

Richiardi, Matteo; Bronka, Patryk; van de Ven, Justin

Working Paper

Attenuation and reinforcement mechanisms over the life course

CeMPA Working Paper Series, No. CeMPA WP 2/24

Provided in Cooperation with:

University of Essex, Centre for Microsimulation and Policy Analysis (CeMPA)

Suggested Citation: Richiardi, Matteo; Bronka, Patryk; van de Ven, Justin (2024) : Attenuation and reinforcement mechanisms over the life course, CeMPA Working Paper Series, No. CeMPA WP 2/24, University of Essex, Centre for Microsimulation and Policy Analysis (CeMPA), Colchester

This Version is available at:

<https://hdl.handle.net/10419/306907>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

CeMPA Working Paper Series

CeMPA WP 2/24

Attenuation and reinforcement mechanisms over the life course

Matteo Richiardi

Patryk Bronka

Justin van de Ven

February 2024



Attenuation and reinforcement mechanisms over the life course

Matteo Richiardi[†], Patryk Bronka[†] and Justin van de Ven[†]

Revised October 2024

Abstract

We analyse the complex dynamic feedback effects between different life domains over the life course, providing a quantification of the direct (not mediated) and indirect (mediated) effects. To extend the analysis in scope and time beyond the limitations of existing data, we use a rich dynamic microsimulation model of individual life course trajectories parameterised and validated to the UK context. We interpret findings in terms of the implied attenuation or reinforcement mechanisms at play, and discuss implications for health and economic inequalities.

Keywords: Dynamic microsimulation, feedback mechanisms, life course analysis, life events, mediation analysis.

JEL codes: C15, C63, C99, D30, R20.

[†] Centre for Microsimulation and Policy Analysis, University of Essex. Email for correspondence: matteo.richiardi@essex.ac.uk

1 Introduction

Individual life domains are highly interconnected, and events occurring in one domain generally affect other domains.¹ This interconnectedness determines whether the impact of an event gets amplified or dampened when attention is broadened from the event in isolation to the life course as a whole. In turn, the interplay between attenuation and reinforcement mechanisms determines individual resilience and vulnerability, with important bearings on economic and health inequalities. While studies abound that take into account the role of mediators in the determination of outcomes, they often lack an integrated approach that considers *all* factors under analysis as both potential mediators and outcomes, in a dynamic context.

Most studies perform mediation analysis on *observed* data. The increasing use of computational models (specifically: in the social sciences and public health) opens up new research possibilities that take advantage of the potential to experiment with interacting causal pathways. Protocols and procedures for manipulating causal pathways and analysing counterfactuals in computational models are however not yet established.

In this paper, we propose a computational procedure to isolate direct and indirect (mediated) effects in a microsimulation framework. The paper explores the feedback loops between health, family and labour market outcomes, and the associated implications for income and health inequalities over alternative time horizons. To construct the relevant counterfactuals, we use a rich dynamic microsimulation model parameterised for the United Kingdom, which projects individual life course trajectories over the three inter-related domains of work, family and health. The model is linked to an underlying tax-benefit calculator, which provides a realistic description of the impact of taxes and benefits at both the individual and population level.

The structure of the model accommodates dynamic interactions between all simulated variables. Given this underlying structure, we offer a model-based decomposition of the overall effect of specific events in terms of their *direct* effects – the un-mediated, self-reflective impacts of (changes of) one variable on the future evolution of the same variable – *cross-effects* (the impacts on the future evolution of other variables), and *indirect* effects (the mediated impacts on the same variable, coming from the impact on other variables). Our primary focus of interest is the comparison of the direct and indirect effects, which we use to construct a synthetic indicator of the complex interactions that shape individual trajectories. This in turn permits quantification of the importance of attenuation and reinforcement mechanisms over time.

As an illustration of the methodology, we consider the impact of two different exogenous events: a partnership dissolution, and a sudden health deterioration. These scenarios have been selected as two polar cases in our empirical and modelling context, with partnership status having implications that extend to all individual life domains, while health effects being more limited to the health sphere.

To the best of our knowledge, this paper is the first empirical study investigating the feedback loops between multiple life domains and their implications for income and health inequalities over the life cycle. The detailed calibration of the model to a real-world setting allows us to explore quantitatively multiple sources of inequalities, at an individual level and over multiple time horizons.

¹ Just think of the last time you came home after a frustrating day at work.

Our results indicate that partnership status has significant effects on other life domains, and highlight attenuation mechanisms that facilitate bouncing back to a partnered status following a union dissolution. On the other hand, health is found to have fewer connections to other life domains, with limited feedback that attenuate or exacerbate the effects of an adverse shock.

The structure of the paper is as follows. Section 2 frames our work in the context of life course analysis. Section 3 describes the microsimulation approach. Section 4 introduces the counterfactual analysis that underpins our decomposition. Section 5 presents our microsimulation model. Section 6 explains the two conceptual experiments. Sections 7 and 8 present the results for the two experiments in turn. Section 9 discusses implications for inequality and resilience. Section 10 summarises and concludes.

2 The impact of shocks

Figure 1 shows the possible links (direct acyclic graph, DAG) between variables of interest – for simplicity, the diagram includes only two variables at different observational times – for illustration, say socio-economic status (X_1), and health (X_2). Each variable possibly has an un-mediated impact on its future values and on the future values of the other variables, as well as mediated effects.

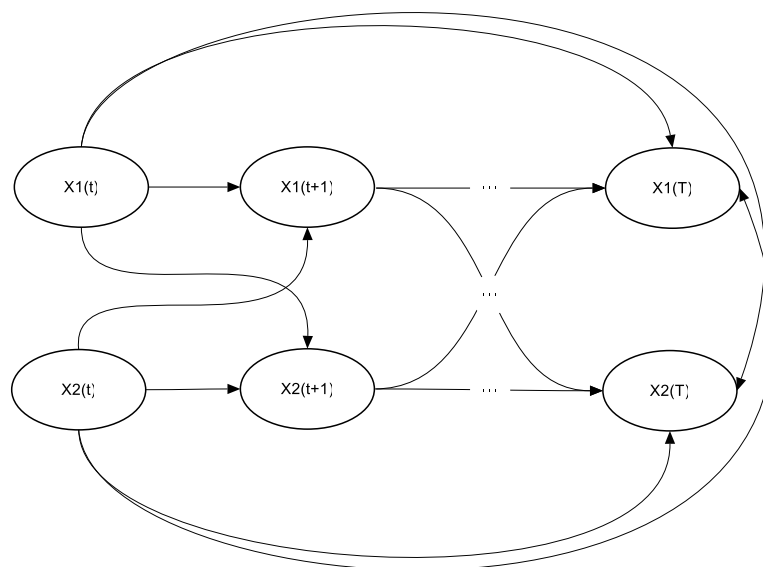


Figure 1: Dynamic determination of individual outcomes

The figure helps classifying the related literature. We focus in particular on studies that have looked at the impact of specific events – economic events, health events, family-related events – rather than specific individual characteristics. This is because events – also referred to as *shocks* – can sometimes lead to quasi-natural experiments, facilitating identification of the associated effects. Examples of shocks examined in the literature include job displacement (often following mass layoffs), acute hospital admissions (due to road accidents or strokes, for example), partnership dissolution and divorce. The literature is vast.

For example, some studies of job displacement consider the implications for prospective labour market circumstances (e.g. Farber, 2017); the effects of $X_1(t)$ on $X_1(T)$ in Figure 1. Other studies explore the influence of job displacement on domains beyond the labour market, such as cardiovascular health (Black et al., 2015), mental health (Paul et al., 2018), and fertility (Huttunen and Kellokumpu, 2016); the effects of $X_1(t)$ on $X_2(T)$ in Figure 1. When exploring family composition (X_1 variables measured at t), Preetz (2022) investigates the effects of partnership dissolution on life satisfaction and mental health, while Glaser et al. (2008) look at the impact on support in later life, and Barbuscia et al. (2022) focus on a number of health conditions, including self-rated health, depressive mood, and sleep disorder (X_2 variables measured at T). Some studies consider the influence of modifiers (and potential mediators), e.g. how family structure has a bearing on job displacement and subsequent recovery (Attewell, 1999); $X_2(t)$ on $X_2(T)$, mediated by $X_1(t+\tau)$.

Public health studies focus more on cross-effects of health shocks on other domains (the effects of $X_2(t)$ on $X_1(T)$ in Figure 1). For instance, García-Gómez et al. (2013) consider the impact of acute hospital admissions on employment and income, while Lenhart (2019) analyses the impact of declines in self-reported health status and the onset of health conditions on subsequent labour market outcomes, and Bonekam and Wouterse (2023) study the impact of hospital admissions on wealth. Most of these studies are based on longitudinal panel surveys, with some using cohort data (e.g. Griffiths et al., 2021, or Wörn et al., 2023, consider the effects of job loss on mental and physical health during the Covid-19 pandemic) and others using administrative data (e.g. Fadlon and Nielsen, 2021, exploring family labour supply responses to severe health shocks).

Our objective is to broaden the view presented by the extant literature, allowing the distinction between determinants and outcomes to blur, and all state variables to co-evolve. To facilitate that analysis, we turn to dynamic microsimulation.

3 Microsimulation as data enrichment

Empirical studies are constrained by the available data. Household panel surveys provide rich information on a number of individual characteristics, but their longitudinal dimension is limited (e.g. EU-SILC, the main survey for the European Union, has a rotational structure of only 4 years). Even long-standing surveys such as the PSID, introduced in 1968, report complete life histories only for a highly selected (older, native, less mobile etc.) sub-sample of the US population. Cohort surveys share similar limitations to panel surveys, subject to further limitations imposed on the segment of the population for whom data are reported. While administrative data sometimes help to fill gaps of respondent surveys, public access to such resources is often restricted, and they generally describe a relatively limited set of characteristics, thus precluding the broad analysis that is required to disentangle complex feedback effects between alternative life domains.

Dynamic microsimulation (O'Donoghue, 2014; O'Donoghue and Dekkers, 2018) offers a way to integrate empirical evidence potentially derived from multiple sources into a coherent and consistent framework, allowing extrapolation of the implied dynamics beyond the temporal limits of the observational data.

In microsimulation, the state of micro units (for example, individuals, households, firms) is modified starting from some initial configuration, on the basis of biological, institutional or behavioural rules. Examples of biological rules are ageing and death. Examples of institutional rules are tax and benefits systems. Examples of behavioural rules are any choices that the units can make, for instance – in the case of individuals – related to education, household composition, fertility, labour supply, lifestyle and health behaviour, retirement.

Microsimulations can be considered as synthetic databases containing detailed information about a population of interest. The 'initial configuration' of a microsimulation database is typically augmented by projecting new variables not described by the initial configuration, or extending the reported time horizon. An example of the former of these cases is when disposable household income is projected by combining market income described by the initial configuration with a description of the corresponding tax-benefit rules. Where data are projected through time, then dynamic microsimulations involve the simulation of panel data that can be subject to subsequent empirical analysis.

Mathematically, dynamic microsimulation models are Markov chains, where at each time t an agent $i \in \{1, \dots, N\}$ is fully described by some set of state variables $\mathbf{x}_{i,t} \in \mathbb{R}^K$. When the model is cast in discrete time (i.e. sampled at regular intervals, for instance yearly) the evolution of her (vector of) state variables is specified by the difference equation:

$$\mathbf{x}_{i,t+1} = \mathbf{f}_i(\mathbf{x}_{i,t}, \mathbf{x}_{-i,t}, \boldsymbol{\theta}, \mathbf{P}_t, \boldsymbol{\xi}_{i,t}) \quad (1)$$

where $\boldsymbol{\theta}$ is a vector of behavioural parameters, \mathbf{P}_t are time-varying environmental parameters (potentially including past, present, and anticipated future policies), and $\boldsymbol{\xi}_{i,t}$ are stochastic disturbances. Individual outcomes can also depend on the state variables of other agents $\mathbf{x}_{-i,t}$, for instance their partners or children.

Structural modelling, in this context, refers to the parameters $\boldsymbol{\theta}$ governing behaviour – for instance those describing utility functions – being policy invariant. Expectations about the future are accommodated in the notation as they can be expressed as a function of the state variables \mathbf{x} and the policy parameters \mathbf{P} . Realism in the policy description requires \mathbf{P} to be a consistent reflection of the “real-world” environment. Finally, the notation can be generalised from partial equilibrium approaches – where there are only specific types of agents in the economy (say, individuals but not firms) – to general equilibrium approaches – where there are more agent types $\{i, j, h, \dots\}$ each defined by their own state variables $\mathbf{x}_{i,t}, \mathbf{x}_{j,t}, \mathbf{x}_{h,t} \dots$ possibly depending on the state variables of all other agents of any type (as in an agent-based setting).

In this context, interaction between different life domains is simply defined as variables pertaining to one domain having a causal impact on the evolution of other domains. Consider for instance health (h) and employment (e), and suppose their respective laws of motion are specified as follows:²

$$h_{i,t+1} = h(h_{i,t}, e_{i,t}, \dots, \boldsymbol{\theta}_h, \mathbf{P}_t, \boldsymbol{\xi}_{i,t}) \quad (2)$$

² In terms of Figure 1, employment is X_1 and health X_2 . The example easily generalises to more domains, and other variables.

$$e_{i,t+1} = e(e_{i,t}, h_{i,t}, \dots, \theta_e, P_t, \xi_{i,t}) \quad (3)$$

Health status at time t affects both health and employment outcomes at time $t+1$, and similarly for employment status at time t . The structure is similar to micro-level dynamic factor models (Altonji et al., 2022; Barigozzi and Pellegrino, 2023), with the added flexibility associated to the algorithmic nature of the simulation approach.

Suppose we are interested in health outcomes at time T , and wish to evaluate the impact of a health event at time 0. In the model, there are two causal pathways: one goes directly from health at time t to health at time $t+1$, for all $t = 0, \dots, T$; the other pathway is mediated by employment outcomes.

This modelling framework can be confronted with a reductionist approach, which would entail estimation of the following specifications, *in isolation*:

$$h_{i,t+1} = h'(h_{i,t}, \dots, \theta'_h, P_t, \epsilon_{i,t}) \quad (2')$$

$$e_{i,t+1} = e'(e_{i,t}, \dots, \theta'_e, P_t, \epsilon_{i,t}) \quad (3')$$

If the time span is sufficiently long, indirect effects would be captured by the lagged dependent variable. For example, in eq. 2' the coefficient on the lagged health variable in the reduced-form specification would pick up the effect on employment, and the subsequent effect of employment on future health. The reductionist approach would produce *on average* the same outcomes as the multidimensional approach, provided the estimators are well-behaved. However, in the reductionist specification the coefficient of the lagged health status would suffer from an omitted variable bias (employment), leading to a mis-representation of the true persistency effect of the health shock. This in itself could lead to incorrect policy implications.³

Moreover, a reductionist approach is by construction blind to what happens in other life domains. For instance, by using equations 2'-3', it would not be possible to predict the impact of an economic shock on health status, or the impact of increasing levels of education on future population health, with related implications for analysis of income inequality.

4 Analytical strategy

Dynamic microsimulation is generally used to project population aggregates, based on individual simulated outcomes. Here, we use it to construct differentiated individual counterfactuals that allow us to quantify how different causal pathways dynamically contribute to outcomes.

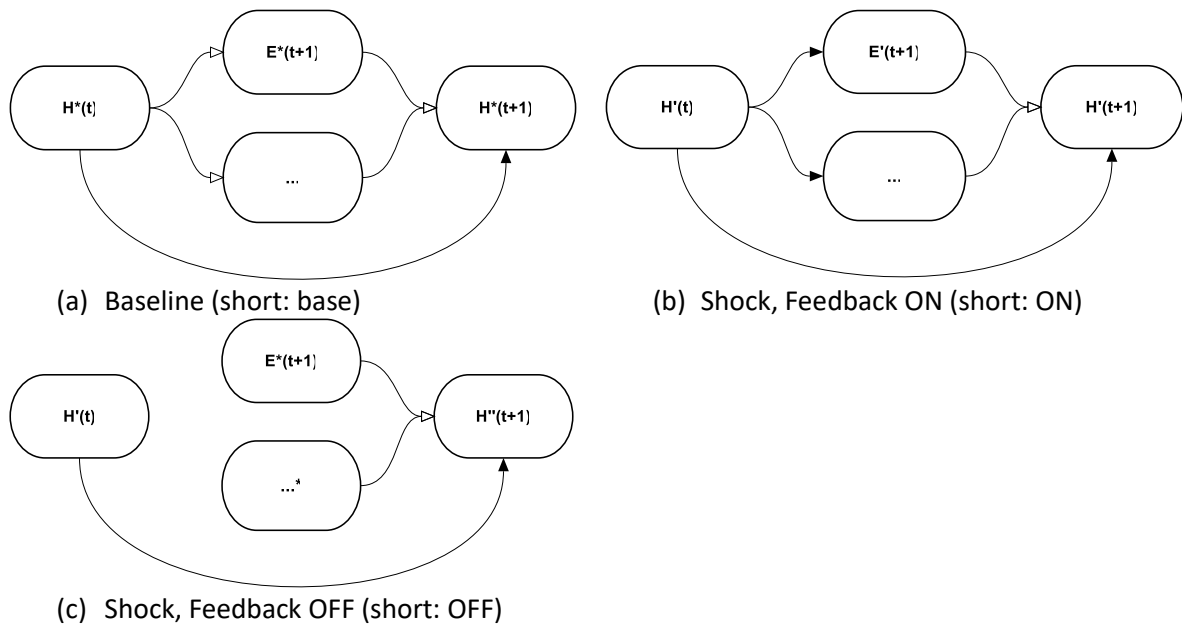
Our analytical strategy entails running three sets of simulations. The first simulation provides the *baseline* (short name: 'base'), with default parameterisation, and without any artificially imposed shock to initial conditions. This is stylised in Figure 2, panel (a), with reference to the evolution of a single variable of interest, say health (H). Baseline values are identified with an asterisk. Health at time t has a direct impact on health at time $t+1$, and a mediated effect through its effects on employment (E) and other variables (not shown in the figure).

The second set of simulations entail shocking the initial conditions, for instance by decreasing the level of initial health. This is shown in panel (b), and referred to as '*Shock, Feedback ON*' (short name: 'ON'). The new values of the variables are indicated with a 'prime' sign.

³ If, for example, a policy influenced health in a way that altered the coincident relationship with employment.

Finally, panel (c) depicts counterfactual simulations where the initial shock is only allowed to have a direct impact on the future values of the shocked variable itself (health in our example), while the evolution of all the other variables is taken from the baseline ('Shock, Feedback OFF', short name: 'OFF').⁴ This is indicated in panel (c), by the asterisks on " $E^*(t+1)$ " and other variables, "...*", whereas the shocked value of $H'(t)$ is depicted as feeding through to a "feedback off" shocked value $H''(t+1)$.

Figure 2: Counterfactuals



The ON vs. base comparison answers the question: "How would a shock in a given life domain broadly affect life trajectories?" The OFF vs. base comparison answers the question: "How would a shock in a given life domain affect life trajectories, if it did not spill over to other life domains?"

As already discussed, the overall effect of the shock involves a direct effect on the same domain where the shock occurred (health in the figure), a cross-effect on other domains (employment, etc.), and an indirect effect from the other domains back to the shocked domain.

In this framework, the direct (un-mediated) effect can be measured by comparing the 'Feedback OFF' scenario with the baseline, with respect to the evolution of the shocked variable. The cross-effect can be measured by comparing the 'Feedback ON' scenario with the baseline, with respect to the evolution of the other variables of interest. The indirect (mediated) effect can be measured, following a diff-in-diff approach, by contrasting differences between the 'Feedback ON' scenario and the baseline with differences between the 'Feedback OFF' scenario and the baseline.

It is not a priori clear whether the impact of the shock on the future evolution of the shocked variable itself should be greater under the ON or OFF scenarios. The case where the difference with respect to the baseline is higher in the 'Feedback ON' than in the 'Feedback OFF' scenario, that is when the total

⁴ This requires matching the simulated individuals in the 'Feedback ON' scenario with their counterparts in the baseline.

effect is higher than the direct effect, implies *reinforcement* mechanisms; the opposite indicates *attenuation* mechanisms.

We can then construct a feedback indicator as follows:

$$F_{x,t} = \frac{\text{total effect}}{\text{direct effect}} = \frac{(ON_{x,t} - base_{x,t})}{(OFF_{x,t} - base_{x,t})} \quad (4)$$

where x is the variable being considered. Values of $F > 1$ reveal reinforcement mechanisms, while $F < 1$ indicates attenuation mechanisms.

This indicator is related to the ‘proportion mediated’ (PM) indicator in mediation analysis (Ditlevsen, 2005; Ananth, 2019):

$$PM_{x,t} = \frac{\text{indirect effect}}{\text{total effect}} = \frac{\text{total effect} - \text{direct effect}}{\text{total effect}} = 1 - \frac{1}{F_{x,t}} \quad (5)$$

or $F_{x,t} = \frac{1}{1 - PM_{x,t}}$. Our preference for the F indicator is due to its more straightforward interpretation in terms of the dominance of reinforcement vs attenuation mechanisms as discussed above.

5 The model

For this study, we use the SimPaths dynamic microsimulation model developed at the Centre for Microsimulation and Policy Analysis at the University of Essex (Bronka et al., 2023), estimated on UK data.⁵ SimPaths implements a hierarchical architecture where individuals are structured in benefit units (for fiscal purposes), and benefit units are structured in households.⁶ The model runs at a yearly frequency, consistent with the yearly frequency of the survey data on which the different processes are estimated. The model is composed of seven different modules: (i) Demography, (ii) Education, (iii) Health, (iv) Household composition, (v) Non-labour income, (vi) Labour supply, and (vii) Consumption. Each module is in turn composed of different processes or sub-modules; for example, the demographic module contains an ageing process and a process for leaving the parental home, and the labour supply module includes a wage setting process together with a process determining the number of hours of work supplied.

Simulated modules and processes are organised as displayed in Figure 3. In each simulated year, agents are first subject to an ageing process (involving age and year specific probabilities of dying), followed by a population alignment process.

Population alignment adjusts the simulated population to match population projections produced by the Office for National Statistics (ONS). Specifically, the ONS reports population estimates for the UK distinguished single year of age and gender for 12 geographic regions for each year between 2011 and 2023 inclusive. The ONS also reports projections for the same disaggregated population subgroups for each year between 2024 and 2043.

In each simulated year, the alignment process begins by evaluating the population short-fall/excess associated with each age, gender and region category, relative to ONS population

⁵ The model is coded in Java using the JAS-mine simulation library (Richiardi and Richardson, 2017).

⁶ A benefit unit is comprised of a single adult or adult couple and their dependent children. There can be households comprised of a single benefit unit, and benefit units comprised of a single individual.

estimates/projections. The model then simulates internal migration by moving benefit units with baby girls from regions with an excess of females aged 0 to those with a short-fall until all short-fall or excess regions are exhausted. Any net differences that remain between the simulated and ONS reported numbers of females aged 0 by region are resolved by removing benefit units to reflect (implicit) international emigration, or cloning benefit units to reflect international immigration.

Having matched ONS estimates/projections for the number of females aged 0, the alignment process proceeds to consider males aged 0. The model allows for up to one birth each year, which ensures that no benefit unit includes both a female and male aged 0. This means that the same process as described for females aged 0 can be applied to align the simulated number of males aged 0, without risking distortion to the previously matched numbers for females.

Subsequent gender and age categories are considered in turn, where benefit unit migration is limited to the set of units in which the youngest member corresponds to the gender/age category under consideration. Closure of this procedure is facilitated by the fact that the incidence of benefit units comprised of a single individual increases as the gender/age category under consideration proceeds to higher ages.

Following alignment, the education module determines whether students should remain in education, or – for individuals who are no longer in education – re-enter education. Students are assumed not to work and therefore do not enter the labour supply module. Individuals who leave education have their level of education re-evaluated (for those who returned to education, their level of education can only go up) and can enter the labour market.

The health module calculates an individual's continuous health score, a measure of mental distress, and evaluates whether the individual is long-term sick or disabled (in which case, he / she is not at risk of work).⁷

The household composition module projects cohabiting relationship formation and dissolution. This aspect of the model is the principal source of interactions between simulated agents. When a relationship forms, the partners are selected via a matching process that is designed to reflect correlations observed in survey data. Females in couples can give birth to a (single) child in each simulated year, as determined by a fertility process. Fertility is modelled at the individual level, and is aligned to fertility rates implied by official population projections.

The labour supply module projects potential wages for each simulated adult in each year using a wage equation with parameters estimated using a Heckman-corrected regression on contemporary survey data. Given potential wages, hours of work supplied by all adult members of a benefit unit are evaluated by identifying the utility-maximising number of discrete hours of work, in a random utility model framework.⁸ This calculation involves identifying disposable income for each feasible labour alternative, which is imputed from a detailed description of the contemporary UK tax and benefit system, as described in van de Ven *et al.* (2022).⁹

⁷ The status of long-term sick / disabled is reversible though.

⁸ A'la van Soest (1995). The structural labour supply module is replaced by a simpler probabilistic transition module for the Covid-19 years (2020 and 2021), during which it is considered that households were less able to choose their preferred level of hours worked.

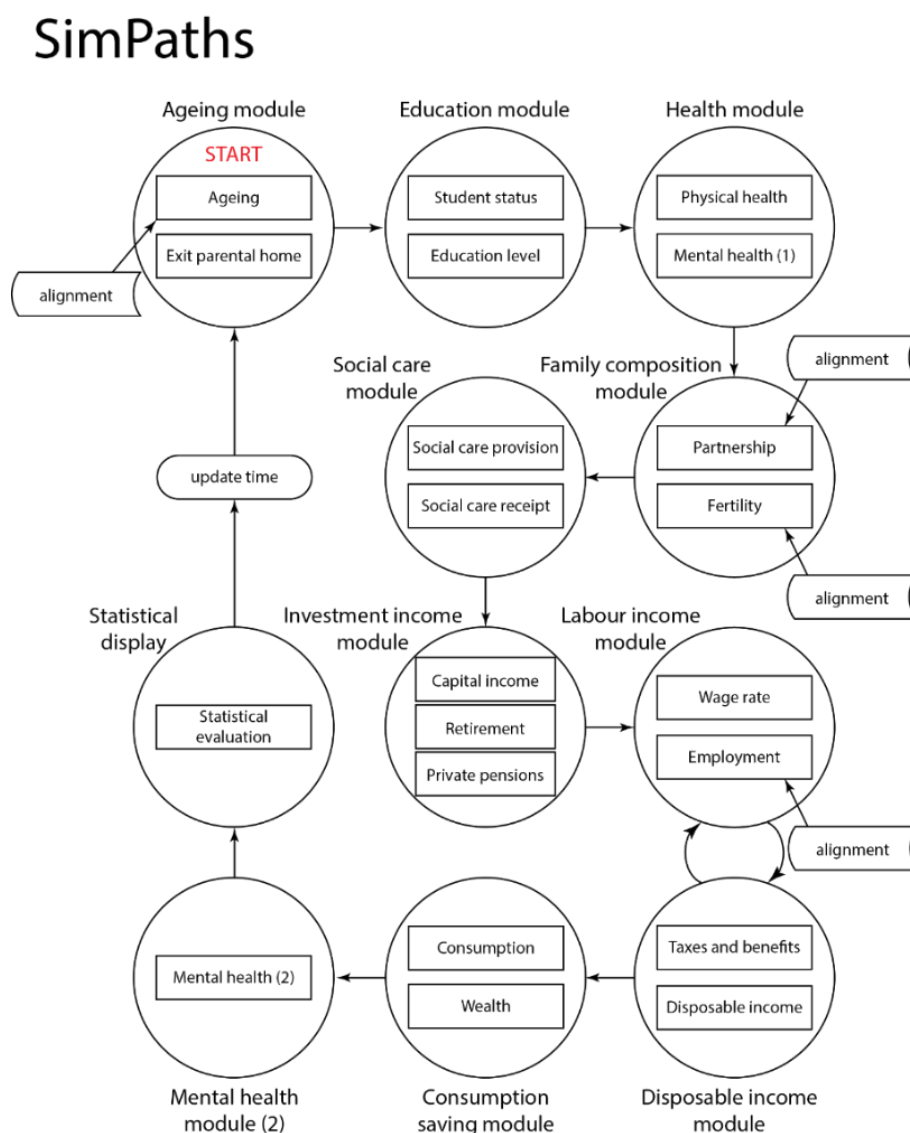
⁹ Imputations are based on data derived from UKMOD, a tax-benefit calculator for the UK; see Richiardi *et al.* (2021).

Finally, a simple consumption module transforms disposable income into consumption by applying an homogenous saving rate, calibrated to the data. The same saving rate is also used when calculating capital income.

Simulations can be initialised in any year between 2011 and 2017 – they start in 2011 for this study – based on a representative cross-section of the UK population in the respective year, and can run until 2060. The period of overlap with existing data is used for validation purposes.

The model structure, as well as the estimated parameters based on the UK Household Longitudinal Survey (UKHLS) and Family Resources Survey (FRS) data and validation to historical time series for the period 2011-2020, are described in Bronka et al. (2023), and summarised in Appendix 1.

Figure 3: Structure and order of processes modelled in SimPaths



6 The experiments

We consider two experiments: a partnership dissolution, and a health shock, applied to the cohort of men aged 30 in the initial year of the simulation (2011). Simulations are run until 2050, when the simulated individuals reach the age of 69.¹⁰ In the first scenario, all partnerships involving men aged 30 in the initial year of the simulation are dissolved. In the subsequent periods, these men might decide to re-partner, thus entering the market for partnership, where they might (or might not) find a suitable partner. The comparison group in the baseline is therefore composed of all partnered men aged 30 in 2011 – the same group of men, in a world in which the shock did not occur.

In the second scenario, the self-rated health status of all men aged 30 in the initial year is reduced to 1 (“poor”, on a five point self-reported scale varying from “poor” to “excellent”).¹¹ The comparison group in the baseline comprises all men aged 30 in 2011 (irrespective of their partnership status) – again, the same group of men, in a world in which the shock did not occur.

Table 1 reports the sample size for each experiment (number of individuals shocked in 2011). Coming from the UKHS survey data, the individuals selected for our experiments are representative of the respective segments of the UK population (see Appendix 2 for a comparison of the characteristics of the two samples with alternative survey data).

Table 1: Sample sizes

Experiment	Shocked individuals
(a) Partnership dissolution	8,401
(b) Health shock	25,232

Note: Each of the initially shocked individuals is simulated from 2011 to 2050.

To be noted, the focus on partnership dissolution requires a change with respect to the standard version of SimPaths described in Section 5, as population alignment to official demographic projections must be switched off. This is because population alignment in the model depends on household structure (see Bronka et al., 2023 for more details), which is obviously impacted by the partnership dissolution. Retaining population alignment would then imply that the simulated populations in the different scenarios are not the same, preventing us from matching individuals from the baseline in the ‘Feedback OFF’ scenario. For coherence, population alignment is switched off also for the health shock experiment.

What this implies is that the experiments are run on a closed and constant population. Without population alignment, the initial cohort of men remains representative of the UK male population aged 30 in 2011 (see Appendix 2), and their partners remain representative of the partners of the UK males aged 30 in 2011. However, no immigration or emigration is allowed.

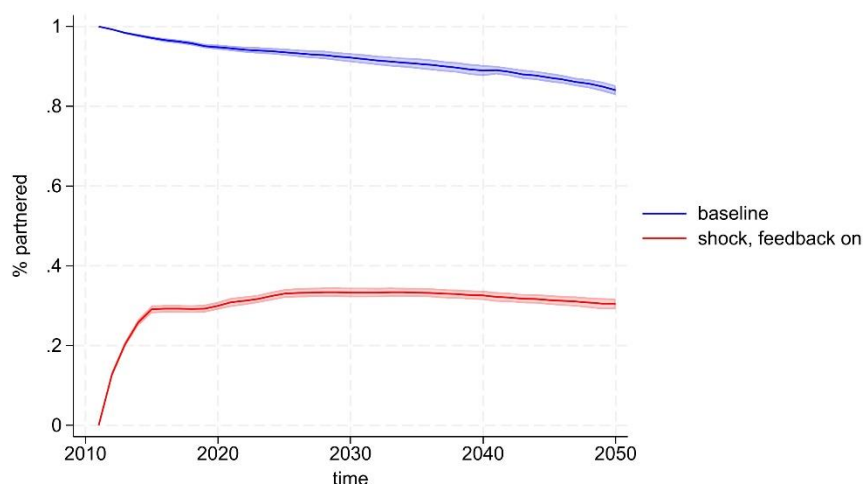
¹⁰ Focussing on a specific cohort allows a better understanding of the simulated dynamics. Moreover, when shocking relationship status, the overall “market for partnership” is affected only marginally.

¹¹ In the sample, approximately 19% of men aged 30 in 2011 report excellent health, 53% very good, 21% good, 6% fair, 1% poor. The shock hits all of them and brings down their health status to ‘poor’.

7 Results: Partnership dissolution

Figure 4 shows the evolution of the partnership rate in the baseline and feedback ON scenarios, for the affected individuals (partnered men aged 30 in 2011). The share of partnered men declines over time in the baseline, mostly due to a regression to the mean (the sample is positively selected to start with). In the counterfactual, it takes about 3 years for the partnership rate to increase and reach an equilibrium level of about 30%.¹²

Figure 4. Partnership dissolution: Total effects



Note: 90% confidence intervals are shown as shaded areas. Confidence intervals are based on 40 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details. Sample: Partnered men aged 30 in 2011.

The difference between the baseline and the feedback ON scenario is a measure of the *total* effect of the shock. Becoming unpartnered at age 30 increases the probability of being single at age 65 by almost 60 percentage points. This may contrast with the observation that partnership breakdowns are a common occurrence at all ages, and in particular for young adults, and they do not seem to lead to such drastic long-term consequences, in real life. This is because most individuals re-partner quite quickly and would therefore not be classified as ‘single’ in a survey, despite having gone through a partnership dissolution. In other words, as is well known, stock sampling leads to length-time bias, with a higher likelihood that the short duration spells will be omitted from the sample. Our experiment should therefore be interpreted as putting individuals in an un-partnered spell that is long enough to be recorded in a survey, that is in a long-term single status, with potentially larger long-term consequences.

Cross-effects on other variables are explored in Figure 5. The partnership dissolution at 30 has a small negative effect on health (panel (a)) until approximately the age of 50. The effect on employment

¹² The partnership rate in the counterfactual scenario levels off at around 30%. For comparison, age-specific partnership rates in the baseline are much higher for prime-age men, remaining approximately constant at around 80% between 30 and 65 years of age. The difference is explained by the fact that most of the transitions from single to partnered happen before the age of 30.

(panel (b)) is more pronounced, with a decrease in the probability of being employed of around 5 percentage points, again until the age of 50. After that, employment rates in the baseline drop. This is because estimated labour supply for men (and women) in couples is reduced after the age of 50, something that is also observed in the survey data. The drop in the baseline therefore reflects the higher percentage of partnered individuals.¹³ The same composition issue (a higher percentage of partnered men) explains why employment rates in the baseline fall below those of the counterfactuals at older ages.

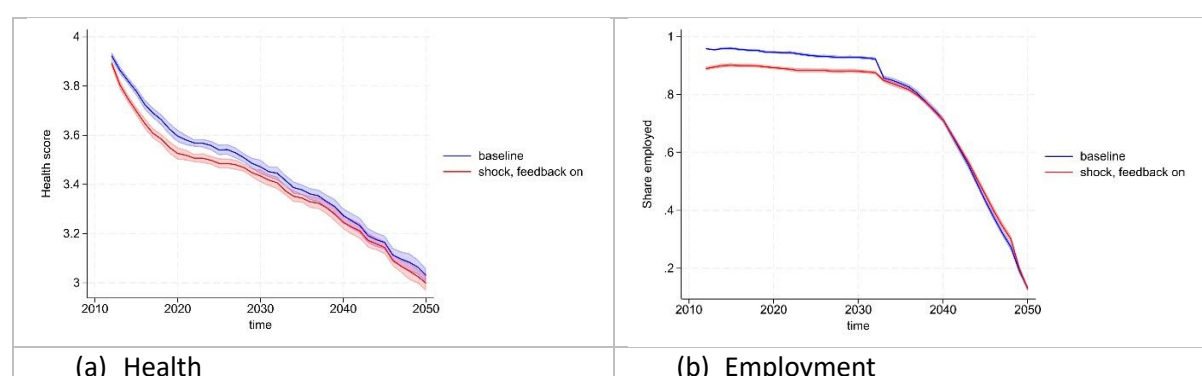
Panels (c) and (d) show the effects on income. Gross income (panel (c)) is higher in the baseline, reflecting higher employment rates and longer work hours – see Appendix 4 for more details. The same pattern is found for equivalised disposable income (panel (d)). The effects of the shock on equivalised disposable income however differ depending on household structure. Specifically, the simulations assume that children follow their mother when a relationship dissolves. In the period immediately after the separation, two effects are consequently at play: a mechanical change to the equivalisation factor, and the incidence of maintenance payments if the couple has children.

The equivalisation factor assumed for analysis is the modified OECD scale. This scale assigns a value of 1 to the first adult; 0.5 to the second and each subsequent person aged 14 and over; and 0.3 to each child aged under 14. A partnership dissolution will tend to reduce the equivalence scales of divorced men, which works to increase equivalised disposable income.

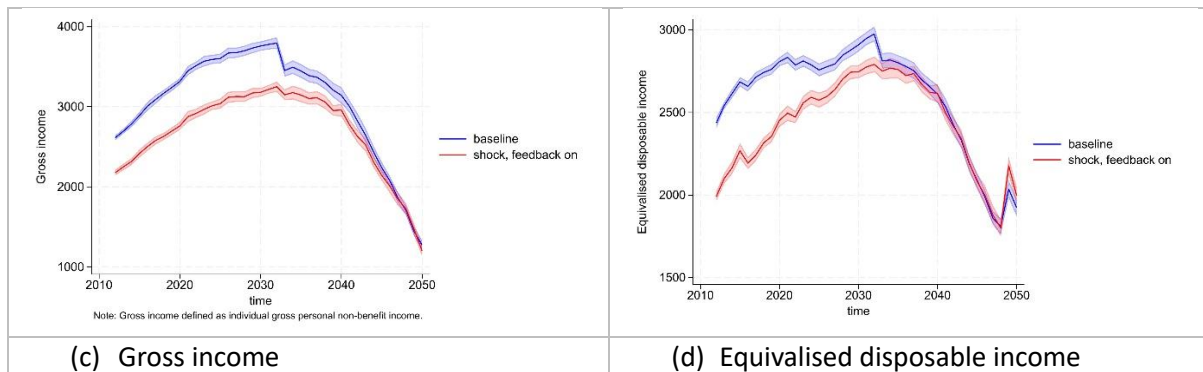
In contrast, maintenance payments work to decrease the equivalised disposable incomes of men following relationship dissolution (where children are involved). The simulations project maintenance payments based on the rules in place from 2012 onwards (Child Maintenance Service, 2024). On average, maintenance payments reduce equivalised disposable income for men in the simulation, by 12% (1st quartile: 9%; 3rd quartile: 15%).

On balance, panel (d) of Figure 5 indicates that the influence of weaker employment and maintenance payments on average dominate reduced equivalence scales, resulting in lower measures of equivalised disposable income following the simulated shock to partnership status.

Figure 5. Partnership dissolution: Cross-effects



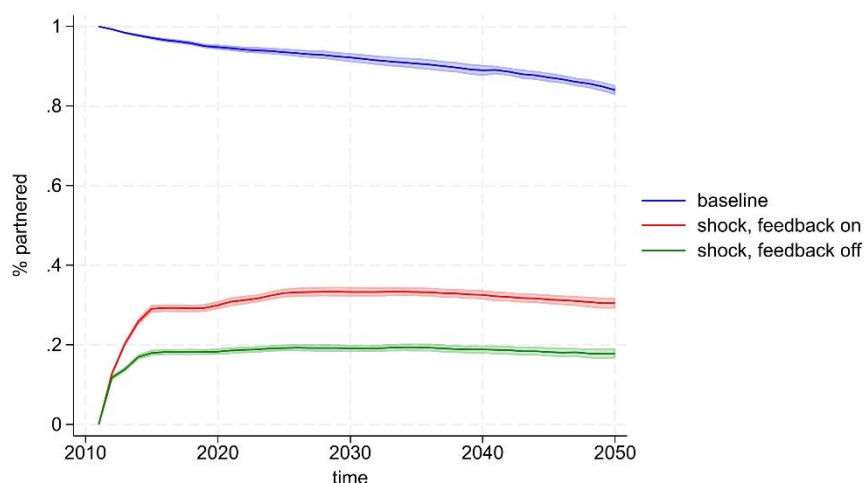
¹³ All individuals in the baseline start as partnered, and all individuals in the scenario start as single. However, over time the initial partnerships might “naturally” (that is, as simulated by the behavioural equations rather than artificially imposed in the experiment) come to an end, while single individuals might re-partner.



Note: 90% confidence intervals are shown as shaded areas. Confidence intervals are based on 40 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details. Sample: Partnered men aged 30 in 2011.

To understand the role of mediated effects, we bring in the 'Feedback OFF' scenario. Figure 6 is the equivalent to Figure 4, with the 'Feedback OFF' scenario added.

Figure 6. Partnership dissolution: Total and direct effects

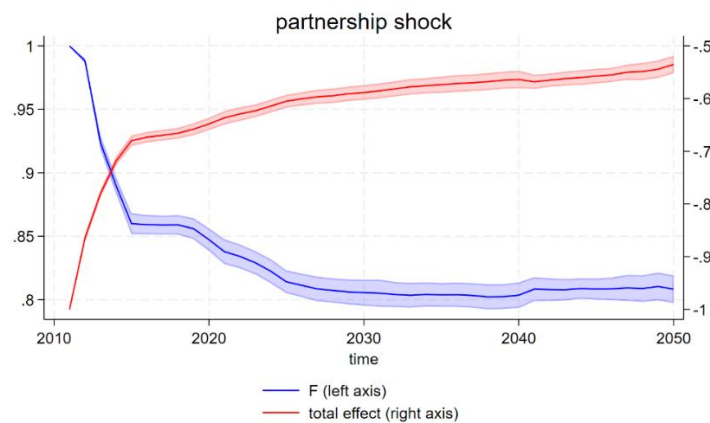


Note: Base vs. ON = total effect; Base vs. OFF = direct effect. 90% confidence intervals are shown as shaded areas. Confidence intervals are based on 40 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details. Sample: Partnered men aged 30 in 2011.

Comparing the blue ('baseline') to the green ('feedback OFF') series displayed in Figure 6 indicates the influence of direct effects on projections. The figure indicates that the direct effect in isolation would see partnership increase from zero to 20%, compared to more than 30% when all the mediator effects are factored in ('Feedback OFF', red line).

Finally, Figure 7 reports evolution of the F index, computed as per eq. (4). The F index (blue line) measures the ratio between the total and the direct effect of the shock on the shocked variable itself. The index starts at 1, as in the initial period the direct effect is the only one at work. The index then swiftly declines, reaching a plateau slightly above 80%. This means that the complex dynamic interactions between life domains compensate for around 20% of the initial impact of the shock. The figure also displays the total size of the effect (red line), to help contextualising the increased relative importance of the mediated effects.

Figure 7. Partnership dissolution: F index



Note: The F index (eq. 4) measures the ratio of the total to the direct effect of the shock. Values above 1 indicate reinforcement mechanisms are at work, while values below 1 indicate attenuation mechanisms. 90% confidence intervals are shown as shaded areas. Confidence intervals are based on 40 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details. Sample: Partnered men aged 30 in 2011.

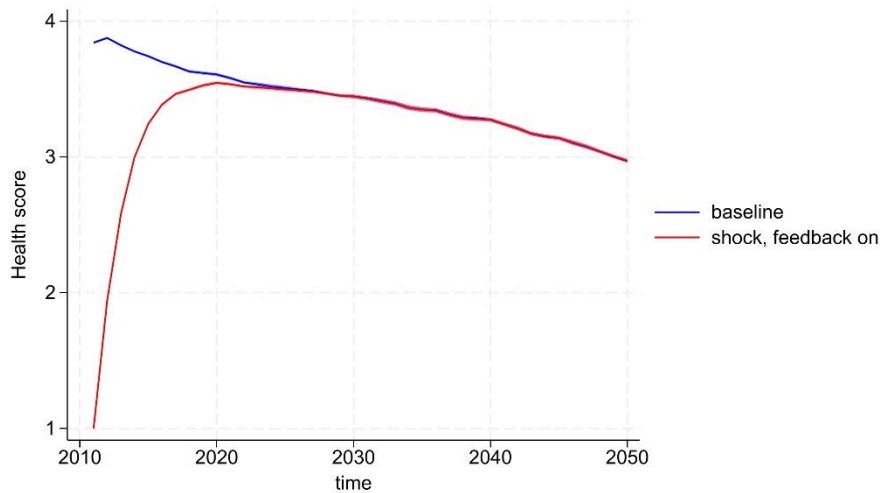
8 Results: Health shock

We contrast the results on the effects of a partnership dissolution with a second experiment, where we reduce the self-rated health score of all men aged 30 in the initial year of the simulation to 1.¹⁴ Figure 8 shows the evolution of health in the baseline and ‘Feedback ON’ scenario, our measure for the total effect of the shock. It takes approximately 10 years for the shock to be absorbed, on average, with big health gains obtained during the first 3-5 years after the shock. The main explanation for this simulated response is that the individuals receiving the shock are young. Despite some persistency in

¹⁴ Qualitative measures of self-reported health are associated with well-known problems of comparability and interpretation. The measures are nevertheless included for analysis as they are analytically convenient and strongly related to other health-related characteristics reported by the UKHLS. For example, the measures of self-reported health relate closely to associated measures for disability reported below. Furthermore, the average number of Activities of Daily Living (ADLs) that individuals aged 65 and over report needing help with in the UKHLS increases from 0.294 for those reporting Excellent health, 0.323 (Very Good), 0.614 (Good), 1.580 (Fair), to 3.835 for those reporting Poor health (averaged over waves “g”, “i” and “k” of the survey).

the process determining health (see Bronka et al., 2023 for details), their other determinants typically point to good health, hence increasing the chances of a recovery.¹⁵

Figure 8. Health shock: Total effects



Note: 90% confidence intervals are shown as shaded areas. Confidence intervals are based on 40 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details. Sample: Men aged 30 in 2011.

The health shock has a large impact on the probability of being disabled (Figure 9, panel (a)). Disability is simulated as a severe health condition but not an absorbing state; individuals can move into and out of disability, with probabilities estimated on survey data. These probabilities depend, amongst other things, on age, gender, education and socio-economic position (see Table A2 in Appendix 1). This explains why, as health gradually recovers, the probability of being disabled also returns to the simulated baseline.¹⁶

The model assumes that disabled people are not available to work.¹⁷ Given that prime age men are typically observed to work despite poor health, it is the simulated increase in the disability rate that explains the small negative effects on employment projected for the health (panel (b) in the figure).

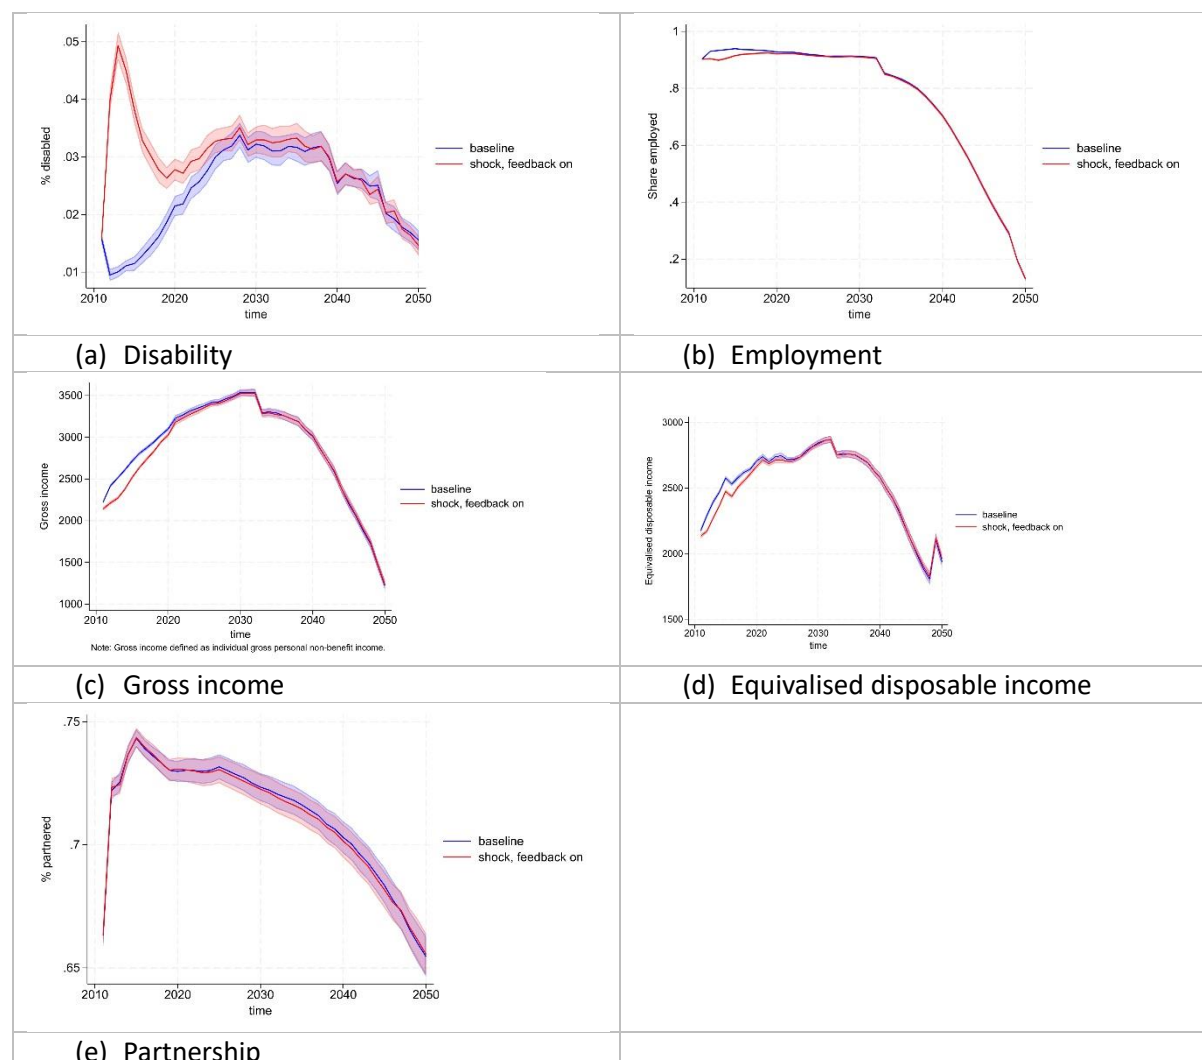
¹⁵ For contrast, we also experimented with a simulated health shock to the cohort of men aged 50 in the initial year of the simulation. Given that the average health score for men aged 50 is lower than for men aged 30, the simulated shock is on average smaller. Despite this, and coherently with the intuition, we observe that it takes more for the shock to be absorbed, and for average health to return to the (declining) trajectory that is observed in the baseline. Appendix 5 describes the results of this exercise in more detail.

¹⁶ The figure points to a counterintuitive lack of a clear age gradient in disability rates, in the baseline. This is because the disability variable is based on the 'long-term sick or disabled' response to the question related to current economic activity ('jbstat') in the UKHLS data. At older ages an increasingly proportion of the sample reports 'retired' as their economic status, even if they previously indicated 'long-term sick or disabled'.

¹⁷ 'long-term sick or disabled' in the 'jbstat' classification is an alternative state to employment.

The effects on employment carry over to effects on gross and net incomes (panels (c) and (d)). Interestingly, the health shock is not projected to influence partnership (panel (e)).

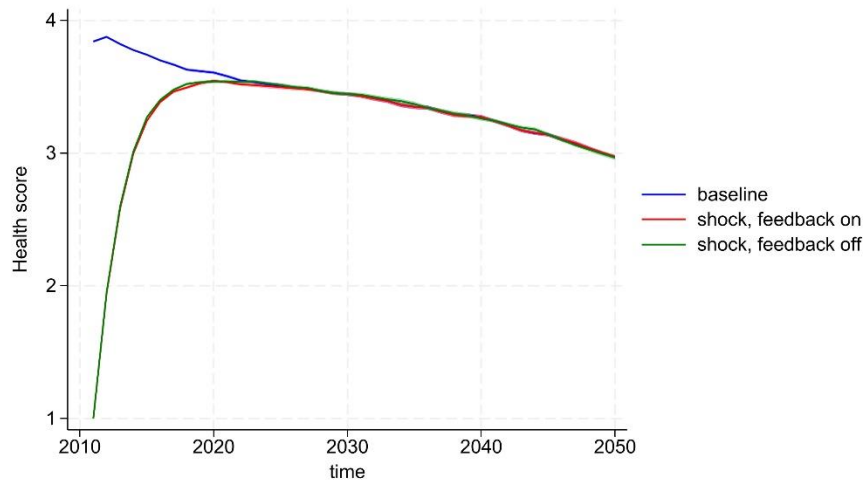
Figure 9. Health shock: Cross-effects



Note: 90% confidence intervals are shown as shaded areas. Confidence intervals are based on 40 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details. Sample: Men aged 30 in 2011.

Given the limited spillovers to other domains, it is no surprise that the mediated effect is also limited, as implied by Figure 11: the total effect substantially coincides with the direct effect.

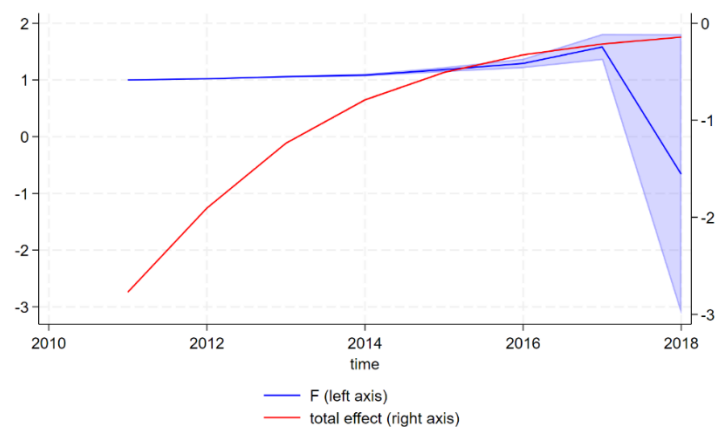
Figure 11. Health shock: Total and direct effects



Note: Base vs. ON = total effect; Base vs. OFF = direct effect. 90% confidence intervals are shown as shaded areas. Confidence intervals are based on 40 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details. Sample: Men aged 30 in 2011.

This can also be seen in Figure 12, which reports the evolution of the F index up to 2018: the index is mostly projected to be close to 1.

Figure 12. Health shock: F index

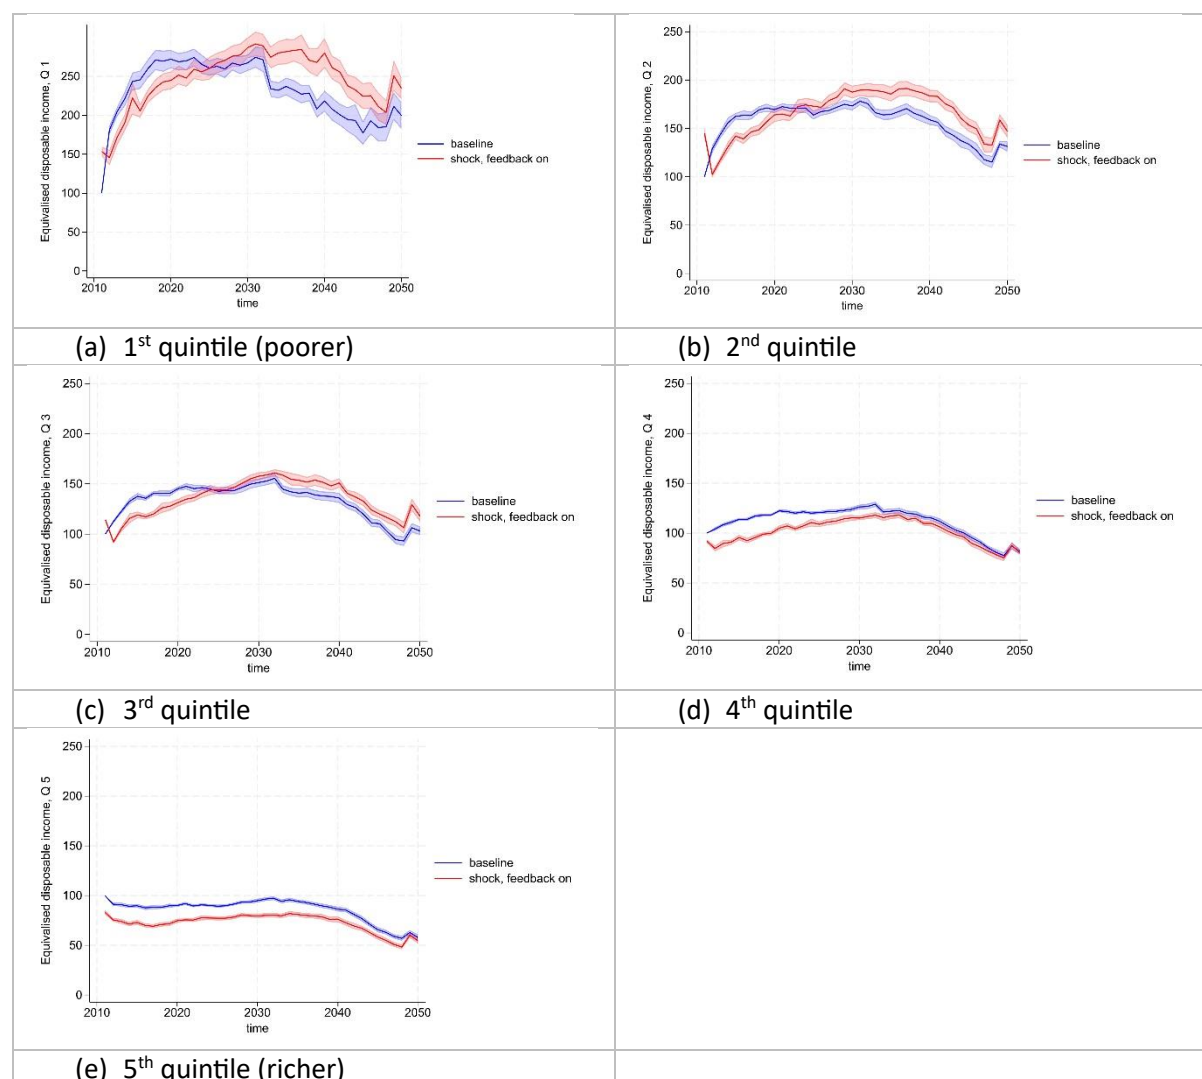


Note: The F index (eq. 4) measures the ratio of the total to the direct effect of the shock. Values above 1 indicate reinforcement mechanisms are at work, while values below 1 indicate attenuation mechanisms. 90% confidence intervals are shown as shaded areas. Confidence intervals are based on 40 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details. Sample: Men aged 30 in 2011.

9 Distributional Implications

Understanding how direct and indirect effects play out over multiple time horizons sheds new light on how inequalities unfold over the life course, and individual resilience to adverse events. For instance, we can analyse the cross-effects of a partnership dissolution on income distinguished by the initial socio-economic position.¹⁸ Figure 13 is a replica of Figure 5d, by quintiles of equivalised disposable income in 2011. In each quintile, income is normalised to 100 in the initial year of the simulation to facilitate comparisons of associated distributional effects.

Figure 13: Partnership dissolution: Cross-effects on equivalised disposable income, by income quintile in the initial year of the simulation



Note: Panels refers to different quintiles of equivalised disposable income in 2011, normalised to 100. 90% confidence intervals are shown as shaded areas. Confidence intervals are based on 40 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts

¹⁸ Given its more limited effect on variables other than health, we omit here a discussion of the distributional effects of the health shock.

- separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details. Sample: Partnered men aged 30 in 2011.

There are three interesting things to notice in Figure 13. First, incomes tend to grow more, in percentage terms, for lower quintiles. Second, simulations show a cross-over, for quintiles 1-3, between the baseline and the scenario later in life, with average equivalised disposable income starting lower in the shocked scenario, but surpassing that of the baseline after around 15 years. Third, the impact of the partnership dissolution is more negative for higher earners.

To explain these dynamics, we begin discussion at the moment when partnerships are dissolved, in 2011. As Table 2 reports, the proportion of shocked men whose partner was working, and their partner's earnings, increase with disposable income. This is partly by construction, as disposable income is computed at the benefit unit level, thus including earnings from both partners (as well as benefits that accrue to both the individual partners and the household). On the other hand, the average number of children – with the exception of the first quintile – decreases with income.

Table 2: Distributional characteristics of the partnership shock sample, 2011

Equivalised disposable income	Partners' employment rate	Partners' gross employment income (£/month, 2015 prices) (std. dev. in parenthesis)	Number of Children
1st Quintile	0.251	635.79 (362.56)	1.14
2nd Quintile	0.626	836.65 (477.02)	1.62
3rd Quintile	0.843	1175.16 (615.43)	1.32
4th Quintile	0.902	1645.88 (838.22)	0.90
5th Quintile	0.954	2429.90 (1383.04)	0.45

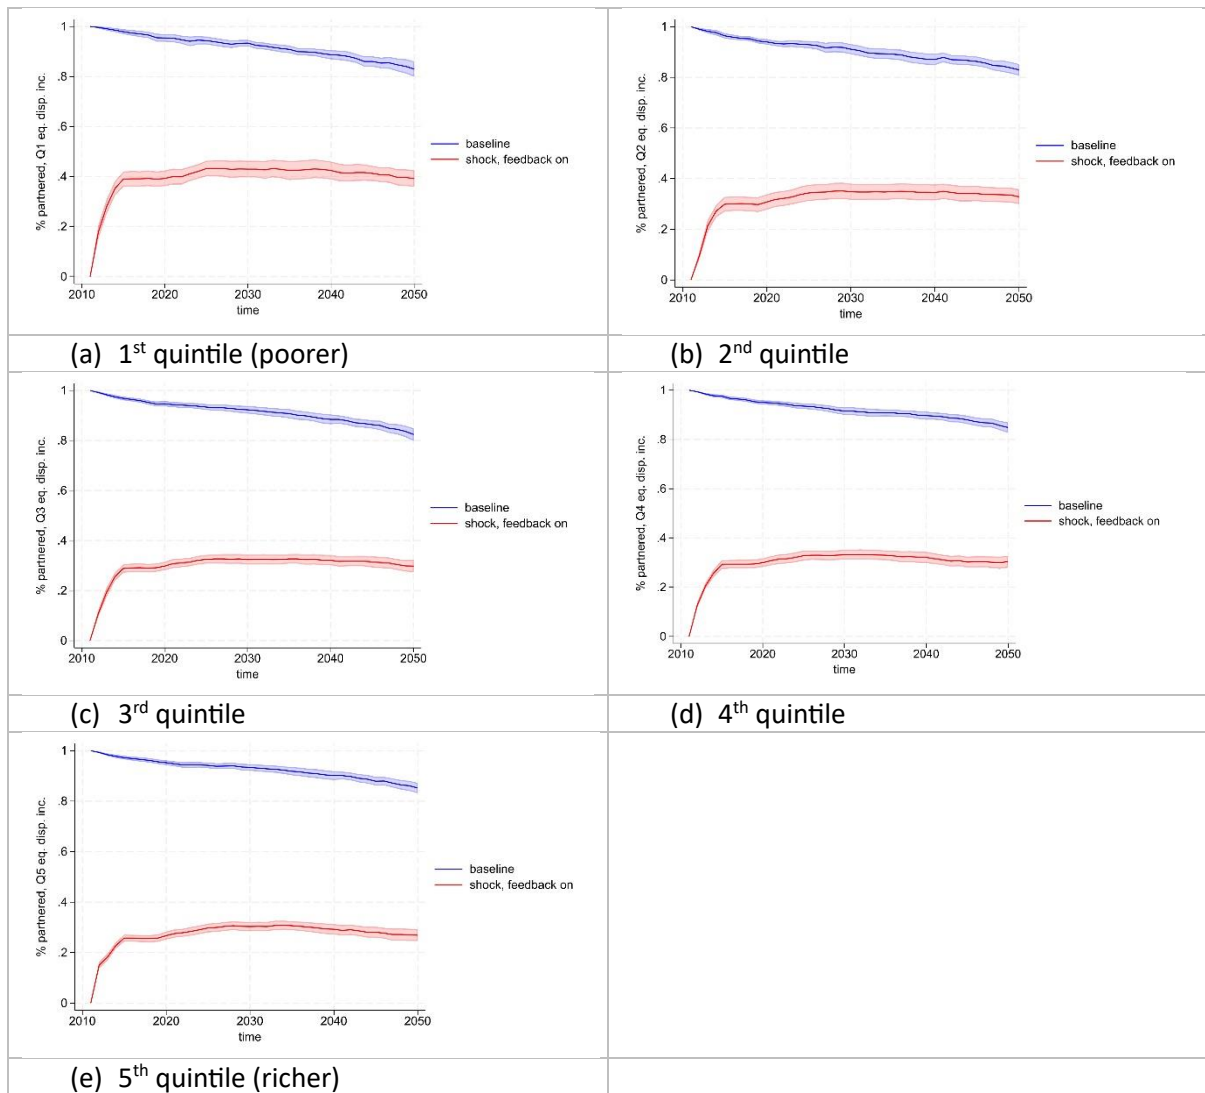
This explains why the shock to partnership affects men in the upper quintiles more, as we have explained when presenting Figure 5d: lost partner earnings combined with smaller reductions in equivalence scales of men toward the top of the distribution dominate the relatively lower maintenance payments that they must typically pay, in relation to their income.¹⁹

Over time moreover, we observe a lower probability of re-partnering as we move up along the income distribution (Figure 15). Given that being partnered is associated, on average, with higher equivalised disposable income, a lower probability of being re-partnered implies a more negative impact on income.²⁰

Figure 15: Partnership dissolution: Direct effects on partnership rates, by income quintile in the initial year of the simulation

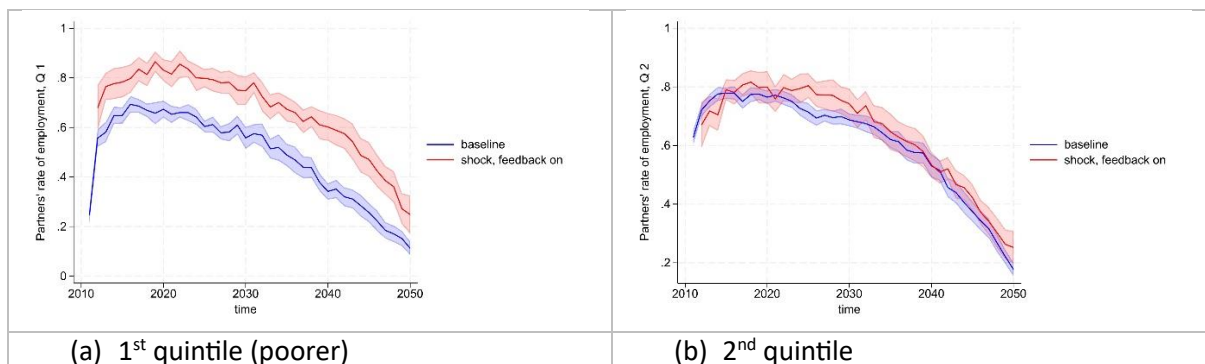
¹⁹ Richer individuals are charged a smaller fraction of their gross weekly income as maintenance payment. Moreover, they get fewer benefits, and have fewer children.

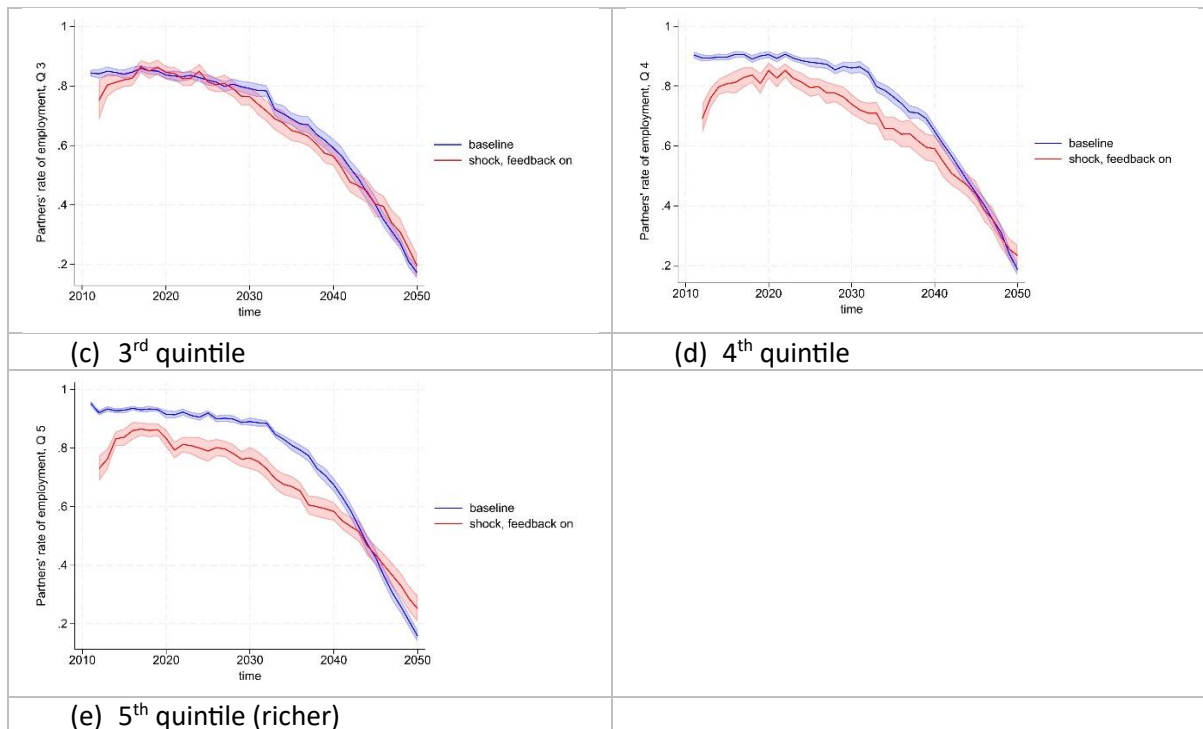
²⁰ The effect vanishes at very old ages as surviving individuals are more likely to be widowed, if previously partnered.



Furthermore, there is a negative relationship in the simulated baseline between income quintile and the probability that a partner is employed (if one exists, Figure 16). As noted previously, income quintiles are computed based on equivalised disposable income in 2011 and are thus lower by construction, *ceteris paribus*, if the partner is not working.

Figure 16: Partnership dissolution: Cross effects on partnership rates of employment, by income quintile in the initial year of the simulation





Note: Panels refers to the fraction of partners who are in employment, for different quintiles of equivalised disposable income in 2011. 90% confidence intervals are shown as shaded areas. Confidence intervals are based on 40 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details. Sample: Partnered men aged 30 in 2011.

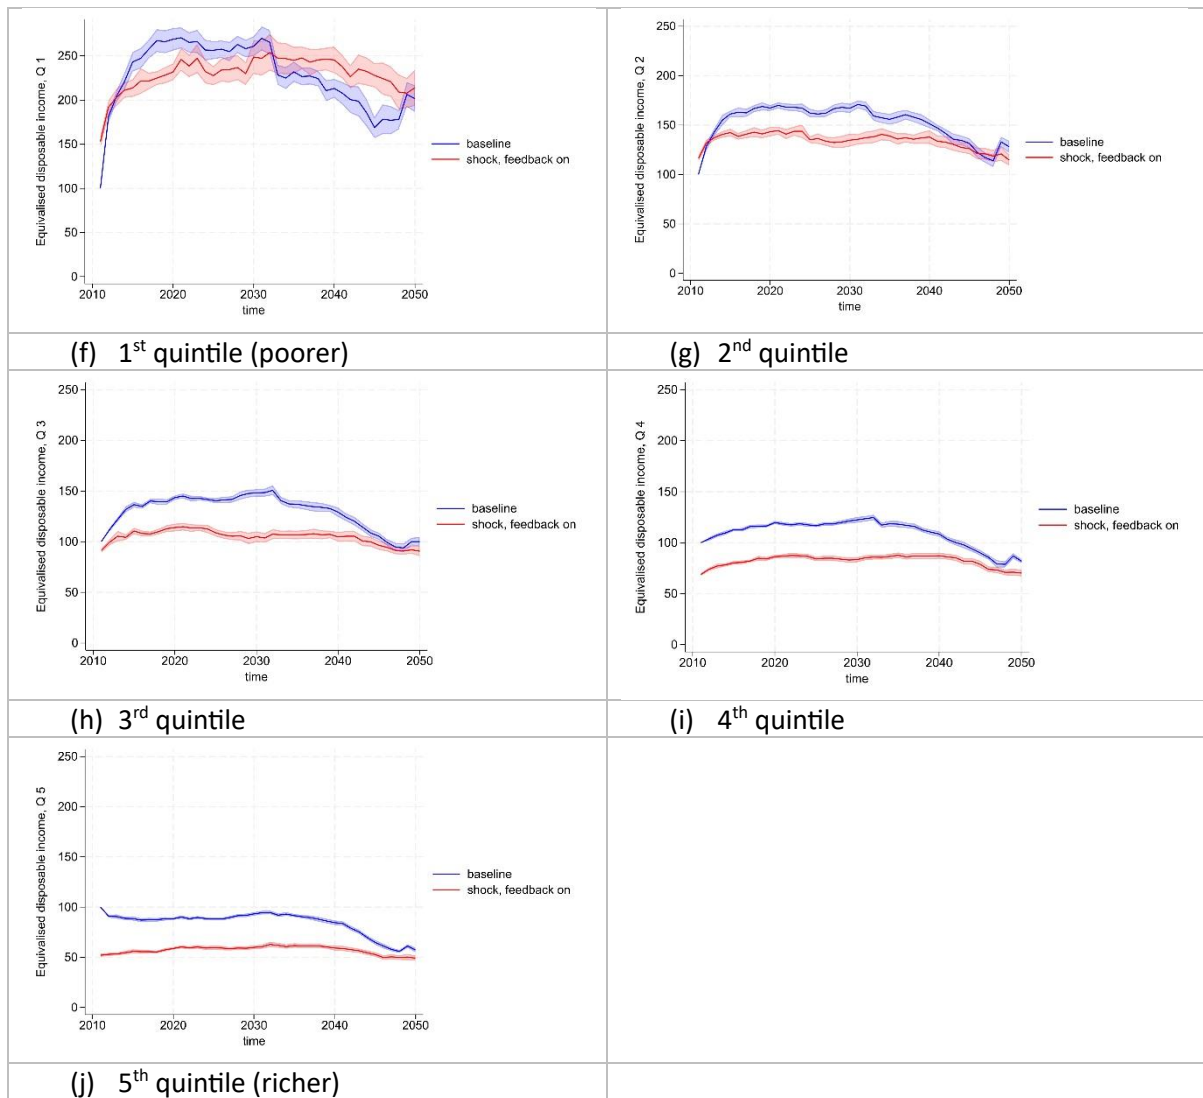
The fact that lower quintiles have a higher probability of re-partnering and a higher increase in the probability that the partner is working, with respect to the baseline, explains why their equivalised disposable income grows more. With a smaller initial loss and a higher rate of growth, equivalised disposable income in the lower quintiles surpasses that of the baseline.

The effects reported above for men are comparable to those of their female partners (Figure 17).²¹ The negative impact of the partnership dissolution is however larger for women, reflecting the gender pay gap.²²

Figure 17: Partnership dissolution: Cross-effects on equivalised disposable income, by income quintile in the initial year of the simulation, for the female partners.

²¹ The female partners affected by the experiment are obviously not constrained to be aged 30.

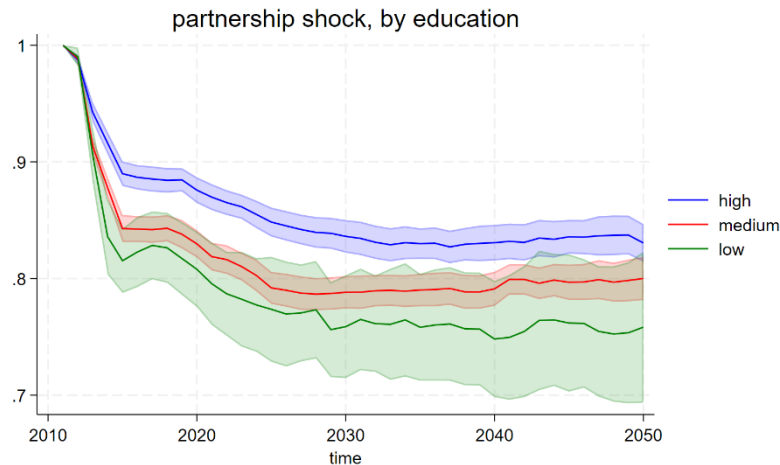
²² Even more than for men, at very old ages only few women are partnered, so, the impact of the initial shock becomes negligible.



Note: Panels refers to different quintiles of equivalised disposable income in 2011, normalised to 100. 90% confidence intervals are shown as shaded areas. Confidence intervals are based on 40 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details. Sample: Female partners of men aged 30 in 2011.

In addition to studying resilience to a specific shock, this analytical framework could also be used to determine an overall score of resilience for the population of interest, by considering the effects of multiple shocks weighted by the likelihood of their occurrence (as estimated in the data). To be noted, the result that a partnership dissolution lowers on average equivalised disposable income takes into account that adjustments are made on different life domains. However, the proportion of the total effect that is mediated is higher for individuals with lower level of education, suggesting that attenuation mechanisms – including the safety net provided by the welfare state – are stronger for this group (Figure 18).

Figure 18: Partnership dissolution: F index, by level of education



Note: The F index (eq. 4) measures the ratio of the total to the direct effect of the shock. Values above 1 indicate reinforcement mechanisms are at work, while values below 1 indicate attenuation mechanisms. Lines refers to different level of education. 90% confidence intervals are shown as shaded areas. Confidence intervals are based on 40 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details. Sample: Partnered men aged 30 in 2011.

10 Conclusions

In this paper we have illustrated a new approach to the study of the complex interactions between life domains, which allows researchers to move beyond the limitations of existing data sources. The approach relies on a structural model projecting life trajectories over time, with a consideration of the heterogeneity of individual characteristics and experiences. This allows to investigate the overall impact of specific life events on any of the outcomes included in the model, as well as the construction of specific counterfactuals to disable individual causal pathways. Exploiting this feature, we have derived a framework for characterising feedback between life domains in terms of their attenuating or reinforcing mechanisms. An illustrative application to young adult men in the UK shows that partnership status is closely linked to most other life domains, with attenuating mechanisms that absorb around 20% of the total effect of a shock. On the other hand, health has fewer connections to other life domains, and shocks to health on average do not get attenuated or reinforced by the web of complex interactions governing life trajectories.

References

- Ananth C.V. (2019), Proportion mediated in a causal mediation analysis: how useful is this measure?, *BJOG: an International Journal of Obstetrics & Gynaecology*, 126: 983-983, <https://doi.org/10.1111/1471-0528.15691>.
- Altonji J.G., Martins A.P., Siow A. (2002), Dynamic factor models of consumption, hours and income, *Research in Economics*, 56(1): 3-59.
- Attewell P. (1999), The Impact of Family on Job Displacement and Recovery, *The ANNALS of the American Academy of Political and Social Science*, 562(1): 66-82, <https://doi.org/10.1177/000271629956200105>.

- Barigozzi M., Pellegrino F. (2023), Multidimensional Dynamic Factor Models, paper available at <https://arxiv.org/pdf/2301.12499.pdf>.
- Barbuscia A., Cambois E., Pailhé A., Comolli C.L., Bernardi L. (2022), Health after union dissolution(s): Cumulative and temporal dynamics. *SSM Popul Health*, 7(17):101042, doi: 10.1016/j.ssmph.2022.101042.
- Bijlsma M.J., Wilson B. (2020), Modelling the Socio-Economic Determinants of Fertility: A Mediation Analysis using the Parametric g-formula, *Journal of the Royal Statistical Society Series A: Statistics in Society*, 183(2): 493–513, <https://doi.org/10.1111/rssa.12520>.
- Black S.E., Devereux P.J., Salvanes K.G. (2015), Losing Heart? The Effect of Job Displacement on Health, *ILR Review*, 68(4): 833-861, <https://doi.org/10.1177/0019793915586381>.
- Bonekamp J., Wouterse B. (2023), Do different shocks in health matter for wealth?, *Journal of Health Economics*, 87:102719, <https://doi.org/10.1016/j.jhealeco.2022.102719>.
- Brand J.E. (2015), The Far-Reaching Impact of Job Loss and Unemployment, *Annual Review of Sociology*, 41:359–75.
- Bronka P., van de Ven J., Kopasker D., Katikireddi V.S., Richiardi M. (2023), SimPaths: An open source dynamic framework for life course analysis, *CeMPA Working Paper* 6/23.
- Child Maintenance Service (2024), How we work out child maintenance – A step-by-step guide, <https://assets.publishing.service.gov.uk/media/660be57967958c001f365a21/cmsb007gb-how-we-work-out-child-maintenance.pdf> (accessed 17/09/2024).
- Ditlevsen S., Christensen U., Lynch J., Damsgaard M.T., Keiding N. (2005), The Mediation Proportion: A Structural Equation Approach for Estimating the Proportion of Exposure Effect on Outcome Explained by an Intermediate Variable, *Epidemiology*, 16(1): 114-120, doi: 10.1097/01.ede.0000147107.76079.07
- Fadlon I., Nielsen T.H. (2021), Family Labor Supply Responses to Severe Health Shocks: Evidence from Danish Administrative Records, *American Economic Journal: Applied Economics*, 13(3): 1-30.
- Farber H.S. (2017), Employment, Hours, and Earnings Consequences of Job Loss: US Evidence from the Displaced Workers Survey, *Journal of Labor Economics*, 35(S1): S235-S272.
- García-Gómez P., van Kippersluis H., O'Donnell O., van Doorslaer E. (2013), Long Term and Spillover Effects of Health Shocks on Employment and Income. *J Hum Resour*, 48(4):873-909, doi: 10.1353/jhr.2013.0031.
- Glaser K., Stuchbury R., Tomassini C., Askham J. (2008), The long-term consequences of partnership dissolution for support in later life in the United Kingdom, *Ageing and Society*, 28(3):329-351, doi:10.1017/S0144686X07006642.
- Griffiths D., Sheehan L., van Vreden C. et al. (2021), The Impact of Work Loss on Mental and Physical Health During the COVID-19 Pandemic: Baseline Findings from a Prospective Cohort Study, *J Occup Rehabil* 31: 455–462, <https://doi.org/10.1007/s10926-021-09958-7>.
- Huttunen K., Kellokumpu J. (2016), The Effect of Job Displacement on Couples' Fertility Decisions, *Journal of Labor Economics* 34(2): 403-442.

- Lenhart O. (2019), The effects of health shocks on labor market outcomes: evidence from UK panel data. *Eur J Health Econ* 20: 83–98, <https://doi.org/10.1007/s10198-018-0985-z>.
- O'Donoghue C. (2014), *Handbook of Microsimulation Modelling*, Emerald Group Publishing Limited.
- O'Donoghue C., Dekkers G. (2018), Increasing the impact of dynamic microsimulation modelling, *International Journal of Microsimulation*, 11: 61-96.
- ONS, Office for National Statistics (2018), Birth characteristics in England and Wales: 2021, Statistical Bulletin, released 19 January 2023, ONS website
- Paul K.I., Hassel A., Moser K. (2018), Individual Consequences of Job Loss and Unemployment, in Klehe U.-C., van Hooft E. (eds), *The Oxford Handbook of Job Loss and Job Search*, <https://doi.org/10.1093/oxfordhb/9780199764921.013.028>.
- Preetz R. (2022), Dissolution of Non-cohabiting Relationships and Changes in Life Satisfaction and Mental Health. *Front Psychol*, 4(13):812831, doi: 10.3389/fpsyg.2022.812831.
- Richiardi M., Collado D., Popova D. (2021), UKMOD – A new tax-benefit model for the four nations of the UK, *International Journal of Microsimulation*, 14(1): 92-101.
- Richiardi M., Richardson R.E. (2017), JAS-mine: a new platform for microsimulation and agent-based modelling, *International Journal of Microsimulation*, 10: 106-134.
- van de Ven J., Bronka P., Richiardi M. (2022), Dynamic simulation of taxes and welfare benefits by database imputation, *CeMPA Working Paper* 3/22.
- Van de Ven J., Popova D. (2024), UKMOD country report 2021-2027, *CeMPA Working Paper* 4/24.
- van Soest A. (1995), Structural Models of Family Labor Supply: A Discrete Choice Approach, *Journal of Human Resources*, 30(1): 63–88, <https://doi.org/10.2307/146191>.
- Wörn J., Reme B.A., Skirbekk V. (2023), Job loss and psychological distress during the COVID-19 pandemic: a national prospective cohort study, *BMC Public Health*, 23(1):1447, doi: 10.1186/s12889-023-16303-5.

Appendix 1. Model structure and validation

A detailed description of the model – including regression results and validation statistics – is available in Bronka et al. (2023). Here we provide an overview of the various processes comprising the model, organised in modules as per Figure 3. We also provide a map of all relationships between state variables in the model (Table A2), and some key validation statistics (Figures A1-A15).

The list of all processes that form a simulation cycle is reported in Table A1.

Table A1: List of modules and estimated processes

Module	Process
Ageing	Age increases.
	Probability of leaving the parental home for those who have left education. (Students stay in the parental home).
Education	Probability of remaining in education for those who have always been in education without interruptions.
	Probability of returning to education for those who had left school.
	Level of education for those leaving education.
Health	Self-rated health status for those in continuous education.
	Self-rated health status for those not in continuous education (out of education or returned having left education in the past).
	Probability of becoming long-term sick or disabled for those not in continuous education.
	(Mental Health (1)) Level of psychological distress on GHQ-12 Likert scale and binary case-based indicator of psychological distress.
	(Mental Health (2)) Effect of exposure to employment-state transitions, household income change, and poverty for individuals aged 25 – 64 on psychological distress.
Family composition	Probability of entering a partnership for those in continuous education.
	Probability of entering a partnership for those not in continuous education.
	Probability of partnership break-up.
	Probability of giving birth to a child.
Social care	Probability of needing care for individuals over an age threshold.
	Probability of receiving care for individuals under an age threshold with a disability or long-standing illness or over the age threshold, distinguished by formal, partner, son, daughter, and other providers.
	Hours of care for those in receipt of care, and financial cost for those receiving formal care.
	Probability of providing informal social care.
	Hours of informal social care, among those providing care.
Investment income	Probability of retiring for single individuals.
	Probability of retiring for partnered individuals.
	Probability of receiving capital income for those in continuous education.
	Probability of receiving capital income for those not in continuous education.
	Amount of capital income for those in continuous education.
	Amount of capital income for those not in continuous education.
Labour income	Amount of pension income for those who are retired and were not retired in the previous year.
	Heckman corrected wage equation; females not employed last period.

	Heckman corrected wage equation; males not employed last period.
	Heckman corrected wage equation; females employed last period.
	Heckman corrected wage equation; males employed last period.
	Hours worked, single males.
	Hours worked, single females.
	Hours worked, single male adult children.
	Hours worked, single female adult children.
	Hours worked, males with dependent partner.
	Hours worked, females with dependent partner.
	Hours worked, couples.
Disposable income	Benefit reciprocity indicator.
	Amount of disposable income.
Consumption & saving	Consumption.
	Home ownership.
	Savings and assets.
Statistical display	Evaluate summary statistics for simulated population.

Following a standard approach in dynamic microsimulation modelling, each of the processes included in the model and described by eq. (1) is estimated separately in the data. Table A2 summarises all the interdependencies in the model.

Table A2: Relationship between variables in the model

Control variable	Dependent variable																			
	student status	education level	health status	mental health	disability status	partnership status	fertility	childcare cost	home owner	retirement status	pension income (£)	capital income (£)	low wage offer	potential wage (£)	hours worked	need social care	receive social care	type of care received	amount of care received	provide social care
gender	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c
age	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c	c
education	l	l	c	c	c	c,l	c	c	c	c	c	c	c	c	c	c	c	c	c	c
maternal education	c	c		l										c						
paternal education	c	c												c						
partnership status			l	c,l	l	l	c,l	c	l	c	l	l			c	c	c	c	c	c
number of children	l		l	l	l	l	l	c	l		l	l			c					
age of children	l					l	l	c												
health status			l	l	c,l	c,l	c		l		l	l	c	c		c	c		c	c
mental health				l																
disability status			l	l	l					l				c				c		
need social care																l				
receive social care																	l			
type of care received																		l	c	
amount of care received																				
provide social care																				l
amount of care provided																				c
activity status	l		l	c,l		l	l	c	l	l	l	l	l	c						
hours worked														c	c,l					
disposable income (£)			l	c,l	l	l			l	l					c					
employment income (£)												l								
benefit income (£)				c																
capital income (£)						l			l			l,l2								
pension income (£)						l			l		l,l2	l,l2								
potential wage (£)						l			l		l			l						
home owner				c					l											
region	c	c	c	c	c	c	c	c	c	c	c	c	c	c		c	c	c	c	c
year	c	c	c	c	c	c	c	c	c	c	c	c	c	c						

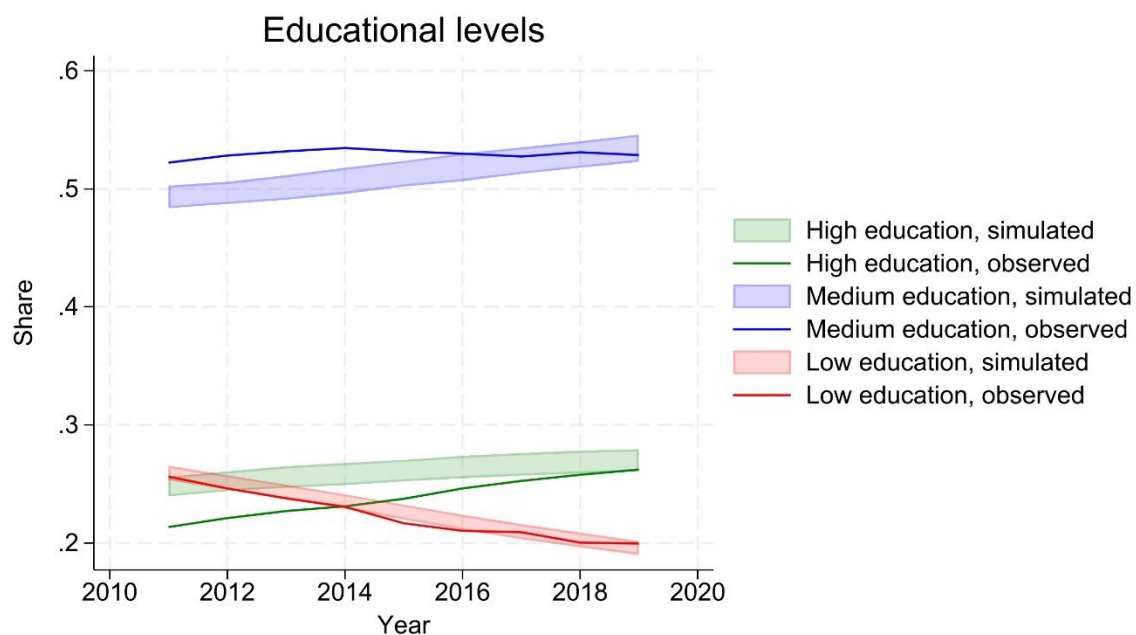
(c) denotes covariate reported in same period as projected characteristic. (l) denotes covariate lagged one period relative to projected characteristic. (l2) denotes covariate lagged two periods relative to projected characteristic.

All models are probability transition models, with the exception of the labour supply module, employing a structural random utility model (RUM). Bronka et al. (2023) provides more details on specifications and estimates.

The model is validated cross-sectionally over the period 2011-2019. Validation measures include the share of students by age, educational attainments, health score by age and gender, psychological distress by age and gender (score and caseness), partnership status, number of children, activity status, employment rate by age and gender, hourly wages (averages and distribution), hours worked, gross income by age and sources, net income by age, equivalised disposable income, poverty rate, and various inequality measures. We also validate pairwise correlations between all the variables in the model.

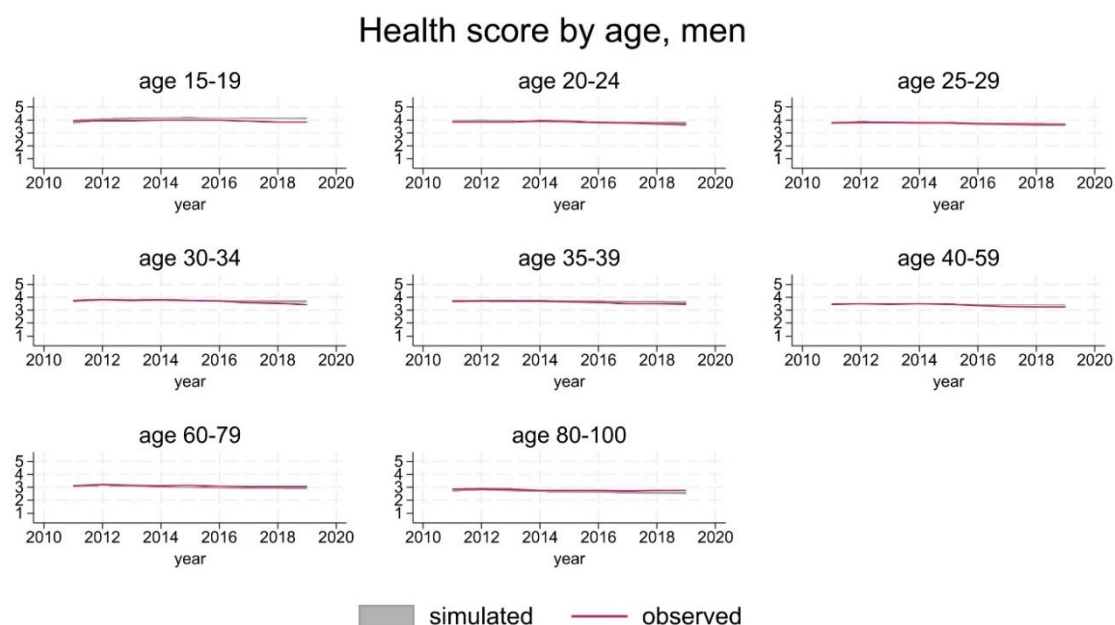
Selected validation graphs are reported below – the reader is referred again to Bronka et al. (2023) for the full set of indicators, as well as a discussion of the performances of the model.

Figure A1: Educational attainment



Note: 95% confidence intervals are shown as shaded areas. Confidence intervals are based on 20 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details.

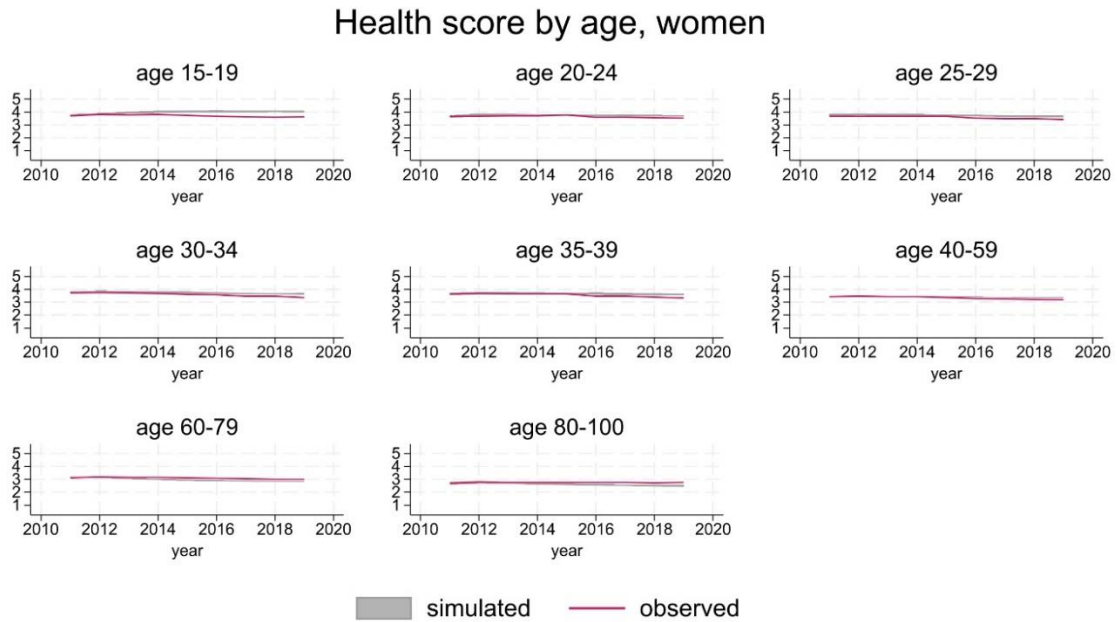
Figure A2: General health score, men



Note: 95% confidence intervals are shown as shaded areas. Confidence intervals are based on 20 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation

starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details.

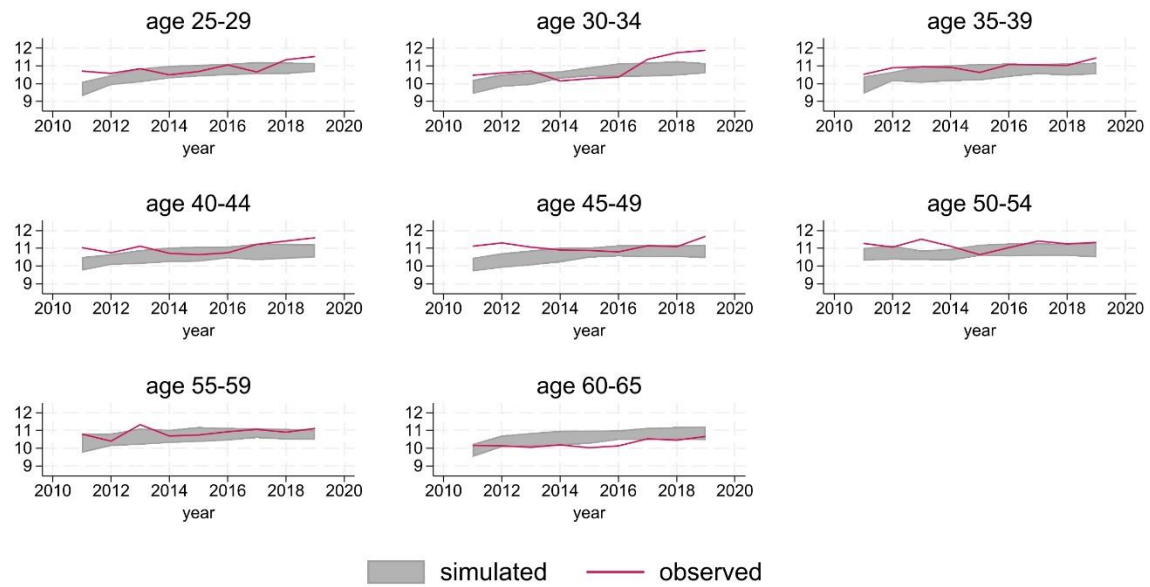
Figure A3: General health score, women



Note: 95% confidence intervals are shown as shaded areas. Confidence intervals are based on 20 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details.

Figure A4: Psychological distress, men

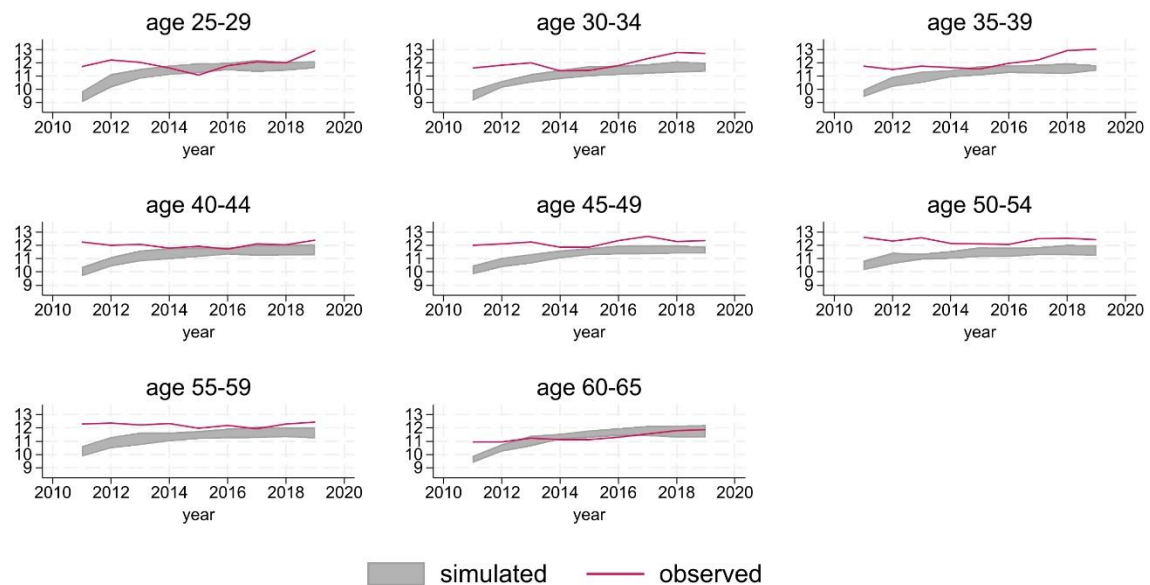
Psychological distress score by age, men



Note: 95% confidence intervals are shown as shaded areas. Confidence intervals are based on 20 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details.

Figure A5: Psychological distress, women

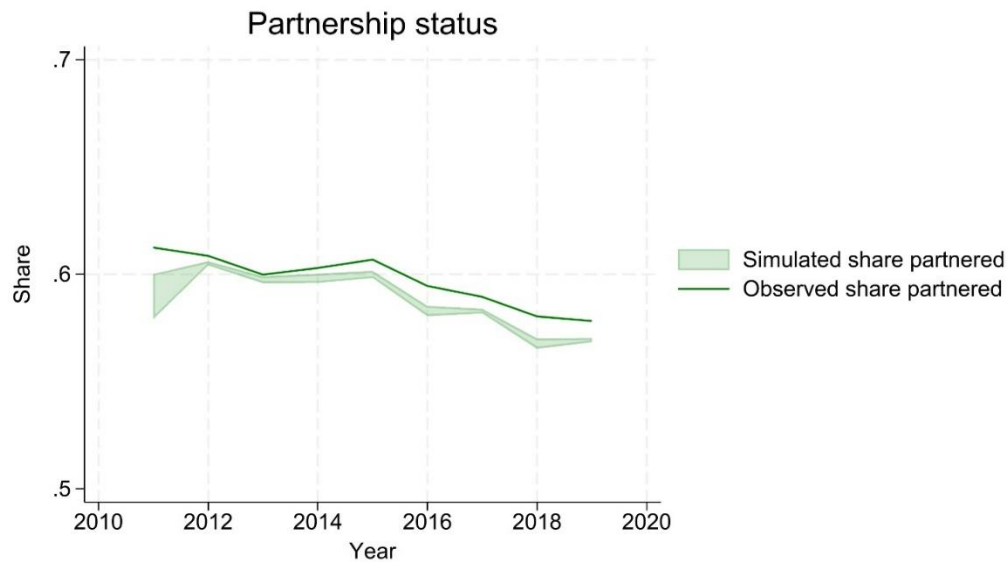
Psychological distress score by age, women



Note: 95% confidence intervals are shown as shaded areas. Confidence intervals are based on 20 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation

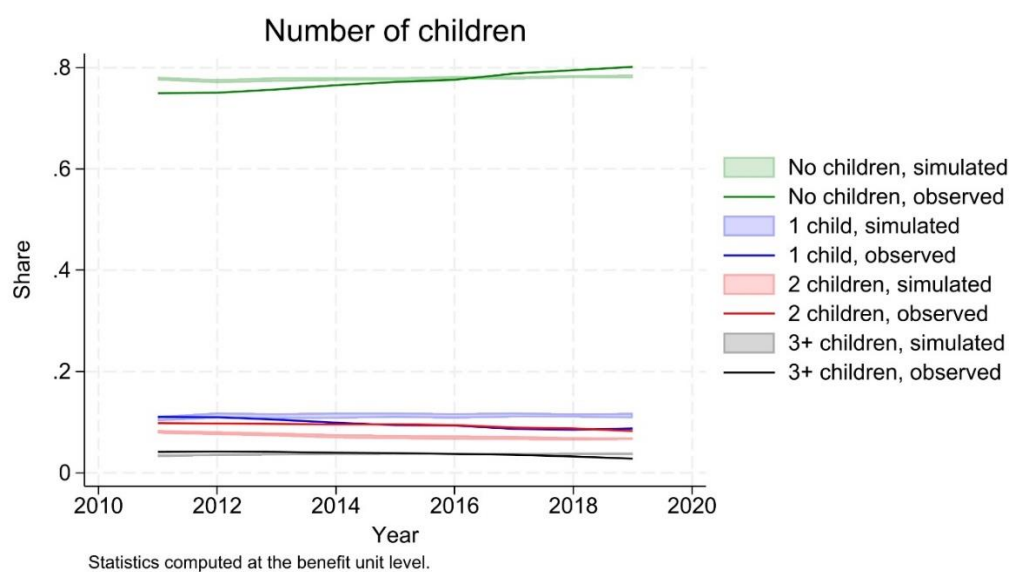
starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details.

Figure A6: Partnership status



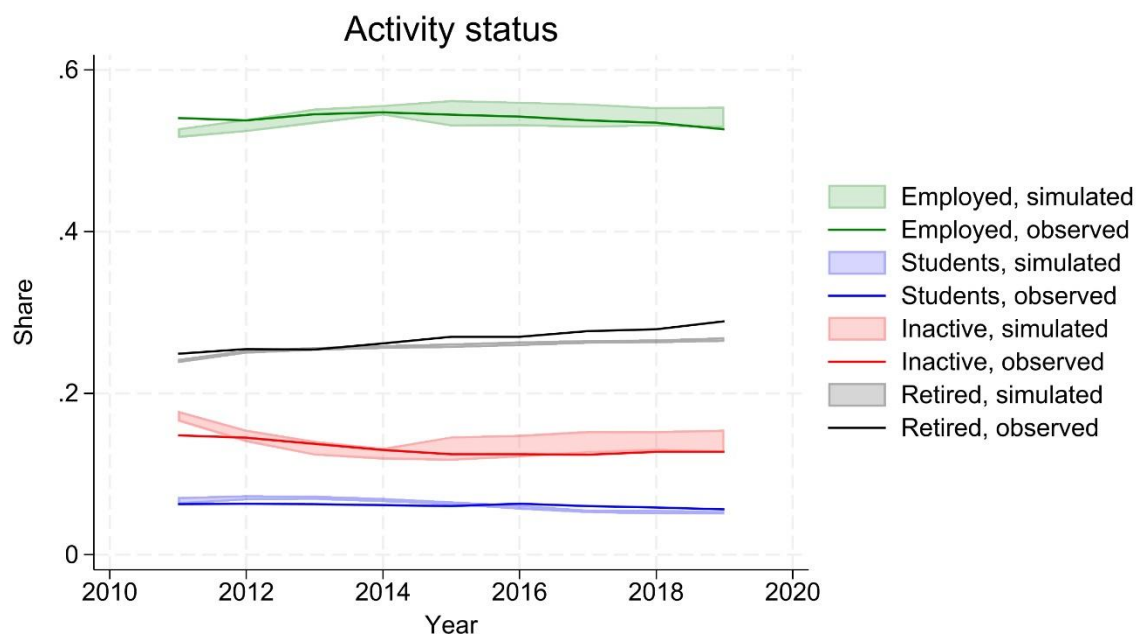
Note: 95% confidence intervals are shown as shaded areas. Confidence intervals are based on 20 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details.

FigureA7: Number of children



Note: 95% confidence intervals are shown as shaded areas. Confidence intervals are based on 20 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details.

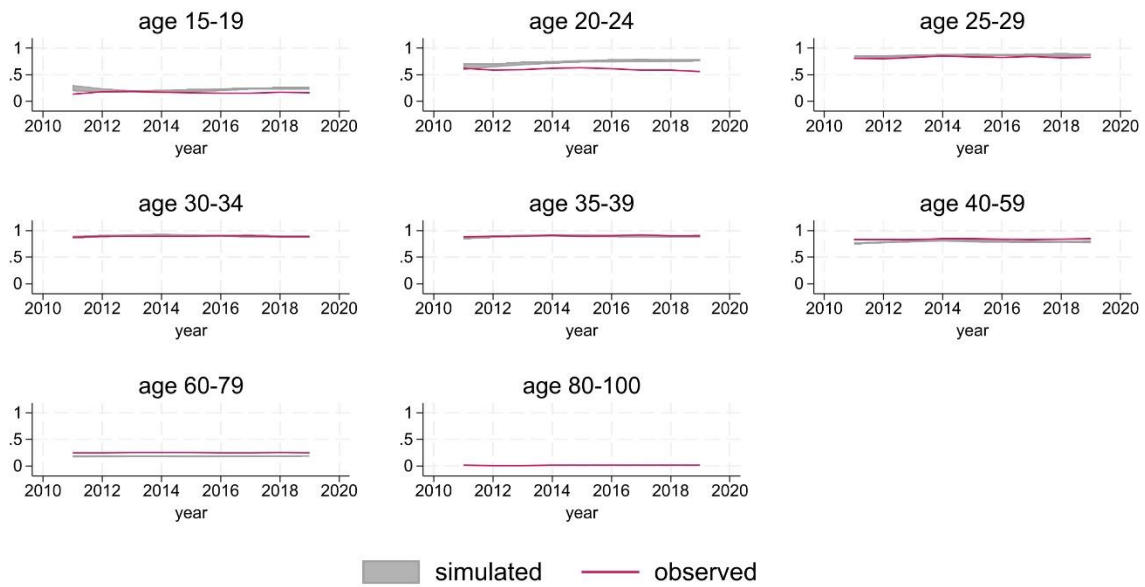
Figure A8: Activity status



Note: 95% confidence intervals are shown as shaded areas. Confidence intervals are based on 20 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details.

Figure A9: Employment rates, men

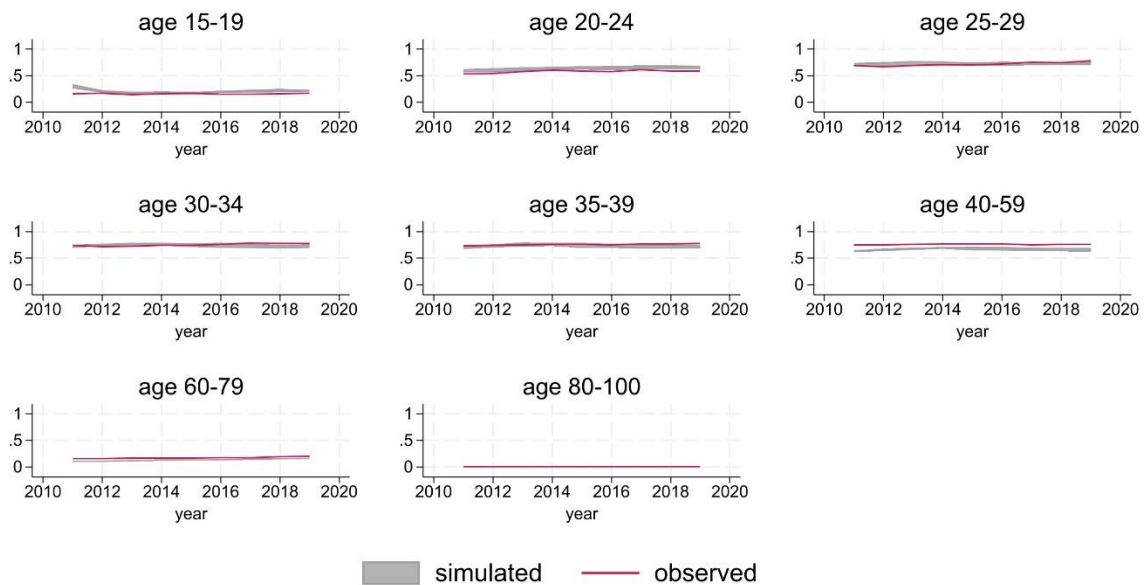
Employment rate by age, men



Note: 95% confidence intervals are shown as shaded areas. Confidence intervals are based on 20 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details.

Figure A10: Employment rates, women

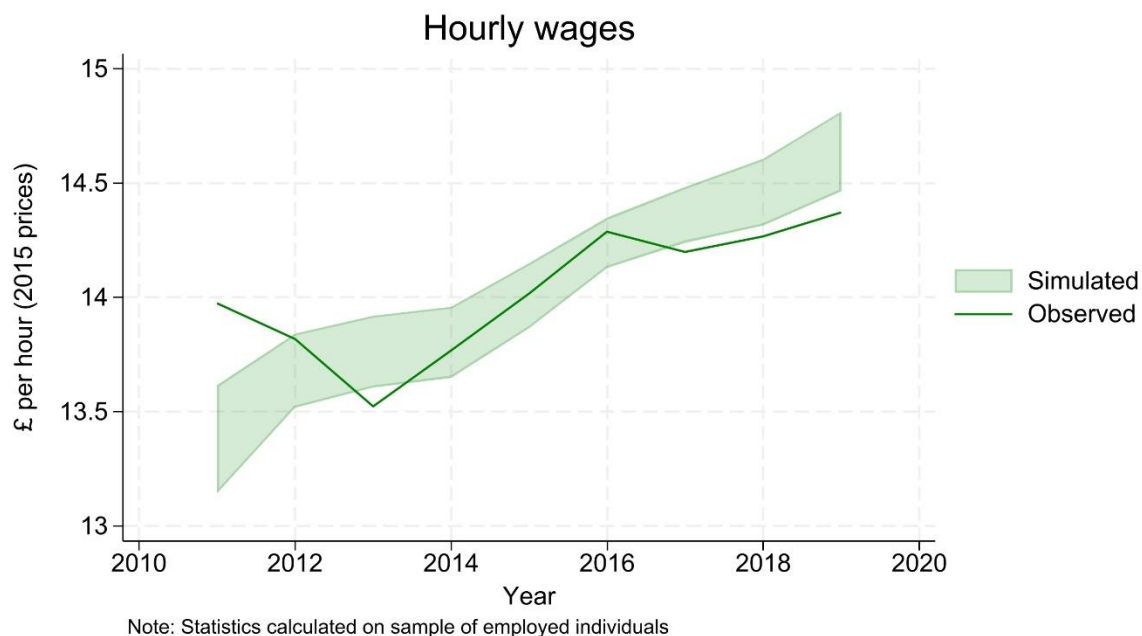
Employment rate by age, women



Note: 95% confidence intervals are shown as shaded areas. Confidence intervals are based on 20 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation

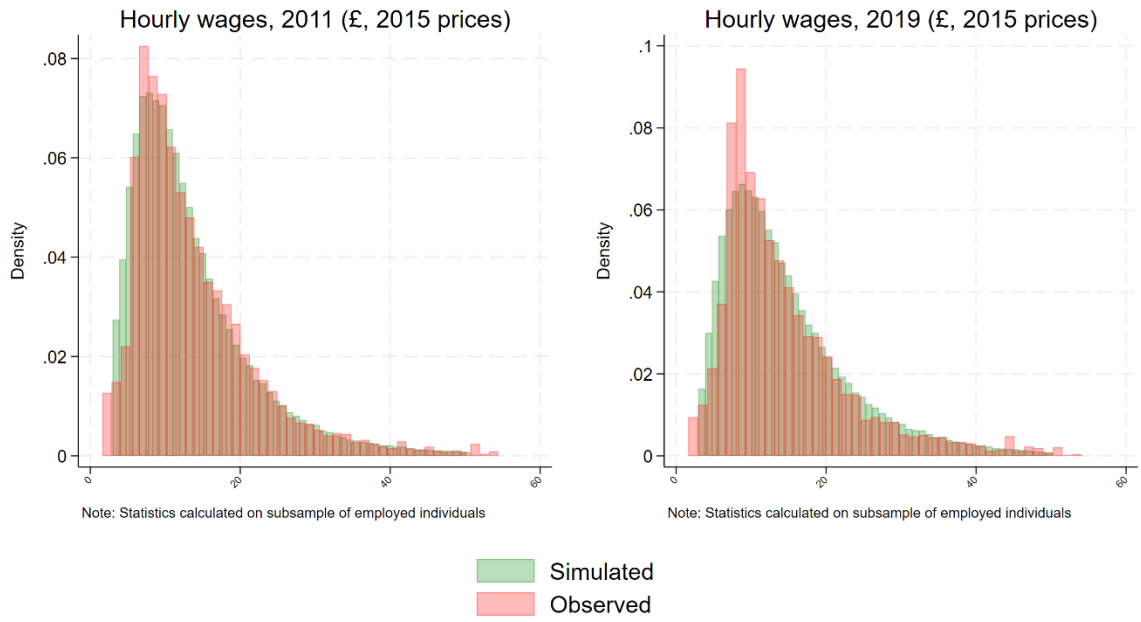
starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details.

Figure A11: Real wages, trend

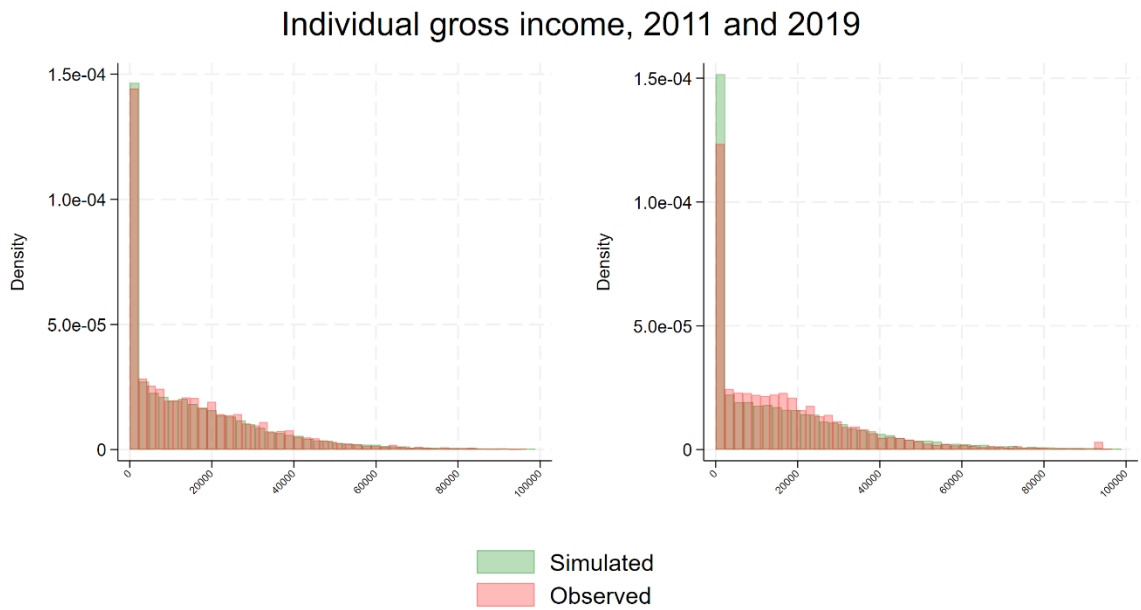


Note: 95% confidence intervals are shown as shaded areas. Confidence intervals are based on 20 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details.

Figure A12: Real wages, distribution

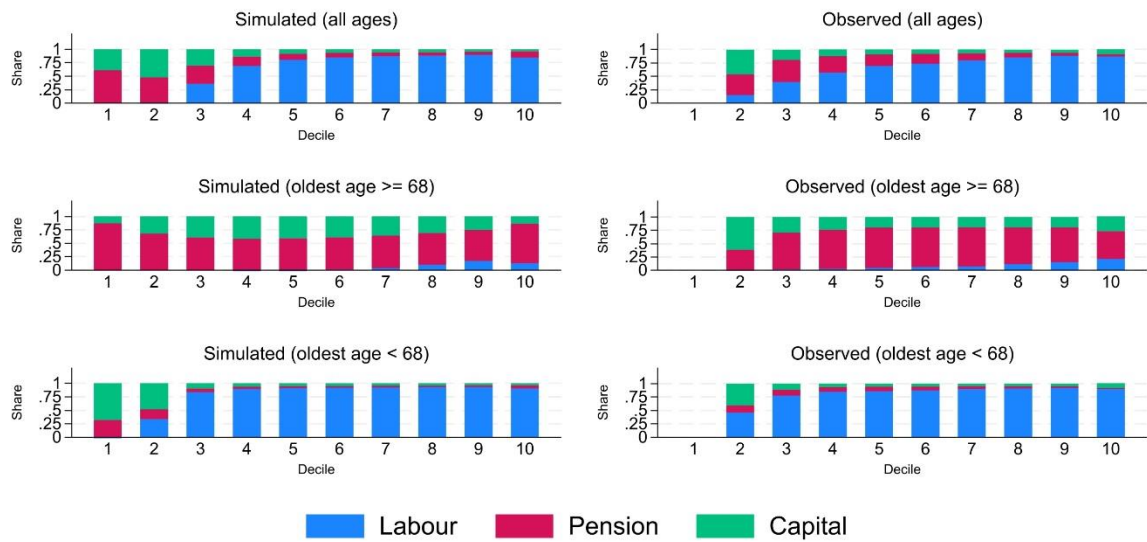


FigureA13: Gross income, distribution



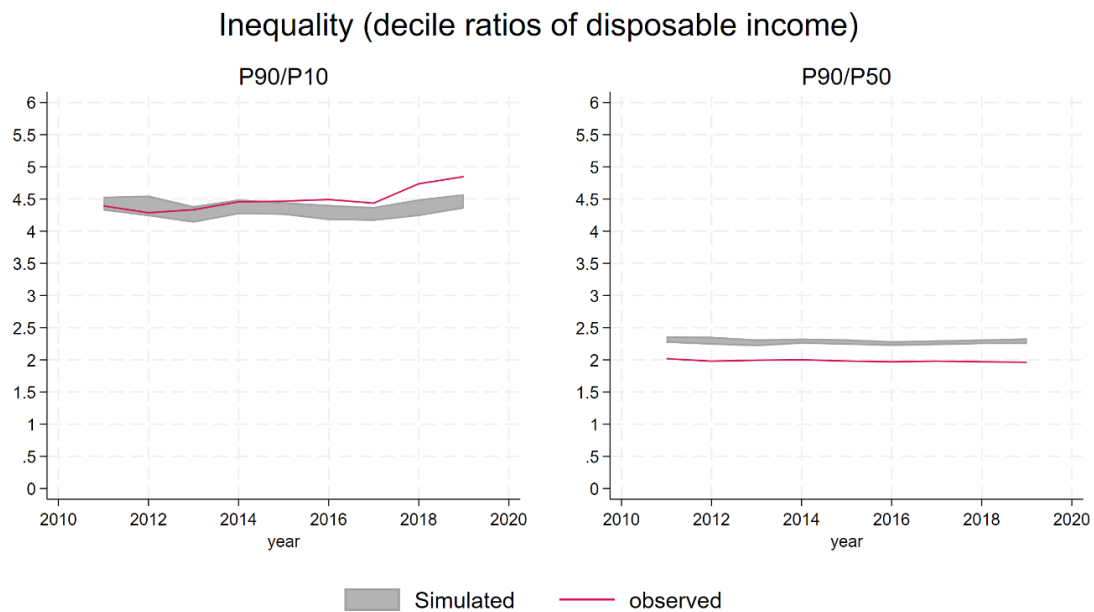
FigureA14: Income sources, share

Gross income shares, by age and decile



Statistics computed at the benefit unit level.
Based on values in £ per year (2015 prices).

Figure A15: Inequality



Note: 95% confidence intervals are shown as shaded areas. Confidence intervals are based on 20 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details.

Appendix 2. Descriptive statistics

In this Appendix we compare the characteristics of our samples (men aged 30 observed as partnered in 2011 for the shock to partnership, and all men aged 30 for the health shock) with external statistics. The initial population for the simulation is drawn from UKHLS. For comparison, we use data coming from the Family Resources Survey (FRS) – more specifically the FRS-based input data for UKMOD (Richiardi et al., 2021; van de Ven and Popova, 2024).

Table A3: Characteristics of the partnership shock sample²³

Variable	Simulated Data 2011	UKMOD Data 2011
Observations	8401	431
Age	30 (0)	30 (0)
Share partnered	1	1
High education	0.395	0.411
Medium education	0.534	0.558
Low education	0.071	0.031
Sex	1	1
Share disabled	0.006	0.010
Weekly hours of work	37.80 (11.79)	39.18 (13.42)
Share employed	0.943	0.936
Share students	0.018	0.010
Share not employed	0.039	0.054
Gross income (2015 £)	2417.08 (1632.04)	2361.79 (1895.48)

Note: Standard deviations (for continuous variables) in parenthesis.

Table A4: Characteristics of the health shock sample

Variable	Simulated Data 2011	UKMOD Data 2011	Simulated Data 2019	UKMOD Data 2019
Observations	25232	657	24622	841
Age	30 (0)	30 (0)	38 (0)	38 (0)
Share partnered	0.663	0.613	0.730	0.786
High education	0.378	0.379	0.389	0.380
Medium education	0.543	0.583	0.532	0.577
Low education	0.079	0.037	0.078	0.043
Sex	1	1	1	1
Share disabled	0.016	0.032	0.019	0.037
Weekly hours of work	35.85 (14.07)	36.21 (15.76)	37.02(12.59)	36.01 (15.20)

²³ Note that due to the cross-sectional nature of the FRS and UKMOD data, a comparison of how the partnership shock sample has evolved between 2011 and 2019 is not possible, as in the simulation individuals are selected conditional on having a partner in 2011. Statistics are reported using cross-sectional weights, except for the number of observations which is reported unweighted.

Share employed	0.903	0.885	0.931	0.899
Share students	0.020	0.021	0.005	0.002
Share not employed	0.077	0.092	0.064	0.099
Gross income (2015 £)	2224.65 (1637.80)	2162.51 (1861.34)	3023.66 (2168.92)	2715.37 (2572.45)

Note: Standard deviations (for continuous variables) in parenthesis.

Appendix 3. Uncertainty in microsimulation models: Tunnels vs. Funnels

Uncertainty regarding a model's projections can arise from a variety of reasons (Bilcke et al., 2011; Creedy et al., 2007). In particular, sources of uncertainty are generally distinguished in (i) input data, for instance due to sampling errors in the initial population; (ii) model structure, that is the validity of the general modelling approach used (also called "methodological uncertainty"); (iii) model specification, which concerns the choice of the covariates and the functional forms used, and in particular the crucial assumption that any regularity observed in the data will not break up in the future; (iv) model parameters, pointing to the imprecision of the estimates and/or externally provided parameters; (v) Montecarlo variation of the model output, which originates from the fact that the simulated aggregate quantities are also imprecise estimates of the theoretical aggregate quantities that the model implicitly defines; and finally (vi) coding errors, pointing to "bugs" in the code used for estimation / simulation of the model. All the above sources are common to all empirically based, stochastic models used for projections.

Generally speaking, source (i) should be limited, due to the use of appropriate input data and sampling weights. Source (ii) is often left unexplored, by making the common assumption that the model is well specified. Insights on the importance of this source of uncertainty are then generally left to meta-analyses of different modelling approaches. Source (iii) is commonly dealt with by performing sensitivity and robustness analysis on model specification. Quantification of this source of uncertainty remains however elusive, and researchers often limit themselves to reporting measures of goodness of fit, to corroborate their choices. Montecarlo variation of the model outcome (source v) can be brought down to negligible by appropriately scaling up simulated population size, and/or performing multiple simulation runs and averaging results.²⁴ Source (vi) – errors in programming – is generally ignored. Bugs are however hammered out over time, especially for sustained modelling efforts which lead to a solid research infrastructure used across a range of different projects, as is the case for SimPaths.

²⁴ Increasing population size is statistically equivalent to running the simulation multiple times if the model is ergodic. On the other hand, if the model is non-ergodic (e.g. there are multiple equilibria) this is generally not the case. However, dynamic microsimulation models are mostly concerned with a time-limited evolution of the system under analysis (e.g. 50 iterations at a yearly frequency), starting from initial conditions that are generally taken from the data. In such a situation, the long-term properties of the stochastic dynamic system are of lesser importance (and more difficult to check), especially considering that some inputs of the model (e.g. population projections) keep changing the implicit underlying long-term statistical equilibrium of the system (see Grazzini and Richiardi, 2013 for a discussion). However, it is generally the case that the estimated processes rely on the ergodic assumption - as this is common in econometrics - and the resulting model is therefore ergodic. After increasing the simulated population size until enough statistical power for the statistics of interest is obtained, there remains little to be gained by running the simulation multiple times. Recurring to multiple simulations runs with smaller simulated population sizes is however a good strategy when running large simulations requires too much time, possibly exceeding memory. Note that the above discussion assumes fixed parameters – source (iv) considers the case when the parameters themselves are uncertain.

The remaining source of uncertainty that is amenable to statistical analysis is parameters uncertainty stemming from sampling errors in estimation (source iv). This is generally left unexplored in microsimulation studies, although this is recognised and criticised (see for instance Mitton et al., 2000; Goedemé et al, 2013).

There are two approaches that can be used to deal with this uncertainty (Creedy et al., 2007). The first is what we might label “brute force”, and prescribes bootstrapping the coefficients of the estimated equations from their estimated distribution (e.g. multivariate normal in case of multinomial probit regressions) with mean equal to the point estimate, and variance-covariance matrix equal to the estimated variance-covariance. Bootstrapping needs to be performed only once, at the beginning of the simulation: the entire simulation is then performed with the bootstrapped values of the coefficients. The second approach provides an approximation by assuming from the onset a normal distribution for the resulting confidence intervals, requiring fewer draws from the parameter distribution.²⁵

A value added of the JAS-mine simulation platform is that it allows for a simple implementation of the “brute-force” approach, by exploiting the bootstrapping feature of its Regression library within a multi-run implementation: the simulation is run many times, each using a different set of coefficients. The result is a distribution of model outcomes around the central projections obtained with the point estimates of the coefficients. Confidence intervals can then be computed based on the estimated distribution of model outcomes.

Typically, these confidence intervals do not increase significantly over time, except for a relatively brief initial period. This is also the case for most of the outcomes reported in this paper. There are a number of reasons that explain this counterintuitive feature. As discussed in the text (see eqs. 1-3) most processes in a microsimulation model share an autoregressive form of the type

$$y_{i,t+1} = \lambda y_{i,t} + \beta x_{i,t} + u_{i,t} \quad (\text{A2.1})$$

where y is the outcome of interest and u is noise. Our bootstrapping procedure entails running multiple simulations, each with a different set of coefficients β and, crucially, λ . However, the coefficients are generally estimated with low enough errors, implying that outcomes of different simulation runs diverge only slowly over time. Moreover, the persistency parameter λ is generally below 1, implying that the dynamic stochastic process y converges to an equilibrium – the quicker the lower the value of λ . The equilibrium itself is different across simulation runs, given that β is different, but the differences remain bounded. Finally, even in an hypothetical case where λ was above 1, divergence of aggregate outcomes at the population level would be limited. This is because at an individual level each process is replicated, with characteristics x that typically either remain constant (e.g. sex, education) or change only slowly with time (e.g. age, income), only for a limited number of periods, after which the individual is either not at risk anymore (e.g. for the education, fertility, employment) or is removed from the population. New individuals might enter the simulation, but starting from similar initial conditions. This implies that, at a population level, different simulations with different values of the coefficients converge to different equilibria irrespective of the value of λ , rather than keep diverging over time.²⁶

²⁵ Note that bootstrapping requires multiple simulation runs with different sets of coefficients, differently from the multiple simulations runs with constant coefficients considered for source (v).

²⁶ This effect dampens the impact of uncertainty over λ even when λ is below 1.

This can be contrasted with the typical 'funnel' charts produced by uncertainty over the growth rate of a variable (e.g. GDP), where in different scenarios the outcome grows at different exponential rates and therefore quickly diverges. Here, the persistency parameter is above 1. Moreover, there are no population effects where outcomes are continuously reset and averaged. In dynamic microsimulation, with limited persistency, small confidence intervals, and continuous changes in the simulated units, uncertainty takes the form of 'tunnels' rather than 'funnels'.

Appendix 4. Effects of partnership on employment outcomes

We mentioned in the text that the negative impact of a partnership dissolution on gross income (Figure 5) is due to decreased employment and hours worked. Comparing partnered and unpartnered men in the baseline supports this claim (Table A5). 94% of partnered men are employed, against 82% of unpartnered men. When working, partnered men work slightly longer hours (40.1 hours per week against 38.6) and earn higher hourly wages (£14.5 against £12.6).

Table A5: Partnered vs. Unpartnered men aged 30 in 2011: Descriptive statistics

Variable	Obs.	Mean	Std. Dev.	Min/Max
Partnered men				
Employed	8,401	0.94		0 / 1
Weekly hours worked (if employed)	7,918	40.1	7.4	6 / 48
Hourly earnings (2015 £)	7,918	14.5	8.4	1.5 / 110.9
Single men				
Employed	4,199	0.82		0 / 1
Weekly hours worked (if employed)	3,444	38.6	10.0	6 / 48
Hourly earnings (2015 £)	3,444	12.6	7.7	1.5 / 150

These differences however, reflect in part a selection into partnership, in the baseline.

Indeed, our earning equation does not control for partnership status, so moving into a single status does not *per se* affect wages, in the simulations. However, partnership status matters for labour supply. Our random utility model (RUM) of labour supply is estimated separately for singles and couples – selected coefficients are reported in Tables A6 and A7.

Table A6: Labour supply estimates (utility function) – Single men

REGRESSOR	COEFFICIENT
IncomeDiv100	-0.026
IncomeSqDiv10000	0.000
MaleLeisure	0.181
MaleLeisureSq	-0.001
MaleLeisure_IncomeDiv100	0.000
FixedCost_Male	-1.922

Note: The specification also controls for other characteristics – see Bronka et al. (2023) for more details.

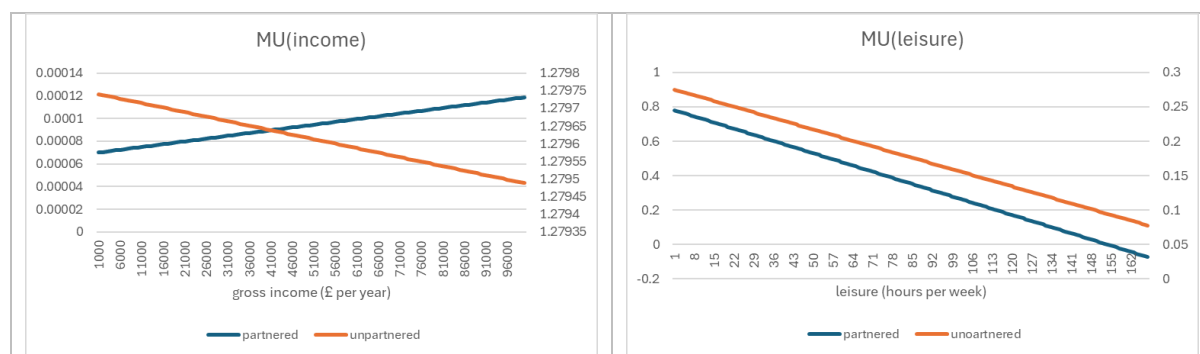
Table A7: Labour supply estimates (utility function) – Couples

REGRESSOR	COEFFICIENT
IncomeDiv100	-0.05109015
IncomeSqDiv10000	2.45749E-06
MaleLeisure	0.643
FemaleLeisure	2.200
MaleLeisureSq	-0.003
FemaleLeisureSq	-0.007
MaleLeisure_FemaleLeisure	0.000
MaleLeisure_IncomeDiv100	0.000
FemaleLeisure_IncomeDiv100	0.000
MaleLeisure_MaleAgeDiv100	0.064
MaleLeisure_MaleAgeSqDiv10000	0.078
FemaleLeisure_FemaleAgeDiv100	-0.047
FemaleLeisure_FemaleAgeSqDiv10000	0.095
FixedCost_Male	-1.581
FixedCost_Female	-3.734

Note: The specification also controls for other characteristics – see Bronka et al. (2023) for more details.

Based on those coefficients, we compute marginal utilities of income and leisure, for some standard characteristics (Figure A16).²⁷

Figure A16: Marginal utilities of income and leisure.



The marginal utilities of (male) leisure are decreasing with leisure – that is, increasing with hours worked) – both for couples and singles. On the other hand, the marginal utility of (male) income is increasing for couples (reflecting the utility that the partner receives from additional male income), and decreasing for singles. Hence, *ceteris paribus* men in couples supply more labour.

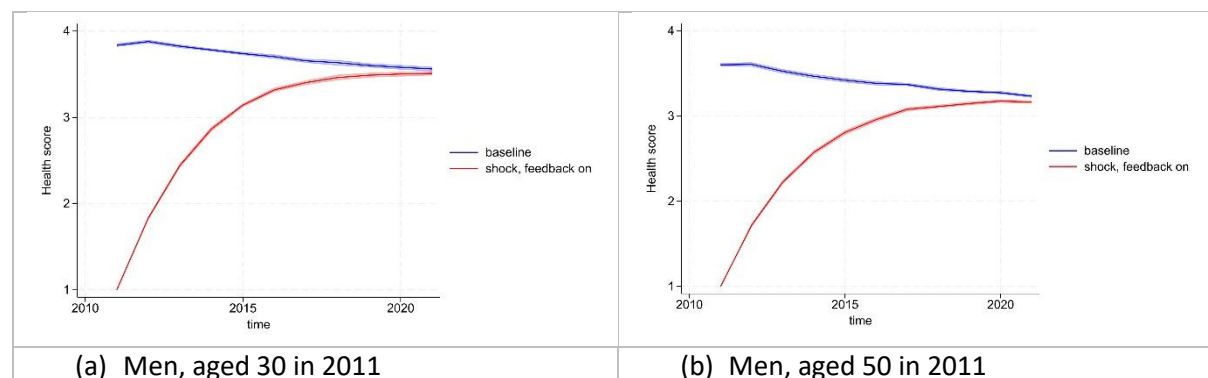
²⁷ Men aged 30 working 40 hours per week, with a spouse (if partnered) working 30 hours per week; total income of the household (when partnered): £50,000 per year. Patterns are robust to changes in those reference characteristics. Note that RUMs are unitary models of labour supply and assume that household maximise the joint utility of the partners.

Appendix 5. Effects of a health shock in later life

As discussed in the text, the rapid absorption of the health shock presented in Figure 8 is broadly explained by the young age of shocked individuals (30 in the initial year). To confirm the robustness of our results, we conduct additional experiments in which we shock men of older ages. In this Appendix we report the results of an experiment for men aged 50 in the initial year in the simulation. Coherently with our experimental setup for men aged 30, the health shock is implemented as a reduction of the self-rated health score to 1 in the initial year. Older men start from a lower average health score, so the initial shock is of lower relative intensity.

Figure A17 contrasts the evolution of the simulated self-rated health score for men aged 50 with that of men aged 30. As in Figure 8, we report the average health score in the baseline and in the ‘Feedback ON’ scenario (the difference between the two scenarios being our measure for the total effect of the shock) for the two groups of men.

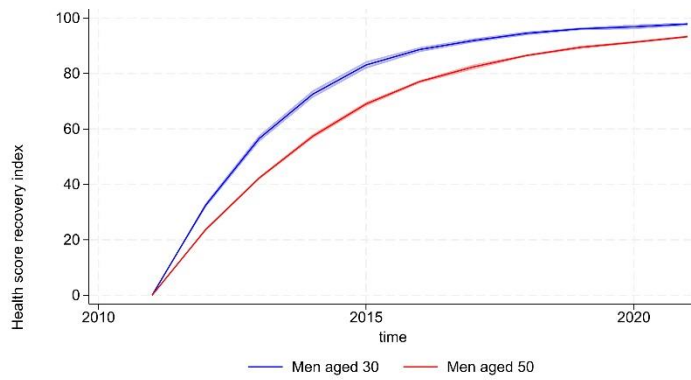
Figure A17. Health shock, comparison of total effects for men aged 30 and 50 in 2011



Note: 90% confidence intervals are shown as shaded areas. Confidence intervals are based on 40 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details.

The shock for the older cohort takes a bit longer to be absorbed, and it is still visible when the affected individuals reach the age of 60. The natural deterioration in health in the baseline however implies that at older ages differences become less and less appreciable. A comparison of the recovery trajectories from the initial shock is reported in Figure A18, where we normalised the health score in the ‘Feedback ON’ scenario with respect to the health score in the baseline (as already noted, the size of the shock is bigger for men aged 30).

Figure A18. Recovery from the health shock, cohorts of men aged 30 and 50 in 2011 compared



Note: The graph shows the average health score in the 'Feedback ON' scenario, normalised to the average health score in the baseline, representing the fraction of the original shock that has been absorbed in each period. 90% confidence intervals are shown as shaded areas. Confidence intervals are based on 40 simulation runs. In each run, coefficients are bootstrapped from their estimated variance-covariance matrix before the simulation starts - separately for each process - and kept constant for the entire duration of the simulation. See Appendix 3 for more details.