history information collected in SHARELIFE Waves 3 and 7 to relate lifetime occupational exposure to later-life health outcomes while controlling for early-life background.²

The resulting sample contains 64,795 unique individuals with information on at least one occupation. Table 1 shows the sample distribution by country.

² For respondents interviewed in SHARELIFE Wave 7, health outcomes are measured in Wave 7 and retrospective employment histories are used up to the corresponding interview age. For respondents interviewed in SHARELIFE Wave 3, health outcomes are taken from the last available regular SHARE interview (Wave 1 or 2), and retrospective employment histories are correspondingly truncated at the age reached in that interview. This ensures that occupational exposures are measured before the observed health outcome and avoids using job information realised after outcome measurement. The results are robust to restricting the analysis to respondents observed in SHARELIFE Wave 7 only.

2.2 Occupations and occupational stressors

A central feature of our analysis is the possibility of attaching comparable measures of occupational demands to each job held over the life course. To do so, we use the occupational histories recorded in SHARELIFE, which report occupations using the International Standard Classification of Occupations (ISCO-08) (ILO 2012), and link them to external occupation-level measures of job demands.

2.2.1 Occupations

SHARELIFE provides detailed information on respondents' lifetime occupations, including the sequence of jobs held over the life course. Each job episode is coded according to the International Standard Classification of Occupations (ISCO-08), which classifies jobs on the basis of the tasks and duties they involve.³ ISCO-08 is organised in four hierarchical levels, ranging from ten broad major groups, such as managers and elementary occupations, to 436 detailed 4-digit unit groups.

In the analysis, we rely on the 4-digit ISCO classification. This level of detail is important because substantial heterogeneity in tasks and working conditions remains even within the same major group. As illustrated in Table A.5 in Appendix Section A.5 for the major group "Elementary occupations," occupations within the same broad category can involve very different tasks. For example, "Crop Farm Labourers" (ISCO 9211) and "Fast Food Preparers" (ISCO 9411) both belong to elementary occupations, yet differ substantially in their task content. Using 4-digit occupations therefore allows us to attach occupational stressors more precisely and to measure lifetime occupational exposure more accurately.

2.2.2 Occupational stressors

We construct measures of lifecycle occupational exposure by combining the retrospective employment histories recorded in SHARELIFE Waves 3 and 7 with external occupation-level indicators of job demands.

For each individual, SHARELIFE records all job spells together with their start year, end year, and associated 4-digit ISCO occupation code. We use this occupation code to assign the occupational stressor measures described below to each spell. Lifecycle occupational exposure is then computed as the duration-weighted average of the stressor intensity assigned across all job spells over the working life. This measure therefore captures the average level of occupational demands experienced over the career, weighting each occupation by the number of years spent in it. Accordingly, the analysis estimates the association of assigned occupation-level stressor indices based on harmonised 4-digit ISCO codes, rather than directly observed individual historical strain. Appendix Section A.6 provides further details on the construction of these measures.

Our main measure of occupational physical strain is the metabolic equivalent of task (MET), constructed by Deyaert et al. (2017). We focus on this measure because

³ See the International Labour Organization documentation on ISCO-08.

it provides a task-based and externally defined indicator of the physical demands of work, offering a more granular measure of occupational physical strain than broad occupational categories or self-reported job conditions commonly used in the literature. MET captures the ratio of the metabolic rate associated with a given activity to the standard resting metabolic rate and therefore measures how physically demanding an activity is relative to rest. For example, sitting quietly corresponds to 1 MET, whereas walking at a very low speed in an office environment corresponds to around 2 METs. Higher MET values therefore indicate more physically demanding occupations.

To construct occupation-level MET values, Deyaert et al. (2017) match the job-specific tasks listed in the ISCO classification to the corresponding energy expenditure values reported in the Compendium of Physical Activities by Ainsworth et al. (2011), and then average these values across tasks within each 4-digit occupation. These MET values are defined at the activity level and therefore do not capture individual-specific differences in energy expenditure. For example, “University and higher education teachers” (ISCO-08 code 2310) are assigned an average MET value of 1.58. An illustrative mapping example is reported in Table A.1 in Appendix Section A.1.

Beyond physical strain, we also consider psychosocial stress and exposure to carcinogenic agents in the workplace. To measure these dimensions, we rely on the Job Exposure Matrix (JEM) developed by Kroll (2015), which is derived from a large-scale representative survey conducted in 2006 by the German Federal Institute for Vocational Education and Training (BIBB). The psychosocial stress metric, defined on a scale from 0 to 10, combines three dimensions of workplace pressure: mental stress (feeling overwhelmed, disappointed, and having low fault tolerance at work), social pressure (conflicts with colleagues or superiors and lack of control options), and time pressure (time pressure, shift work, and excessively long working hours). Similarly, the carcinogenic agents exposure metric, also defined on a scale from 0 to 10, assigns each 4-digit ISCO occupation to a decile of exposure to toxins, gases, and climatic stress.

We also account for the routine content of job-specific tasks. The repetition of manual tasks in poor ergonomic conditions may affect health regardless of the level of physical effort required. Likewise, easy and repetitive cognitive tasks may affect mental health. To capture this dimension, we use the measures developed by Mihaylov and Tijdens (2019), which assign to each 4-digit ISCO occupation both the manual and the cognitive routine content of tasks. Starting from the ISCO classification, they classify tasks as either manual or cognitive routines depending on their susceptibility to replacement by computer-controlled technology. They then compute, for each 4-digit ISCO occupation, the share of manual routine tasks and the share of cognitive routine tasks.

A limitation of these occupational stressor measures is that they are time-invariant and reflect occupational characteristics measured around the first decade of the 2000s. They therefore do not capture changes in job content over time, for example due to technological change or automation. We return to this issue in Section 4.3.3, where we assess the sensitivity of the results to potential cohort differences in the measurement of occupational exposures.

2.3 Health outcomes and background variables

2.3.1 Health outcomes

As our main outcome, we use self-perceived poor health (*sphus*). We focus on this measure because it is widely used in the literature as a summary measure of overall health status (Fletcher et al. 2011; Ravesteijn et al. 2018). In SHARE, respondents rate their health on a five-point scale ranging from 1 (*excellent*) to 5 (*poor*). Higher values therefore indicate worse health. We complement this main measure with a broader set of health outcomes to assess whether the baseline results are robust across alternative and more specific dimensions of health. These include limitations in activities of daily living (ADL), limitations in instrumental activities of daily living (IADL), limitations with usual activities, mobility limitations, maximum grip strength, the number of chronic conditions, body mass index (BMI), disability due to injury, and an indicator for whether health limits work.⁴

SHARE also records information on a broad set of diagnosed diseases, which allows us to explore which specific conditions are associated with lifetime occupational exposure. We consider the following diagnoses: heart attack, hypertension, high cholesterol, stroke, diabetes, chronic lung disease, cancer, stomach or duodenal ulcer, Parkinson's disease, cataracts, hip fracture, other fractures, Alzheimer's disease, emotional disorders, rheumatoid arthritis, osteoarthritis or other rheumatism, and chronic kidney disease.

2.3.2 Additional background variables

SHARE contains rich information on respondents' background characteristics over the life course.⁵ First, it provides standard socio-demographic characteristics, including gender, age at interview, year of birth, country of residence, and marital status.

Another important feature of SHARE is the availability of retrospective information on childhood circumstances between ages 10 and 15. These variables include the occupation of the main breadwinner in the household, measured as a categorical variable based on the 1-digit ISCO classification, as well as the number of rooms in the home, the number of books available in the household, and respondents' relative ability in mathematics and language at age 10. The survey also records a broad set of childhood health conditions, which we summarise into an index of early-life health using principal component analysis.

SHARELIFE also provides detailed information on respondents' working lives. This includes the total number of years worked up to the interview, the number of job changes over the career, and several characteristics of the main job, including the NACE sector, whether it was part-time and whether the respondent worked as a civil servant, in a private company, or as self-employed.

Finally, SHARE also contains information on a number of health-related behaviours, which may be relevant for health over the life cycle (Mahler and Yum 2024). In par-

⁴ Appendix Section A.2 provides further details on the construction of the health outcomes.

⁵ Appendix Section A.3 provides further details on the construction of the control variables.

ticular, we observe whether respondents ever smoked daily, whether they report being physically inactive in their leisure time, and the frequency with which they consume legumes or eggs, meat, and fruit.

2.4 Descriptive statistics

Table 2 reports descriptive statistics for the main outcome and the occupational stressors used in the baseline analysis; descriptive statistics for the remaining health outcomes and control variables are reported in Appendix Sections A.2 and A.3, respectively.

Figure 1 provides preliminary descriptive evidence on the relationship between lifetime occupational physical strain and later-life health. It shows a clear gradient in self-perceived health: respondents reporting worse health tend to display a MET distribution shifted upwards. This purely descriptive pattern shows that higher lifetime occupational physical strain is associated with worse health at older ages, and helps motivate the empirical analysis that follows.

Figure 2 documents substantial heterogeneity in lifetime occupational physical strain across demographic, socioeconomic, sectoral, geographic, and childhood dimensions. Panel a shows marked cross-country differences in the distribution of MET, with higher values in countries such as Romania, Poland, and Spain, and lower values in countries such as Switzerland and Israel. Panel b highlights equally pronounced sectoral variation: construction, agriculture, forestry and fishing, and mining display the highest MET values, whereas finance is by far the least physically demanding sector. Panel c shows a clear educational gradient, with respondents holding higher levels of education concentrated in substantially less physically demanding occupations, while panel d documents a gender gradient, with men displaying a higher MET distribution than women. Panels e–h further show that lifetime occupational exposure

Table 2 Descriptive statistics for the main outcome and occupational stressors

Variable	Obs.	Mean	Std. dev.	Min	Max
<i>Main outcome</i>					
Self-perceived poor health	64,795	3.19	1.07	1	5
<i>Occupational stressors</i>					
Metabolic equivalent of task (MET)	64,795	2.33	0.79	1.30	6.23
Psychosocial stress (OSI)	64,795	5.19	2.74	0	10
Carcinogenic exposure (CAI)	64,795	5.38	2.91	0	10
Routine manual	64,795	0.10	0.19	0	1
Routine cognitive	64,795	0.22	0.26	0	1

Notes: This reports descriptive statistics for the main outcome and the occupational stressors used in the baseline analysis. Self-perceived poor health is measured on a five-point scale ranging from 1 (excellent) to 5 (poor), with higher values indicating worse health. MET measures occupational physical intensity. OSI and CAI denote psychosocial stress and carcinogenic exposure, respectively. Routine manual and routine cognitive capture the routine task content of occupations. Occupational stressors are constructed as duration-weighted averages over the working life using retrospective job histories linked to external occupation-level measures




Fig. 1 Distribution of occupational physical intensity by self-perceived health status. *Notes:* This reports the distribution of lifetime occupational physical intensity, measured in MET values, across categories of self-perceived health in the analysis sample. The boxplots display the median and interquartile range. The whiskers extend to the lower and upper adjacent values

is related to childhood background. Respondents with better values of the childhood health index tend to display higher lifetime MET exposure, and MET is also higher among those who grew up in households with fewer books and among those reporting weaker relative mathematical and language ability at age 10. Taken together, these descriptive patterns suggest that lifetime occupational physical strain is already shaped by early-life circumstances.

3 Empirical strategy

We estimate the following linear specification by OLS:

$$y_i = \beta_0 + \beta_1 \mathbf{S}_i + \beta_2 \mathbf{X}_i + u_i, \quad (1)$$

where y_i denotes the health outcome of individual i . $\mathbf{S}_i$ is the vector of lifetime occupational stressors (Metabolic Equivalent of Task (MET), psychosocial stress, carcinogenic exposure, routine manual intensity, and routine cognitive intensity). $\mathbf{X}_i$ collects the covariates described in Section 2, including socio-demographic characteristics, childhood background variables, job-related characteristics, and, where available, health-related behaviours. u_i is the error term. Standard errors are two-way clustered by 4-digit ISCO occupation and country, allowing for correlation in the residuals within occupations and within countries. All specifications use SHARE sample weights to account for the survey design, including unequal sampling probabilities and differential non-response.

Two endogeneity concerns are relevant in this setting. The first is self-selection into occupations: individuals with poorer underlying health or less healthy behaviours may be more likely to sort into jobs with less desirable working conditions. The




Fig. 2 Descriptive patterns in lifetime occupational physical intensity. *Notes:* This reports lifetime occupational physical intensity, measured in MET values. **a** and **b** show boxplots by country and industry; **c** reports MET values by education; **d** shows boxplots by gender; and **e–h** relate MET exposure to childhood health, books in the household at age 10, and self-reported mathematical and language ability at age 10

second is reverse causality: health and behaviour may also influence later occupational trajectories, for instance, if workers in worse health move out of physically demanding jobs or are less able to remain in better positions over time.

Our control variables include unusually detailed information that is often unavailable in similar studies, such as childhood circumstances, working conditions, and health-related behaviours. Nevertheless, our empirical strategy does not support a definitive causal interpretation. Accordingly, we interpret the coefficients in β_1 as associations between lifetime occupational stressors and later-life health outcomes. Section 4.2 shows that the results remain robust across a range of sensitivity checks.

A further consideration in interpreting the empirical strategy concerns the age structure of the SHARE sample. Because SHARE covers individuals aged 50 and over, our analysis is necessarily limited to those who survive to at least age 50 and are observed in the survey. This may generate non-random selection by occupational exposure. A natural concern is that adverse working conditions may increase mortality risk or reduce the probability of survey participation, leading the estimation sample to be selected towards healthier survivors among the exposed. To the extent that adverse working conditions reduce survival or survey participation among less healthy exposed individuals, the estimated associations may be attenuated. Nevertheless, the direction of survivor bias cannot be established definitively.

4 Empirical results

4.1 Baseline results

Table 3 reports the estimates of Eq. 1. In the most parsimonious specification, lifetime occupational physical intensity (MET) is positively and statistically significantly associated with worse self-reported health at older ages. The coefficient in column 1 is 0.174 and remains statistically significant throughout all specifications. Some of the other job characteristics are also statistically significant in the less saturated specifications. However, as additional controls are introduced, these associations lose statistical significance. Once we account for socio-demographic characteristics, childhood conditions, and job-related factors, only MET remains statistically significant. In column 4, which we treat as our baseline specification, the estimated coefficient on MET is 0.073. This indicates that the association between occupational physical intensity and health at older ages is robust to the inclusion of a rich set of confounding factors.⁶

To provide an intuitive benchmark for the magnitude of the coefficients, we first express them in terms of equivalent years of ageing.⁷ The baseline estimate in column 4 implies that a one-unit increase in MET corresponds to a deterioration in self-perceived health equivalent to about 2.6 additional years of ageing. In our data, the

⁶ As a robustness check, Appendix Section A.9.1 re-estimates the baseline specification after standardising all job stressors and, alternatively, transforming them into deciles. The results remain robust.

⁷ Specifically, following Ravesteijn et al. (2018), we estimate the association between age and self-perceived poor health in a separate regression of the same outcome on age, which yields a coefficient of 0.0276. We then divide the coefficient of interest by this age gradient, thereby benchmarking the effect of working conditions against the effect of ageing.

Table 3 Main results: occupational stressors and self-perceived poor health

	(1) sphus	(2) sphus	(3) sphus	(4) sphus	(5) sphus
MET	0.174*** (0.022)	0.109*** (0.021)	0.093*** (0.021)	0.073*** (0.023)	0.073** (0.027)
Psychosocial stress	−0.008 (0.007)	−0.007 (0.005)	0.000 (0.004)	0.003 (0.005)	0.000 (0.004)
Carcinogenic exposure	0.008 (0.008)	0.004 (0.007)	0.011* (0.006)	0.009 (0.006)	0.007 (0.007)
Routine manual	0.353*** (0.077)	0.263*** (0.069)	0.133* (0.075)	0.112 (0.087)	0.144 (0.090)
Routine cognitive	0.045 (0.089)	0.024 (0.063)	0.055 (0.051)	0.038 (0.052)	0.032 (0.058)
Constant	2.751*** (0.094)	3.043*** (0.153)	1.170** (0.424)	1.104** (0.496)	0.639** (0.286)
Socio-demo	No	Yes	Yes	Yes	Yes
Child	No	No	Yes	Yes	Yes
Job	No	No	No	Yes	Yes
Habits	No	No	No	No	Yes
Observations	64,795	55,908	55,908	55,765	43,355
R-squared	0.027	0.060	0.165	0.175	0.191

Notes: This reports OLS estimates of the association between lifetime occupational stressors and self-perceived poor health. The dependent variable is self-perceived poor health. MET measures occupational physical intensity, while the remaining variables capture psychosocial stress, carcinogenic exposure, routine manual tasks, and routine cognitive tasks. Columns progressively add socio-demographic characteristics, childhood circumstances, job-related controls, and health-related habits. Standard errors, reported in parentheses, are two-way clustered by 4-digit ISCO occupation and country. SHARE sample weights are applied.

* $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$

25th percentile of MET is approximately 1.7, whereas the 75th percentile is about 2.73, so a one-unit increase closely corresponds to an interquartile shift in occupational physical intensity. In practical terms, this is comparable to moving from a relatively less physically demanding occupation, such as driving instructors (MET = 1.75), to a more physically intensive one, such as shelf fillers (MET = 2.71).

Finally, column 5 shows that the coefficient on MET remains very similar once health behaviours are added, although this comes at the cost of a reduction in sample size.

4.2 Sensitivity analysis

4.2.1 Selection into occupations

A first concern is that lifetime occupational exposure may partly reflect sorting into occupations. Individuals from less advantaged backgrounds, or with different initial health conditions and behaviours, may be more likely to enter jobs with less desirable

working conditions. If these same factors are also related to health at older ages, the baseline estimates may partly capture pre-existing differences across individuals rather than only differences in job demands accumulated over the life course.

Previous literature shows that family background and childhood conditions are important predictors of lifetime socioeconomic and health outcomes (Case et al. 2002, 2005; Currie 2009). Our setting is particularly informative in this respect because the data include a rich set of childhood variables that are often unavailable in related studies. In particular, we observe the occupation of the main household breadwinner, which helps capture the family labour-market environment and the intergenerational transmission of job-related stressors. We also include indicators of childhood housing conditions, which proxy the socioeconomic circumstances in which individuals grew up. In addition, the number of books in the household and relative mathematical and language ability at age 10 capture early cognitive and educational resources that are likely to shape subsequent educational attainment, first job placement, and broader occupational opportunities (Brunello et al. 2017). We also include a PCA-based index of childhood health, which captures pre-labour-market health endowments that may influence both later health and the ability to enter, or avoid, specific jobs. Together, these variables absorb several key dimensions of initial selection into occupations.

To assess the sensitivity of the estimates to omitted variable bias, we implement the Oster procedure. The Oster approach builds on the idea that the movement in the coefficient of interest and in the R^2 across specifications can be informative about the potential role of unobserved confounders (Altonji et al. 2005; Oster 2019). Intuitively, if the coefficient remains stable once a rich set of observables is included, and the explanatory power of the regression rises substantially, then selection on unobservables would need to be relatively strong to overturn the result. This approach is particularly useful in our setting because the controlled specification includes not only childhood circumstances, but also detailed socio-demographic characteristics and job-related variables such as career length, job mobility, sector, and employment arrangements. Nevertheless, the Oster exercise is best interpreted as an informative sensitivity analysis rather than conclusive evidence, as residual confounding may still reflect unobserved characteristics.

Formally, the Oster procedure computes a value δ such that the estimated coefficient β would be reduced to zero. The parameter δ measures how strong selection on unobservables would need to be relative to selection on observables in order to fully eliminate the estimated effect. A value of $\delta = 1$ means that unobservables would need to be as important as the observed controls to explain away the result, while larger values imply that even stronger selection on unobservables would be required.

A key implementation choice concerns the maximum attainable R^2 , $R_{\max}$, which bounds the explanatory power that could be achieved if both observed and unobserved factors were included in the regression. Because $R_{\max}$ is not directly observed, it must be calibrated. Following Oster (2019), we set $R_{\max} = m \cdot \tilde{R}^2$, where $\tilde{R}^2$ denotes the R^2 of the baseline specification. We use $m = 1.3$ as our benchmark, in line with Oster's suggestion for social science applications, and examine the sensitivity of the results to alternative values.

Table 4 reports the implied values of δ from the Oster procedure for the baseline specification. Under the benchmark calibration $R_{\max} = 1.3\tilde{R}^2$, the implied value is

Table 4 Oster procedure: implied values of δ^*

m	$R_{\max}$	δ^*
1.1	0.193	2.377
1.2	0.210	1.239
1.3	0.228	0.838
1.4	0.246	0.633
1.5	0.263	0.509

Notes: This reports the implied values of Oster's δ^* , that is, the relative degree of selection on unobservables compared with observables required to reduce the estimated coefficient on MET to zero. For each value of m , $R_{\max}$ is defined as $m \cdot \bar{R}^2$, where $\bar{R}^2 = 0.175$ is the R^2 from the baseline specification. The baseline specification includes socio-demographic characteristics, childhood circumstances, and job-related controls

$\delta^* = 0.838$. This means that, in order to drive the estimated MET coefficient to zero, selection on unobservables would need to be of roughly comparable magnitude to selection on observables. The corresponding values range from 2.377 when $m = 1.1$ to 0.509 when $m = 1.5$, showing that the result varies mechanically with the choice of $R_{\max}$.

Taken together, the Oster procedure indicates that, given the rich set of relevant controls included in the baseline specification, omitted variable bias would not trivially account for the estimated relationship under standard calibrations. At the same time, the results also indicate that unobserved confounding of a non-negligible magnitude could in principle attenuate the estimated effect. Nevertheless, the Oster exercise remains informative by providing a quantitative gauge of how important unobserved factors would need to be relative to the observed determinants already accounted for in the baseline specification.

4.2.2 Reverse causality and the timing of job mobility

A second concern is reverse causality. If individuals with different health conditions or health-related behaviours face different probabilities of remaining in, leaving, or moving into better jobs over time, the observed relationship between occupational stressors and later-life health may partly reflect the effect of health on occupational trajectories, rather than only the effect of occupational exposure on health. For example, individuals with poorer health may face greater difficulty retaining better jobs, while worsening health may also reduce access to occupations requiring a minimum health threshold and thereby induce moves towards less demanding work.

To examine this issue, we begin by showing that lifetime occupational mobility is relatively limited over the course of the career. Table 5 shows that around two-thirds of individuals held no more than two job spells over their entire working life, more than 84% experienced fewer than four distinct 4-digit ISCO occupations, and more than 44% remained within the same 1-digit ISCO major group throughout. The timing of job changes is also informative. Table 6 (a) shows that job mobility is concentrated early in adulthood: more than half of all job changes occur before age 30, and roughly three quarters take place before age 40. By contrast, Table 6 (b) shows

Table 5 Job mobility over the working life

(a) Number of jobs			
N jobs	Freq	%	Cum %
1	24,612	37.98	37.98
2	17,958	27.72	65.70
3	10,795	16.66	82.36
4	5827	8.99	91.35
5 or more	5603	8.65	100.00
Total	64,795	100.00	
(b) Unique occupations, 4-digit ISCO			
N jobs (4-digit ISCO)	Freq	%	Cum %
1	23,997	37.04	37.04
2	19,030	29.37	66.40
3	11,540	17.81	84.21
4	5714	8.82	93.03
5 or more	4514	6.97	100.00
Total	64,795	100.00	
(c) Unique occupations, 1-digit ISCO			
N jobs (1-digit ISCO)	Freq	%	Cum %
1	28,751	44.37	44.37
2	21,246	32.79	77.16
3	10,390	16.04	93.20
4	3335	5.15	98.34
5 or more	1073	1.66	100.00
Total	64,795	100.00	

Notes: This reports the distribution of job mobility over the working life in the analysis sample. Panel (a) reports the total number of jobs held over the career. Panels (b) and (c) report the number of unique occupations held over the career, defined at the 4-digit and 1-digit ISCO level, respectively

that the first reported diagnosis is concentrated much later in life, with only about 16% of individuals reporting a first disease before age 40, while more than 65% report their first disease after age 50 or no disease at all.

Taken together, these patterns suggest that occupational exposure is typically shaped by a relatively small number of occupational positions, most of which are determined early in working life. This makes it less likely that lifetime exposure is driven mainly by health-induced job changes at older ages. We now provide several diagnostics and sensitivity checks to investigate the role of reverse causality.

We start from Table 7, which compares MET values in the first and last job by childhood health conditions, where better childhood health corresponds to the 0–74th percentiles and worse childhood health to the 75th–100th percentiles. Individuals with poorer childhood health both start and end their careers in less physically demanding jobs: the average MET in the first job is 2.322 in the worse-health group, compared with 2.430 among those with better childhood health, while in the last job the corresponding values are 2.234 and 2.340. These differences are statistically significant in both cases. At the same time, the decline in physical job demands over the career is very similar

Table 6 Timing of job changes and first disease diagnosis

(a) Age at job change			
Age	Job changes	%	Cum %
10–20	11,603	13.18	13.18
21–30	33,377	37.92	51.11
31–40	20,946	23.80	74.90
41–50	14,803	16.82	91.72
51–60	6716	7.63	99.35
> 60	568	0.65	100.00
Total	88,013	100.00	
(b) Age at first diagnosis			
Age	First disease	%	Cum %
10–20	2198	3.48	3.48
21–30	3060	4.84	8.32
31–40	5001	7.91	16.23
41–50	10,695	16.91	33.14
51–60	14,072	22.26	55.40
> 60	11,239	17.78	73.17
No disease	16,963	26.83	100.00
Total	63,228	100.00	

Notes: This reports the timing of job changes and first disease diagnosis in the analysis sample. Panel (a) shows the age distribution of job changes over the working life. Panel (b) shows the age at first disease diagnosis; the category “No Disease” identifies respondents who do not report any disease diagnosis

across the two groups. The average difference between the first and last job is 0.090 for individuals with better childhood health and 0.087 for those with worse childhood health, and the difference between the two is small and statistically insignificant.

Further evidence on reverse causality is reported in Table 8. To assess whether the baseline relationship is partly driven by health-induced occupational adjustments, we implement two complementary diagnostics. First, we restrict the sample to individuals with very stable careers, namely, those who held only one job over the entire working

Table 7 MET in the first and last job by childhood health

	Better Childhood health		Worse Childhood health		Difference	<i>t</i> -test
	Mean	Std error	Mean	Std error		
MET first job	2.430	(0.004)	2.322	(0.007)	0.109	13.166***
MET last job	2.340	(0.004)	2.234	(0.006)	0.106	13.420***
Difference (first–last)	0.090	(0.0037)	0.087	(0.0064)	0.003	0.390

Notes: This reports mean MET values in the first and last job by childhood health group. Better childhood health includes respondents in the 0–74th percentiles of the childhood health index, whereas worse childhood health includes respondents in the 75th–100th percentiles. The last two columns report the difference in means across groups and the corresponding *t*-test. The final row reports the within-group difference between MET in the first and last job

Table 8 Robustness: stable careers and early-career occupational exposure

	Job mobility (one job) sphus (1)	Occupational mobility (one occ., 4-digit) sphus (2)	Early exposure (up to age 40) sphus (3)
MET	0.063** (0.023)	0.066** (0.030)	0.058*** (0.016)
Psychosocial stress	−0.002 (0.005)	−0.004 (0.005)	0.003 (0.004)
Carcinogenic exposure	0.005 (0.004)	0.008* (0.004)	0.009 (0.006)
Routine manual	0.080 (0.071)	0.114* (0.056)	0.070 (0.078)
Routine cognitive	−0.001 (0.077)	0.016 (0.085)	−0.037 (0.061)
Constant	1.698*** (0.538)	1.496*** (0.517)	1.320** (0.480)
Observations	17,017	22,016	54,314
R-squared	0.232	0.216	0.167

Notes: This reports robustness checks for the baseline specification relating lifetime occupational stressors to self-perceived poor health (sphus). The baseline specification includes socio-demographic characteristics, childhood circumstances, and job-related controls. Column 1 restricts the sample to respondents with exactly one job over the working life. Column 2 restricts the sample to respondents with one unique 4-digit ISCO occupation. Column 3 measures occupational exposure using only jobs held up to age 40. Standard errors, reported in parentheses, are two-way clustered by 4-digit ISCO occupation and country. SHARE sample weights are applied. * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$

life or remained within a single occupation throughout the career, defined at the 4-digit ISCO level. Second, we reconstruct occupational exposure using only jobs held up to age 40, so that MET is measured well before the stage of the life cycle in which most health deterioration is observed.

The estimates provide limited support for reverse causality as the main explanation for the baseline result. In the one-job sample, the MET coefficient remains positive and statistically significant at 0.063, compared with 0.073 in the baseline specification. When the sample is restricted to individuals who remain within a single occupation at the 4-digit level, the corresponding coefficient is 0.066, again very close to the baseline estimate. When exposure is measured using only jobs held up to age 40, the estimated coefficient remains positive and is precisely estimated at 0.058. Relative to the baseline estimate, these coefficients are between roughly 80% and 90% of the original magnitude, suggesting that only a limited share of the baseline relationship could plausibly be attributed to reverse causality.

4.2.3 Lewbel instrumental variable approach

Occupational sorting is likely shaped by many of the same factors that also affect health later in life. This makes conventional exclusion restrictions difficult to defend in our

setting, since variables that influence lifetime exposure to physically demanding work are also likely to affect later-life health directly. To address this concern, we implement the heteroskedasticity-based identification strategy proposed by Lewbel (2012). We interpret this exercise as a complementary robustness check to the baseline estimates, rather than as a standalone causal design.

The Lewbel approach constructs internal instruments from higher-moment conditions in the data. Intuitively, identification comes from systematic differences in the dispersion of the endogenous regressor across observable groups. In our application, the endogenous regressor is lifetime occupational physical intensity (MET).⁸ The internal instruments are generated as products of demeaned variables in a subset of observables, denoted by Z , and the first-stage residual from the regression of MET on the full set of controls.

Formally, consider the system

$$Y_1 = X'\beta_1 + Y_2\gamma_1 + \varepsilon_1, \quad (2)$$

$$Y_2 = X'\beta_2 + \varepsilon_2, \quad (3)$$

where Y_1 denotes later-life health, Y_2 is lifetime occupational physical intensity (MET), and X includes the exogenous controls in the benchmark specification. Following Lewbel (2012), the internal instruments take the form $(Z - \bar{Z})\hat{\varepsilon}_2$, where $\hat{\varepsilon}_2$ is the residual from the first-stage regression of MET on X . Identification requires two conditions:

$$\text{Cov}(Z, \varepsilon_2^2) \neq 0, \quad (4)$$

$$\text{Cov}(Z, \varepsilon_1\varepsilon_2) = 0. \quad (5)$$

The first condition requires that the variance of the first-stage residual vary systematically with the variables in Z , so that the data contain the heteroskedasticity needed to generate relevant instruments. The second condition requires that, although unobserved factors may jointly affect occupational selection and later-life health, their covariance does not vary systematically with the variables in Z , conditional on the full set of controls. In this sense, the Lewbel design does not require the absence of selection on unobservables, but rather a restriction on how this selection co-varies with the instrument-generating variables.

In our application, the variables used to generate the internal instruments include age, gender, birth cohort, country indicators, childhood housing conditions (rooms per capita), childhood socioeconomic indicators (number of books at home, self-reported ability in mathematics and language, and parental occupation), and the childhood health index. These variables are chosen because they are plausible sources of group-level heteroskedasticity relevant for the first condition above, while the second condition is more credible when Z is restricted to predetermined characteristics

⁸ The benchmark specification treats MET as the only endogenous regressor. We do not include the other occupational stressors because they may also be jointly determined with occupational sorting and later-life health, making them difficult to isolate in a single-endogenous-regressor Lewbel design. Appendix A.10 reports sensitivity checks in which all five occupational stressors are treated as endogenous, with broadly robust results.

realised before later-life health is measured. In substantive terms, this means that while unobserved traits such as preferences for physically demanding work may affect both occupational choice and health, the covariance between the unobserved determinants of occupational choice and later-life health is assumed to be stable across the groups defined by Z , once the full set of controls is included. The benchmark specification generates 35 excluded instruments.

The data are consistent with these requirements. A Breusch–Pagan-type test applied to the first-stage residuals strongly rejects homoskedasticity (BP LM = 2116.000, $p = 0.000$), supporting the relevance condition $\text{Cov}(Z, \varepsilon_2^2) \neq 0$. The Kleibergen–Paap rk LM statistic rejects underidentification (KP LM = 98.890, $p = 0.000$), while the corresponding rk Wald F statistic is large (10431.950), indicating strong first-stage relevance. Weak-IV-robust inference based on the Anderson–Rubin test also rejects the null of no effect (AR $p = 0.000$). The Hansen overidentification test does not reject the joint validity of the excluded instruments ($p = 0.133$). While this does not establish instrument validity, it indicates that the overidentifying restrictions implied by the moment conditions are not contradicted by the data.

The Lewbel estimate is close to the baseline OLS result (Table 9). When MET is instrumented using the heteroskedasticity-based procedure, the estimated coefficient is 0.087 (s.e. 0.026), compared with 0.073 (s.e. 0.023) in the baseline OLS specification. The estimate therefore remains positive, precisely estimated, and economically similar in magnitude.

Appendix A.10 provides additional evidence on the sensitivity of the Lewbel design. In particular, it shows that the MET coefficient remains stable across alternative definitions of the instrument-generating set, with results that are most stable when country indicators are included. The appendix also reports specifications in which all five occupational stressors are treated as endogenous.

4.3 Additional evidence and robustness

We provide additional evidence on the scope of the baseline results and their robustness. We first extend the analysis to a broader set of health outcomes and disease-specific measures, and then assess the sensitivity of the findings to measurement concerns and alternative specifications.

4.3.1 Other health outcomes

We examine whether the relationship between occupational physical intensity and health holds across alternative outcome measures. Our baseline specification uses self-perceived poor health, one of the most widely used measures in the literature. However, occupational physical intensity may affect other dimensions of health as well, and not necessarily in the same way. Table 10 therefore reports estimates for a range of additional health outcomes, including ADL, IADL, GALI,⁹ mobility limitations, objectively measured maximum grip strength, chronic conditions, BMI, disability injury, and work-limiting health problems. The estimates are statistically significant for

⁹ GALI is a dummy variable indicating any limitation in daily activities.

Table 9 Lewbel IV estimates: MET only

	Self-perceived poor health (1)
Occupational physical intensity (MET)	0.087*** (0.026)
Observations	55,765
Number of endogenous variables	1
Number of excluded instruments	35
BP LM (MET)	2116.000
BP p-value (MET)	0.000
KP rk LM	98.890
KP rk LM p-value	0.000
KP rk Wald F	10,431.950
AR F	6.090
AR p-value	0.000
Hansen p-value	0.133

Notes: This reports Lewbel IV estimates of the association between lifetime occupational physical intensity and self-perceived poor health, treating only MET as endogenous. The specification corresponds to the benchmark model including socio-demographic characteristics, childhood circumstances, and job-related controls. Coefficient standard errors, reported in parentheses, are two-way clustered at the 4-digit ISCO and country level. The reported IV diagnostic tests are instead computed using one-way clustering at the 4-digit ISCO level, because the corresponding test statistics are not straightforward to implement under two-way clustering in our setting. BP denotes the Breusch–Pagan test from the auxiliary heteroskedasticity regression used to generate the Lewbel instruments. KP denotes the Kleibergen–Paap underidentification and weak-instrument diagnostics, AR denotes the Anderson–Rubin weak-instrument-robust test, and Hansen reports the overidentification test. The benchmark specification generates 35 excluded instruments using age, gender, birth cohort, country indicators, childhood housing conditions, childhood socioeconomic indicators, and the childhood health index. SHARE sample weights are applied. * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$

nearly all outcomes and display a pattern that is consistent with the baseline findings. This is also true for grip strength, which is objectively measured and therefore not subject to reporting bias in the way self-assessed health may be.

The estimates also reveal heterogeneity by gender, with effects generally larger for women across most outcomes. One possible explanation is intrinsic biological differences in physical capacity, such as average differences in muscle mass and strength, which may make physically demanding work more burdensome for women (Holtermann et al. 2010). In this sense, a given level of occupational intensity may translate into a higher physiological cost for female workers. An alternative, though not mutually exclusive, explanation is that women and men may sort into physically demanding jobs with different task content. Even within occupations characterised by similar average physical intensity, the specific activities performed, working conditions, or combinations of tasks may differ by gender, potentially leading to different health consequences. At the same time, for outcomes based on self-reports, these differences should be interpreted with some caution, as part of the observed gender heterogeneity may also reflect differences in reporting behaviour.

Table 10 Occupational physical intensity and alternative health outcomes

	sphus (1)	adl (2)	iadl (3)	gali (4)	mobility (5)	maxgrip (6)	chronic (7)	bmi (8)	Disability injury (9)	Health limits work (10)
All										
MET	0.073*** (0.023)	0.021 (0.014)	0.060*** (0.015)	0.040*** (0.007)	0.200*** (0.031)	-0.321*** (0.093)	0.084*** (0.017)	0.282*** (0.081)	0.019*** (0.006)	0.044*** (0.008)
Observations	55,765	55,765	55,765	55,765	55,765	51,886	55,765	55,765	55,758	47,400
R-squared	0.175	0.098	0.202	0.127	0.220	0.652	0.180	0.068	0.045	0.108
Women										
MET	0.118*** (0.021)	0.027 (0.020)	0.073** (0.027)	0.054*** (0.010)	0.257*** (0.053)	-0.331** (0.120)	0.095*** (0.031)	0.386*** (0.134)	0.008 (0.010)	0.041*** (0.007)
Observations	30,316	30,316	30,316	30,316	30,316	27,996	30,316	30,316	30,313	25,756
R-squared	0.200	0.138	0.249	0.145	0.241	0.293	0.213	0.089	0.049	0.112
Men										
MET	0.030 (0.033)	0.020* (0.010)	0.049*** (0.016)	0.025*** (0.009)	0.166*** (0.037)	-0.313** (0.128)	0.079*** (0.017)	0.237*** (0.068)	0.020*** (0.006)	0.043*** (0.013)
Observations	25,449	25,449	25,449	25,449	25,449	23,890	25,449	25,449	25,445	21,644
R-squared	0.174	0.079	0.160	0.127	0.178	0.337	0.159	0.067	0.064	0.137

Notes: This reports OLS estimates of the association between lifetime occupational physical intensity, measured in MET values, and alternative health outcomes. Each column corresponds to a different outcome, and the specification matches column 4 of Table 3, including socio-demographic characteristics, childhood circumstances, and job-related controls. Panels report estimates for the full sample, women, and men. Standard errors, reported in parentheses, are two-way clustered by 4-digit ISCO occupation and country. SHARE sample weights are applied. * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$

4.3.2 Diseases

We next examine whether occupational physical intensity is associated with specific disease profiles over the life course. Using retrospective information on lifetime diagnoses, we analyse both (i) the probability of ever being diagnosed with a given condition and (ii) the timing of disease onset using duration models.

Table 11 reports the results for disease incidence based on linear probability models. In the full sample, a one-unit increase in MET raises the probability of heart attack and stroke by approximately 1 percentage point, and is also positively associated with lung disease. For other conditions, including hypertension, diabetes, cancer, and neurodegenerative diseases, the estimated effects are small and statistically indistinguishable from zero.

To complement the incidence analysis, we estimate Cox proportional hazard models of age at first diagnosis (Table 12). While the linear probability models capture whether an individual has ever been diagnosed with a given condition by the time of the interview, the Cox model exploits information on the timing of first diagnosis and relates occupational physical intensity to the hazard of disease onset over the life cycle. This approach also accounts for right-censoring, since individuals who have not yet experienced the disease by the interview date still contribute information up to their observed age.

The estimates are reported as hazard ratios: values above one indicate that higher MET exposure is associated with a higher rate of first diagnosis at a given age, conditional on not having been diagnosed previously. For example, a hazard ratio of 1.10 means that, among individuals who have not yet developed the condition, those with one additional unit of MET face a 10% higher rate of first diagnosis at each age. By incorporating the timing of disease onset, this approach allows us to distinguish between earlier and later health deterioration, providing a more informative measure of the relationship between occupational exposure and health over the life cycle.

The Cox estimates broadly confirm and strengthen the patterns observed in the incidence analysis. Higher occupational physical intensity is associated with a significantly higher hazard of several conditions, particularly cardiovascular, metabolic, and respiratory diseases. In the full sample, a one-unit increase in MET raises the hazard of heart attack, hypertension, stroke, diabetes, and lung disease, and is also associated with higher hazards of stomach disease, emotional disorders, and rheumatoid arthritis. These magnitudes imply economically meaningful increases in the likelihood of earlier disease onset.

The results also point to gender differences in the disease profile associated with occupational physical intensity. Among women, higher MET exposure is associated with significantly higher hazards of heart attack, hypertension, diabetes, and lung disease, with additional positive associations for Alzheimer's disease and rheumatoid arthritis. Among men, significant associations emerge for heart attack, hypertension, diabetes, lung disease, and stomach disease, as well as for other fractures and osteoarthritis. Overall, these findings suggest that physically demanding work may be associated with different patterns of earlier disease onset for men and women.

A further caveat is that the disease variables are based on self-reported doctor diagnoses. As a result, the estimated associations may reflect not only differences in

Table 11 Occupational physical intensity and ever diagnosis of specific diseases

<i>Panel A</i>										
	Heart attack	Hypertension	Cholesterol	Stroke	Diabetes	Lung	Cancer	Stomach	Parkinson	
MET – All	0.013*** (0.004)	0.013 (0.008)	–0.001 (0.006)	0.008*** (0.002)	0.007 (0.004)	0.010*** (0.003)	0.002 (0.002)	0.004* (0.002)	0.000 (0.001)	
MET – Women	0.017*** (0.008)	0.024** (0.011)	–0.003 (0.009)	0.007** (0.003)	0.027*** (0.007)	0.007 (0.005)	0.002 (0.006)	0.003 (0.005)	0.000 (0.002)	
MET – Men	0.010** (0.005)	0.007 (0.008)	0.004 (0.006)	0.008** (0.004)	–0.004 (0.008)	0.011* (0.005)	0.000 (0.002)	0.004 (0.004)	0.000 (0.001)	
Observations (All)	55,591	55,591	55,591	55,591	55,591	55,591	55,591	55,591	55,591	
Observations (Women)	30,236	30,236	30,236	30,236	30,236	30,236	30,236	30,236	30,236	
Observations (Men)	25,355	25,355	25,355	25,355	25,355	25,355	25,355	25,355	25,355	
<i>Panel B</i>										
	Cataracts	Hip fracture	Other fractures	Alzheimer	Emotional disorders	Rheumatoid arthritis	Osteoarthritis	Kidney		
MET – All	–0.002 (0.003)	–0.002 (0.002)	0.005 (0.003)	–0.001 (0.001)	0.007 (0.008)	–0.002 (0.005)	0.013 (0.008)	0.001 (0.002)		
MET – Women	–0.003 (0.005)	–0.004 (0.003)	0.003 (0.003)	–0.001 (0.002)	–0.001 (0.006)	0.002 (0.011)	0.017 (0.011)	0.004 (0.005)		
MET – Men	–0.005 (0.004)	–0.001 (0.002)	0.009* (0.004)	0.000 (0.001)	0.014 (0.010)	0.000 (0.005)	0.010 (0.008)	–0.002 (0.003)		
Observations (All)	55,591	55,591	54,446	54,446	37,135	37,135	37,135	37,135		
Observations (Women)	30,236	30,236	29,640	29,640	20,368	20,368	20,368	20,368		
Observations (Men)	25,355	25,355	24,806	24,806	16,767	16,767	16,767	16,767		

Notes: This reports linear probability model estimates of the association between lifetime occupational physical intensity, measured in MET values, and the probability of ever being diagnosed with each disease. Entries report the coefficient on MET. The specification includes socio-demographic characteristics, childhood circumstances, and job-related controls. Panel A reports cardiovascular, metabolic, respiratory, and other major conditions, while Panel B reports additional chronic conditions. Rows labelled “All,” “Women,” and “Men” report estimates separately for the full sample and by gender. Standard errors, reported in parentheses, are two-way clustered by 4-digit ISCO occupation and country. SHARE sample weights are applied. * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$

Table 12 Occupational physical intensity and the hazard of first diagnosis

<i>Panel A</i>		Heart attack	Hypertension	Cholesterol	Stroke	Diabetes	Lung	Cancer	Stomach	Parkinson
MET – all		1.130*** (0.031)	1.056*** (0.018)	1.000 (0.019)	1.107** (0.056)	1.098*** (0.024)	1.232*** (0.044)	1.011 (0.039)	1.090** (0.045)	0.881 (0.070)
MET – women		1.217*** (0.054)	1.090*** (0.031)	0.998 (0.032)	1.131 (0.102)	1.200*** (0.044)	1.195*** (0.064)	1.039 (0.062)	0.987 (0.065)	0.940 (0.169)
MET – men		1.106*** (0.034)	1.061*** (0.018)	1.014 (0.023)	1.103 (0.067)	1.070*** (0.026)	1.228*** (0.058)	0.986 (0.052)	1.161** (0.073)	0.888 (0.099)
Diagnosed cases (all)		6165	20,784	12,228	1986	6393	2804	2510	2152	390
Diagnosed cases (women)		2675	11,364	6773	876	3103	1408	1366	1172	175
Diagnosed cases (men)		3490	9420	5455	1110	3290	1396	1144	980	215
Observations (all)		54,367	51,410	53,455	55,199	54,412	55,147	55,122	55,132	55,531
Observations (women)		29,652	27,885	29,027	30,043	29,610	29,998	29,970	29,987	30,203
Observations (men)		24,715	23,525	24,428	25,156	24,802	25,149	25,152	25,145	25,328

Table 12 continued

<i>Panel B</i>		Cataracts	Hip fracture	Other fractures	Alzheimer	Emotional	Rheumatoid	Osteoarthritis	Kidney
MET – all		0.998	1.032	1.057	1.103	1.122**	1.167***	1.049	0.930
		(0.039)	(0.063)	(0.045)	(0.074)	(0.050)	(0.048)	(0.035)	(0.077)
		1.053	0.904	1.026	1.285**	–	1.246***	1.073	1.128
MET – women		(0.050)	(0.077)	(0.058)	(0.137)	–	(0.070)	(0.049)	(0.143)
		0.923	1.104	1.111*	–	–	1.064	1.088*	0.712***
		(0.056)	(0.099)	(0.065)	–	–	(0.064)	(0.048)	(0.085)
Diagnosed cases (all)		4013	872	2214	692	1983	2661	4996	421
		2481	516	1248	361	1373	1772	3305	217
		1532	356	966	331	610	889	1691	204
Observations (all)		54,950	55,441	53,509	54,255	36,385	35,779	34,594	36,665
		29,849	30,148	29,122	29,535	19,834	19,437	18,573	20,103
		25,101	25,293	24,387	24,720	16,551	16,342	16,021	16,562

Notes: This reports hazard ratios from Cox proportional hazard models relating lifetime occupational physical intensity, measured in MET values, to the hazard of first diagnosis for each condition. Diagnosed cases denote the number of individuals ever diagnosed with the condition, while non-diagnosed individuals are treated as right-censored at interview age. The specification includes socio-demographic characteristics, childhood circumstances, and job-related controls. Standard errors, reported in parentheses, are clustered at the 4-digit ISCO level. Estimates that do not converge are omitted. * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$

underlying disease incidence, but also differences in the probability and timing of diagnosis. In particular, access to healthcare, health-seeking behaviour, and diagnosis rates may vary across individuals and socioeconomic groups. To the extent that these factors are correlated with occupational histories, they may affect the measured timing of first diagnosis. For this reason, the disease-specific results should be interpreted as associations between occupational exposure and self-reported diagnosed disease incidence or onset, rather than as direct estimates of the biological onset of disease.

4.3.3 Cohort heterogeneity

The occupational stressor indices are constructed using externally harmonised occupation–task mappings benchmarked to early-2000s conditions and applied to retrospective career histories spanning several decades. A natural concern is that these benchmarks may not fully capture occupational exposure for older cohorts whose working lives predate the measurement window. This concern is particularly relevant because the content of occupations may have evolved over time as mechanisation, automation, and technological progress altered the physical demands of many jobs. If earlier versions of the same occupations were more physically demanding, applying early-2000s benchmarks to earlier job spells would tend to understate exposure and attenuate the estimated relationship. To assess this possibility, we examine heterogeneity by birth cohort using alternative thresholds. The estimated effect of MET remains positive across all cohort splits and tends to be larger for younger cohorts, whose working lives overlap more closely with the benchmark period. Under the median split (1948), the coefficient is 0.055 for older cohorts and 0.085 for younger cohorts. Under the 75th percentile split (1955), it rises from 0.065 to 0.097. Under the 90th percentile split (1960), the difference becomes more pronounced, increasing from 0.063 for older cohorts to 0.142 for younger cohorts. Overall, this pattern is consistent with greater attenuation among older cohorts rather than an overstatement of the baseline estimates. Detailed results are reported in Appendix Table A.6.

4.3.4 Flexible estimation: Causal Forest

We re-estimate the main specification using a Causal Forest within the Generalised Random Forest framework (Athey et al. 2019). The identifying assumption remains the same as in the baseline regressions—conditional independence given observed covariates—and therefore this approach does not address selection on unobservables. Instead, it relaxes the linearity and constant-effect assumptions imposed by OLS by allowing for flexible non-linearities and interactions in the relationship between occupational stressors and health outcomes. In addition, the tree-based algorithm can accommodate missing values in covariates during the splitting procedure, allowing the estimation to retain observations that would otherwise be excluded under listwise deletion. As a result, the Causal Forest is estimated on a sample that is approximately 15% larger than in the baseline OLS specification.

The forest is estimated using the same covariate set as in the baseline specification, with hyperparameters selected via the built-in cross-validation procedure of the GRF

algorithm. Appendix A.11.1 reports additional details on the estimation procedure, tuning parameters, and diagnostic evidence.

Table A.14 reports the average treatment effects implied by the forest. The results are qualitatively consistent with the corresponding OLS estimates. While point estimates are somewhat larger in magnitude, the direction and statistical significance of the effects remain unchanged. Overall, the findings are robust to allowing for a more flexible relationship between occupational exposure and health.

5 Task-level evidence: LASSO and Job Strain Intensity

In this section, we introduce a novel approach based on the LASSO estimator (Tibshirani 1996) and granular task information at the 4-digit ISCO level. We first use task-level information to predict the tasks that are most strongly associated with poor later-life health. We then aggregate this information into a new measure of *Job Strain Intensity* (JSI), which captures harmful occupational conditions at the 4-digit ISCO level, including dimensions of work that are difficult to observe directly. This approach serves two related purposes: it provides a more granular characterisation of the link between job tasks and health, and it allows us to test whether the estimated effect of physical job demands remains robust after controlling for a broader, data-driven measure of harmful working conditions.

5.1 Granular job-specific tasks and LASSO estimator

To analyse the task content of occupations, we use the harmonisation proposed by Deyaert et al. (2017), which groups the job-specific tasks listed by ILO (2012) for each 4-digit ISCO code into homogeneous task categories.¹⁰ This classification gives us 302 harmonised task categories that can be compared across occupations and that form the basis of the analysis. Section A.12.1 in the Appendix reports the full list.

To identify which job-specific tasks performed in the main job are most strongly associated with poor self-perceived health in old age, we estimate a LASSO model (Tibshirani 1996). To account for selection into occupations, we first residualise both the outcome and the task indicators with respect to childhood characteristics, and then apply the LASSO to the residualised variables:

$$\min_{\gamma} \left(\frac{1}{N} \sum_{i=1}^N (\tilde{y}_i - \beta_0 - \sum_{k=1}^q \tilde{z}_{ik} \gamma_k)^2 + \lambda \sum_{k=1}^q |\gamma_k| \right) \quad (6)$$

where i indexes individuals, k indexes tasks, and q denotes the total number of harmonised tasks. The variable $\tilde{y}_i$ represents the residualised health outcome for

¹⁰ For instance, using the ILO (2012) job-specific tasks, Deyaert et al. (2017) assign tasks such as “serving on government administrative boards or official committees” (1111 – Legislators), “participating in departmental and faculty meetings and in conferences and seminars” (2310 – University and higher education teacher), and “advising clients on legal matters” (3411 – Legal and related associate professionals) to the harmonised task category “Meetings.”

individual i , obtained by partialling out the childhood characteristics described in Section 2.¹¹ Similarly, $\tilde{z}_{ik}$ denotes the residualised task indicators for the main job of individual i . The parameter λ is the regularisation (shrinkage) parameter.

Intuitively, the aim is to assess which tasks performed in the main job are most predictive of poor self-perceived health in old age. Residualising the outcome and task indicators with respect to childhood characteristics helps mitigate selection concerns, although the resulting estimates should still be interpreted as associations. The LASSO approach is well suited to this setting because the task matrix is high-dimensional and sparse. Since our focus is on tasks associated with worse health, we impose a non-negativity constraint on the coefficients γ_k . The penalty parameter λ is chosen through 10-fold cross-validation. The resulting estimates retain 67 tasks as positively associated with poor self-perceived health (see Table 13).

The tasks reported in Table 13 span a wide range of activities, highlighting substantial heterogeneity in the types of job content associated with worse health outcomes. They include, for example, relatively white-collar activities such as advising and client interaction, alongside clearly manual tasks such as construction, electrical work, and machinery operation. This heterogeneity is important: even within occupations that appear similar at an aggregate level, the underlying tasks can differ markedly in terms of physical intensity, exposure conditions, and organisational demands.

Standard indicators of occupational stressors are unlikely to fully capture this variation, as they typically summarise complex job content into a limited set of predefined dimensions. In addition, such measures often rely on external codings that may overlook relevant aspects of working conditions or introduce measurement error, particularly when based on stylised descriptions of occupations rather than detailed task information. By contrast, the data-driven approach adopted here allows the analysis to directly leverage the richness of task-level information and to uncover the dimensions of job content that are most strongly associated with later-life health outcomes.

5.2 Job Strain Intensity

We next construct a data-driven indicator of occupational strain, which we label *Job Strain Intensity* (JSI). The index is based on the harmful tasks selected by the LASSO procedure and is defined, for each 4-digit ISCO occupation, as the share of harmful tasks among all harmonised tasks associated with that occupation.¹² Higher values of JSI therefore indicate that an occupation is more heavily concentrated in task dimensions associated with poor later-life health.

To assess the validity of the measure, Fig. 3 reports average JSI by 1-digit ISCO group. The pattern is clear and consistent with economic intuition: blue-collar occupations exhibit substantially higher JSI values than white-collar occupations.

¹¹ $\tilde{y}_i = y_i - \hat{y}_i(X_{child})$, where $\hat{y}_i(X_{child})$ is the predicted value of y_i based on childhood characteristics.

¹² This construction implicitly assumes that all tasks within a 4-digit ISCO occupation are performed for equal amounts of time.

Table 13 Tasks selected by the LASSO procedure

Adjusting maintaining equipment	Leveling floor	Supervising operational programmes
Advising (financial, legal)	Machining, sheet metal, welding	Sweeping, vacuum-cleaning, polishing..
Advising on selling goods..	Maintenance repairs	Tailoring by hand
Arts and crafts	Making bed, cleaning bathroom	Tailoring light
Assembling/modifying/installing..	Masonry concrete	Training animals
Carpentry light	Monitoring equipment plant	Truck driving
Carpet cleaning	Operating/coordinating control room	Unloading truck
Chemical treatment	Operating auxiliary plant	Vehicle maintenance/repair
Cleaning kitchen, bathroom, toilet	Operating heavy duty equipment	Walking, carrying heavy object
Cleaning stone walls/heavy outdoor	Painting varnishing coating	Washing by hand, moderate
Collecting mineral samples (mine)	Patrolling	Weighing goods and issues
Construction and restoration	Physical care (bathing, dressing,...)	Farming
Coordination at workplace and site	Preparing medicine/dosage	Fishing
Delivering postal mail	Preparing wall areas	Repair maintenance
Designing/preparing blueprints	Preparing, cooking, serving meals	Raking lawn or leaves
Digging ditches	Preparing/maintaining medical aids	Receiving and verifying payments
Driving car/van	Purchasing ordering (electronic)	Researching/verifying client status
Electrical work	Routine lab tests	Scanning/recording/wrapping goods
Establishing contact stakeholders	Security check maintaining order	Selecting/modifying products
Feeding felled trees into processor	Sending/receiving information	Stacking and loading logs
Gardening general	Standing arts and crafts	Stimulating discussion
Grading/mixing produce	Stock control	Stone quarrying
Heavy cleaning		

Notes: This reports the 67 harmonised task categories retained by the constrained LASSO procedure as positively associated with poor self-perceived health. The LASSO is estimated on residualised task indicators and a residualised health outcome, after partialling out childhood characteristics. The penalty parameter is chosen by tenfold cross-validation




Fig. 3 Average Job Strain Intensity by 1-digit ISCO occupation. *Notes:* This reports average job strain intensity by 1-digit ISCO occupation in the analysis sample. The horizontal axis lists the nine broad ISCO occupation groups, and the vertical axis reports the corresponding average value of the job strain intensity measure

5.3 Estimation results controlling for Job Strain Intensity

We next assess whether the estimated effect of physical job demands (MET) remains once broader dimensions of occupational strain are taken into account. To this end, we augment the baseline specification by including the Job Strain Intensity (JSI) index, which summarises task-level characteristics associated with poor health outcomes in a data-driven way.

A potential concern is that constructing JSI using the same sample and outcome variable employed in the regression may induce mechanical correlation or overfitting. To address this issue, we implement an out-of-sample procedure. Specifically, we randomly split the sample into two halves: the LASSO procedure is applied in the first half to identify harmful tasks and construct the JSI index at the 4-digit ISCO level, while the regression analysis is conducted exclusively in the second half. This ensures that the variation in JSI used in the estimation is independent of the outcome data.

Table 14 reports the results. In column 1, without JSI, the coefficient on MET is 0.087. When JSI is included (column 2), the MET coefficient declines slightly to 0.080, while JSI enters positively and is statistically significant. The attenuation of the MET coefficient is modest, and the effect remains both economically and statistically meaningful. This indicates that the relationship between physical job demands and later-life health is not driven by other dimensions of occupational strain captured by JSI.

Because the construction of JSI depends on the random train–test split, we further assess the stability of the results by repeating the procedure across 1,000 alternative

Table 14 Out-of-sample results including Job Strain Intensity

	sphus (No JSI) (1)	sphus (with JSI) (2)
Job Strain Intensity		0.097** (0.046)
MET	0.087*** (0.025)	0.080*** (0.025)
Psychosocial stress	−0.002 (0.005)	−0.002 (0.005)
Carcinogenic exposure	0.013 (0.008)	0.010 (0.008)
Routine manual	0.062 (0.111)	0.037 (0.113)
Routine cognitive	0.085 (0.056)	0.100* (0.055)
Constant	0.935** (0.367)	0.988*** (0.348)
Socio-demo	Yes	Yes
Child	Yes	Yes
Job	Yes	Yes
Habits	No	No
Observations	27,623	27,623
R-squared	0.183	0.184

Notes: This reports OLS estimates of the association between occupational stressors and self-perceived poor health in the out-of-sample analysis. Column 1 excludes job strain intensity, while column 2 additionally includes it. Both specifications include socio-demographic characteristics, childhood circumstances, and job-related controls. Standard errors, reported in parentheses, are two-way clustered by 4-digit ISCO occupation and country. SHARE sample weights are applied. * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$

random partitions of the data. The results, reported in Appendix A.12.2, show that the estimated coefficients remain stable across splits.

5.4 Tasks and jobs most predictive of poor health

A task-level perspective allows us to move beyond aggregate occupational measures and isolate the specific activities that are most predictive of poor later-life health. Rather than relying on the composite JSI index, we therefore return to the disaggregated task indicators and use the LASSO framework to discuss a parsimonious set of tasks with the strongest predictive content.

To obtain an interpretable ranking, we depart from the standard cross-validation criterion and instead select the penalty parameter λ so as to retain five tasks. This yields a concise summary of the task dimensions most strongly associated with poor self-perceived health. The LASSO procedure selects the following five tasks: *cleaning kitchen, bathroom, and toilet; identifying, reporting, and documenting hazards; operating heavy-duty equipment; farming; and repair and maintenance.*

The stability of this selection is supported by the LASSO solution path reported in Appendix Table A.19. Tasks enter the model in a monotonic way as the penalty is relaxed and remain selected thereafter, indicating that the ranking is not driven by a knife-edge choice of λ . In particular, *repair and maintenance* enters first, followed by *farming*, then *cleaning kitchen, bathroom, and toilet*, and, as the fourth selected task, *identifying, reporting, and documenting hazards; operating heavy-duty equipment* enters only in the specification retaining five tasks.

These tasks map naturally into economically meaningful dimensions of job strain. They capture, respectively, repetitive physical effort, exposure to risk and safety environments, high-intensity mechanical work, agricultural labour, and maintenance activities involving sustained physical and technical demands. Importantly, they span both traditionally blue-collar activities and tasks present in broader occupational settings, highlighting that harmful job content is not confined to a single occupational category.

More generally, this exercise illustrates that JSI captures dimensions of job strain that are only imperfectly proxied by standard indicators. Measures of routineness, such as Mihaylov and Tijdens (2019), primarily reflect automatable production tasks and may miss other forms of strain arising in non-automatable but physically or ergonomically demanding activities. For example, cleaning tasks involve substantial repetitive motion and postural stress despite not being classified as routine manual work, while operating heavy machinery may entail peak exertion and exposure risks not fully captured by average energy expenditure measures such as MET. In this sense, JSI can be interpreted as a broader proxy for occupational strain, encompassing ergonomic, environmental, and organisational dimensions of work.

Finally, Appendix Tables A.20–A.23 show the occupations associated with these five selected tasks, as well as the occupations in the top decile of the JSI distribution. The substantial overlap between the two indicates that the task-level analysis and the occupation-level JSI measure identify a similar set of high-strain jobs.

6 Discussion and conclusion

Using SHARE data from European countries, this paper examines how lifetime exposure to occupational stressors is associated with health later in life. The results show that physical job intensity stands out as an occupational stressor that remains consistently associated with poorer health once a rich set of childhood circumstances, socio-demographic characteristics, job-related factors, and health-related behaviours is taken into account. By reconstructing individuals' full occupational histories, the paper relates later-life health to exposure over the entire working trajectory rather than to a limited phase of the career. The results on disease incidence and timing point in the same direction, indicating that physically demanding jobs are associated with earlier health deterioration, particularly through cardiovascular, metabolic, and respiratory conditions. While the setting does not allow a fully causal interpretation, the results are supported by a wide range of complementary checks, including rich controls, analyses of selection and reverse causality, Lewbel IV estimates, alternative

health outcomes, disease-specific evidence, cohort heterogeneity, and flexible Causal Forest estimation.

The paper also shows how the analysis can move beyond broad occupational stressors and examine the specific task content of jobs. The task-based analysis points to specific tasks that are more strongly associated with poor health and uses this information to construct a broader measure of Job Strain Intensity. For instance, the five tasks most strongly predictive of worse health in the data-driven LASSO analysis are *cleaning kitchen, bathroom, and toilet, identifying, reporting, and documenting hazards, operating heavy-duty equipment, farming, and repair and maintenance*. This complements the main findings by showing that the relationship between work and health is shaped not only by broad occupational categories, but also by the concrete content of jobs.

More broadly, the findings suggest that working conditions may help explain how labour-market disparities are related to inequalities in health. Differences in jobs shape not only wages and careers, but also the distribution of health later in life. In this respect, the more granular approach developed here may be useful not only for predicting which dimensions of work are most harmful, but also for informing a range of policy discussions, including targeted prevention, workplace safety, job redesign, and the design of social insurance and retirement rules. In particular, if workers in more physically demanding occupations face worse health and shorter healthy working lives, uniform institutional arrangements may impose unequal burdens across jobs, in line with recent concerns about actuarial fairness and occupational differences in life expectancy (Deeg et al. 2021; Marcus et al. 2022).

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s00148-026-01207-6>.

Acknowledgements We thank the Editor, Klaus F. Zimmermann, and two anonymous referees for their constructive comments and guidance, which substantially improved the paper. We also thank Beatrix Eugster, Lorenzo Cappellari, and Agar Brugiavini for helpful comments and discussions.

Funding Open access funding provided by University of Innsbruck and Medical University of Innsbruck. This study was funded by the European Union – Next Generation EU, Mission 4 Component 2, within the NRRP (National Recovery and Resilience Plan) project “Age-It – Ageing well in an ageing society” (PE0000015 – CUP H73C22000900006). Petru Crudu gratefully acknowledges financial support from the Austrian Science Fund (FWF), project grant P36845-G (DOI: 10.55776/P36845). The views and opinions expressed are only those of the authors and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

Data availability The SHARE data used in this study are available to registered researchers through the Survey of Health, Ageing and Retirement in Europe (SHARE), subject to SHARE’s terms of use. The occupation-level measures used to construct occupational stressors are obtained from the external sources cited in the paper, including Deyaert et al. (2017), Kroll (2015), and Mihaylov and Tijdens (2019). Replication code is available upon request.

Declarations

Conflict of interest The authors declare no competing interests.

Declaration of generative AI in scientific writing During the preparation of this work, the authors used ChatGPT to check and improve the fluency and clarity of the writing. After using this tool, the authors reviewed and edited the content as needed and took full responsibility for the content of the published article.

Competing interests The authors declare that they have no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ainsworth BE, Haskell WL, Herrmann SD, Meckes N, Bassett DR Jr, Tudor-Locke C, Greer JL, Vezina J, Whitt-Glover MC, Leon AS (2011) 2011 Compendium of Physical Activities: a second update of codes and MET values. *Med Sci Sports Exerc* 43(8):1575–1581
- Altonji JG, Elder TE, Taber CR (2005) An evaluation of instrumental variable strategies for estimating the effects of Catholic schooling. *J Human Res* 40(4):791–821
- Athey S, Tibshirani J, Wager S (2019) Generalized random forests. *Ann Stat* 47(2):1148–1178
- Börsch-Supan A, Brandt M, Hunkler C, Kneip T, Korbmacher J, Malter F, Schaan B, Stuck S, Zuber S (2013) Data resource profile: the Survey of Health, Ageing and Retirement in Europe (SHARE). *Int J Epidemiol* 42(4):992–1001
- Brugiavini A, Orso CE, Genie MG, Naci R, Pasini G (2019) Combining the retrospective interviews of Wave 3 and Wave 7: the third release of the SHARE Job Episodes Panel. Munich Center for the Economics of Aging (MEA), Munich, p 36
- Brunello G, Weber G, Weiss CT (2017) Books are forever: early life conditions, education and lifetime earnings in Europe. *Econ J* 127(600):271–296
- Case A, Deaton AS (2005) Broken down by work and sex: how our health declines. In: *Analyses in the economics of aging*. University of Chicago Press, pp 185–212
- Case A, Fertig A, Paxson C (2005) The lasting impact of childhood health and circumstance. *J Health Econ* 24(2):365–389
- Case A, Lubotsky D, Paxson C (2002) Economic status and health in childhood: the origins of the gradient. *Am Econ Rev* 92(5):1308–1334
- Cillekens B, Huysmans MA, Holtermann A, van Mechelen W, Straker L, Krause N, van der Beek AJ, Coenen P (2022) Physical activity at work may not be health enhancing. A systematic review with meta-analysis on the association between occupational physical activity and cardiovascular disease mortality covering 23 studies with 655 892 participants. *Scandinav J Work Environ Health* 48(2):86
- Currie J (2009) Healthy, wealthy, and wise: socioeconomic status, poor health in childhood, and human capital development. *J Econ Literat* 47(1):87–122
- Deeg DJ, De Tavernier W, de Breijl S (2021) Occupation-based life expectancy: actuarial fairness in determining statutory retirement age. *Front Sociol* 6:675618
- Deyaert J, Harms T, Weenas D, Gershuny J, Glorieux I (2017) Attaching metabolic expenditures to standard occupational classification systems: perspectives from time-use research. *BMC Public Health* 17(1):1–10
- Eurostat (2021) Duration of hours worked per week
- Eurostat (2022) Duration of working life
- Fletcher JM, Sindelar JL, Yamaguchi S (2011) Cumulative effects of job characteristics on health. *Health Econ* 20(5):553–570
- Holtermann A, Mortensen OS, Burr H, Sogaard K, Gyntelberg F, Suadicani P (2010) Physical demands at work, physical fitness, and 30-year ischaemic heart disease and all-cause mortality in the Copenhagen Male Study. *Scandinav J Work Environ Health* 357–365

- ILO (2012) International Standard Classification of Occupations. ISCO-08. Structure, group definitions and correspondence tables. Geneva, Switzerland. <http://www.ilo.org/public/english/bureau/stat/isco/docs/publication08.pdf>
- Krause N, Brand RJ, Arah OA, Kauhanen J (2015) Occupational physical activity and 20-year incidence of acute myocardial infarction: results from the Kuopio Ischemic Heart Disease Risk Factor Study. *Scandinavian J Work Environ Health* 124–139
- Kroll LE (2015) Job exposure matrices (JEM) for ISCO and KIdB (version 2.0). Data file version 1.0.0. <https://doi.org/10.7802/1102>
- Lewbel A (2012) Using heteroscedasticity to identify and estimate mismeasured and endogenous regressor models. *J Bus Econ Stat* 30(1):67–80
- Mahler L, Yum M (2024) Lifestyle behaviors and wealth-health gaps in Germany. *Econometrica* 92(5):1697–1733
- Marcus E, Anders A, Per G, Karin M (2022) Occupational differences in mortality and life expectancy persist after retirement and throughout life. *Scandinavian J Public Health* 14034948221081628
- McEwen BS (2000) Allostasis and allostatic load: implications for neuropsychopharmacology. *Neuropsychopharmacology* 22(2):108–124
- Mihaylov E, Tijdens KG (2019) Measuring the routine and non-routine task content of 427 four-digit ISCO-08 occupations
- Oster E (2019) Unobservable selection and coefficient stability: theory and evidence. *J Bus Econ Stat* 37(2):187–204
- Ravesteijn B, Kippersluis Hv, Doorslaer Ev (2018) The wear and tear on health: what is the role of occupation? *Health Econ* 27(2):e69–e86
- Samitz G, Egger M, Zwahlen M (2011) Domains of physical activity and all-cause mortality: systematic review and dose-response meta-analysis of cohort studies. *Int J Epidemiol* 40(5):1382–1400
- Tibshirani R (1996) Regression shrinkage and selection via the lasso. *J R Stat Soc Ser B Stat Methodol* 58(1):267–288

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.