The Journal of the Economics of Ageing 32 (2025) 100590

Contents lists available at ScienceDirect

The Journal of the Economics of Ageing

journal homepage: www.elsevier.com/locate/jeoa

The macroeconomic impact of chronic disease in the United Kingdom ☆

Yannick Schindler [a,b,*], Andrew J. Scott [c,b,d]

[a] London School of Economics, United Kingdom
[b] London Business School, United Kingdom
[c] Ellison Institute of Technology Oxford, United Kingdom
[d] CEPR, United Kingdom

ARTICLE INFO

JEL classification:
I15
J11
J21
O40

Keywords:
Aging population
Chronic disease
Economic growth
Health
Labor market transitions

ABSTRACT

This paper examines the macroeconomic impact of chronic disease in the United Kingdom (UK). We use individual-level data to estimate how diagnoses of six major diseases affect labor market transitions and combine these with a tractable growth model with age-specific productivity and labor force participation to quantify the impact of chronic disease on UK economic growth. Using a novel machine learning approach to classify National Health Service (NHS) cost data, we also provide new estimates of disease-specific treatment costs. Our findings indicate that a 20% reduction in disease incidence would increase annual GDP by 0.99% after five years and 1.51% after ten years. Most of the gains are due to increased participation in the labor force, especially among workers aged 50 to 65 years. Reductions in mental health conditions and musculoskeletal conditions contribute the most to these effects. Our analysis points to three important features of preventative health policies: (1) the potential welfare gains are substantial and manifest themselves in terms of both improved population health and increased output growth, (2) only around 40% of long-term effects appear after five years, and (3) the 50–65 age group experiences the largest labor force participation gains. This last feature is due to two factors: improved health at those ages prevents transitions into health-related inactivity and a larger share of workers reaches this age band as a result of reduced transitions into inactivity at earlier ages. This compounding effect underscores the importance of targeting prevention efforts at earlier ages.

Introduction

A large body of work has considered the links between health and economic growth.[1] Although not uncontroversial, this literature generally points to a positive relationship (Barro, 1996; Arora, 2001; Bhargava et al., 2001; Aghion et al., 2010; Barro, 2013; Bloom et al., 2014b). It is this positive relationship that underpins economic fears of an aging society (Lee and Mason, 2017; Goodhart and Pradhan, 2020). As the population ages and the proportion of older individuals increases, a shift in the disease burden is occurring towards chronic long-lasting diseases (GBD 2019 Diseases and Injuries Collaborators, 2020). The consequence is a substantial projected increase in health costs and lower GDP as an increased incidence of disease lowers labor supply and rising health costs divert funds from capital accumulation. The substantial negative magnitude of these impacts has been demonstrated for Asian countries (Bloom et al., 2014a; Chen and Bloom, 2019; Bloom et al., 2020), for Central and South America (Bloom et al., 2018), and for the United States (US) (Chen et al., 2018).

In this paper, we consider the impact of chronic disease on the UK economy. The UK is an interesting example as in recent years it has significantly increased its health expenditure but has witnessed a large increase in health-related economic inactivity and a deterioration in health, life expectancy, and GDP trends relative to many other countries (see OBR (2024b) Chapter 3).

Given these trends, the UK faces mounting pressure to recalibrate its healthcare strategy (Ara Darzi, 2024). One particular recommendation, is to switch emphasis from *ex-post* treatment of chronic diseases towards prevention, which currently accounts for only 5% of total UK health expenditure (ONS, 2024). Such a shift represents a delicate fiscal challenge. An increase in spending on preventative health would need to come at the expense of ex post treatment of diseases given

☆ We thank Sir John Bell, Tamsin Berry, John Deanfield, Ines Hassan, Roshni Joshi, and Naomi Lee for multiple helpful conversations as well as participants at the London Business School "Health and Growth" conference in April 2024 for helpful comments. Thanks are also due to Katerina Nikalexi for excellent research assistance. We gratefully acknowledge funding from ESRC Research, United Kingdom Grant T002204.

* Corresponding author at: London School of Economics, United Kingdom.
E-mail address: y.schindler@lse.ac.uk (Y. Schindler).

[1] See Bloom et al. (2004) for a survey.

https://doi.org/10.1016/j.jeoa.2025.100590

Available online 30 June 2025

constraints on public finances. However, if preventative healthcare generates positive economic returns through improved population health and productivity, this could potentially resolve the apparent fiscal dilemma. This possibility motivates our central research question: what would be the impact on UK economic growth if targeted investments in preventative health measures successfully reduced the incidence of the six most prevalent chronic diseases?

To answer this question we pay particular attention to the link between health outcomes and labor market participation by using individual level data from the UK Labour Force Survey (LFS) and the UK Household Longitudinal Study (UKHLS). An extensive literature has shown that adverse health shocks have sizable impacts on individual transitions from employment into retirement (Bound et al., 1999; Disney et al., 2006; Hagan et al., 2009; Jones et al., 2010; Christensen and Kallestrup-Lamb, 2012) into unemployment (García-Gómez et al., 2010; Christensen and Kallestrup-Lamb, 2012), or out of the labor force (Cai and Kalb, 2006; García-Gómez, 2011). In this paper we estimate the impact of six chronic diseases on the probability of transitioning out of the labor force and from full-time to part-time employment using the estimator proposed by Dube et al. (2024), which accounts for biases that arise in staggered treatment contexts when using traditional two-way fixed effects (TWFE) estimators (de Chaisemartin and D'Haultfœuille, 2020; Goodman-Bacon, 2021; Borusyak et al., 2024).

We combine these estimates of how disease impacts labor market transitions with an augmented (Solow, 1956) growth model (similar to Chen et al., 2018) in which labor force participation and productivity vary by age. In the model, changes in disease incidence affect the rate at which workers transition from employment to inactivity—and from full-time to part-time employment—leading to changes in aggregate employment and labor productivity. We then simulate how improvements in population health, through reduced disease incidence, would affect both short-run and long-run economic growth by altering labor market transition patterns and the aggregate labor supply.

The structure of the paper is as follows: Section "UK health and labor market trends" provides a brief summary of UK health and its recent trends; Section "Model" outlines our model; Section "Calibration" discusses the data we use and our model calibration; Section "Estimating the impact of disease on labor force participation" estimates the impact of disease on labor force participation; Section "Results" analyzes our results; and the final section concludes.

UK health and labor market trends

The UK, like many countries, is experiencing population aging. The share of the population aged over 50 has risen from 36% in 1970 to 43% today, and United Nations projections indicate it will reach 51% by 2070. The proportion aged over 65 has similarly increased from 18% in 1970 to 24% today, with projections showing 32% by 2070. The most dramatic proportional change is among those aged over 80, where the share has risen from 4% in 1970 to 8% today and is projected to reach 15% by 2070 (UN, 2024).

This increase in the proportion of older individuals has implications for the disease burden. As shown in Fig. 1, the overall prevalence of disease rises strongly with age, particularly for cardiovascular disease (CVD), chronic obstructive pulmonary disease (COPD), cancer, and diabetes. As the proportion of people reaching advanced ages increases, the burden of disease rises and shifts toward chronic non-communicable diseases (NCDs). In 2019, it is estimated that there were 6.6 million people in the UK with a major illness, of whom 3.6 million were over 70. By 2040, this figure is predicted to reach 9.1 million, of whom 5.6 million will be over 70 (Toby Watt et al., 2023). Furthermore, because the disease burden increases with age, so does health expenditure. Health expenditure rises sharply with age: while annual spending per person averages £2300 at age 50, it more than doubles to £4800 by age 65 and nearly triples again to £13,400 by age 90 (OBR, 2024b). This steep age gradient in healthcare costs means that population aging places significant upward pressure on total health expenditure.

As well as this shifting age structure, the UK is also experiencing a slowdown in health and life expectancy improvements. While life expectancy has risen dramatically over the past century—from 55 years in 1922 to over 81 in 2022—the pace of improvement has recently faltered. For most of the past 70 years, life expectancy increased steadily by about 1.8 years each decade. However, in the last decade, it rose by just 0.7 years—a more pronounced slowdown than seen in most other high-income countries. This slower life expectancy growth has been accompanied by declining healthy life expectancy, which fell by over one year between 2011–13 and 2020–22. While there are some positive signs of health improvements at older ages (Old and Scott, 2023), the overall picture is concerning: a combination of demographic aging and rising disability rates among younger people (OBR, 2024b) has led to a deterioration in population health.

The impact of health on employment becomes increasingly significant with age. Data from the UK Labour Force Survey shows that the prevalence of work-limiting health conditions rises sharply across age groups: from 10% among those aged 30–35, to 19% among those aged 50–55, and 26% for those aged 60 and above (ONS, 2024c). As the population ages and more people work into their later years, this age-related rise in health conditions becomes increasingly consequential for the overall workforce. Moreover, the deterioration of health among younger age groups has compounded this problem. The combined effect has led to a substantial increase in health-related economic inactivity (Haskel and Martin, 2022): by July 2023, 2.8 million working-age people (30% of the total) were economically inactive due to ill health, up from 2.3 million (24%) a decade earlier.

Model

Key to our analysis is an augmented Solow–Swan growth model (Solow, 1956, Swan 1956) similar to that used in (Bloom et al., 2020). We assume output is produced according to a Cobb–Douglas aggregate production function:

$$Y_t = A_t K_t^{\alpha} L_t^{1-\alpha} \tag{1}$$

where Y_t is aggregate output, A_t is total factor productivity, K_t is the capital stock, L_t is the effective labor supply, and α is the capital share of output. The subscript t denotes time. Capital accumulates according to the law of motion:

$$K_{t+1} = (1-\delta)K_t + s(1-\omega)Y_t \tag{2}$$

where δ is the depreciation rate, s the savings rate, and ω the share of output spent on healthcare. The implicit assumption in (2) is that a fraction s of any reductions in healthcare spending is invested into the physical capital stock, while the remaining fraction is consumed.

A key channel through which chronic illness affects aggregate output is through the aggregate labor supply. To capture this mechanism and to reflect that labor supply varies with age, we decompose effective labor supply into :

$$L_t = \sum_{a=0}^{A} N_t(a) \times \mu_t(a) \times E_t(a) \times H_t(a) \tag{3}$$

where a indexes age, A is maximum age of agents, $N_t(a)$ is total population, $\mu_t(a)$ is worker productivity, $E_t(a)$ is the share of the labor force that is working (i.e. the labor force participation rate), and $H_t(a)$ is a variable that reflects hours worked. Hours worked is determined by the fraction of workers who are working full-time or part-time:

$$H_t(a) = f_t(a) + \left[1 - f_t(a)\right] \lambda \tag{4}$$

where $f_t(a)$ is the fraction working full-time and λ is the fraction of full-time hours provided by part-time workers.

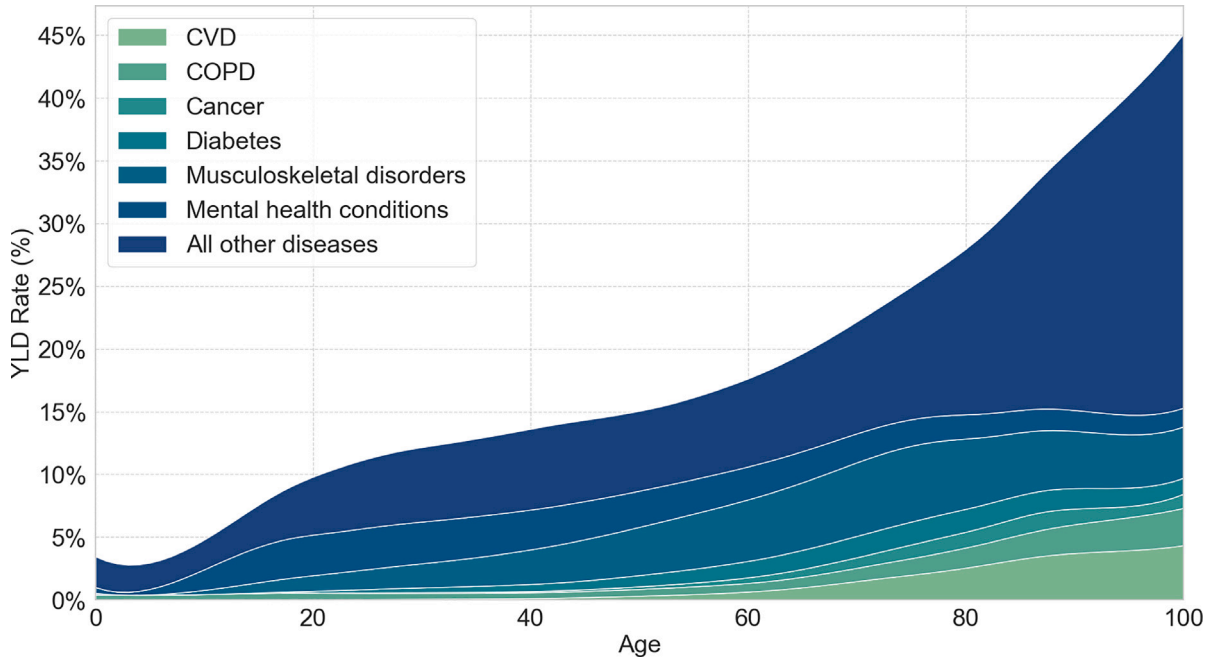

Fig. 1. Years Lived with Disability (YLD) rates by age.
Note: This figure illustrates the Years Lived with Disability (YLD) rates by age for six leading chronic diseases—cardiovascular disease (CVD), chronic obstructive pulmonary disease (COPD), cancer, diabetes, musculoskeletal conditions, and mental health conditions—along with all other diseases. Data are sourced from the 2021 Global Burden of Disease (GBD) dataset (IHME, 2021).

Population dynamics are governed by fertility, migration, and age-specific mortality rates so that the population aged a in time t is given by:

$$N_t(a) = N_{t-1}(a-1)\left(1 - \gamma_{t-1}(a-1)\right) + I_t(a) \tag{5}$$

where $\gamma_t(a)$ is the age-specific mortality rate and $I_t(a)$ is the number of immigrants aged a in time t. The population at $a = 0$ is given by the number of births in time t: $N_t(0) = B_t$.

Agents in the model begin work at age $\tilde{a}$. Labor force participation varies by age and is endogenously determined by a set of age-specific transition probabilities that dictate the likelihood of moving between employment, unemployment, and inactivity. The entire three-state system can be expressed in matrix form as:

$$\begin{bmatrix} E_t(a) \\ U_t(a) \\ N_t(a) \end{bmatrix} = \begin{bmatrix} \tau_{t,E\to E}(a) & \tau_{t,U\to E}(a) & \tau_{t,N\to E}(a) \\ \tau_{t,E\to U}(a) & \tau_{t,U\to U}(a) & \tau_{t,N\to U}(a) \\ \tau_{t,E\to N}(a) & \tau_{t,U\to N}(a) & \tau_{t,N\to N}(a) \end{bmatrix} \begin{bmatrix} E_t(a) \\ U_t(a) \\ N_t(a) \end{bmatrix}$$

where $E_t(a)$, $U_t(a)$, and $N_t(a)$ are the share of employed, unemployed, and inactive workers of age a in time t respectively, and $\tau_{t,i\to j}(a)$ is the transition probability from state i to state j for workers aged a in time t. The transition probability depends on a, reflecting the fact that the likelihood of moving between employment, unemployment, and inactivity varies by age and the dependence on t reflects the evolution of transition probabilities over time. The main counterfactual we perform in the model is to reduce the incidence of a set of chronic diseases which in turn affects the transition probabilities in the matrix above.

The distribution of employed workers across full-time and part-time working status is modeled in a similar fashion. Conditional on being employed, a worker can transition between full-time and part-time. The fraction of employed workers working full-time at age a in time t is given by:

$$f_t(a) = \tau_{FT\to FT} f_{t-1}(a-1) + \tau_{PT\to FT}(1 - f_{t-1}(a-1)) \tag{6}$$

where $\tau_{FT\to FT}$ is the probability of remaining full-time employed conditional on being full-time employed in the previous period and $\tau_{PT\to FT}$ is the probability of transitioning from part-time to full-time employment.

Calibration

We follow Bloom et al. (2020) and set the capital stock depreciation rate δ to 5%. The capital share of income α is set to 60% in accordance with the Office for National Statistics (ONS) Labour Costs and Labour Income Statistical Bulletin ONS (2022a).

We consider $t = 0$ to be the period in which the reduction in disease incidence is realized. We therefore abstract from any issues related to the practical implementation of a prevention policy and any lags between the start of the policy and its impact on disease. We set $A_0 = 1$ and calibrate total factor productivity (TFP) growth such that the growth trajectory of our baseline model matches that of the March 2024 Office for Budget Responsibility (OBR) Economic and Fiscal Outlook report (OBR, 2024a).

In our simulations, we assume that in $t = 0$ the economy starts from a steady-state level of capital. With δ, α, and A_0 calibrated, we set the savings rate s such that the model's capital–output ratio in $t = 0$ matches the ONS estimate of the UK's capital–output ratio of 3.9 (ONS, 2017). This calibration results in an estimate for s of 20.3%, close to the UK gross fixed capital formation rate of 18%, as reported by the World Bank (World Bank, 2024).

Treatment costs

To calibrate how changes in disease incidence affect the share of output spent on healthcare, we use the National Health Services (NHS) National Cost Collection (NCC) data for the fiscal year 2021–2022.[2] The NCC process aims to provide a comprehensive overview of costs associated with delivering various healthcare services across NHS providers. Healthcare costs are broken down by Healthcare Resource Group (HRG), which are standardized groupings of clinically similar treatments (e.g., heart transplant or hand fracture). In the NCC process, costs are allocated across HRGs by matching NHS outlays to specific patient

[2] Fiscal year 2021–2022 is the latest time period for which publicly available data in accessible format is available.

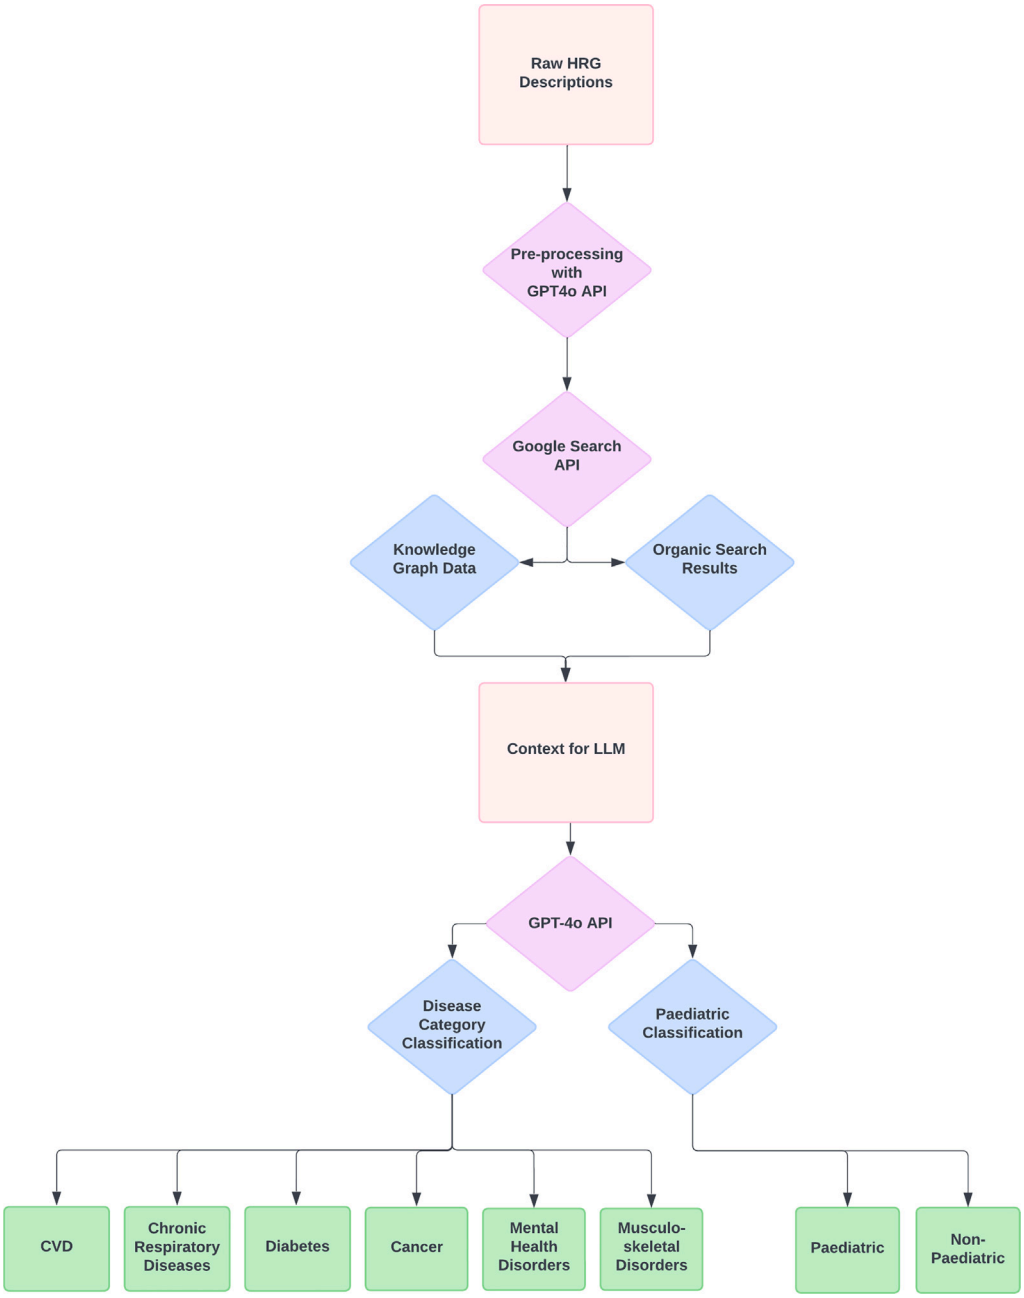

Fig. 2. Flowchart of AI classification pipeline.
Note: This flowchart illustrates the steps of the classification pipeline used to categorize HRG codes into one of six disease categories. The pipeline consists of three main steps: data preprocessing, context extraction, and classification. The data preprocessing step involves cleaning the HRG descriptions and removing unnecessary details. The context extraction step entails gathering additional information about the HRG codes from web searches for each HRG description. The classification step utilizes a large language model (LLM) to categorize each HRG into one of six disease categories based on the cleaned HRG description and the extracted context.

episodes within the NHS (e.g., a patient visiting the A&E department after a heart attack). Critically, the cost data reflect not just the direct costs of treatment but also the indirect costs of treatment, such as wages for administrative staff or other overhead.[3] One limitation of the NCC data is that it only includes NHS costs incurred by NHS trusts—which provide secondary, tertiary, and community care—and does not include costs incurred by primary care providers such as general practitioners (GPs).

The NCC cost data is broken down across 2825 HRGs. To estimate how the incidence of a particular disease affects NHS healthcare spending, we develop a classification pipeline that leverages recent advances in natural language processing (NLP). The classification process starts with a large language model (LLM) pre-processing the HRG descriptions, stripping away unnecessary details like complication and comorbidity (CC) scores, as these are not pertinent to the classification task. We then utilize the Google Search API to gather comprehensive context for each cleaned HRG string, including the Google Knowledge Graph entry and relevant snippets from the top five search results. The Google Knowledge Graph is a knowledge database that provides information alongside Google's online search results. As of May 2020, it contained 500 billion facts about 5 billion entities. Table 1 shows

[3] For a more detailed description of the NCC process, see Amies-Cull et al. (2023).

Table 1
Top 25 Base URLs with cumulative percentage.

Base URL	Count	Cumulative %
ncbi.nlm.nih.gov	4080	15.37
england.nhs.uk	2718	25.61
assets.publishing.service.gov.uk	1913	32.82
elht.nhs.uk	1282	37.65
nhs.uk	723	40.37
sciencedirect.com	602	42.64
nice.org.uk	581	44.83
my.clevelandclinic.org	460	46.56
digital.nhs.uk	439	48.21
emedicine.medscape.com	422	49.80
en.wikipedia.org	332	51.05
mayoclinic.org	332	52.31
hopkinsmedicine.org	286	53.38
dhcw.nhs.wales	281	54.44
alderhey.nhs.uk	206	55.22
cancerresearchuk.org	201	55.97
uhbristol.nhs.uk	179	56.65
ambulatoryemergencycare.org.uk	165	57.27
mtw.nhs.uk	159	57.87
thelancet.com	157	58.46
salisbury.nhs.uk	136	58.97
sbuhb.nhs.wales	123	59.44
webmd.com	119	59.89
ahajournals.org	117	60.33
cddft.nhs.uk	105	60.72

Note: This table shows the unique base URLs from which information about each HRG code was extracted to generate additional contextual information before the LLM classification task. The table is sorted by the total count of times that information from a particular base URL was extracted. A base URL is the root domain of a website, excluding the protocol (e.g. https://) and any subdirectories. For example, `nhs.uk` is the base URL of https://www.nhs.uk/conditions/laryngeal-cancer/.

the top 25 base URLs from which information about the HRG codes was extracted to generate additional contextual information before the LLM classification task. In the final step, we pass each HRG description with its respective text context to an LLM, which uses the provided context to categorize each HRG into our six predefined chronic disease categories: cancer, diabetes, respiratory disorders, cardiovascular disease, musculoskeletal disorders, and mental health disorders.

To ensure the accuracy of our classification approach, we conducted a validation exercise using a random sample of 450 HRG codes (approximately 16% of the total 2825 HRGs). We then evaluated our LLM-based classification against these expert assessments using standard performance metrics including precision, recall, and F1 scores. The full validation methodology and detailed results are presented in Appendix "Validation of LLM-based healthcare cost classification". Fig. 2 provides a visual representation of the classification pipeline.

Table 2 summarizes how our review of the NCC cost data apportions NHS costs across our six chronic diseases. We find that approximately a quarter (£12.89 bn) of all NHS Trust expenditures (£53.5 bn) in the 2021–2022 period can be explicitly allocated to the treatment of these chronic diseases. One limitation of our analysis is that it does not account for *general* healthcare categories, such as NHS expenditures on ambulance services that may be associated with chronic disease. For example, ambulance transport to a hospital for someone suffering from a heart attack is not counted as a CVD cost, whereas expenditures on performing coronary angioplasties do count. For this reason, our estimates are likely conservative.

To arrive at an estimate of total UK healthcare expenditure on our six chronic diseases, we apportion the UK's total healthcare spending for 2023 as reported in the UK Health Accounts (£292 bn) into the six disease based on Table 2. As a last step, we divide these estimated healthcare expenditures by UK GDP to arrive at an estimated fraction of output spent on each disease.[4] Our results compare favorably with existing studies of NHS healthcare expenditures of £12 billion for mental health (Department of Health & Social Care, 2023), £7.4 billion for CVD (Raleigh et al., 2022), £5.81 billion for cancer (Aggarwal and Sullivan, 2014), £4.9 billion for COPD and asthma (NHS England, 2024a), and £4.76 billion for MSK conditions (Office for Health Improvement & Disparities, 2022).

For diabetes, our estimate is substantially lower than most figures reported by public health authorities and existing studies. We believe there are at least two reasons that explain this discrepancy: (1) existing studies (e.g. Hex et al. (2024)) include treatment costs for diseases that are well-established co-morbidities of diabetes, such as coronary heart disease, which we apportion exclusively to our CVD disease category, and (2) much of the treatment costs for diabetes likely fall to GP care in the UK, which is not included in the NHS NCC data. To the extent that expenditure on diabetes appears in other disease categories, our overall estimates of the impact of GDP are valid; however, the relative importance of cost savings per disease will be incorrect.

Labor force participation

Agents in our model can either be employed, unemployed, or inactive. The transition probabilities between these states are estimated from the ONS LFS Longitudinal Data (ONS, 2024c).[5] To reduce noise, we employ locally weighted smoothing to generate smoothed values of transition probabilities by age. The LFS data contain information on workers up until the age of 69, so we extrapolate transition rates from age 70 through age 100. Fig. 3 shows the estimated and extrapolated transition probabilities. By age 21, around 27% of females and 36.3% of males are working in full-time education (ONS, 2024d), enabling us to estimate reliable transition probabilities from that age. We use the age-specific transition probabilities to infer the steady-state LFP-by-age curve and use this to set values for $LFP_0(a)$ for all a. Fig. 4 shows that the steady-state LFP-by-age curve implied by the age-specific transition probabilities matches well with the empirically observed LFP-by-age curve in the LFS dataset. We use the same methodology to calibrate the transition probabilities between full-time and part-time work.

Productivity by age

To calibrate how productivity varies with age we use data from UKHLS to estimate a Mincer-type equation, regressing individual wages on age while controlling for other relevant characteristics that influence earnings. The UKHLS is a large-scale, longitudinal survey collecting annual data from households across the United Kingdom. More details about this dataset are given in Section "Estimating the impact of disease on labor force participation". The UKHLS extract we use covers the period from 2009 through 2019 and the dataset comprises 86,096 individuals.

To estimate the wage-age profile we estimate:

$$\ln(\text{Wage}_{i,t}) = \sum_{a=20}^{80} \beta_a D_a + \alpha_i + \delta_t + \mathbf{X}'_{i,t}\boldsymbol{\theta} + \varepsilon_{i,t} \tag{7}$$

where $\text{Wage}_{i,t}$ is an individual's total monthly gross labor income (variable `fimnlabgr_dv` in UKHLS), D_a is a dummy variable for each age 20 through 80 (we use age 19 as the reference period), α_i is an individual fixed effect, δ_t is a year fixed effect, $\mathbf{X}_{i,t}$ is a vector of additional fixed effects and control variables, $\boldsymbol{\theta}$ is the corresponding vector of coefficients, and $\varepsilon_{i,t}$ is an error term. The vector $\mathbf{X}_{i,t}$ includes fixed effects for industry, education level, labor force status (i.e. employed vs. unemployed vs. inactive), and working mode (i.e. full-time vs. part-time). We control for these additional fixed effects as individuals

[4] For UK GDP we use the ONS chain volume measure of GDP for 2022 (ONS, 2024a)

[5] Estimates of the annual transition rates are obtained from LFS Five-Quarter Longitudinal Datasets from January 2013 through December 2019.

Table 2
Healthcare costs by disease (in £millions).

Disease	Inpatient	Outpatient	Other	Total cost	Percentage of NHS trust spend	Estimated total UK health spend	Percentage of UK GDP (2022)
Mental health	797	39	10,117	10,953	12.32%	35,974	1.58%
CVD	4101	831	310	5242	5.90%	17,228	0.76%
Cancer	3503	990	678	5171	5.82%	16,994	0.75%
COPD	2960	430	221	3611	4.06%	11,855	0.52%
Musculoskeletal	1451	1362	618	3430	3.86%	11,271	0.50%
Diabetes	208	239	157	604	0.68%	1986	0.09%
Total	13,020	3891	12,101	29,011	32.64%	95,300	4.20%

Note: This table shows the costs recorded in the NHS National Cost Collection (NCC) data for the fiscal year 2021–2022 for each of the six disease categories, broken down by inpatient (HRG), outpatient, and other costs. The table also displays the percentage of the total NHS Trust spend (£88.89 billion), the estimated total UK health spend based on the proportion of NHS Trust spend (total UK health spend £292 billion), and the percentage of the UK GDP (£2,271 billion in 2022) that each disease category represents.

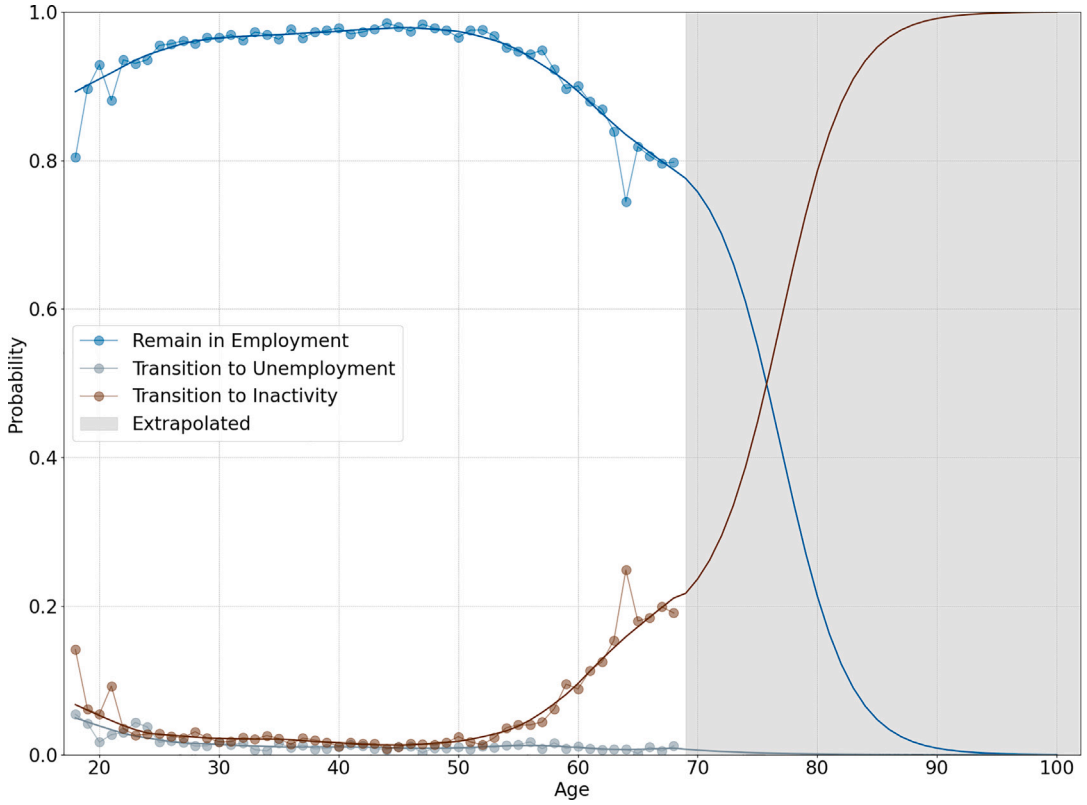

Fig. 3. Labor force transition probabilities from employment by age.
Note: This figure shows the estimated probabilities of labor force transitions from employment for ages 18–100, using the ONS Labour Force Survey (LFS). The plot displays three transition probabilities: remaining employed, transitioning to unemployment, and transitioning to inactivity. Solid lines represent smoothed estimates using LOWESS (Locally Weighted Scatterplot Smoothing), while markers show the raw data points. The shaded area indicates extrapolated predictions for ages beyond the observed data. Data for ages 18–69 are directly estimated from the LFS; transitions for ages 70–100 are extrapolated using sigmoid and exponential decay functions. The smoothing parameter used in the LOWESS estimation is 0.2.

change the industry in which they work, their level of education, their labor force status, and their working mode as they age.

Our specification differs from the traditional Mincer approach in two key ways. First, we use education level fixed effects rather than years of schooling to better reflect the UK education system, where qualifications (e.g., A-levels, vocational certificates, university degrees) matter more for earnings than the simple duration of schooling. Second, we replace the standard quadratic experience terms with age fixed effects to allow for a more flexible age-earnings relationship that can capture potential non-linearities in the lifecycle earnings profile. While we refer to our regression as "Mincer-type", our primary goal is to estimate the wage-age profile of UK workers while controlling for a rich set of covariates. This empirical specification generates age-earnings profiles that we can directly use to calibrate our macroeconomic model.

The resulting profile of estimates is shown in Fig. 6. Because very few individuals in UKHLS are in paid work after the age of 80 it is difficult to obtain accurate estimates of the association between age and wage at older ages. For our simulations, we therefore extrapolate the estimated earnings through to age 100 using an exponential decay function, with earnings set to zero at age 100.

Disease incidence, population, and mortality

We calibrate disease incidence by age for our six chronic diseases from the Global Burden of Disease (GBD) dataset (IHME, 2021), plotted in Fig. 5. Population and mortality statistics are obtained from the ONS (ONS, 2021, 2020).

Estimating the impact of disease on labor force participation

Our analysis uses longitudinal data from UKHLS, following 86,063 individuals between 2009 and 2019. The average individual is observed

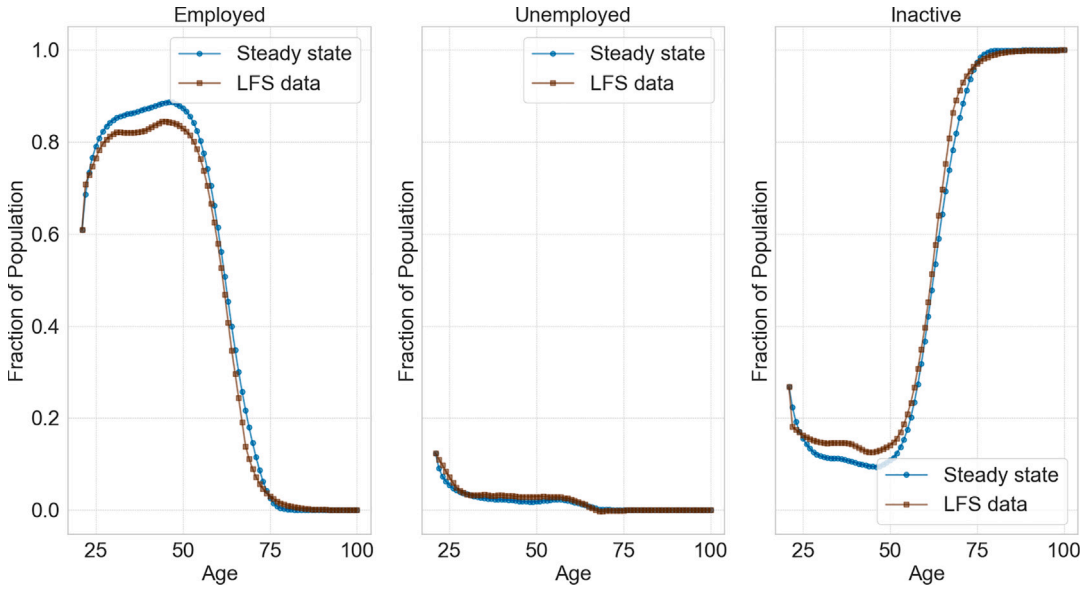

Fig. 4. Labor force participation by age: Data vs Steady-state.
Note: This figure compares the observed and steady-state labor force participation (LFP) rates across ages 21–100. The blue (round marker) lines represent the steady-state LFP curves, which are derived by iterating forward the age-specific labor force transition probabilities until the age-LFP distribution converges. The brown (square marker) lines represent the smoothed estimates from the January 2013 through December 2019 Labour Force Survey (LFS) Five-Quarter Longitudinal Datasets. Data were smoothed using the LOWESS (Locally Weighted Scatterplot Smoothing) technique.

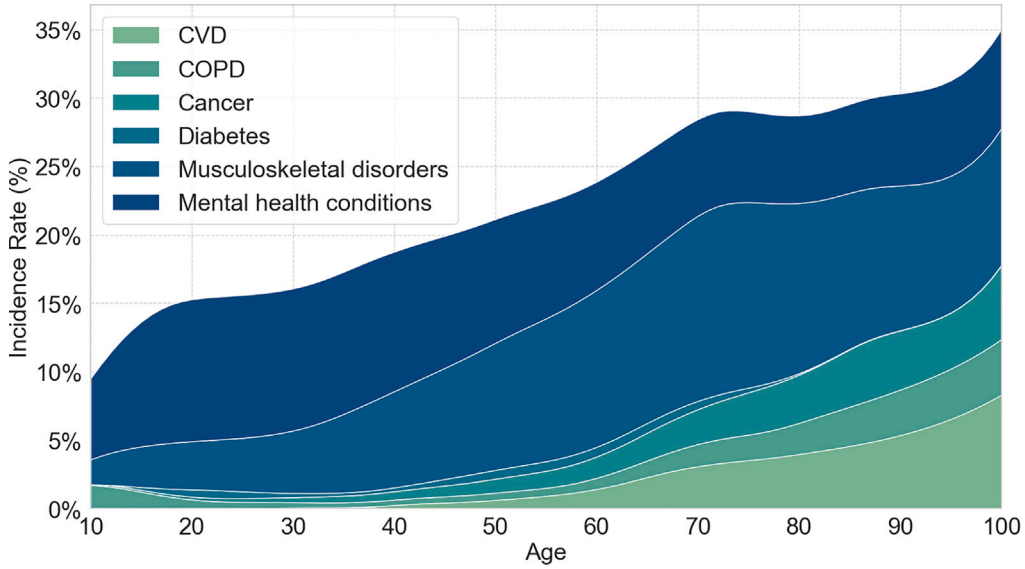

Fig. 5. Disease incidence rate by age.
Note: This figure shows the annual disease incidence rates for the six chronic diseases we study by age. Data are taken from the 2021 Global Burden of Disease (GBD) dataset (IHME, 2021). The data aggregate disease incidence by age brackets; therefore, we interpolate the individual age-specific disease incidence rates.

for 4.8 waves (years) of the survey, with a median observation length of 4 waves. Disease incidence varies substantially across conditions: 48.0% of individuals experience a musculoskeletal condition during the sample period, while 27.0% report a mental health episode. The incidence of other chronic conditions is lower, with 4.0% developing respiratory conditions, 3.7% cardiovascular disease, 3.3% diabetes, and 2.9% cancer. Many individuals develop multiple conditions over time—while 44.7% of the sample never reports a chronic condition, 28.5% develop one condition, 21.1% develop two conditions, and 5.6% develop three or more conditions during the observation period. This pattern of multiple diagnoses underscores the importance of accounting for comorbidities, a limitation of our current analysis that we discuss in Section "Impact on growth".

Disease classification and variable construction

For cancer, diabetes, COPD, and CVD, we identify onset through respondents' reports of formal medical diagnoses. However, for musculoskeletal and mental health conditions, which typically develop gradually rather than having clear onset points, we take a different approach that better captures when these conditions begin to affect work capacity. For musculoskeletal conditions, we identify onset as the point in time when an individual transitions from reporting that

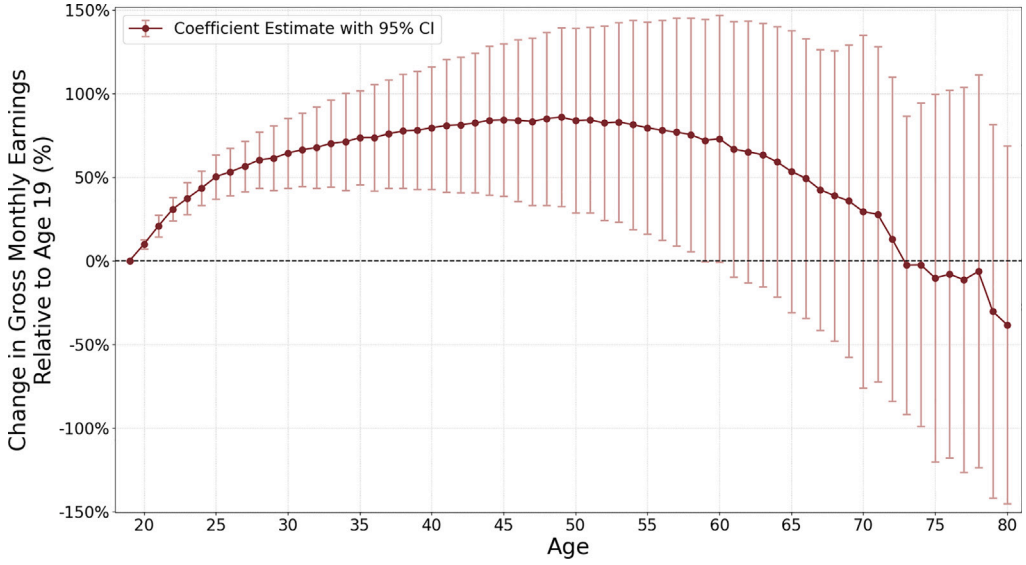

Fig. 6. Estimated wage-age profile of UK workers.
Note: This figure shows estimated percentage changes in gross monthly earnings relative to age 19 for ages 19–100, using UKHLS data from 2009–2019 for employed individuals. A Mincer-type regression is estimated with log gross monthly earnings as the dependent variable, including age fixed effects with age 19 as the reference. Controls include individual, year, education, and industry fixed effects. For age, the variable `age_dv` was used. For gross monthly earnings, the variable `fimnlabgrs_dv` was used. Total observations: 216,173 on 51,965 unique individuals.

Table 3
Disease diagnosis in UK Household Longitudinal Study (UKHLS)

Disease	Disease diagnosis in UKHLS
Cancer	Cancer or malignancy
Diabetes	Diabetes
Respiratory	Asthma, Chronic bronchitis, Emphysema
Cardiovascular Disease	Congestive heart failure, Coronary heart disease, Angina, Heart attack or myocardial infarction, Stroke
Musculoskeletal (MSK)	Respondent reports pain interferes with work
Mental Health (MH)	GHQ-12 score ≥ 3

Note: This table presents the mapping of diseases to their corresponding diagnoses in the UK Household Longitudinal Study (UKHLS).

pain does not interfere with their work to reporting that it does. For mental health conditions, we identify onset when an individual's GHQ-12 score (a validated screening tool for psychological distress) crosses the clinical threshold of 3. The GHQ-12 is an established screening tool for detecting psychological distress and minor psychiatric disorders in general populations (Goldberg, 1988). This measurement strategy focuses on capturing meaningful changes in mental health status that affect work capacity rather than relying solely on formal medical diagnoses, which may lag behind actual changes in a person's mental health. Table 3 provides an overview of how the onsets of different chronic diseases are defined for the purpose of our estimation exercise.

Sample construction

Our analysis follows several steps. For each disease, we create a panel that includes all individuals who report a diagnosis of the given disease (treated individuals) as well as individuals who never report a diagnosis of the given disease (control individuals). To balance the panel across treated and untreated individuals, we construct matching cells by first dividing individuals into groups based on their characteristics: we assign individuals to five-year age bands based on their median age across all observations, and to pay categories based on their median earnings when employed (using deciles for the employed, plus a separate category for those predominantly inactive). Each possible combination of these age bands and pay categories forms a matching cell. We then ensure that each cell contains an equal number of treated and control individuals.

We then create indicators of employment status, converting the original categorical variables in UKHLS into binary indicators for economically active, unemployed, and economically inactive. An individual is classified as economically active if they report being in paid employment (either full-time or part-time) or self-employed. Economic inactivity encompasses those who report being retired, long-term sick or disabled, taking care of home or family, or "doing something else".

We only include treated individuals who remain in the sample for at least three years post-diagnosis. This restriction mitigates potential bias that could arise from individuals leaving the sample due to mortality, which would otherwise distort our transition probability estimates. Additionally, we restrict our treated sample to individuals who were economically active in the wave immediately preceding diagnosis, as our primary interest lies in estimating the effect of chronic disease diagnosis on labor market transitions among the working population. Similarly, when estimating the effect of a chronic disease diagnosis on full-time-to-part-time transitions, we restrict the treated sample to individuals who were working full-time in the wave immediately preceding diagnosis.

Empirical strategy

Our estimation employs the local projection difference-in-differences (LP-DiD) method proposed by Dube et al. (2024) to address challenges associated with staggered treatment timing, as individuals experience the onset of disease at different points in time. Traditional TWFE estimators can be biased in such settings because they may inadvertently compare newly treated individuals to those previously treated who might still be experiencing treatment effects (de Chaisemartin and D'Haultfœuille, 2020; Goodman-Bacon, 2021; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Athey and Imbens, 2022; Borusyak et al., 2024). The LP-DiD approach accounts for this source of bias by employing local projections to estimate dynamic treatment effects while implementing a "clean control" condition. This condition ensures that only observations unaffected by prior treatments are included in the control group, thereby avoiding the bias.

Our baseline regression takes the form:

$$y_{i,t+h} - y_{i,t-1} = \beta^h \Delta D_{i,t} + \boldsymbol{\gamma}^h \mathbf{X}_{i,t} + \delta_t^h + e_{i,t}^h \tag{8}$$

where $y_{i,t}$ is the outcome variable of interest (transition to inactivity or transition to part-time employment), $D_{i,t}$ is a binary treatment variable indicating disease onset, $\mathbf{X}_{i,t}$ is a vector of controls including age polynomials, lagged employment status, and industry indicators, δ_t^h is a time fixed effect, and $e_{i,t}^h$ is the error term. The subscript i indexes individuals, t indexes survey waves, and h indexes the treatment time horizon relative to disease onset.

Identification assumptions

The key assumptions required for the LP-DiD estimator to identify the effect of chronic disease onset on labor market outcomes are that individuals are not able to anticipate their diagnosis (no anticipation) and that the labor market outcomes of treated individuals would have continued to evolve in parallel with the labor market outcomes of untreated individuals in the absence of disease onset (parallel trends). No anticipation is more plausible for sudden-onset conditions like cancer and certain components of CVD—such as stroke—where diagnoses often occur without prior warning. This assumption is less likely to hold for diagnoses such as mental health and MSK conditions. Nevertheless, the absence of significant effects in the pre-treatment periods—even for mental health and MSK conditions—suggests that treated individuals are not changing their labor force participation any differently than untreated individuals, at least at the annual frequency, providing some support for the assumption.

A threat to the parallel trends assumption arises in two ways. First, there may be unobserved factors that simultaneously affect both the probability of a diagnosis and the probability of transitioning from employment to inactivity. For instance, if high-stress jobs both increase the risk of health conditions and lead to earlier retirement, this would violate our identifying assumptions. We address this concern by including a comprehensive set of controls in our estimation: individual fixed effects to account for time-invariant individual characteristics, calendar year fixed effects to capture economy-wide trends, age polynomials to flexibly control for life-cycle patterns, lagged employment status to accounts for path dependence in labor market outcomes, and industry indicators to control for sector-specific working conditions. Second, the direction of causality might run from employment status to health diagnosis rather than vice versa. This reverse causality is particularly concerning for mental health conditions, where job loss itself might trigger or lead to the diagnosis of depression or anxiety. While we see no evidence that individuals systematically switch into inactivity prior to any of the health diagnoses, we would require higher-frequency data to more substantively rule out the possibility of reverse causality for mental health diagnoses.

Estimation results

Our empirical results suggest that the onset of chronic diseases significantly increases the likelihood of workers transitioning from employment to inactivity. For instance, a cancer diagnosis increases the probability of exiting employment by 11.5 percentage points in the diagnosis year. Similarly, CVD diagnoses lead to an immediate increase of 9.1 percentage points in the likelihood of becoming inactive, with the effect remaining elevated at 4.1 percentage points two years after diagnosis. Mental health episodes consistently raise the probability of exiting the workforce by 1.3 to 2.9 percentage points in the years following diagnosis. In contrast, the effects of COPD on labor force exit are less pronounced and not statistically significant beyond the diagnosis year.

Fig. 8 presents scaled estimates of the impact of a CVD diagnosis on labor force participation, expressed as relative risks across age groups. For individuals aged 40–49, a CVD diagnosis is associated with an 11-fold increase in the relative risk of employment exit during the year of diagnosis. This elevated risk persists, with affected individuals still experiencing a sixfold increase in the relative likelihood of employment exit two years post-diagnosis.

It is worth noting that the hump-shaped effect pattern observed in Fig. 7 may partly reflect a selection mechanism. In particular, patients with more severe symptoms might exit employment relatively quickly after diagnosis, leading to an early, pronounced effect. Over time, as these more severely affected individuals leave the sample, the remaining group largely comprises workers with milder symptoms, which can result in a diminished treatment effect in later periods.

In addition to influencing exits from employment, chronic diseases affect transitions from full-time to part-time work. As shown in Table 5, mental health conditions increase the probability of transitioning to part-time work by 2.5 percentage points in the diagnosis year, with effects persisting up to three years post-diagnosis (1.1 percentage points). MSK conditions demonstrate a similar pattern, with a 2.9 percentage point increase in transition probability during the diagnosis year; however, the persistence of these effects is less pronounced. Fig. 9 illustrates the dynamic evolution of these effects, underscoring that health shocks induce labor supply adjustments not only through complete withdrawal but also through the intensive margin of hours worked. While we attempted to estimate similar effects for the other four chronic conditions, the resulting estimates proved too imprecise to support reliable inference.

Limitations of analysis and implications for causal interpretation

A key limitation of our approach is that we generally treat the timing of diagnosis as equivalent to disease onset. This assumption is more reasonable for acute conditions with sudden onset, such as cancer, than for gradually developing conditions like mental health disorders or musculoskeletal problems. To address this issue for gradual-onset conditions, our empirical strategy attempts to capture meaningful changes in health status rather than formal medical diagnoses. For musculoskeletal conditions, we identify "onset" as the point in time when an individual transitions from reporting that pain does not interfere with their work to reporting that it does. This approach likely captures the actual onset of work-limiting musculoskeletal problems more accurately than formal diagnostic records, as we are directly measuring when these conditions begin to affect work capacity. Similarly, for mental health conditions, we identify onset based on when an individual's GHQ-12 score crosses the clinical threshold indicating psychological distress, rather than relying on formal psychiatric diagnoses. This approach better reflects the emergence of mental health challenges that affect work capacity.

Nevertheless, some caution in causal interpretation remains warranted. While our event study estimates show no significant pre-trends in labor market transitions before these health status changes, this

Table 4
Effect of chronic disease on likelihood of transitioning from employment into inactivity.

	Cancer	COPD	CVD	Diabetes	Mental	Musculoskeletal
Diagnosis Year +3	0.017	0.003	−0.004	−0.009	0.013***	0.011
	(0.024)	(0.009)	(0.008)	(0.010)	(0.001)	(0.006)
Diagnosis Year +2	0.028	0.010	0.041***	−0.004	0.014**	0.015**
	(0.021)	(0.009)	(0.010)	(0.007)	(0.005)	(0.005)
Diagnosis Year +1	0.032	0.008	0.049***	0.021	0.025***	0.019**
	(0.021)	(0.011)	(0.013)	(0.011)	(0.006)	(0.006)
Diagnosis Year	0.115***	0.024**	0.091***	0.076***	0.029***	0.023***
	(0.009)	(0.009)	(0.021)	(0.011)	(0.003)	(0.004)
Diagnosis Year −1 (Omitted)	0	0	0	0	0	0
Diagnosis Year −2	0.002	0.002	0.012	0.007	0.001	−0.002
	(0.009)	(0.007)	(0.015)	(0.005)	(0.003)	(0.002)
Diagnosis Year −3	−0.007	0.001	−0.002	−0.006	−0.001	−0.003
	(0.023)	(0.005)	(0.008)	(0.006)	(0.002)	(0.002)
Diagnosis Year −4	−0.006	0.003	0.006	−0.001	0.005*	0.003
	(0.018)	(0.009)	(0.010)	(0.004)	(0.002)	(0.004)
Control for age + age^2	✓	✓	✓	✓	✓	✓
Control for LF status	✓	✓	✓	✓	✓	✓
Control for industry	✓	✓	✓	✓	✓	✓
N with diagnosis	358	638	319	444	7,165	12,605
N without diagnosis	358	638	319	444	7,165	12,603

Standard errors in parentheses: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Note: This table presents the results from the LP-DiD estimation for different time horizons and across six chronic diseases. The outcome variable is the probability of transitioning from employment to inactivity, and the reported coefficient estimates are to be interpreted as percentage point (in decimal form) increases in the likelihood of transitioning from employment to inactivity. For example, the "Diagnosis" coefficient for cancer is 0.115, meaning that workers with a cancer diagnosis are 11.5 percentage points more likely to transition from employment to inactivity following their diagnosis compared to workers who do not have such a diagnosis. Standard errors are clustered at the individual and (annual) survey level. Estimates and 95% confidence intervals are plotted in Fig. 7.

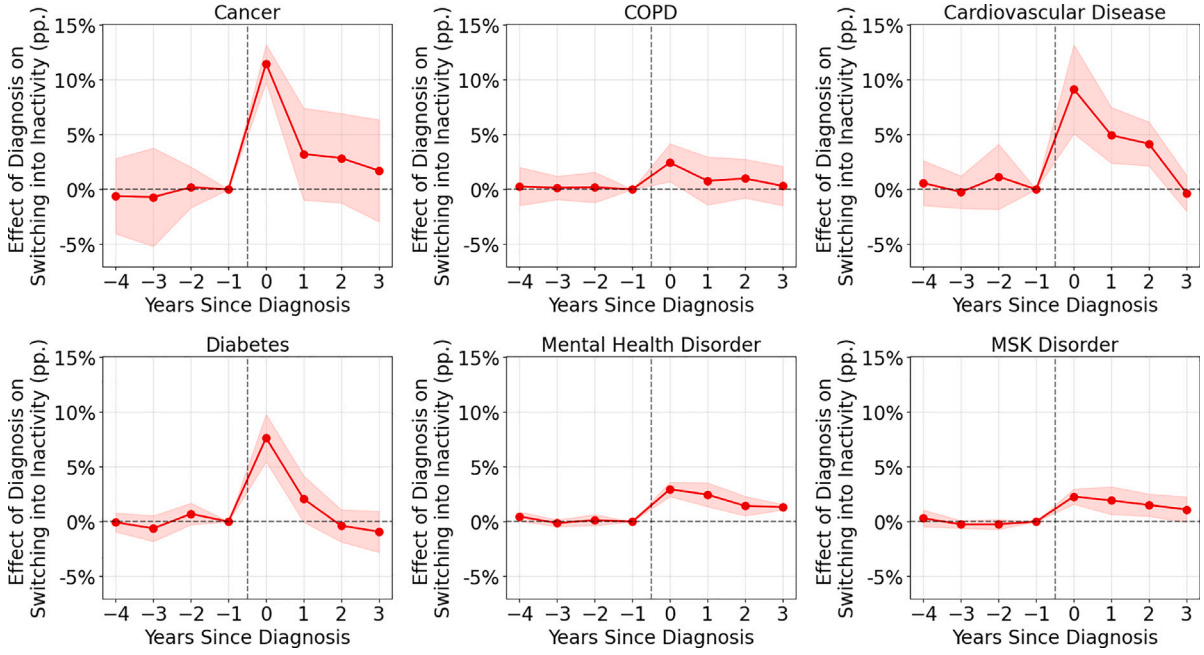

Fig. 7. Effect of chronic disease diagnosis on likelihood of transitioning into inactivity.
Note: This figure presents estimates of the percentage point change in the likelihood of transitioning from employment to inactivity following a chronic disease diagnosis. Estimates are obtained from using the LP-DiD estimator to estimate the specification in Eq. (8) using a panel that is balanced across age and income. To obtain estimates for the effect on individuals who received their diagnosis while in employment, the pool of treated individuals is restricted to include only those who receive a given disease diagnosis while in employment (i.e., those who report being actively in the labor force in the survey wave prior to the wave in which they report their diagnosis). The shaded areas indicate 95% confidence intervals around the point estimates.

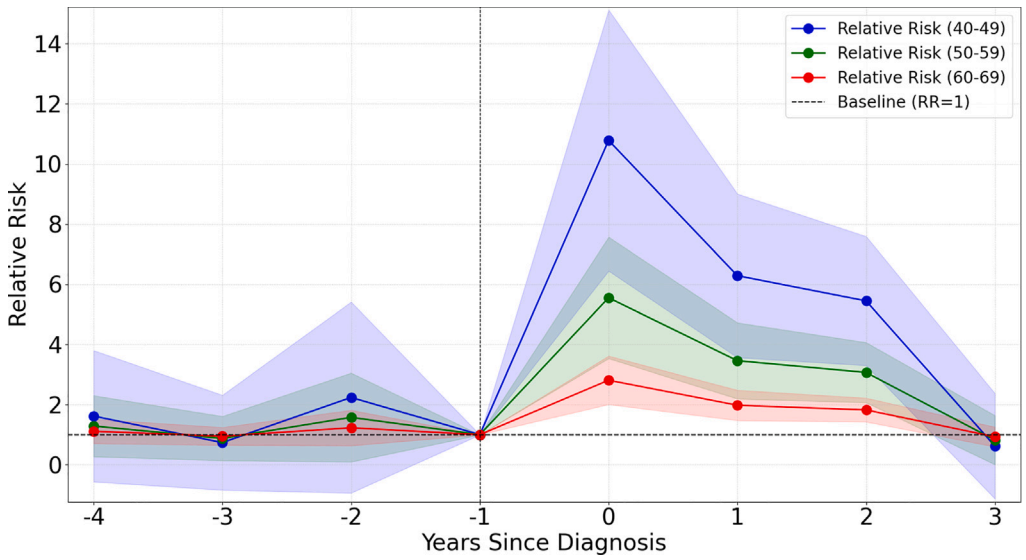

Fig. 8. Relative Risk of CVD diagnosis on likelihood of transitioning into inactivity (by age group).
Note: This figure displays the estimated Relative Risk (RR) of transitioning from employment to inactivity following a chronic disease diagnosis, using the LP-DiD estimation method across different time horizons and age groups. The RR is calculated by scaling the estimated coefficients with the average baseline transition rates specific to each age group (40–49, 50–59, and 60–69 years old). An RR greater than 1 indicates an increased likelihood of transitioning to inactivity compared to individuals without a diagnosis. For example, in the case of cardiovascular disease (CVD) for the 60–69 age group, a Relative Risk of 2 at Year +1 means that diagnosed individuals are twice as likely to exit employment one year after diagnosis compared to their healthy counterparts. The shaded areas represent the 95% confidence intervals around the RR estimates for each age group.

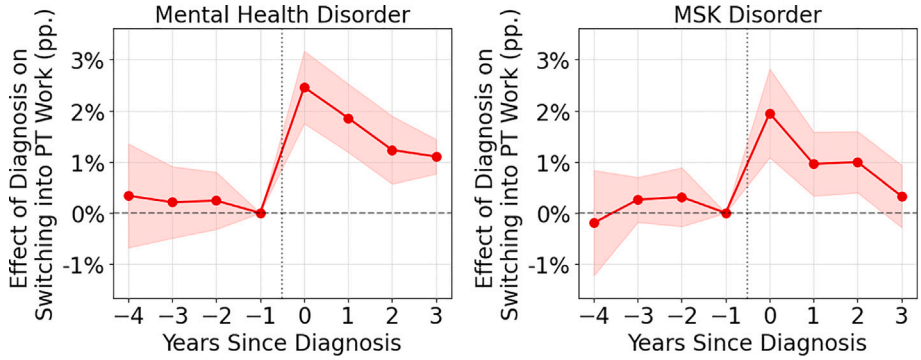

Fig. 9. Effect of chronic disease diagnosis on likelihood of transitioning into part-time employment.
Note: This figure presents estimates of the percentage point change in the likelihood of transitioning from full-time to part-time employment following a chronic disease diagnosis. Estimates are obtained from using the LP-DiD estimator to estimate the specification in Eq. (8) using a panel that is balanced across age, sex, and income. To obtain estimates for the effect on individuals who received their diagnosis while in employment, the pool of treated individuals is restricted to only include those who receive a given disease diagnosis while in full-time employment (i.e. those who report working full-time in the survey wave prior to the wave in which they report their diagnosis). The shaded areas indicate 95% confidence intervals around the point estimates.

could reflect the annual frequency of our data rather than true absence of pre-onset effects. Additionally, for mental health conditions in particular, the relationship with labor market outcomes may be bidirectional—deteriorating mental health may affect work capacity, but job loss or work stress could also trigger mental health episodes.

In Appendix "Calculation of counterfactual transition rates" we explain how we translate our estimates of disease impacts on labor market outcomes into changes in aggregate transition rates between employment states in our model. The aggregate transition rate from employment to inactivity at each age is a weighted average of transition rates for healthy individuals and those with each disease, where the weights are determined by disease incidence. A reduction in disease incidence therefore reduces these transition rates proportionally to three factors: the additional transition risk associated with each disease, the size of the reduction in disease incidence, and the initial incidence rate of the disease at each age.

Results

To quantify the effect of chronic disease on UK output we use our model to analyze the impact of a 20% reduction at all ages in the incidence of cancer, CVD, COPD, diabetes, MSK, and mental health conditions. To be clear, in our counterfactual scenario, we model a permanent 20% reduction in disease incidence that begins in period $t = 0$ and persists throughout the simulation horizon. This reduction means that in every period from $t = 0$ onward, individuals face a 20% lower probability of being diagnosed with any of the six chronic diseases we study.

The choice to focus on a 20% reduction in incidence is to some extent arbitrary, reflecting a balance between ambition and feasibility. Clinical evidence from Cholesterol Treatment Trialists (CTT) suggests that statin therapy can reduce the risk of a major vascular event in patients by 22% within one year of the therapy's start. Clark (2018) finds

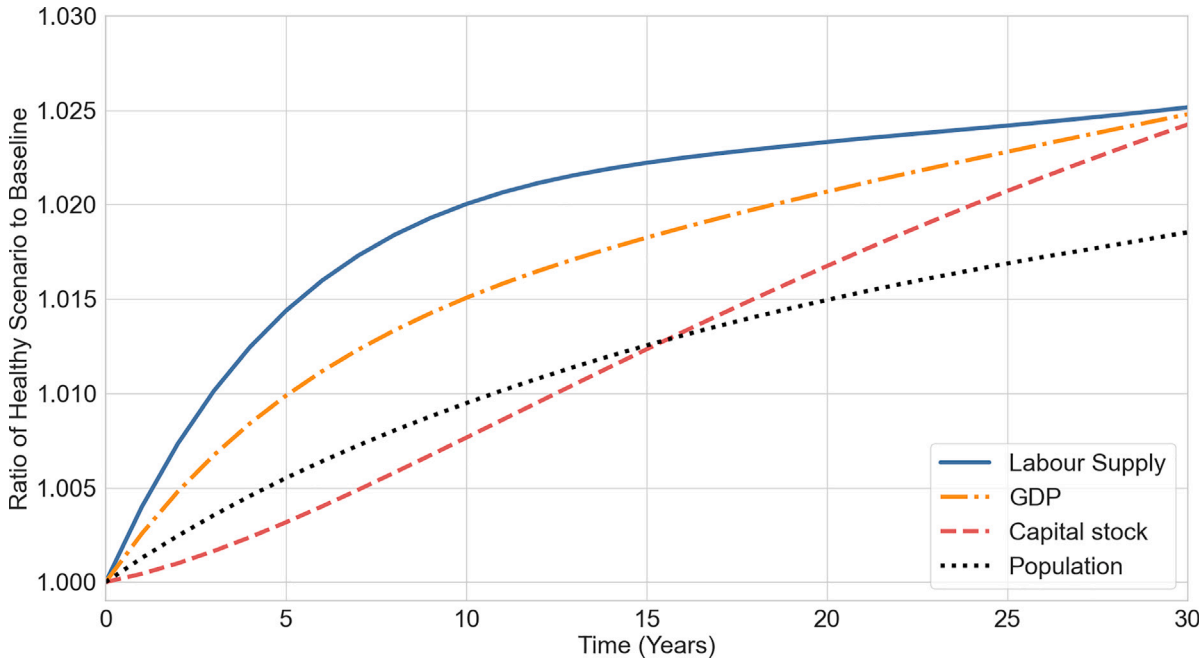

Fig. 10. Effect of 20% reduction in chronic diseases on UK macroeconomic aggregates.
Note: This figure shows how the model's aggregate labor supply, capital stock, output, and population evolve in the healthy scenario relative to the baseline scenario. Each line represents the ratio of the given aggregate in the healthy scenario to the same aggregate in the baseline.

Table 5
Effect of chronic disease on likelihood of transitioning from full-time to part-time work.

	Mental health	Musculoskeletal
Diagnosis Year +3	0.011***	0.003
	(0.002)	(0.003)
Diagnosis Year +2	0.012**	0.010**
	(0.003)	(0.003)
Diagnosis Year +1	0.019***	0.010**
	(0.003)	(0.003)
Diagnosis Year	0.025***	0.020***
	(0.004)	(0.004)
Diagnosis Year −1 (Omitted)	0	0
Diagnosis Year −2	0.002	0.003
	(0.003)	(0.003)
Diagnosis Year −3	0.002	0.003
	(0.004)	(0.002)
Diagnosis Year −4	0.003	−0.002
	(0.005)	(0.005)
Control for age +age^2	✓	✓
Control for LF status	✓	✓
Control for industry	✓	✓
N with diagnosis	5,246	8,850
N without diagnosis	5,246	8,850

Standard errors in parentheses: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Note: This table shows the LP-DiD estimates for the effect of a chronic disease diagnosis on switching from full-time to part-time work, focusing on two conditions: mental health and musculoskeletal disorders. The outcome variable is the probability of transitioning from full-time to part-time work. The coefficient estimates are interpreted as percentage point (in decimal form) increases in the likelihood of transitioning from full-time to part-time work. For example, a coefficient of 0.025 implies a 2.5 percentage-point increase in the probability of switching to part-time work following a diagnosis, compared to individuals without that diagnosis. Standard errors are clustered at the individual and year levels.

that a variety of therapies for treating mental health conditions offered in the UK—ranging from counselling to mindfulness-based cognitive therapy—were able to significantly improve anxiety and depression symptoms in over half of enrolled patients. Friedson et al. (2021) provide evidence that a one US dollar increase in cigarette taxes is associated with an 8% reduction in adult smoking leading to a 6% reduction in mortality from heart disease and lung cancer. However, our simulations are not intended to assess the impact of any specific prevention program but to estimate a "growth-health" elasticity for the UK economy.

In interpreting our results, it is important to bear in mind three key factors. Firstly, as our results are broadly linear in incidence, the choice of 20% is not particularly limiting. The impact of a 10% reduction is approximately half of the impact reported in this section. Secondly, we are focusing on the economic consequences of a given reduction in disease incidence, not the effects of any particular treatment or prevention program. Different prevention programs will have varying timelines to achieve reductions in disease incidence, which would introduce additional lags to our results that we do not consider. Thirdly, we do not make any attempt to identify the interventions that would achieve these hypothetical reductions; our focus is purely on the magnitude of gains from any successful intervention.

Impact on growth

Fig. 10 shows the differential evolution of aggregate labor supply, output, capital stock, and population levels in the healthy scenario versus the baseline scenario. Our analysis suggests that a 20% reduction in chronic disease would have substantial effects on output growth, with annual output being 0.99 percentage points higher after 5 years, 1.51 percentage points higher after 10 years, and 2.07 percentage points higher after 20 years. These magnitudes align well with findings from similar studies in other economies. When we model a 20% reduction in the same five chronic diseases examined by Chen et al. (2018) for the US economy, we find that the cumulative output gain over 30 years would amount to 56.1% of current UK output, comparable to their projected gain of 51.2% of US output. Our results are also consistent with Bloom et al. (2020)'s analysis of Asian economies, where a 20% reduction in disease prevalence yields cumulative output gains of 25% for China, 12% for Japan, and 18% for South Korea relative to current output. These comparisons should be interpreted cautiously, given

differences in methodology. Bloom et al. (2020), for instance, focuses on reductions in disease prevalence rather than incidence.

Broader comparisons with the existing literature are challenging for several reasons. Many studies focus on different countries with distinct demographic profiles, healthcare systems, and labor market institutions. Others examine different combinations of diseases or measure disease burden using alternative metrics such as disability-adjusted life years or healthcare expenditure rather than output growth. Studies also vary in their methodological approaches, with some using cost-of-illness frameworks that, as discussed below, can overstate economic impacts by not accounting for labor market adjustment mechanisms. These differences in scope, context, and methodology make direct comparisons of magnitudes across studies difficult and potentially misleading.

Our estimates of the economic impact of chronic disease are lower than those typically found in cost-of-illness studies (see e.g. McKinsey, 2020). This difference likely reflects two key methodological differences in our approach. First, cost-of-illness studies often treat income losses from labor force exit as permanent and complete, whereas our analysis permits workers to return to the labor force after exit. Second, while cost-of-illness studies typically treat healthcare costs as a pure economic burden and assume that any reduction in these costs would directly boost GDP, our approach takes a more nuanced view. In our model, when disease prevention reduces healthcare costs, these savings are partially reinvested into the economy's capital stock at the national savings rate, while the remainder contributes to consumption. This better reflects the reality that when society needs to spend less of the economy's output on healthcare, this frees up resources for other productive uses—and additional consumption—rather than translating one-for-one into higher GDP.

As shown in Fig. 11, the majority of the output impact arises from changes in the aggregate labor force participation rate and in the proportion of full-time workers—together accounting for more than three-quarters of the output boost. Increases in aggregate labor supply due to reduced mortality account for 9% of the 10-year effect while changes in the rate of capital accumulation—stemming both from an endogenous adjustment of the capital stock to increased labor supply and treatment cost savings being re-invested—account for 11% of the 10-year effect.

Breakdown by disease type and age brackets

Fig. 12 provides a breakdown of the output effect into the six different chronic diseases considered. Mental health and MSK conditions account for the majority of the output boost in the improved-health scenario. This is because the incidence rates of these conditions are significantly higher than those of other chronic diseases. Although the immediate effects of the onset of mental health and MSK conditions on labor force participation are less severe than those for cancer, CVD, or diabetes (see Table 4), the larger incidence of mental health and MSK conditions results in a greater total impact. Furthermore, the higher incidence of mental health and MSK conditions at earlier ages means that intervening early boosts future labor force participation over many years and enhances labor force participation at ages when workers are most productive (see Fig. 6).

Reduced disease incidence lowers the likelihood that workers of any age will transition from employment to inactivity, leading to increases in the aggregate labor force participation rate. Fig. 13 shows how this effect varies with age and how it builds over time, demonstrating that the labor force participation rate is hump-shaped across age and peaks between the ages of 50 and 65. The employment effects of disease reduction are determined by the interaction between disease incidence and labor force participation across age groups. Prior to age 50, the relatively low incidence of disease limits the impact of disease reduction, despite high labor force participation rates. Beyond age 65, the diminishing size of the active workforce constrains the potential pool of workers who could be prevented from health-related labor market exit, notwithstanding elevated disease rates. This interaction produces maximum employment effects in the 50–65 age range, where both disease incidence and labor force participation are substantial.

Compounding effects of improved transition probabilities

It is important to note that in Fig. 13, the effect on labor force participation for the 50–65 age band is nearly four times greater five years after the intervention compared to the effect size in the first year after the intervention. Even after five years, only around 60% of the full effect has been realized. The reason for this is a key compounding channel associated with preventative health measures: prevention reduces ill-health-related exits at a given age while also increasing the share of individuals employed at that age due to reduced exit rates at earlier ages. To illustrate this, recall that employment at age a and time t is given by:

$$E_t(a) = \tau_{t,E\to E}(a)\, E_{t-1}(a-1) + \tau_{t,U\to E}(a)\, U_{t-1}(a-1) + \tau_{t,N\to E}(a)\, N_{t-1}(a-1) \tag{9}$$

where $E_t(a)$ denotes the share of employed individuals at age a in period t, $U_t(a)$ and $N_t(a)$ denote the share of unemployed and non-participating individuals at age a and time t, $\tau_{t,E\to E}(a)$ is the probability that an individual who was employed at age $a-1$ in period $t-1$ remains employed at age a in period t. Similarly, $\tau_{t,U\to E}(a)$ and $\tau_{t,N\to E}(a)$ represent the probabilities that an unemployed or non-participating individual, respectively, transitions into employment.

When the incidence of chronic disease is reduced at $t = 0$, $E_1(a)$ increases solely through the higher retention rate $\tau_{t,E\to E}(a)$ applied to the predetermined labor force participation rate $E_0(a-1)$. In the next period, $E_2(a)$ benefits from both the higher retention rate and the elevated labor force participation rate $E_1(a-1)$; this rate is higher because these workers have already benefited from improved retention when aging from $a-2$ to $a-1$ in Period 1. The process continues in future periods until $E_t(a-1)$ reaches a new steady state. This simple chain of reasoning demonstrates why a small improvement in the probability of remaining employed across the age distribution can have a significant effect on labor force participation rates several periods into the future.

It is also worth noting that the gains in labor force participation shown in Fig. 13 cannot be achieved simply by targeting preventative health measures at workers aged 50–65, even though that group exhibits the largest labor force participation response. While the labor force participation impact is indeed greatest for the 50–65 age band, these gains would be substantially lower if prevention were confined solely to older workers. In fact, the full benefits of a reduction in chronic disease incidence arise only when prevention is also targeted at younger workers. Early intervention improves health at younger ages and increases the probability that they remain employed, giving rise to the compounding effect described above. By the time these workers reach ages 50–65, the earlier improvements in labor force participation have accumulated, leading to much larger overall gains.

The impact on health

So far the analysis has focused only on the increases in economic growth arising from better health. But from a welfare perspective, health is extremely valuable in its own right beyond any impact on earnings capacity (see Topel and Murphy, 2006). Table 6 shows how

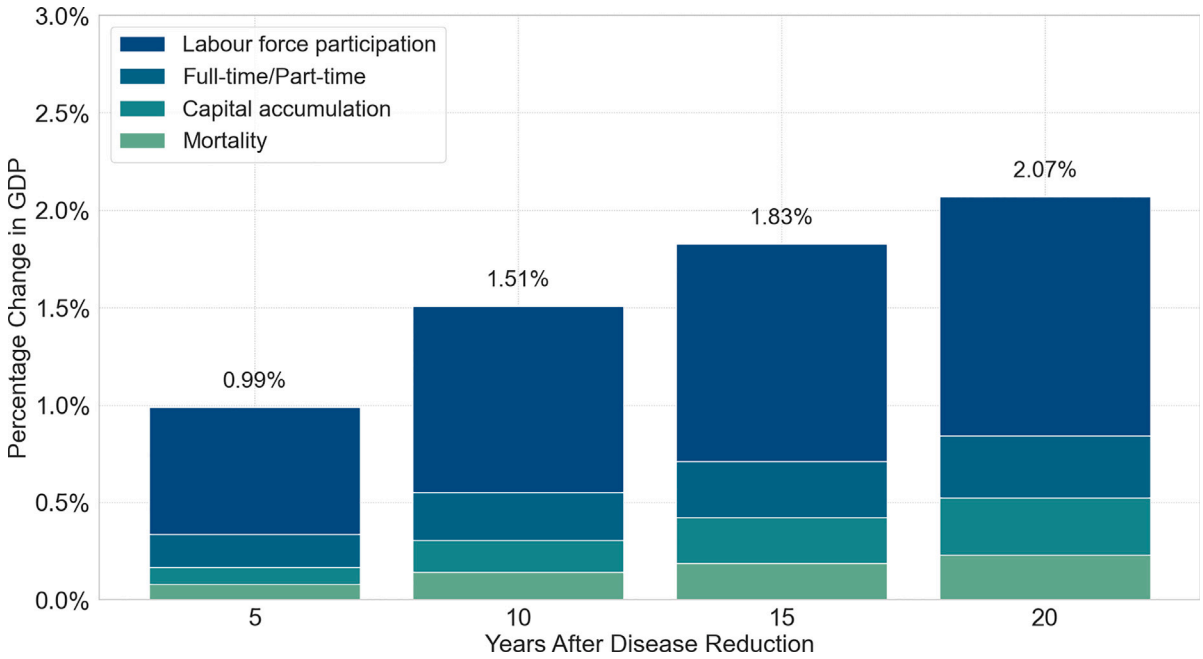

Fig. 11. Effect of 20% reduction in chronic diseases on UK output.
Note: This figure shows the effect on the model's aggregate output of a 20% reduction in disease incidence across six major chronic diseases and all age groups. The effects are decomposed into four main channels. To apportion the total effect to each of these channels, we run the healthy scenario simulation four times, activating an additional channel in each iteration.

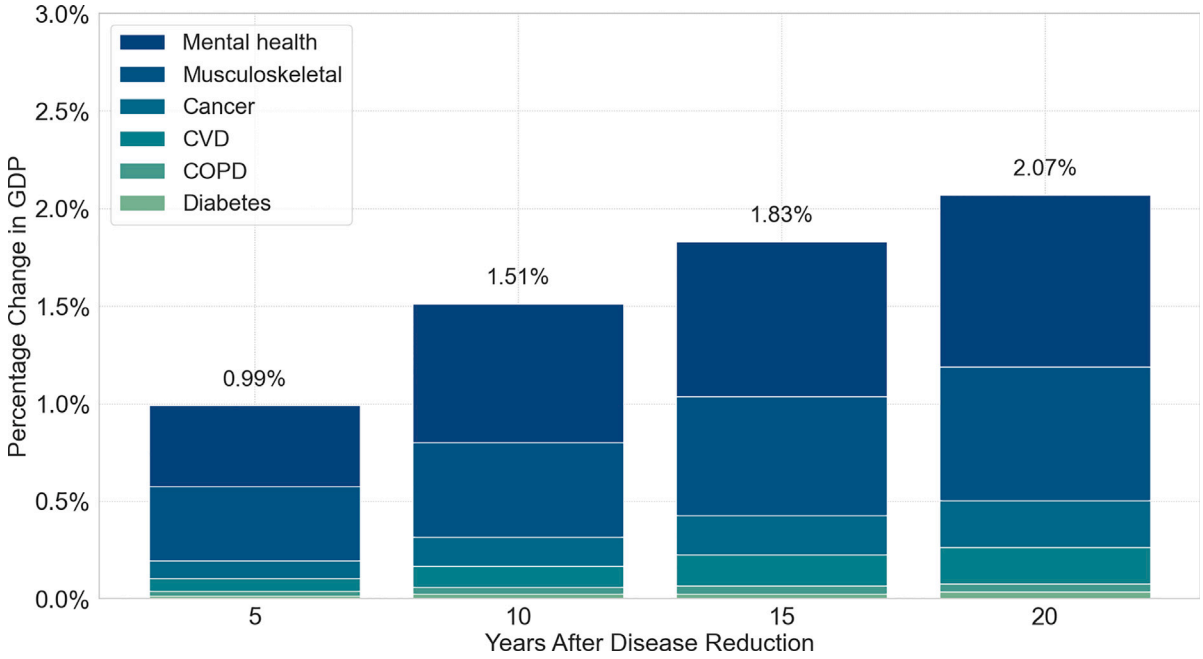

Fig. 12. Effect of 20% reduction in chronic diseases on UK output (by disease).
Note: This figure illustrates the cumulative effect on GDP at 5, 10, 15, and 20 years following a 20% reduction in the incidence of six major chronic diseases. The stacked bars represent the contribution of each disease to the total output effect at different time horizons after the reduction in disease incidence has taken effect. The reduction in disease incidence is assumed to occur instantaneously at Year 0.

various dimensions of health and longevity are affected by our hypothetical 20% reduction in chronic disease prevalence.[6]

Table 6 shows that, alongside improvements in economic output, there are substantial gains in a variety of health measures. Most notably, there is an increase in life expectancy of 1.23 years and a rise in healthy life expectancy of 2.13 years, resulting in a compression of morbidity. As emphasized in Goldman et al. (2013) and Scott et al. (2021), there are substantial welfare gains from raising life expectancy; however, those from achieving a compression of morbidity are even greater. Fig. 15 shows how a 20% reduction in chronic disease would

[6] In this section we base our analysis on the long-run case when a 20% reduction in disease incidence has fed through into a 20% reduction in prevalence.

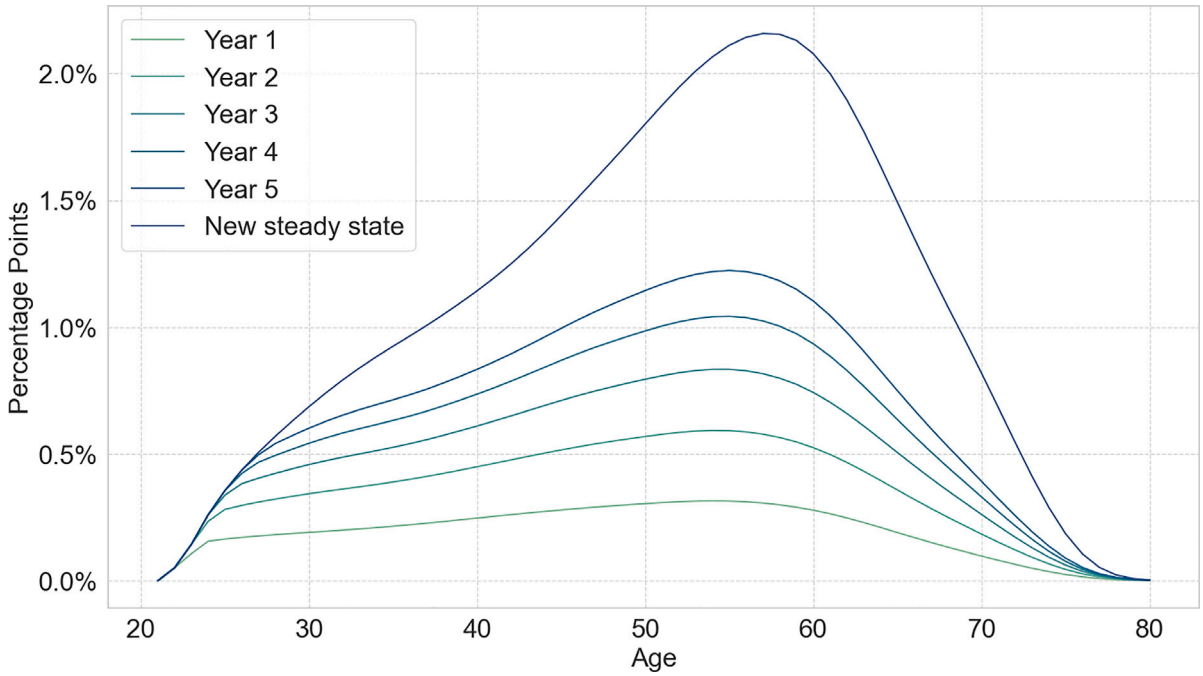

Fig. 13. Effect of 20% reduction in chronic diseases on labor force participation rate (by age).
Note: This figure shows the effect on the model's labor force participation by age following a 20% reduction in the incidence of six major chronic diseases (CVD, cancer, COPD, diabetes, mental health conditions, and MSK conditions) across age. The colored lines represent the percentage point increase in the age-specific labor force participation rate at different years after implementing the disease reduction (i.e., the Year 5 line corresponds to five years after the 20% reduction in disease incidence has gone into effect). The disease incidence reduction is assumed to occur instantaneously at Year 0.

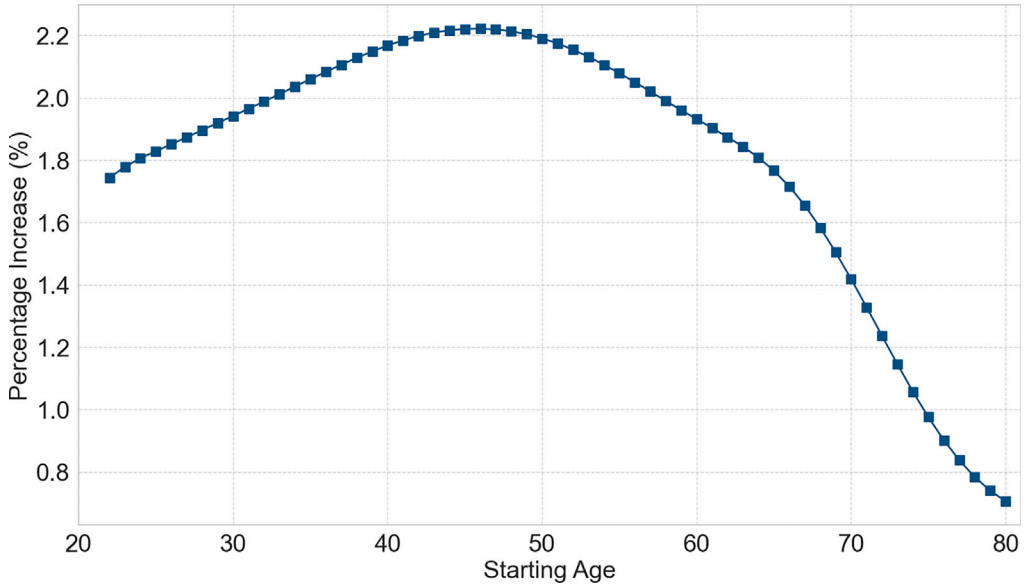

Fig. 14. Percentage increase in expected remaining working life in healthy scenario (by starting age).
Note: This figure shows the percentage increase in expected remaining working life by age following a 20% reduction in the incidence of six chronic diseases (CVD, cancer, COPD, diabetes, mental health conditions, and MSK conditions). Expected working life is calculated as the expected number of years an individual will spend in employment from their current age onward, taking into account age-specific probabilities of transitioning between employment, unemployment, and inactivity. The calculation assumes individuals are employed at the time of the disease reduction. The *y*-axis shows the percentage increase in expected working life relative to the baseline scenario without disease reduction. Appendix "Calculating expected working life" provides additional details on how the expected remaining working life for each age is calculated.

affect the Years Lived With Disabilityy (YLD) rate at different ages in the UK.

Furthermore, as the burden of chronic disease is reduced and fewer workers are forced into inactivity due to illness, average working lives increase. According to the model, a 20% reduction in the incidence of chronic diseases extends expected working life by 0.66 years. However, this extension is predicated on workers beginning their careers after the disease reduction has been implemented, thereby benefiting from lower morbidity risks throughout their entire working lifecycle. As illustrated in Fig. 14, the magnitude of this benefit exhibits substantial

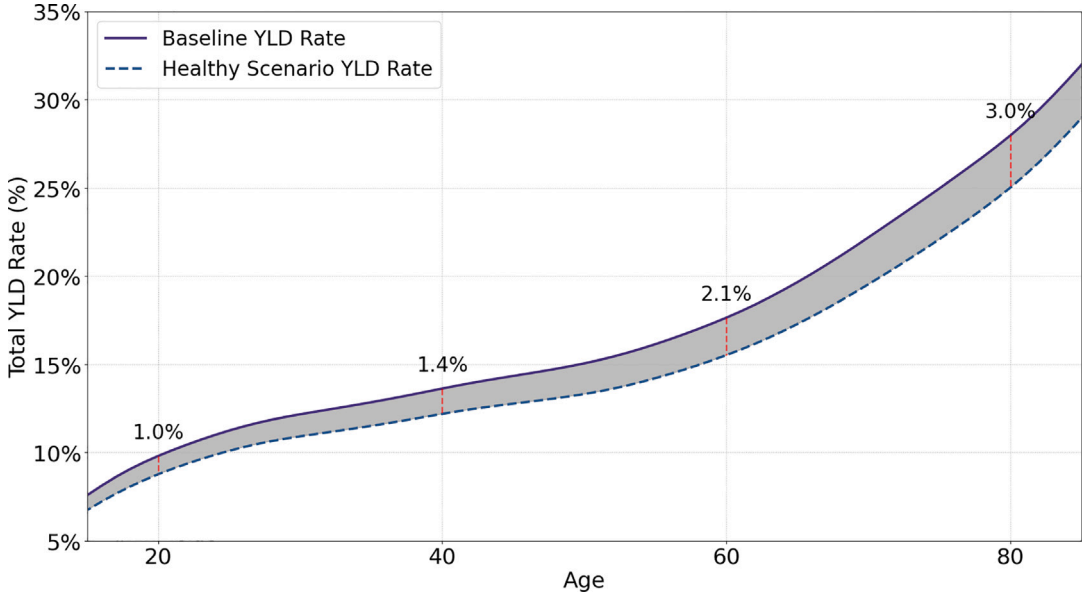

Fig. 15. Reduction in Life Years Lived With Disability (YLD) rate.
Note: This figure shows the baseline and healthy scenario total Years Lived With Disability (YLD) rates across age, illustrating the effect of a 20% reduction in our selected chronic diseases on the proportion of life years lived with disability. The solid line represents the baseline YLD rate, while the dashed line represents the healthy scenario YLD rate after the chronic disease reduction. Vertical lines at ages 20, 40, 60, and 80 indicate the percentage point reduction in YLD rates at those specific ages.

heterogeneity across age cohorts, with peak gains of approximately 2.2% for workers in their mid-forties. The observed hump shape reflects the interplay between mechanical and behavioral factors: the shortening of remaining working life as workers age initially amplifies the percentage impact of additional working years, but this effect is eventually overwhelmed by the reduced importance of health interventions among older cohorts, where workforce detachment has already occurred through standard retirement processes. This age-dependent variation in benefits underscores the differential economic impacts of health interventions across the workforce.

Table 6 also shows the improvement in the total years of life lost to death and disease across the whole population and over time (DALY). In the UK, the National Institute for Health and Care Excellence (NICE) assesses health expenditure as cost-efficient if a Quality Adjusted Life Year (QALY) can be achieved for less than £30,000. Based on this threshold—and assuming an approximate equivalence of DALY and QALY measures of life years—the value of the improvements in health from reducing the incidence of chronic diseases by 20% is worth £433 billion after five years and £1.58 trillion after ten years. Given that our model projects a cumulative increase in GDP over five years of £64 billion and over ten years of £254 billion, it is clear that health improvements on their own are the main source of welfare gains. Nevertheless, the main focus of this paper has been on the GDP gains from improved health because these gains provide a funding mechanism for preventative care.

Table 6
Change in health metrics.

Health metric	Change due to healthy scenario
Life expectancy	+1.23 years
Healthy life expectancy	+2.13 years
Expected working life at age 22	+0.66 years
Average reduction in YLD across all ages	−10.09%
Cumulative DALYs Saved:	
After 5 Years	14 million (£432.9 bn)
After 10 Years	53 million (£1,575.9 bn)
After 20 Years	202 million (£6,069.1 bn)
After 30 Years	453 million (£13,595.0 bn)
After 40 Years	807 million (£24,213.5 bn)

Note: The impact on life expectancy is obtained by first calculating the proportion of deaths caused by chronic diseases at different ages, then determining the overall reduction in mortality when that proportion is decreased by 20% across all age groups. Healthy life expectancy is calculated by first calculating the share of the population living without chronic disease at each age and then using these shares to calculate the number of years a newly born person can expect to live without chronic disease. Working lives are calculated through an iterative process using the approach outlined in Appendix A. YLD stands for Years Lived with Disability and YLD rates are obtained from the Global Burden of Disease dataset (IHME, 2021). The healthy scenario YLD rate is calculated by first determining the share of YLD due to chronic diseases at each age, then reducing that share by 20% across all age groups. DALYs are the sum of YLD and YLL where YLL stands for Years of Life Lost. YLLs are calculated as the number of deaths at a given age multiplied by remaining life expectancy at that age, summed across all ages. The value of a single DALY is taken from the National Institute for Health and Care Excellence (NICE) QALY threshold of £30,000 (National Institute for Health and Care Excellenc, 2022).

Policy discussion

Our results suggest a misalignment between the economic benefits of prevention and current fiscal policy constraints. While preventative health investments deliver substantial economic growth benefits, only around 40% of these benefits accrue within five years. This creates a challenge under the UK's current fiscal framework, where the government must demonstrate that the current budget will balance by 2029/30 and that net financial debt will fall as a share of GDP in that year.

Several institutional mechanisms could help align political incentives with these valuable long-term health investments. Outcome-based contracting offers one promising approach. Social Impact Bond structures, for instance, allow governments to transfer the upfront costs and risks of prevention programs to private or philanthropic investors, who are repaid only when defined health improvements are achieved.[7]

[7] See, for example, Tortorice et al. (2020) for more background and examples of Social Impact Bonds.

This model reduces the short-term budgetary burden while maintaining accountability for results. Additionally, dedicated prevention funds or "Health Investment Trusts" can provide stable, ring-fenced financing that transcends individual political terms. Cross-party agreements to maintain multi-year prevention budgets could also be effective, as they make it politically costly for successive administrations to dismantle existing prevention commitments. A target for the health system to achieve specific gains in healthy life expectancy would also help with a focus on prevention and offset the demands of fiscal policy.

The framing of prevention programs also matters for political feasibility. While our model shows 40% of long-term GDP impacts emerging in the first five years of disease reduction, prevention initiatives can be designed to deliver observable improvements within a single political term. Early indicators like reduced emergency room visits or declining obesity rates allow implementing administrations to demonstrate progress, even if the full economic benefits emerge later. Starting with local pilot programs that generate quick feedback on effectiveness can build the evidence base and political momentum for broader implementation.

The magnitude of potential gains from reduced chronic disease underscores the importance of finding workable solutions to these implementation challenges. The challenge lies less in demonstrating the value of prevention but in designing institutional frameworks that can sustain commitment to these investments over appropriate time horizons.

Limitations

A limitation of our analysis is that we do not account for disease comorbidities. Chronic conditions often cluster within individuals, and their combined effect on labor market outcomes may exceed the sum of their independent effects. While explicitly modeling disease interactions would be valuable, data limitations and computational complexity make this challenging in our current framework.

Another limitation of our analysis is that we do not fully capture general equilibrium effects that could arise from large-scale improvements in population health. In practice, significant increases in labor supply might affect wage levels, which could in turn influence both labor force participation decisions and firms' input choices. Additionally, structural changes in healthcare demand could affect the sectoral composition of employment and output. Our model also does not capture how improved life expectancy might change individuals' labor force participation decisions across their lifecycle—for instance, people might choose to extend their careers given expectations of a longer, healthier life. These adjustment channels are beyond the scope of our current modeling framework, which focuses on capturing the first-order effects of improved population health on aggregate labor supply and the induced capital accumulation response.

Conclusion

Our results show that substantial gains in UK output can be realized within five years of achieving reductions in the incidence of chronic disease. We find that these gains primarily stem from keeping people in full-time employment, with the most significant conditions driving this result for the UK being mental health and musculoskeletal (MSK) conditions. In addition to providing a substantial boost to output, we quantify how reductions in chronic disease incidence lead to an increase in life expectancy and an even greater increase in healthy life expectancy, resulting in a compression of morbidity.

The largest labor force participation effects occur among workers aged 50–65; yet, achieving these gains requires prevention across the entire age distribution of workers rather than targeting older workers alone. Early intervention creates two reinforcing channels: it improves the probability of remaining employed at each age (through higher retention rates) while simultaneously increasing the share of individuals employed at that age (through accumulated retention from previous periods). These channels compound over time, generating substantially larger participation gains at older ages than would be possible through late-life intervention alone.

This compounding channel is key to the magnitude of the impact but has the important implication that the majority of economic gains only materialize in the long run. This issue is exacerbated by our focus on the economic impacts arising from successfully reducing incidence. There will be further lags introduced due to the gap between interventions and their impact on disease incidence. These lags will act as a deterrent to governments with medium-term fiscal constraints, such as the UK.

Our analysis demonstrates the substantial macroeconomic potential of chronic disease prevention in the context of population aging. While we find significant output gains within five years, primarily through increased full-time employment retention, the compounding nature of prevention creates larger long-term effects. Although implementation lags and medium-term fiscal constraints may pose political challenges, our results indicate that successful prevention could significantly offset the economic pressures of population aging. While important questions remain regarding specific interventions, their costs, financing mechanisms, and institutional requirements, we view our analysis as evidence that preventative health investment represents a meaningful policy lever for addressing demographic challenges to economic growth.

CRediT authorship contribution statement

Yannick Schindler: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Project administration, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Andrew J. Scott:** Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Resources, Software, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Calculating expected working life

Calculating expected working lives at any age is done using recursion. In the model, individuals can be in one of three states at any age t: Employed (E), Unemployed (U), or Not in the Labor Force (N). The transition probabilities between these states are denoted as follows:

$$P_{ij}(t)$$

where $i, j \in \{E, U, N\}$ represent the states an individual transitions between from age t to $t+1$. For example, $P_{EE}(t)$ represents the probability of remaining employed, while $P_{EU}(t)$ represents the probability of transitioning from employed to unemployed.

Let $V_E(t)$ denote the expected total years of employment from age t onwards, given that an individual is employed at age t (i.e. expected working life). Similarly, $V_U(t)$ and $V_N(t)$ denote the expected years of employment starting in the unemployed and not in the labor force states, respectively. These expected working life values are defined recursively as follows.

For individuals starting in the employed state at age t:

$$V_E(t) = 1 + P_{EE}(t) \cdot V_E(t+1) + P_{EU}(t) \cdot V_U(t+1) + P_{EN}(t) \cdot V_N(t+1)$$

For individuals starting in the unemployed state at age t:

$$V_U(t) = P_{UE}(t) \cdot V_E(t+1) + P_{UU}(t) \cdot V_U(t+1) + P_{UN}(t) \cdot V_N(t+1)$$

For individuals starting in the not in labor force state at age t:

$$V_N(t) = P_{NE}(t) \cdot V_E(t+1) + P_{NU}(t) \cdot V_U(t+1) + P_{NN}(t) \cdot V_N(t+1)$$

The recursive process starts at the maximum age $T = 100$, where it is assumed that no further working life remains, i.e.,

$$V_E(101) = V_U(101) = V_N(101) = 0$$

With the recursive relationships defined above and this boundary condition, $V_E(t)$ can be obtained for any $t \in [0, 100]$.

Appendix B. Validation of LLM-based healthcare cost classification

This appendix provides details of our validation of the Large Language Model (LLM) classification pipeline, which we used to categorize Healthcare Resource Groups (HRGs) from the NHS National Cost Collection (NCC) data into six major disease categories. The validation exercise addresses whether the LLM-based classification produces reliable disease categorizations when compared to expert human judgment.

Validation methodology: To validate our classification approach, we performed a systematic evaluation using a random sample of 450 HRG codes drawn from the LLM-classified records set in the NCC dataset. This sample represents approximately 16% of the total of 2825 HRG codes.

The validation process involved two primary steps. First, we categorized each of the 450 randomly drawn HRG codes into the six chronic disease categories. This categorization was then reviewed by two medical experts: Dr. Roshni Joshi, who holds a PhD in Genetic Epidemiology of Cardiovascular Disease from UCL and has extensive experience in health technology assessment at NICE, and Dr. Naomi Lee, who holds an MD from King's College London and has served as Senior Executive Editor at The Lancet. In case of disagreement, consensus was reached through discussion. These manual classifications served as the reference standard against which we evaluated the performance of our LLM-based classification pipeline.

Second, we computed a comprehensive set of classification metrics comparing the LLM's existing categorizations against our manual classifications. These metrics include precision, recall, F1 score, accuracy, Cohen's Kappa, and Matthews Correlation Coefficient (MCC). Precision measures the proportion of HRGs that were actually disease-related among all HRGs that the LLM classified as belonging to that disease. For example, if the LLM classified 100 HRGs as CVD and 97 of these were confirmed as CVD by the medical experts, the precision would be 0.97. Recall measures the proportion of disease-related HRGs that the LLM successfully identified out of all HRGs that the medical experts classified as belonging to that disease. For instance, if the medical experts identified 100 HRGs as CVD and the LLM correctly identified 67 of these, the recall would be 0.67. The F1 score provides a balanced measure of precision and recall. Accuracy represents the overall proportion of correct classifications. Cohen's Kappa quantifies the agreement between human and LLM classifications while accounting for chance agreement. The Matthews Correlation Coefficient offers a balanced measure of classification quality that is particularly useful for categories with uneven sizes.

Classification performance: Details for the various performance metrics across disease categories are given in Table .7. The validation results demonstrate strong overall performance of the LLM classification pipeline, with a macro-average F1 score of 0.81 across all disease categories.

The relationship between our classification metrics and cost estimation has important implications for interpreting our results. Low precision indicates that the LLM is incorrectly classifying HRGs as belonging to a disease category when they do not (false positives), leading to potential cost overestimation. Conversely, low recall indicates that the LLM is missing HRGs that should be classified in that category (false negatives), suggesting potential cost underestimation.

For CVD, the high precision (0.97) but lower recall (0.67) suggests that our cost estimates are likely conservative, potentially underestimating the true economic burden of CVD. COPD shows balanced performance, with both precision (0.79) and recall (0.88) above 0.75, indicating relatively reliable cost estimates. Cancer similarly demonstrates balanced metrics, with precision at 0.76 and recall at 0.89, suggesting reasonably accurate cost estimation. Mental health presents an interesting case where high recall (1.00) but lower precision (0.72) indicates that we may be overestimating costs by incorrectly including non-mental health HRGs. For musculoskeletal conditions, both precision (0.77) and recall (0.68) are lower, with the low recall suggesting potential cost underestimation. Diabetes shows the strongest overall performance, with perfect precision (1.00) and high recall (0.80), indicating particularly reliable cost estimates; however, the small number of diabetes-related HRGs in our sample suggests that these results should be interpreted cautiously.

Appendix C. Event study analysis with placebo tests

To validate our empirical strategy we conduct an additional analysis consisting of placebo tests using an event study framework.

First, we balance our panel as done for the main analysis described in Section "Estimating the impact of disease on labor force participation". We create cells based on time-invariant characteristics by computing each individual's median age across all observations, assigning them to five-year age bands, and categorizing them by their position in the earnings distribution when employed. Specifically, we assign individuals to earnings deciles if they are observed in employment at least 50% of the time, with a separate category for those who are predominantly inactive. These categorizations combine to create balancing cells defined by the interaction of age band and earnings group.

Within each balancing cell, we ensure that we have an equal number of treated individuals (those who report a diagnosis and have at least three years of post-diagnosis data) and control individuals who never receive a diagnosis. We require treated individuals to be economically active in the wave immediately preceding their diagnosis to focus our analysis on labor market transitions among the working population. When a cell contains more potential controls than treated individuals, we randomly sample from the control pool to achieve equal numbers. In cases with fewer potential controls than treated individuals, we either use control individuals multiple times or reduce the number of treated individuals included in the analysis.

For these balanced treatment and control groups, we then implement our placebo test methodology. We first identify the empirical distribution of diagnosis timing in our sample of treated individuals. We then randomly assign placebo "diagnosis" dates to our control individuals, drawing from this empirical distribution of actual diagnosis timing. This approach ensures that the temporal distribution of placebo events matches that of actual disease diagnoses.

For both treated and control groups, we estimate the following event study specification:

$$y_{i,t} = \delta_t + \sum_{k=-4}^{3} \beta_k D_{i,t}^k + \gamma X_{i,t} + \epsilon_{i,t} \tag{10}$$

where $y_{i,t}$ is an indicator for transitioning from employment to inactivity, δ_t represents time fixed effects, $D_{i,t}^k$ denotes indicator variables for being k periods away from diagnosis (or placebo diagnosis), $X_{i,t}$ is a vector of controls for age and age squared, and $\epsilon_{i,t}$ is the error term. The coefficients of interest are the β_k parameters, which trace out the dynamic response of labor market transitions relative to the diagnosis event. We normalize $\beta_{-1} = 0$, making all effects relative to the period immediately preceding diagnosis.

We restrict our sample to individuals observed for at least seven consecutive waves and with event times between −4 and +3 years relative to diagnosis. Standard errors are clustered at both the individual

Table .7
Classification metrics by disease.

Disease	Precision	Recall	F1	Accuracy	Cohen's Kappa	MCC	% of rows (Manual)	% of rows (LLM)
Cardiovascular Disease	0.97	0.67	0.79	0.97	0.77	0.79	9.33%	6.44%
Chronic Obstructive Pulmonary Disease	0.79	0.88	0.83	0.99	0.83	0.83	3.78%	4.22%
Cancer	0.76	0.89	0.82	0.97	0.80	0.80	8.44%	10.00%
Mental Health	0.72	1.00	0.84	0.97	0.82	0.83	7.33%	10.22%
Musculoskeletal	0.77	0.68	0.72	0.97	0.71	0.71	5.56%	4.89%
Diabetes	1.00	0.80	0.89	1.00	0.89	0.89	1.11%	0.89%
Macro Average	0.83	0.82	0.81	0.98	0.80	0.81	–	–

Note: This table reports the per-disease Precision, Recall, F1, Accuracy, Cohen's Kappa, and Matthews Correlation Coefficient (MCC), as well as the percentage of rows allocated to each disease by manual review and by the LLM. The analysis was conducted on a random sample of 450 rows from the NHS National Cost Collection data. Each row was manually labeled and then classified by the LLM into six categories: Cardiovascular Disease (CVD), Chronic Obstructive Pulmonary Disease (COPD), Cancer, Mental Health Condition, Musculoskeletal (MSK) Condition, and Diabetes. Cohen's Kappa is a measure of agreement between two raters, accounting for chance. The Matthews Correlation Coefficient (MCC) is a correlation coefficient for binary classifications that takes into account true and false positives and negatives.

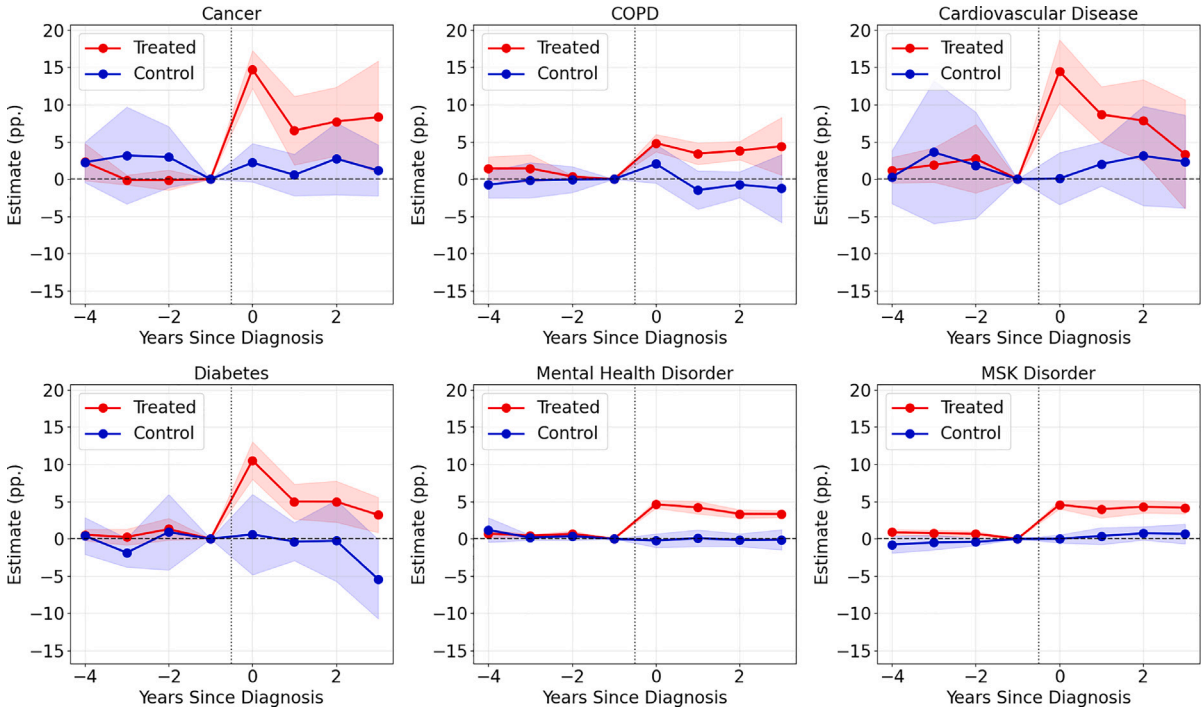

Fig. 16. Event study estimates: Treated vs Control groups.
Note: This figure presents event study estimates for the effect of disease diagnosis on transitions from employment to inactivity. The red lines show estimates for treated individuals around their actual diagnosis dates, while the blue lines show estimates for control individuals around randomly assigned placebo diagnosis dates. Shaded areas represent 95% confidence intervals. The vertical axis measures the percentage point change in the probability of transitioning to inactivity. The vertical dotted line at $t = -1$ indicates the reference period. The sample is restricted to individuals observed for at least seven consecutive waves with event times between −4 and +3 years relative to diagnosis. Standard errors are clustered at both the individual and year levels.

and year levels to account for serial correlation in outcomes within individuals and common shocks within time periods.

Fig. 16 presents the results of this analysis. For each disease category, we plot the estimated β_k coefficients separately for the treated group (red) and control group (blue), along with associated 95% confidence intervals. For all six disease categories, we observe relatively flat pre-trends in both the treated and control groups, with coefficients statistically indistinguishable from zero in the pre-diagnosis period. At the time of diagnosis, we observe sharp increases in transition probabilities for treated individuals, while control individuals with placebo diagnoses show no significant changes in their transition patterns.

These placebo tests provide additional evidence that our estimated effects reflect genuine responses to disease diagnosis rather than underlying differences between treated and control individuals or general trends in labor market transitions.

Appendix D. Calculation of counterfactual transition rates

We can use the estimates obtained in the previous section to infer how transition rates between employment and inactivity for a given age a would change in response to a reduction of δ_j percent in disease incidence for disease j. The total transition rate T_a is a weighted average of the transition rates for healthy individuals and those with each disease, namely:

$$\begin{aligned} T_a &= t_a^H \left(1 - \sum_{j \in J} p_{j,a}\right) + \sum_{j \in J} \left(t_a^H + \Delta t_j\right) p_{j,a} \\ &= t_a^H + \sum_{j \in J} \Delta t_j p_{j,a} \end{aligned} \tag{11}$$

where T_a denotes the total observed transition rate from employment to inactivity at age a, t_a^H denotes the transition rate for healthy individuals (those without any of the six chronic diseases) at age a,

J represent the set of six chronic diseases under consideration, $p_{j,a}$ denotes the prevalence of disease $j \in J$ at age a, and Δt_j denotes the additional probability—in percentage points—that an individual with disease j transitions from employment to inactivity, compared to a healthy individual.

Assume a reduction in the incidence of disease j by a proportion δ_j (e.g., a 20% reduction implies $\delta_j = 0.20$). The new prevalence of disease j at age a becomes $p'_{j,a} = (1 - \delta_j)p_{j,a}$.

The counterfactual total transition rate T'_a after the reduction is:

$$T'_a = t_a^H + \sum_{j \in J} \Delta t_j (1 - \delta_j) p_{j,a}$$

which implies:

$$T'_a - T_a = -\sum_{j \in J} \Delta t_j \delta_j p_{j,a} \tag{12}$$

The above has the simple interpretation that the aggregate transition rate T_a is reduced by $\Delta t_j p_{j,a} \delta_j$ for a δ_j percentage reduction in the incidence of disease j. This means the decrease in the transition rate is directly proportional to (1) the additional transition risk Δt_j associated with each disease j, (2) the proportionate reduction δ_j in disease incidence, and (3) the initial incidence $p_{j,a}$ of a given disease at age a.

References

Aggarwal, Ajay, Sullivan, Richard, 2014. Affordability of cancer care in the United Kingdom – Is it time to introduce user charges? J. Cancer Policy 2 (2), 31–39.

Aghion, Philippe, Howitt, Peter, Murtin, Fabrice, 2010. The relationship between health and growth: When Lucas meets Nelson-Phelps. Rev. Econ. Institutions 2 (1).

Amies-Cull, Ben, Luengo-Fernandez, Ramon, Scarborough, Peter, Wolstenholme, Jane, 2023. NHS reference costs: a history and cautionary note. Heal. Econ. Rev. 13, 54.

Ara Darzi, 2024. Independent investigation of the NHS in England.

Arora, Suchit, 2001. Health, human productivity, and long-term economic growth. J. Econ. Hist. 61 (3), 699–749, Publisher: Cambridge University Press.

Athey, Susan, Imbens, Guido W., 2022. Design-based analysis in difference-in-differences settings with staggered adoption. J. Econometrics 226 (1), 62–79.

Barro, Robert J., 1996. Determinants of economic growth: A cross-country empirical study. NBER working paper.

Barro, Robert J., 2013. Health and economic growth. Ann. Econ. Financ..

Bhargava, Alok, Jamison, Dean T, Lau, Lawrence J, Murray, Christopher J L, 2001. Modeling the effects of health on economic growth. J. Heal. Econ..

Bloom, David E., Cafiero-Fonseca, Elizabeth T., McGovern, Mark E., Prettner, Klaus, Stanciole, Anderson, Weiss, Jonathan, Bakkila, Samuel, Rosenberg, Larry, 2014a. The macroeconomic impact of non-communicable diseases in China and India: Estimates, projections, and comparisons. J. Econ. Ageing 4, 100–111.

Bloom, David E., Canning, David, Fink, Günther, 2014b. Disease and development revisited. J. Political Econ. 122 (6).

Bloom, David E., Canning, David, Sevilla, Jaypee, 2004. The effect of health on economic growth: A production function approach. World Dev. 32 (1), 1–13.

Bloom, David E., Chen, Simiao, Kuhn, Michael, McGovern, Mark E., Oxley, Les, Prettner, Klaus, 2020. The economic burden of chronic diseases: Estimates and projections for China, Japan, and south Korea. J. Econ. Ageing 17, 100163.

Bloom, David, Chen, Simiao, McGovern, Mark, 2018. The economic burden of noncommunicable diseases and mental health conditions: results for Costa Rica, Jamaica, and Peru. Rev. Panam. Salud Pública 1–8.

Borusyak, Kirill, Jaravel, Xavier, Spiess, Jann, 2024. Revisiting event study designs: Robust and efficient estimation. Rev. Econ. Stud..

Bound, John, Schoenbaum, Michael, Stinebrickner, Todd R, Waidmann, Timothy, 1999. The dynamic effects of health on the labor force transitions of older workers. Labour Econ..

Cai, Lixin, Kalb, Guyonne, 2006. Health status and labour force participation: evidence from Australia. Heal. Econ. 15 (3), 241–261, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/hec.1053.

Callaway, Brantly, Sant'Anna, Pedro H.C., 2021. Difference-in-differences with multiple time periods. J. Econometrics 225 (2), 200–230.

Chen, Simiao, Bloom, David E., 2019. The macroeconomic burden of noncommunicable diseases associated with air pollution in China. PLoS ONE 14 (4), e0215663.

Chen, Simiao, Kuhn, Michael, Prettner, Klaus, Bloom, David E., 2018. The macroeconomic burden of noncommunicable diseases in the United States: Estimates and projections. PLOS ONE 13 (11), e0206702. http://dx.doi.org/10.1371/journal.pone.0206702, URL https://dx.plos.org/10.1371/journal.pone.0206702.

Cholesterol Treatment Trialists (CTT) Collaborators, 2012. The effects of lowering LDL cholesterol with statin therapy in people at low risk of vascular disease: meta-analysis of individual data from 27 randomised trials. Lancet 380 (9841), 581–590.

Christensen, Bent Jesper, Kallestrup-Lamb, Malene, 2012. The impact of health changes on labor supply: Evidence from merged data on individual objective medical diagnosis codes and early retirement behavior. Heal. Econ. 21 (S1), 56–100.

Clark, David M., 2018. Realising the mass public benefit of evidence-based psychological therapies: The IAPT program. Annu. Rev. Clin. Psychol. 14, 159–183.

de Chaisemartin, Clément, D'Haultfœuille, Xavier, 2020. Two-way fixed effects estimators with heterogeneous treatment effects. Am. Econ. Rev. 110 (9), 2964–2996.

Department of Health & Social Care, 2023. Progress in improving mental health services in England. Discussion paper, National Audit Office.

Disney, Richard, Emmerson, Carl, Wakefield, Matthew, 2006. Ill health and retirement in Britain: A panel data-based analysis. J. Heal. Econ. 25 (4), 621–649.

Dube, Arindrajit, Girardi, Daniele, Jordà, Òscar, Taylor, Alan, 2024. A local projections approach to difference-in-differences. p. w31184, NBER working paper.

Friedson, Andrew I., Li, Moyan, Meckel, Katherine, Rees, Daniel I., Sacks, Daniel W., 2021. Cigarette taxes, smoking, and health in the long-run. NBER working paper.

García-Gómez, Pilar, 2011. Institutions, health shocks and labour market outcomes across Europe. J. Heal. Econ. 30 (1), 200–213.

García-Gómez, Pilar, Jones, Andrew M., Rice, Nigel, 2010. Health effects on labour market exits and entries. Labour Econ. 17 (1), 62–76.

GBD 2019 Diseases and Injuries Collaborators, 2020. Global burden of 369 diseases and injuries in 204 countries and territories, 1990–2019: a systematic analysis for the global burden of disease study 2019. Lancet 396.

Goldberg, David, 1988. A User's Guide to the General Health Questionnaire. NFER-NELSON.

Goldman, Dana P., Cutler, David, Rowe, John W., Michaud, Pierre-Carl, Sullivan, Jeffrey, Peneva, Desi, Olshansky, S. Jay, 2013. Substantial health and economic returns from delayed aging may warrant a new focus for medical research. Health Aff. 32 (10), 1698–1705.

Goodhart, Charles, Pradhan, Manoj, 2020. The Great Demographic Reversal : Ageing Societies, Waning Inequality, and an Inflation Revival. Palgrave Macmillan.

Goodman-Bacon, Andrew, 2021. Difference-in-differences with variation in treatment timing. J. Econometrics 225 (2), 254–277.

Hagan, Ronald, Jones, Andrew M., Rice, Nigel, 2009. Health and retirement in Europe. Int. J. Environ. Res. Public Heal. 6 (10), 2676–2695.

Haskel, Jonathan, Martin, Josh, 2022. Economic inactivity and the labour market experience of the long-term sick. working paper.

Hex, Nick, MacDonald, Rachael, Pocock, Jessica, Uzdzinska, Barbara, Taylor, Matthew, Atkin, Marc, Wild, Sarah H., Beba, Hannah, Jones, Ross, 2024. Estimation of the direct health and indirect societal costs of diabetes in the UK using a cost of illness model. Diabetic Med. 41 (9), e15326.

IHME, 2021. Global burden of disease (GBD).

Jones, Andrew M., Rice, Nigel, Roberts, Jennifer, 2010. Sick of work or too sick to work? Evidence on self-reported health shocks and early retirement from the BHPS. Econ. Model. 27 (4), 866–880.

Lee, Ronald, Mason, Andrew, 2017. Cost of aging. Discussion paper, IMF.

McKinsey, 2020. Prioritizing health. Discussion paper, McKinsey Global Institute.

National Institute for Health and Care Excellenc, 2022. NICE health technology evaluations: the manual. Publisher: NICE.

NHS England, 2024a. NHS England >> respiratory disease.

OBR, 2024a. Economic and fiscal outlook (March 2024). Discussion paper, Office for Budget Responsibility.

OBR, 2024b. Fiscal risks and sustainability (September 2024). Discussion paper, Office for Budget Responsibility.

Office for Health Improvement & Disparities, 2022. Musculoskeletal health: applying all our health.

Old, Jonathan, Scott, Andrew, 2023. Healthy ageing trends in England between 2002 to 2018: Improving but slowing and unequal. J. Econ. Ageing 26, 100470.

ONS, 2017. Net capital stock to output (ratio): S.1: Asset TOTAL: Ind: F: CP - Office for national statistics.

ONS, 2020. Numbers surviving at exact age (lx), principal projection, UK.

ONS, 2021. Estimates of the population for the UK, England, Wales, Scotland, and Northern Ireland.

ONS, 2022a. Labour costs and labour income, UK. Statistical Bulletin, Office for National Statistics.

ONS, 2024. Healthcare expenditure, UK health accounts - office for national statistics.

ONS, 2024a. Gross domestic product: chained volume measures: Seasonally adjusted £m - office for national statistics.

ONS, 2024c. Labour force survey five-quarter longitudinal dataset.

ONS, 2024d. Milestones: journeying through modern life - office for national statistics.

Raleigh, Veena, Jefferies, Danielle, Wellings, Dan, 2022. Cardiovascular disease in England. Discussion paper, The King's Fund.

Scott, Andrew J., Ellison, Martin, Sinclair, David A., 2021. The economic value of targeting aging. Nat. Aging 1 (7), 616–623.

Solow, Robert M., 1956. A contribution to the theory of economic growth. Q. J. Econ. 70 (1), 65.

Sun, Liyang, Abraham, Sarah, 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. J. Econometrics 225 (2), 175–199.

Swan, T.W., 1956. Economic growth and capital accumulation. Econ. Rec. 32 (2), 334–361.

Toby Watt, Ann Raymond, Laurie Rachet-Jacquet, Anna Head, Chris Kypridemos, Elaine Kelly, Anita Charlesworth, 2023. Health in 2040: projected patterns of illness in England. Discussion paper, The Health Foundation.
Topel, Robert, Murphy, Kevin, 2006. The value of health and longevity. J. Political Econ..
Tortorice, Daniel L, Bloom, David E, Kirby, Paige, Regan, John, 2020. A theory of social impact bonds. NBER working paper.
UN, 2024. World population prospects - population division.
World Bank, 2024. World bank open data.