How do acute health shocks impact individual preferences? Evidence from Employment Decision.

Maria Plakhtieva*

November 2026

Click here for the latest version

Abstract

This paper studies whether acute health shocks - such as cancer, stroke, and heart attack - change the preferences that determine consumption and labour supply decisions. Using a combined British Household Panel Survey and Understanding Society dataset, together with a life-cycle consumption and labour-supply model, we estimate intratemporal, Frisch and Euler-equation conditions recursively in three steps by GMM, comparing treated individuals before and after the shock with an entropy-balanced control group. The analysis is motivated by an event-study, which first documents dynamic changes in labour supply outcomes of the individuals of working age (25 - 65) around the onset of the health shock. We report that male employment falls by 6.3 percentage points at onset and annual hours fall in the

*Plakhtieva: Department of Economics, School of Social Sciences, University of Manchester (maria.plakhtieva@manchester.ac.uk).

following year, while female annual hours start declining one year before the shock occurs. The subsequent structural estimation suggests that health shocks change how individuals trade off consumption and leisure, but not as much how they trade off the present and the future. Amongst women who remain in work, the utility weight on consumption falls by 2.7 pp relative to the control group, while for men the evidence is inconclusive. We find no significant change in the curvature of utility function or in the discount factor. Relative to the comparison group, the annual discount factor falls by 0.016 for women and is unchanged for men. If these shifts reflect preferences, life-cycle models that hold preferences fixed may therefore overstate the welfare loss from reduced work and the employment gains that can be achieved by policies that aim to bring people back to work.

1 Introduction

Health is known to be one of the main determinants of life choices and economic decisions. A large literature documents that health deterioration reduces employment, hours of work and earnings, and that these losses are quite persistent (García-Gómez et al., 2013; Dobkin et al., 2018; Lenhart, 2019; Jones et al., 2020). In England and the US, health deterioration explains up to 15% of the decline in employment between ages 50 and 70 (Blundell et al., 2023). Structural life-cycle models usually interpret these responses as the result of tighter constraints that individuals face. After a health shock, individuals may become less productive and earn lower wages, lose time to illness and treatment, face higher medical expenses and have lower chances of survival (French, 2005; De Nardi et al., 2010; Capatina, 2015). In this literature, the individual after the health shock is

the same person as before, but with a different choice set.

Surprisingly, little is known about whether health shocks also change individual preferences and values in life. Think of an individual who cares about having a good job and high earnings. After receiving a diagnosis, such as cancer, this individual may revise their life priorities, even if after the health shock they do not have any lasting health problems that limit their daily activities. In particular, they may start to attach more importance to leisure, such as time spent outside work or with their family, relative to what their income can buy. Some medical research documents this kind of adjustment: individuals facing a life-threatening or chronic illness change their internal standards and values, and reorder their life goals (Caumette et al., 2023; Sawatzky et al., 2024). If so, part of the reduction in labour supply after a health shock reflects a change in preferences and not simply a change in constraints. Even though it is intuitive that health shocks may affect and shape individual values and preferences, to our knowledge, this relationship has not been estimated from observed labour supply decisions. In a model with fixed preferences, all post-shock changes in behaviour are attributed to constraints, and such a model may therefore overstate the welfare loss from reduced work and the employment gains that can be achieved by policies that aim to bring people back to work.

This paper asks whether acute health shocks, namely cancer, stroke and heart attack, change the preferences that determine consumption and labour supply decisions. We use 19 waves of the British Household Panel Survey (BHPS) and Understanding Society (UKHLS), covering the years 2001–2019, and compare individuals of working age who experience an acute health shock with a control group of individuals without any of these shocks, reweighted by entropy balancing (Hainmueller, 2012). There are two

empirical analyses. In the first, we conduct an event study that documents how five key variables evolve before and after the shock: satisfaction with health, employment, annual hours of work, annual consumption of nondurable goods (measured by food and grocery expenditure) and annual labour income. In the second part, we estimate a life-cycle model of consumption and labour supply, in which health affects the wage, the time available for work and leisure, and the probability of survival. The parameters of the utility function are the within-period weight individuals put on consumption rather than leisure α; the curvature of utility (relative risk aversion) σ; the discount factor δ; and the elasticity of the marginal utility of consumption with respect to health ϕ. We estimate these parameters recursively in three steps by GMM from the intratemporal condition, the Frisch consumption equation and the Euler equation, separately for treated and control individuals before and after the shock, and compare the changes in a difference-in-differences design. In our model, the survival probability enters the effective discount factor separately from δ, which allows us to separate the effect of a health shock on time preferences from its effect on the expected survival.

Our main results suggest that acute health shocks change how individuals trade off consumption and leisure, but we find no evidence that they change how individuals trade off the present and the future. Among treated women who remain in work, the consumption weight α falls by 2.7 percentage points (pp) more than among women in the control group, i.e. after the shock women shift weight from consumption towards leisure more. When we compare treated women only with the small group of controls who never report any health problems, the difference-in-differences is smaller (1.8 pp) and not statistically significant. The estimate for men is similar in size (2.0 pp) but

not statistically significant. The curvature of utility, σ, equals unity and lies inside the 95% confidence intervals for all groups (treated and control, before and after the shock), which suggests that within-period preferences are close to logarithmic. Our difference-in-differences estimates of the discount factor δ, relative to the control group, fall by 0.016 for treated women and rise by 0.004 for treated men, but neither change is statistically significant. The point estimate for women goes in the same direction as the evidence that patients discount the future more heavily in medical decisions after a diagnosis of serious illness (Xu et al., 2024). In terms of annual discount rates, the difference-in differences is 1.6pp for women and -0.6 pp for men. The estimates of the health dependence of the marginal utility of consumption ϕ are sensitive to the choice of the comparison group, and we therefore treat them only as suggestive. Since α is identified from the ratio between the value of leisure and consumption, time absorbed by illness and treatment that is recorded as leisure would look like a shift in preferences. We discuss this and other threats to interpretation in Section 6, and, at this stage, the absence of a significant change in parameters should be read as an absence of evidence rather than as evidence of no change.

The event study shows that the shocks are indeed severe. Self-reported satisfaction with health falls after the health shock by about 0.8 of a category on a seven-point scale for men and by more than one category for women. Men's employment falls by 6.3 pp in the year the shock is first reported and their annual hours fall by about 130 in the following year (roughly 2.5 hours per week), but both recover within two years after the shock. This adjustment is temporary and partly happens at the extensive margin, which the structural sample of individuals who remain in work does not capture, and may explain

why we do not find a significant change in α for men. For women, employment falls by 4.1 pp in the year of the shock, while hours start to decline one year before the shock, which may be a result of symptoms and medical checks that often precede a diagnosis of cancer, the most common shock among women in the sample. Nondurable consumption is largely smoothed through the shock for both men and women, with a slight increase at the onset of a health shock, which is consistent with food being a necessity.

Economic models usually assume that preferences are stable and that differences in behaviour come from differences in prices and incomes (Stigler and Becker, 1977). The idea that preferences change after a health shock is, however, consistent with existing research. For example, using German panel data, Decker and Schmitz (2016) define a health shock as a large decline in measured grip strength, and find that this objective health change raises self-reported risk aversion and that the effect persists for at least four years. Banks et al. (2020) use data on individuals aged 50 and over in 14 European countries, and show that willingness to take financial risks declines with age, and that about half of this decline is explained by life events such as health shocks, retirement and widowhood, particularly for women. The evidence is not unanimous, however. In panel data from rural Paraguay, Chuang and Schechter (2015) find that health shocks do not change experimental measures of risk and time preferences. For time preferences, the link with health is usually thought to go through survival. Yaari (1965) shows that the survival probability enters the effective discount factor, so that a higher mortality risk makes individuals behave as if they were less patient even when their time preferences are unchanged. Becker and Mulligan (1997) go further and make time preference itself endogenous, so that a higher mortality risk can also lower patience, because it reduces

the return to investing in a more future-oriented attitude. Empirically, Chao et al. (2009) survey individuals in an area with high HIV-related mortality in South Africa, and show that people in very poor and very good health, and those with very low or very high expected survival, discount the future most heavily. At the same time, perceptions may adjust after the health shock. Baji and Bíró (2018) compare how subjective survival expectations, objective survival probabilities and self-reported health evolve after cancer, stroke or heart attack. They show that after a stroke or heart attack, subjective survival expectations and self-reported health go back towards their pre-shock levels, even though objective survival remains lower. This matters for our estimation, since we measure survival from adjusted ONS life tables. If individuals expect to live longer than these tables imply, our estimate of δ after the shock is biased upwards, which would hide a fall in patience. We add to this literature by estimating all parameters of a life-cycle utility function rather than a single preference measured in isolation, and by separating the effect of a health shock on time preferences from its effect on survival. We also infer preferences from observed consumption and labour supply choices rather than from survey or experimental measures.

The labour market effects of acute health shocks are well documented. Acute hospital admissions reduce the probability of employment and personal income (García-Gómez et al., 2013); sudden health declines lead to persistent reductions in earnings (Lenhart, 2019); and in the UK, adverse health shocks predict early retirement (Disney et al., 2006). Among older workers in Europe, an acute health shock roughly doubles the probability of leaving the labour market, and for women this response seems to be driven more by preferences for leisure and financial constraints than by health-related disability

(Trevisan and Zantomio, 2016). The closest paper to ours is Jones et al. (2020), who use the same definition of acute health shocks in the UKHLS data and find that labour market participation falls by about 3 pp one year after the shock, with larger effects for women, older workers and individuals with more severe limitations and impairments. We use a longer panel that also includes the BHPS waves, and go beyond the reduced-form effects by estimating the preference parameters of a life-cycle model. Structural life-cycle models give an interpretation to these reduced-form effects. The life-cycle model of consumption and labour supply goes back to Heckman (1974), and MaCurdy (1981) shows how hours respond to anticipated changes in wages through intertemporal substitution. French (2005) adds health to this framework, so that bad health lowers wages, reduces the time available for leisure and lowers the probability of survival, and his model explains a large part of the retirement behaviour of older men. Capatina (2015) models four channels through which health affects individuals, namely productivity, medical expenditure, the time endowment and survival, and finds that the productivity and time channels dominate. Low and Pistaferri (2015) build a life-cycle model in which workers face wage risk and work-limiting health shocks of varying severity and persistence to study the trade-off between insurance and incentives in disability insurance. They find that welfare increases as disability insurance programs become less strict. A separate literature allows the marginal utility of consumption to depend on current health, but the results disagree on the sign. The marginal utility of consumption is lower in poor health in Finkelstein et al. (2013), and higher in Kools and Knoef (2019), while Blundell et al. (2024) also find it lower in poor health and show that this channel explains most of the response of consumption to health shocks at older ages. In all of these models, however,

the preference parameters are fixed over time, and health affects behaviour through the budget constraint, the time constraint, survival and, in some cases, the level of marginal utility. Building on the mechanisms explored in this literature, our paper tests whether the parameters of a life-cycle utility function change after an acute health shock by comparing the same individuals before and after the shock with a weighted control group. Our results suggest that allowing the consumption-leisure trade-off to change after a health shock may matter for the predictions of these models, while there is still little evidence that the discount factor changes.

The rest of the paper is organized as follows. The theoretical model of consumption and labour supply and its analytical solutions are described in Section 2, while Section 3 introduces the data and describes the construction of the treated and control groups. Section 4 presents the event study, and Section 5 reports the structural estimates of the preference parameters. Finally, Section 6 discusses the results, their limitations and the next steps of this research.

2 The Model

To incorporate health shocks into a standard life-cycle model, this section describes the decision of an individual who chooses consumption and allocation of available time between leisure and labour over the life cycle, building on Heckman (1974). Each individual chooses consumption and hours of work (equivalently, leisure-hours) to maximize the expected lifetime utility. They are allowed to save for the future periods but not borrow, and, when making these decisions, they are faced with several forms of uncertainty, re-

lated to survival, health, and wage.

An individual maximizes expected lifetime utility over discrete periods indexed by age $t = t_0, \ldots, T$, where t_0 is the age of entry into employment, and T is the terminal period. Retirement occurs at age $R < T$, after which the individual is no longer employed and consumes out of accumulated assets and non-labour income. An individual is assumed to have a period utility $u_t = u(c_t, h_t, H_t)$, which has the same form at any t and depends on consumption c_t, time spent on work h_t, and health state H_t. The function is strictly concave, twice differentiable in c_t and h_t. In each period, an individual has a fixed amount of time $M_t = M(H_t)$ available so that if he decides to work h_t hours, his leisure time in this period can be expressed as $l_t = M(H_t) - h_t$. Note that total time available for allocation depends on the health state H_t, implying that ill health absorbs time that can be devoted to neither work nor leisure. Let $p_t = p(t, H_t)$ denote the probability of being alive at $t + 1$, conditional on surviving t and on health state H_t, with $p_T = 0$ in a terminal period. And let $P(j, t) = \prod_{k=t}^{j-1} p(k, H_k), \quad P(t, t) \equiv 1$ denote the probability of surviving from t to $j \geq t$, which depends on the realised health profile $(H_t, \ldots, H_{j-1})$ and is therefore stochastic from the perspective of period t. We assume that individual preferences take the following form:

$$u(c_t, h_t, H_t) + \mathbb{E}_t \left[\sum_{j=t+1}^{T} \delta^{j-t} P(j, t)\ u(c_j, h_j, H_j) \right], \tag{1}$$

where δ is the time discount factor.

This setup follows the exponential discounting, additionally separating pure time preferences δ from the survival probability $p(t, H_t)$. The survival-probability channel

was formalised by Yaari (1965), who shows that survival probabilities enter the effective discount factor when lifetime is uncertain. We, thus, incorporate this evidence from the literature in our model and, as $p(t, H_t)$ enters the effective discounting separately, ask whether a health shock changes δ separately from its effect on expected survival.

To estimate the model we follow French (2005) and model within-period utility as the product of two terms. First is the Cobb-Douglas with constant relative risk aversion (CRRA), defined on consumption and leisure. the second is that utility also depends on the health state of the individual, which enters multiplicatively, as introduced by Finkelstein et al. (2013). The within-period utility function is

$$u(c_t, h_t, H_t) = \frac{1}{1-\sigma}\{{c_t}^{\alpha}{l_t}^{1-\alpha}\}^{1-\sigma} \cdot H_t^{\phi}, \text{ where } l_t = M(H_t) - h_t, \tag{2}$$

and σ is the parameter of relative risk aversion, α is the weight on consumption in the Cobb-Douglas, and ϕ measures the elasticity of the marginal utility of consumption with respect to health[1]. We will refer to ϕ as the elasticity with respect to health throughout for simplicity.

The decision of an individual is thus an optimization problem in which he decides on the number of hours he works at each time period t. When working, an individual receives wage $w_t = w(t, H_t)$, which depends on the age-health profile, so that a deterioration in health lowers the return to work. We also assume that an individual has some non-labour income v_t, which may include benefits and transfers. Finally, r_t is the market interest rate, at which agent can save, and a_t is endowment. The choice of time alloca-

[1]Note that in this setup $\frac{\partial u_{c_t}}{\partial H_t} = \phi \frac{u_{c_t}}{H_t}$

tion in the current period and all future periods is then such that it maximizes future discounted utility from labour and consumption (1) until the terminal period T subject to the budget and additional constraints (3).

$$
\begin{aligned}
\text{(BC)} \quad & a_{t+1} = (1 + r_t)a_t + w_t h_t + v_t - c_t, \\
& h_t = M_t - l_t, \quad t = t_0, \ldots, T, \\
& a_{t+1} \geq 0, \\
& 0 \leq h_t \leq M_t, \\
& h_t = 0, \quad \text{for} \quad t \geq R, \\
& c_t \geq 0,
\end{aligned}
\tag{3}
$$

and the terminal condition $a_{T+1} = 0$. Health state therefore affects the employment behaviour through three channels: the marginal utility of consumption, via ϕ; the effective discount rate, via $p(t, H_t)$; and the choice set, via $M(H_t)$ and wage $w(t, H_t)$. First assume that the preference parameters $(\alpha, \sigma, \delta, \phi)$ are constant over the life cycle. Then consider whether an acute health shock changes preferences. To do this, allow them to take different values before and after the shock. Because the shock is irreversible, a change in parameters is a form of state dependence on health history, and is different from the dependence on health state captured by H_t. Section 2.2 allows these parameters to vary among the treated and control individuals before and after the health shock, while Section 5 compares these changes across the groups and time.

2.1 The solution

The sequential representation (1) treats the agent as choosing the entire contingent plan $\{c_t, h_t, a_{t+1}\}_{t=t_0}^{T}$ at $t = t_0$ subject to the period-by-period constraints (3). In the interest of space, we write marginal utilities as $u_c(t) = u_c(c_t, h_t, H_t)$ and $u_l(t) = u_l(c_t, h_t, H_t)$, where $u_l = -u_h$. First order conditions (FOC) provide us with the following system of equations:

$$\delta^{t-t_0} P(t, t_0) u_c(t) = \lambda_t, \tag{4}$$

$$\delta^{t-t_0} P(t, t_0) u_l(t) = \lambda_t w_t, \tag{5}$$

$$\mu_t + \mathbb{E}_t[\lambda_{t+1}(1 + r_{t+1})] = \lambda_t, \tag{6}$$

where λ_t is the Lagrange multiplier measuring the shadow value of income, and μ_t is the multiplier on the no-borrowing constraint. Equations (4) and (5) express marginal utilities of consumption and leisure, respectively, at every age t. The intratemporal FOC (7) below is then derived by dividing (5) by (4), and gives us first moment condition for the estimation. Combining (4) at t and $t+1$ with (6) give the intertemporal Euler equation (8). Finally, the Frisch condition (9) is derived from equations (5), (6), and intratemporal condition (7), and gives us an additional moment condition for the estimation.

$$\frac{u_l(t)}{u_c(t)} = w_t; \tag{7}$$

$$u_c(t) = \delta p_t \mathbb{E}_t[(1 + r_{t+1}) u_c(t+1)] + \frac{\mu_t}{\delta^{t-t_0} P(t, t_0)}; \tag{8}$$

$$\Delta \ln(\delta p_t) + \Delta \ln u_l(t+1) = \Delta \ln \lambda_{t+1} + \Delta \ln w_{t+1}. \tag{9}$$

2.2 Estimation strategy

An estimation strategy for the model is a three-step recursive identification of parameters of interest, using intratemporal condition, the Frisch elasticity condition, and Euler equations as moment conditions. The marginal utilities with the within-period form (2) are

$$\begin{aligned} u_c(t) &= \alpha\, c_t^{\alpha(1-\sigma)-1}\, l_t^{(1-\alpha)(1-\sigma)}\, H_t^{\phi}, \\ u_l(t) &= (1-\alpha)\, c_t^{\alpha(1-\sigma)}\, l_t^{(1-\alpha)(1-\sigma)-1}\, H_t^{\phi}. \end{aligned} \tag{10}$$

Substituting (10) into (7) gives an intratemporal condition

$$\frac{1-\alpha}{\alpha}\frac{c_t}{l_t} = w_t, \tag{IT}$$

which holds period by period whenever hours worked (and thus, leisure) are interior, and identifies the Cobb-Douglas share parameter α. Substituting (10) into (8), setting $\mu_t = 0$, and dividing both sides by $u_c(t)$ yields the estimating Euler equation

$$1 \;=\; \delta\,(1+r_{t+1})\,p_t\,\mathbb{E}_t\left[\left(\frac{c_{t+1}}{c_t}\right)^{\alpha(1-\sigma)-1}\left(\frac{l_{t+1}}{l_t}\right)^{(1-\alpha)(1-\sigma)}\left(\frac{H_{t+1}}{H_t}\right)^{\phi}\right]. \tag{EE}$$

Finally, substituting the marginal utility $u_l(t)$ from (10) into (9) yields the estimating Frisch equation

$$\Delta \ln c_{t+1} = \frac{\ln[\delta\, p_t\,(1+r_{t+1})]}{\sigma} - \frac{(1-\alpha)(1-\sigma)}{\sigma}\Delta \ln w_{t+1} + \frac{\phi}{\sigma}\Delta \ln H_{t+1}. \tag{FE}$$

The three-step estimation is then as follows. We first estimate α parameter from the intratemporal condition (IT), we then take this estimate $\hat{\alpha}$ and use the Frisch equation

(FE) to estimate the risk-aversion parameter σ, and then we utilize the panel data by estimating Euler equation (EE) at one- and two-period horizons to recover the discount factor δ together with the elasticity with respect to health ϕ, taking $\hat{\alpha}$ and $\hat{\sigma}$ as given.

We estimate all parameters separately for men and women. There are four groups g defined by treatment status (control or treated) and period (before and after the health shock). We seek to estimate each structural parameter across these groups to compute the difference-in-difference style estimate for α, σ, δ, and ϕ. All the moments are weighted by the EB weights discussed in Section 3. The intratemporal condition (IT) states that individuals equate the marginal rate of substitution between consumption c_t and leisure l_t in every period to the wage w_t. Under our within-period Cobb–Douglas specification, this implies $\dfrac{w_t l_t}{c_t} = \dfrac{1-\alpha}{\alpha}$, so that α is the share of consumption in full expenditure. This is a static, within-period optimality condition that holds regardless of the dynamic structure and regardless of other parameters of the model, allowing us to identify α entirely from the cross-sectional and time-series relationship between the consumption-leisure ratio and the wage. It requires, however, that hours are freely chosen and that leisure excludes time absorbed by ill health. We estimate the (IT) by GMM using the moment condition

$$\mathbb{E}\left\{ z_{it} \left[\ln\left(\frac{w_{it} l_{it}}{c_{it}}\right) - \ln\left(\frac{1-\alpha}{\alpha}\right) \right] \right\} = 0, \tag{M1}$$

where z_{it} is a vector of group indicators g. Because α is defined from $\ln\left(\dfrac{w_{it} l_{it}}{c_{it}}\right)$ any change in each of these elements around a shock is attributed to a change in α. Identification of the effect of the health shock rests on a parallel-trends assumption: absent the shock, the log ratio of the value of leisure to consumption of the treated would have

evolved as that of the weighted controls. Standard errors are clustered at the individual level and computed by the delta method. Raw means of $\ln\left(\frac{wl}{c}\right)$ for the treated and weighted control groups by time relative to the (pseudo) health shock are shown in Figure 7 in Appendix C.

Second step uses the dynamic information from the panel data to evaluate how consumption responds to the wage growth (through intertemporal substitution)[2]. We estimate the Frisch equation (FE) in first differences,

$$\Delta \ln c_{i,t+1} = \beta_{0,g} + \beta_{w,g} \Delta \ln w_{i,t+1} + \beta_{H,g} \Delta \ln H_{i,t+1} + \beta_{p,g} \ln p_{it} + \eta' X_{it} + \epsilon_{i,t+1}, \quad \text{(M2)}$$

where Δ denotes the first difference over consecutive waves[3], and $\beta_{w,g} = -\frac{(1-\alpha_g)(1-\sigma_g)}{\sigma_g}$ and $\beta_{H,g} = \frac{\phi_g}{\sigma_g}$. The vector X_{it} contains year fixed effects, marital status, a quadratic in age and household size. As in the first step, all slope coefficients are allowed to differ across the four groups g defined by treatment status and period, through a full set of treatment-by-period interactions. We exclude growth rates that span the health shock. Given the group-specific estimates $\hat{\alpha}_g$ from the first step, the structural parameters are recovered as

$$\sigma_g = \frac{1 - \hat{\alpha}_g}{(1 - \hat{\alpha}_g) - \hat{\beta}_{w,g}}, \quad (11)$$

with standard errors, clustered at the individual level, obtained by the delta method.

In the third step we estimate the Euler equation (EE) to recover the discount factor δ and the elasticity of consumption with respect to health ϕ. We take $\hat{\alpha}$ and

[2] While Frisch elasticity equation allows us to estimate σ from the wage coefficient, elasticity parameter ϕ is imprecisely estimated, so we recover it from the Euler equations in Step 3.

[3] $\Delta \ln x_{i,t+1} = \ln x_{i,t+1} - \ln x_{it}$ for any variable x

$\hat{\sigma}$ from the first two steps as given. Taking logs of (EE) gives the system of moment conditions

$$\begin{cases} \mathbb{E}_t\Big\{ z_t^1 \Big[\big(1-\alpha(1-\sigma)\big)\Delta \ln c_{t+1} - (1-\alpha)(1-\sigma)\,\Delta \ln l_{t+1} \\ \qquad\qquad - \ln\delta - \ln(1+r_{t+1}) - \ln p_t - \phi\,\Delta \ln H_{t+1} \Big] \Big\} = 0, \\ \mathbb{E}_t\Big\{ z_t^2 \Big[\big(1-\alpha(1-\sigma)\big)\Delta \ln c_{t+2} - (1-\alpha)(1-\sigma)\,\Delta \ln l_{t+2} \\ \qquad\qquad - \ln\delta - \ln(1+r_{t+2}) - \ln p_{t+1} - \phi\,\Delta \ln H_{t+2} \Big] \Big\} = 0. \end{cases} \tag{M3}$$

which is linear in the unknown parameters. When $\sigma = 1$, the leisure term drops out and (M3) reduces to the Euler equation of the logarithmic Cobb–Douglas utility $u(c_t, l_t) = [\alpha \ln c_t + (1-\alpha)\ln l_t]H_t^{\phi}$. This case is consistent with our second-step estimates, as $\hat{\sigma}$ is not significantly different from one in almost all specifications (see Section 5). The instruments z_t are dated $t-1$ or earlier. The vector z_t^1 contains lagged consumption and leisure, the productivity wage lagged one and two periods, the lagged industry wage, and education. The vector z_t^2 contains the same variables lagged one period further. We estimate this system separately for each gender and group g by two-step efficient GMM, starting from an identity weighting matrix, and we cluster standard errors at the individual level.

3 Data

Empirical estimation is conducted on two groups of individuals - with and without a health shock of interest. We refer to individuals who experienced a health shock as

"treated", and so individuals without a health shock form the "control" group. A health shock is thus considered to be a "treatment" and is assumed to be irreversible. While treated group consists of individuals for whom we observe an onset of cancer, stroke, or heart attack, all control individuals may be further divided into two groups - the group without cancer, stroke, and heart attack but with any other reported health issues such as problems with hearing or sight, which we call "other controls", and a "pure" control group, consisting of individuals without any reported health problems from the list throughout the survey. Pure and other controls do not overlap, and account for 15.4% and 84.6% of the full control group observations. Our main exercise is to estimate the time discounting δ and other parameters of the utility function - α, σ, ϕ - for treated and control groups before and after the health shock separately for men and women. Additionally, we exclude from the sample those individuals, who report having long-term health issues, such as epilepsy, or chronic diseases. These conditions typically develop gradually, so they may affect behaviour before they are first reported and there is no full recovery from such illness.

A health shock is defined by the incidence of cancer, stroke, or heart attack based on data from the British Household Panel Survey[4] (BHPS) and the Understanding Society[5] (UKHLS) data. We use eight waves of the BHPS and eleven waves of the UKHLS resulting in 19 waves of annual data spanning 2001 - 2019; this is because the question concerning cancer and stroke was only introduced in Wave 11 of the BHPS (corresponding year 2001). In this and subsequent waves of the survey, individuals are asked whether they have been diagnosed with any of a list of health conditions, which allows us to con-

[4] British Household Panel Survey - Institute for Social and Economic Research (ISER)
[5] Understanding Society, the UK Household Longitudinal Study.

struct a full health history of acute health shocks. The combination of these three health conditions, which we use to define a health shock, is common in the literature, as they typically occur suddenly and mostly without symptoms. However, while the UKHLS questionnaire asks about 'Heart attack or myocardial infarction', the BHPS asks about 'Heart/high blood pressure or blood circulation problems', which is a broader category, thus, we only use UKHLS observations to define heart attack as an acute health shock.

The demographic and socio-economic controls include age, gender, race, marital status, highest educational qualification, and individual and household labour and non-labour income. For labour market activity, at each wave, respondents are asked about their employment status, type of occupation, as well as the number of hours worked normally and over time (including paid and unpaid), individual labour income and other job characteristics. We focus our analysis on individuals who have experienced a health shock in working age[6], i.e., 25 - 65, excluding proxy individuals (interviews answered by another household member on the respondent's behalf).

To measure individual consumption c_{it} we use the weekly household food and grocery bill, converted to GBP per year, deflated to 2001 prices using the CPI, and normalized by the household composition. We use this measure as proxy for the nondurable consumption. Due to changes in questionnaire design between BHPS and UKHLS food expenditure variables are not directly comparable across the surveys. We rescale food expenditure data from the BHPS to match the level of the data from the first wave of Understanding Society. This yields a single multiplicative factor applied to all BHPS observations to ensure comparability across the BHPS and UKHLS. Hours h_{it} are usual

[6]The UK Pensions Act 1995 legislated equalisation of women's pension age with men's at 65.

weekly hours worked plus reported usual-week overtime, multiplied by number of weeks per year. Leisure is $l_{it} = M - h_{it}$, with a fixed annual time endowment of $M = 5,840$ hours that is not adjusted for health at this stage. Additionally, we measure an employment status from the self-reported current economic activity (employee or self-employed). And while individuals may report no economic activity, their hours worked are positive as they did some paid work in the year of the interview. The wage w_{it} is gross labour income annualized from gross monthly income, deflated to 2001 levels, divided by annual hours. The industry wage used as an instrument is calculated by two-digit SIC 2007 industry and year. The real annual interest rate r_t is taken from the Bank of England website for 2001–2019. Survival p_{it} are taken from ONS life tables by age, sex and year, and are adjusted for the excess mortality associated with cancer, stroke, and heart attack. Health H_{it} is self-reported satisfaction with health measured on a seven-point scale (1 - completely dissatisfied, 7 - completely satisfied), which we rescale linearly to zero-one scale, so that a one-category change corresponds to 0.14.

The estimation sample is restricted to individuals who are observed for at least two consecutive waves and did not enter the survey with a health shock (always-treated). Even though individuals report whether they have ever been diagnosed with a health condition from the list when they are surveyed for the first time, the timing of their health shock is unknown. We end up with the unbalanced panel of 53,168 individuals, comprising 361,205 observations (individual-waves). The number of treated individuals is 1,263, which is not out of line with that of other studies.

Since our main exercise involves comparing both treated and control individuals before and after the treatment, we generate a pseudo treatment date for the control group.

To do so, we use individual characteristics, and construct a health shock timing based on the probability proportional to the fitted hazard model on treated individuals. The hazard rate among treated is estimated based on their age, gender, marital status, race, education level, and calendar year. Table 1 reports raw descriptive statistics for the set of socio-economic covariates in the full control and treated groups prior to the (pseudo) health shock separately for men and women.

Table 1: Descriptive Statistics: Raw Data at $t = -1$

Variable	(1) Treated Mean	(2) Control Mean	(3) Difference	(4) Standardized Diff.	(5) N
Panel A: Men					
Weekly hours worked	41.669	40.739	0.930	0.066	15,470
Real monthly earnings	476.385	1,112.320	-635.935***	-0.541	24,106
Working status	0.275	0.626	-0.350***	-0.752	24,527
Age	63.134	44.468	18.666***	1.086	24,538
Marital status: single	0.168	0.079	0.089***	0.272	24,502
White	0.976	0.818	0.159***	0.541	24,538
Self-employed	0.212	0.190	0.022	0.054	15,765
Year of interview	2003	2006	-3.599***	-0.059	24,538
Work experience (years)	9.564	8.137	1.427***	0.221	19,058
Higher education	0.127	0.297	-0.171***	-0.427	24,538
No education	0.389	0.132	0.257***	0.612	24,538
Number of individuals	553	23,985			
Panel B: Women					
Weekly hours worked	32.118	30.851	1.267	0.087	15,272
Real monthly earnings	344.037	620.261	-276.224***	-0.351	28,301
Working status	0.310	0.506	-0.196***	-0.406	28,620
Age	59.859	44.931	14.928***	0.808	28,630
Marital status: single	0.331	0.170	0.161***	0.378	28,589
White	0.986	0.826	0.159***	0.569	28,630
Self-employed	0.109	0.090	0.019	0.063	15,495
Year of interview	2002	2009	-6.711***	-0.184	28,630
Work experience (years)	9.957	7.547	2.410***	0.385	20,052
Higher education	0.163	0.308	-0.145***	-0.346	28,630
No education	0.413	0.166	0.246***	0.564	28,630
Number of individuals	710	27,920			

Notes: The sample comprises individuals observed at $t = -1$, the last wave before the first (pseudo) health shock, with one observation per individual. Column (3) reports the coefficient on the treatment indicator from a regression of the variable on that indicator, with robust standard errors. Column (4) reports the standardized difference in means, $\frac{\text{mean}(T) - \text{mean}(C)}{\sqrt{[\text{Var}(T) + \text{Var}(C)]/2}}$. Column (5) reports the number of observations in the corresponding regression; counts differ across rows because hours worked, earnings and self-employment are observed only for individuals in work, and because of item non-response on work experience.
Significance level: *** p<0.01, ** p<0.05, * p<0.1

Comparison table shows that the two groups (treated and control) in the raw sample differ on most characteristics. Importantly, treated individuals are substantially older than controls, with a standardized difference in age of 1.09 for men and 0.81 for women, and they are less educated and less likely to be in employment. To improve the balance between the two groups we utilize the Entropy Balancing (EB) method (Hainmueller, 2012), which is applied separately to men and women using characteristics measured in the (pseudo) pre-treatment period ($t = -1$) and we restrict the balanced subsample to individuals who were employed in at least one wave of the survey. Specifically, we balance on age, education level, marital status, and race so that the control group data are reweighted to match the covariate moments in the treatment group. We deliberately do not balance the outcome variables such as hours worked, labour earnings or employment and their pre-shock balance serves as a check on the design. We impose balance on the first and second moments. Table 2 below reports descriptive statistics for the set of socio-economic covariates in the treated and control groups after the EB on individuals who are ever observed in employment in the data, are aged 25 to 65, and do not have missing variables needed for the structural estimation of parameters (consumption, hours worked, and labour income). The restriction is substantial: it retains 143 of 553 treated men and 209 of 710 treated women, and 13,865 of 23,985 control men and 13,532 of 27,920 control women. Within the retained sample, the treated are still older than the controls before reweighting (48.4 against 41.2 years for men and 44.2 against 40.7 for women), less educated, and more likely to be white. After reweighting, all targeted covariates are balanced, with standardized differences no larger than 0.02 in absolute value.

The outcome variables, which are not targeted, are also close to balance after

Table 2: Entropy-Balanced Sample at $t = -1$ on individuals of working age, who worked in at least one wave of the survey: Pre-Treatment Characteristics

Variable	(1) Treated Mean	(2) Control Mean Unweighted	(3) Control Mean EB	(4) Difference (EB)	(5) Standardized Diff. Unweighted	(6) Standardized Diff. EB	(7) N
Panel A: Men							
Weekly hours worked	42.699	42.365	42.662	0.037	0.026	0.003	14,008
Real monthly earnings	1,724.146	1,826.179	1,767.271	-43.125	-0.081	-0.035	13,731
Working status	0.979	0.985	0.991	-0.012	-0.044	-0.102	14,006
Age	48.378	41.167	48.371	0.007	0.626	0.001	14,008
Marital status: single	0.056	0.051	0.056	0.000	0.024	0.000	14,008
White	0.986	0.824	0.984	0.002	0.574	0.017	14,008
Self-employed	0.217	0.186	0.227	-0.010	0.077	-0.024	14,006
Year of interview	2003	2009	2009	-5.748***	-1.320	-1.197	14,008
Work experience (years)	12.797	9.408	9.955	2.842***	0.532	0.443	14,008
Higher education	0.238	0.379	0.238	0.000	-0.309	0.000	14,008
No education	0.175	0.068	0.175	0.000	0.331	0.000	14,008
Number of individuals	143	13,865	13,865				
Effective sample size			7,434				
Panel B: Women							
Weekly hours worked	32.734	32.173	30.884	1.850	0.040	0.133	13,741
Real monthly earnings	1,053.006	1,216.848	1,080.487	-27.481	-0.187	-0.033	13,558
Working status	0.971	0.962	0.975	-0.003	0.050	-0.021	13,738
Age	44.196	40.725	44.191	0.005	0.313	0.000	13,741
Marital status: single	0.115	0.113	0.115	0.000	0.005	0.000	13,741
White	0.990	0.850	0.988	0.002	0.536	0.020	13,741
Self-employed	0.101	0.088	0.087	0.014	0.046	0.049	13,738
Year of interview	2001	2009	2008	-7.026***	-1.925	-1.648	13,741
Work experience (years)	13.244	9.001	9.732	3.512***	0.698	0.566	13,741
Higher education	0.230	0.442	0.230	0.000	-0.460	-0.001	13,741
No education	0.153	0.056	0.153	0.000	0.320	0.000	13,741
Number of individuals	209	13,532	13,532				
Effective sample size			8,783				

Notes: The sample excludes individuals never observed in employment. The control group is reweighted by entropy balancing (Hainmueller, 2012), separately by sex, to match the treated group on the first and second moments of the variables shown in bold, measured at $t = -1$; balance on these moments holds by construction. The remaining variables are not targeted and are reported as checks. Column (2) reports unweighted control means and column (3) entropy-balanced control means in the same sample. Column (4) reports the coefficient on the treatment indicator from a weighted regression, with robust standard errors. Columns (5) and (6) report standardized differences, $\frac{\text{mean}(T) - \text{mean}(C)}{\sqrt{[\text{Var}(T) + \text{Var}(C)]/2}}$, computed without and with the entropy-balancing weights. The effective sample size is $\frac{(\sum_i w_i)^2}{\sum_i w_i^2}$ over the control group. Column (7) reports the number of observations in the corresponding regression.
Significance level: *** p<0.01, ** p<0.05, * p<0.1

reweighting. Employment at $t = -1$ is nearly the same in both groups (0.98 among treated men against 0.99 among weighted controls, and 0.97 against 0.98 for women), with standardized differences of -0.10 for men and -0.02 for women. Earnings differ by -0.04 and -0.03 standard deviations, and weekly hours among those in work by 0.00 for men and 0.13 for women. None of these differences is statistically significant at conventional levels. Work experience, which is not in the balancing set, remains higher among the treated by 0.44 standard deviations for men and 0.57 for women (2.8 and

3.5 years). The largest remaining imbalance is in calendar time: treated individuals are observed on average 5.7 (men) and 7.0 (women) years earlier, with standardized differences of -1.20 and -1.65. However, our estimations include calendar-year fixed effects, so aggregate shocks, macroeconomic conditions and changes in survey design that are common to treated and controls in a given year do not bias the comparison.

Table 8 in Appendix A reports descriptive statistics for the treated and control groups after the EB on the full sample, including never employed throughout the survey.

4 Dynamic Effects of Acute Health Shocks

We first conduct an event-study design (our notation and estimation method follow Sun and Abraham (2021)) to reveal changes in labour market outcomes due to experienced health shocks. In this exercise we focus on five main outcomes of interest: employment status, the number of hours individuals allocate to work, their labour earnings, and individual consumption of nondurable goods, as well as subjective satisfaction with health. This reduced-form analysis is primarily illustrative as it cannot distinguish between changes in preferences and changes in the constraints individuals face, or changes in expected survival. The main exercise of the paper is therefore the structural estimation in Section 5, which recovers the preference parameters (the consumption share α, relative risk aversion σ, the discount factor δ, and the health elasticity ϕ) before and after the shock, and asks whether a health shock alters preferences themselves.

We observe a panel of $i = 1, \ldots, N$ individuals and index time by age a. Treatment status is then denoted as $D_{i,a} \in \{0, 1\}$, such that $D_{i,a} = 1$ if individual i experienced a

health shock by age a and $D_{i,a} = 0$ if not. The age at the first occurrence of a health shock is then denoted with $E_i = \min\{a : D_{i,a} = 1\}$. For never-treated individuals, i.e., if $D_{i,a} = 0$ for all a, we set $E_i = \infty$ implying that a health shock might happen in the long-term future, but we never observe it.

It is standard to estimate dynamic treatment effects with the two-way fixed effects (TWFE) estimator, whereby the following model is estimated by the OLS regression:

$$Y_{i,a} = \gamma_i + \xi_a + \eta_{year} + \sum_{e=-K}^{-2} \mu_e D_{i,a}^e + \sum_{e=0}^{L} \mu_e D_{i,a}^e + \beta X_{i,a} + \nu_{i,a}. \tag{12}$$

Here $Y_{i,a}$ is the outcome of interest for individual i at age a; γ_i, ξ_a, and η_{year} are individual, age and calendar-year fixed effects; $X_{i,a}$ is a vector of controls; and $\nu_{i,a}$ is the error term that captures all unobserved determinants of $Y_{i,a}$ for individual i at age a. We define $D_{i,a}^e = \mathbb{I}\{a - E_i = e\}$ to be an indicator of being e periods away from a health shock. For never-treated individuals $E_i = \infty$, for all e we set $D_{i,a}^e = 0$ at any age a. However, when treatment effects vary across cohorts, each μ_e in (12) is a weighted average of cohort-specific effects that may be biased by effects from other relative periods (Goodman-Bacon, 2021; Sun and Abraham, 2021). As Wooldridge (2021) shows, once the event-time indicators are interacted with cohort indicators, TWFE identifies the cohort-specific average treatment effects on the treated under the same no-anticipation and parallel-trends assumptions as Callaway and Sant'Anna (2021).

We therefore estimate the interacted regression

$$Y_{i,a} = \gamma_i + \xi_a + \eta_{year} + \sum_{s} \sum_{\substack{e=-K \\ e \neq -1}}^{L} \mu_{s,e} \, \mathbb{I}\{E_i = s\} \, D_{i,a}^e + \beta X_{i,a} + \nu_{i,a}, \tag{13}$$

using never-treated individuals as the comparison group. Each $\mu_{s,e}$ identifies the average effect of a health shock e years after it occurs for individuals first treated at age s. Following Sun and Abraham (2021), we aggregate these into event-time effects

$$\hat{\mu}_e = \sum_s \hat{\omega}_{s,e}\, \hat{\mu}_{s,e}, \tag{14}$$

where $\hat{\omega}_{s,e}$ is the sample share of cohort s among treated individuals observed at event time e. Observations are weighted by the EB weights discussed in Section 3, and standard errors are clustered at the individual level. In the current analysis, we set $K = 4$, to allow more visual tests of parallel trends, and $L = 5$, thereby, focusing on four years before the treatment and up to five years after the treatment, where all periods before $-K$ and $+L$ are accumulated into two final estimates. A schematic representation of the data used in this analysis can be found in the Appendix B.

Identification rests on two assumptions: (i) no anticipation, meaning that outcomes do not respond before the shock is first reported; and (ii) parallel trends, meaning that, absent the shock, outcomes of treated and weighted comparison individuals would have evolved in parallel. Health shocks of interest occur suddenly and largely unexpectedly - in the case of stroke and heart attack due to the nature of the condition; in the case of cancer, because it typically becomes known upon diagnosis. Figures 8 - 12 in the Appendix C depict raw means of the five outcomes of interest for the treated and weighted control groups by time relative to the (pseudo) health shock. These means do not account for individual fixed effects, calendar year, age or other controls, so we use them for descriptive purposes, and rely on the event-study estimates to check for the

parallel trends.

Table 3 reports the event-study estimates for subjective satisfaction with health, employment, annual hours worked, nondurable consumption, and annual labour income, separately for men and women in the entropy-balanced sample. Observations outside the event window are accumulated into the endpoint bins. Standard errors are clustered by individual. The estimates of μ_e (14) are plotted in Figures 1 - 5.

Table 3: Event-study estimates for men and women; Entropy-balanced sample; Outcomes in annual levels.

Dep. variable:	Health satisf.		Employment (binary)		Labour hours		Food expenditure (£)		Labour income (£)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	men	women	men	women	men	women	men	women	men	women
μ_{-4}	0.0041	0.040**	-0.0021	0.0042	10.21	19.81	105.3	161.3	-626.0	-637.7
	(0.019)	(0.020)	(0.016)	(0.0088)	(49.60)	(51.44)	(215.4)	(179.3)	(837.7)	(666.6)
μ_{-3}	0.013	0.031	-0.024	0.023*	-34.46	118.4**	161.8	272.2	161.7	-552.0
	(0.026)	(0.020)	(0.015)	(0.014)	(52.57)	(55.65)	(251.0)	(197.4)	(1,025)	(667.8)
μ_{-2}	-0.0053	0.010	-0.012	0.0009	52.00	117.3**	-176.7	-138.5	-657.4	-294.1
	(0.025)	(0.017)	(0.015)	(0.018)	(48.39)	(47.05)	(213.1)	(165.6)	(1,248)	(398.6)
μ_0	-0.118***	-0.199***	-0.063***	-0.041**	-38.50	-13.85	283.0	219.9*	1,116	-378.6
	(0.040)	(0.034)	(0.024)	(0.020)	(42.83)	(40.55)	(196.0)	(132.6)	(832.0)	(459.6)
μ_1	-0.034	-0.040*	-0.0065	-0.019	-130.0*	64.59	224.6	85.42	-95.86	-163.3
	(0.027)	(0.024)	(0.010)	(0.019)	(70.87)	(55.99)	(199.4)	(158.2)	(955.9)	(627.6)
μ_2	0.016	0.027	-0.018	-0.0007	-6.837	-54.28	389.6	56.86	-1,130	-1,054
	(0.027)	(0.023)	(0.016)	(0.010)	(63.32)	(72.93)	(276.7)	(215.3)	(1,090)	(732.7)
μ_3	0.034	0.021	0.0013	0.012	51.62	45.26	201.0	228.0	2,120	1,202
	(0.030)	(0.024)	(0.011)	(0.015)	(71.82)	(63.82)	(285.6)	(234.3)	(1,441)	(1,431)
μ_4	0.047*	0.035	-0.028	-0.015	-65.54	-61.40	370.2	86.09	209.1	-492.3
	(0.027)	(0.023)	(0.023)	(0.020)	(97.76)	(75.88)	(277.5)	(224.6)	(1,767)	(864.5)
μ_5	0.038*	0.042**	-0.0092	-0.022*	130.0**	-47.19	683.7***	-207.3	1,764	-328.3
	(0.020)	(0.019)	(0.0085)	(0.012)	(64.78)	(73.78)	(240.9)	(154.3)	(1,298)	(793.2)
Single	-0.015	0.016	-0.013	0.0041	33.96	-97.39**	3,014***	884.7***	595.4	-440.2
	(0.016)	(0.013)	(0.011)	(0.012)	(49.66)	(48.34)	(639.6)	(205.5)	(1,606)	(550.0)
High education	-0.034*	-0.024	0.0047	0.035*	21.00	129.4**	170.8	472.1***	1,568	2,075
	(0.019)	(0.023)	(0.014)	(0.021)	(86.38)	(62.19)	(485.5)	(154.9)	(1,788)	(1,389)
No education	0.0019	0.032	-0.0054	-0.018	-44.18	-137.9	-206.7	-169.8	447.3	-1,441
	(0.017)	(0.047)	(0.0082)	(0.017)	(68.60)	(126.7)	(198.2)	(169.9)	(1,379)	(1,060)
Self employed	-0.0050	0.016	0.018***	0.037***	54.60	62.91	-274.0	-181.6	-4,887***	-4,414***
	(0.015)	(0.013)	(0.0066)	(0.010)	(41.24)	(88.20)	(246.8)	(111.6)	(1,641)	(914.4)
Observations	103,068	105,589	127,208	122,220	127,208	122,220	123,615	118,327	126,972	122,106
R-squared	0.546	0.504	0.361	0.270	0.639	0.670	0.635	0.543	0.668	0.722
Num. of ind.	12,854	13,193	13,887	13,626	13,887	13,626	13,801	13,522	13,885	13,626

a. Standard errors clustered by individual in parentheses. *** p<0.01, ** p<0.05, * p<0.1

b. Acute health shock happens between periods $e = -1$ and $e = 0$, i.e. μ_0 measures the ATT in the first time period a health shock had been reported.

c. Interaction-weighted estimates Sun and Abraham (2021); treated cohorts are defined by age at the health shock, never-treated individuals form the control group. All specifications include individual, year and age fixed effects.

The health results confirm that the shocks we study are sharp and largely unanticipated. For men, self-reported satisfaction with health shows no pre-trend and falls by

0.118 (about 0.8 of a category) in the year the health shock is first reported, before returning to a pre-shock levels within one year. For women, the drop at the time of a health shock is larger (0.199), and is more than one category. There is also a mild pre-trend: self-reported satisfaction with health is somewhat higher three to four years before the shock than in the year before it ($\mu_{-4} = 0.040$, $\mu_{-3} = 0.031$), indicating a deterioration in self-reported satisfaction with health even before the shock. This is consistent with the composition of shocks by gender. Cancer accounts for 57.6% of shocks among women, compared with 47.2% among men, for whom strokes and heart attacks make up a larger share. Cancer is often preceded by some symptoms and a diagnostic process, whereas strokes and heart attacks typically occur suddenly.

The same mechanism helps interpret the results for hours worked. For women, hours are significantly higher three and two years before the shock than in the reference year (by 118 and 117 hours), so hours fall one period before the shock and stay at the lower level afterwards. Relative to two years before the shock, annual hours at onset are about 130 lower, or roughly 2.5 hours per week. For men, the effect appears after the onset of a health shock: hours fall by about 130 hours in the year following the shock and recover thereafter.

Employment, however, responds mainly in the year of a health shock. For men, the probability of working falls by 6.3 pp at onset, which is about 6% of the pre-shock employment rate of treated men in this sample. This effect does not last long, however, as the coefficient one year after the shock is close to zero (-0.0065), and none of the later coefficients is statistically significant. For women, employment falls by 4.1 pp at onset, without significant effects in the following years, apart from a 2.2 pp decline five years

after the shock. The only significant coefficient before the shock is three years before μ_{-3} among women. It is 2.3 pp higher than in the reference year, which goes in the same direction as the pre-shock decline in their satisfaction with health and hours worked. The drop in employment at onset may partly reflect temporary absence from work rather than dropping out of labour force as employment and hours are measured by different questions (see Section 3). So some individuals who describe themselves as not working in the year of the shock may still report their usual working hours when they are temporarily not working, for example because of sick leave.

Nondurable consumption appears to be largely smoothed through the shock. For women, there are no significant changes after the health shock, at the onset, however, they report spending £220, significant at the 10% level, more on the nondurable goods. For men, nondurable consumption is slightly higher after the shock, yet not statistically different from zero (except $\mu_5 = 683.7$), even as their hours fall. Finally, annual labour income in levels shows no significant response for either men or women.

Appendix D repeats the event-study on all individuals of working age, including never employed throughout the survey. The results are very similar to the baseline specification reported in this Section.

Figure 1: Estimates and 95% CI of μ_e for the subjective satisfaction with health (see Table 3, columns (1) - (2))

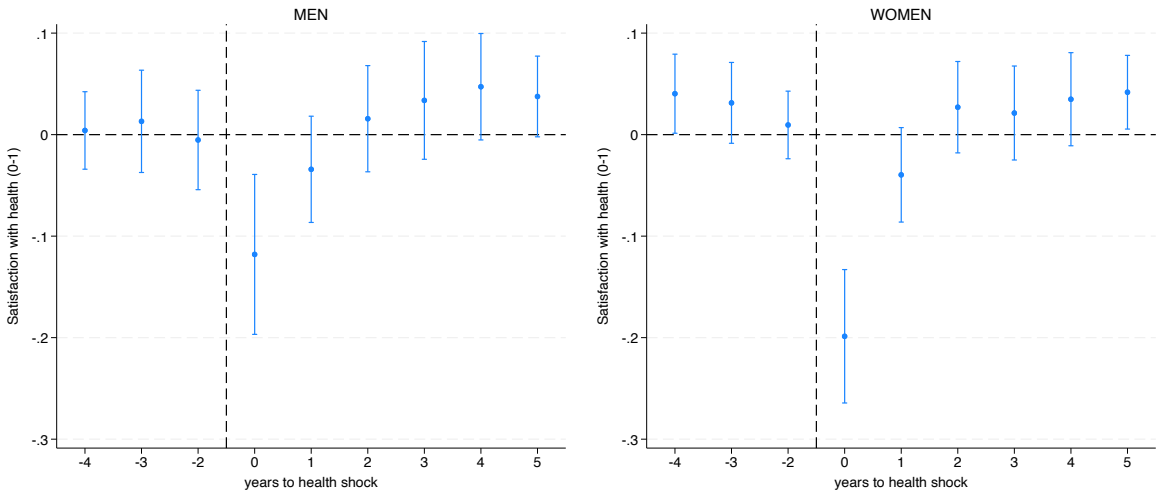


Figure 2: Estimates and 95% CI of μ_e for whether working (see Table 3, columns (3) - (4))

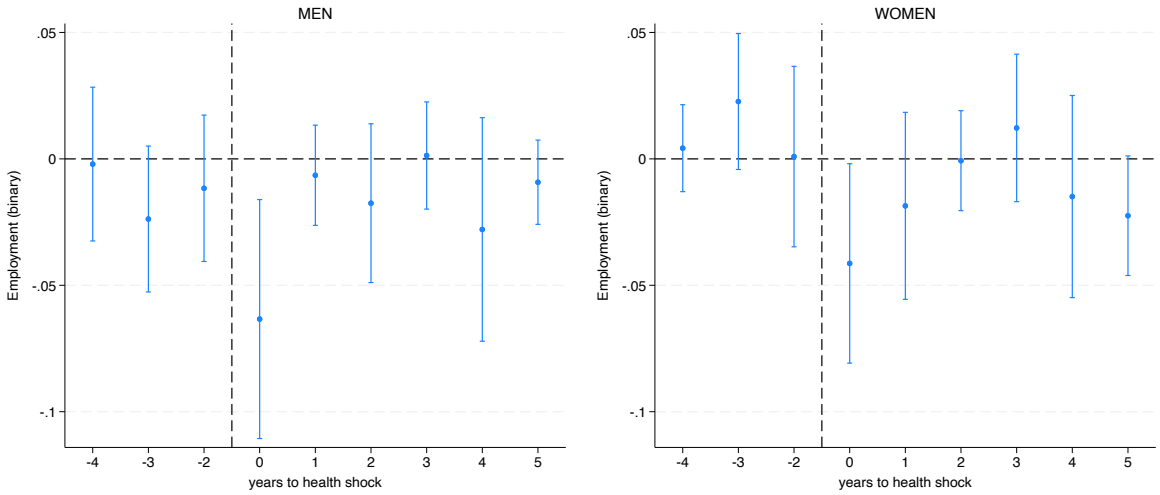


Figure 3: Estimates and 95% CI of μ_e for the number of working hours per year (see Table 3, columns (5) - (6))

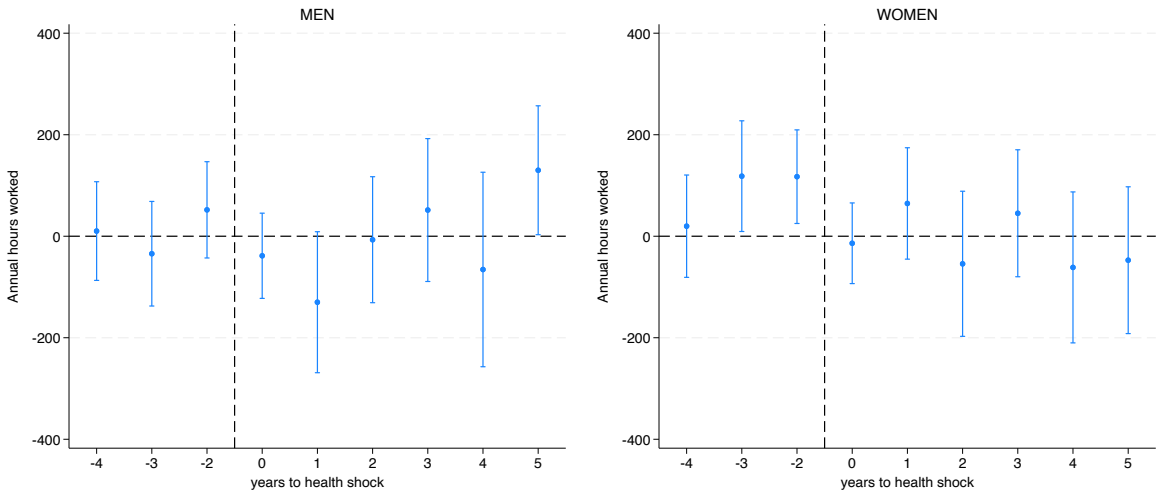


Figure 4: Estimates and 95% CI of μ_e for the annual consumption of nondurable goods (see Table 3, columns (7) - (8))

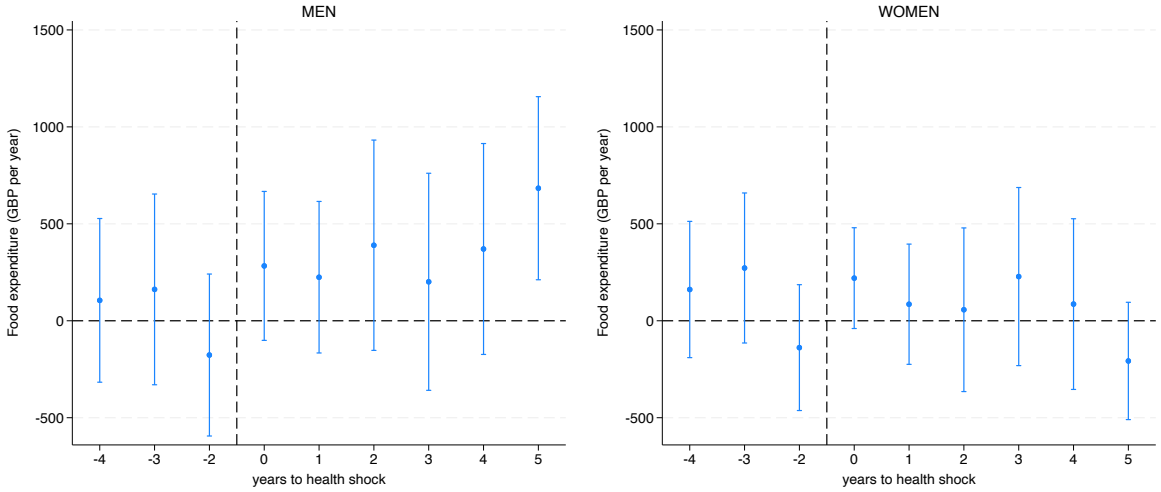
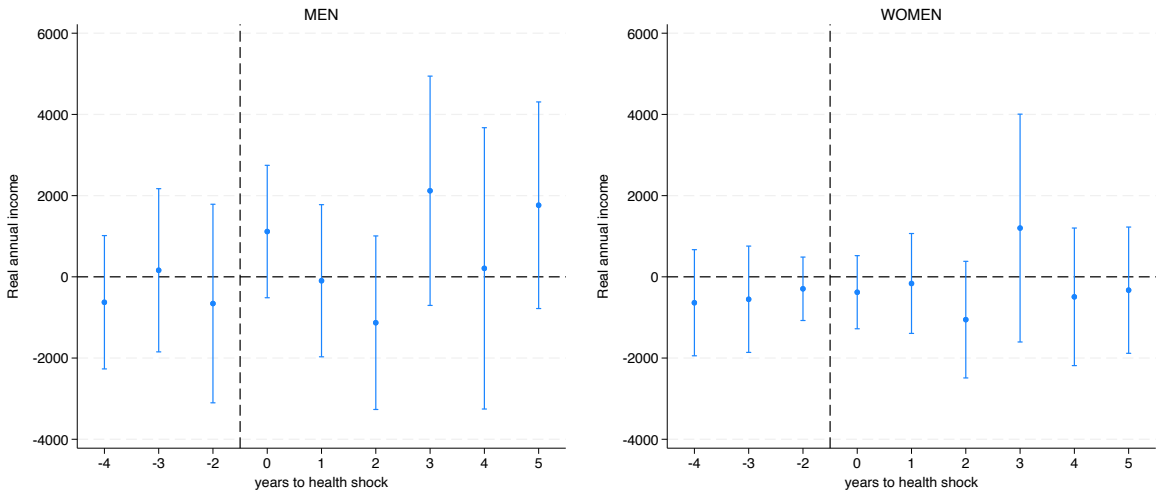


Figure 5: Estimates and 95% CI of μ_e for the annual labour income in levels (see Table 3, columns (9) - (10))



5 Structural Estimation

The structural estimation follows the three steps described in Section 2.2. Step 1 recovers the nondurable consumption share α from the intratemporal first-order condition (IT); Step 2 recovers relative risk aversion σ from (FE); and Step 3 recovers the discount factor δ and the elasticity of marginal utility of consumption with respect to health ϕ from the Euler equations (EE). Each parameter is estimated separately for the treated and control groups, before and after the (pseudo) health shock, for individuals observed in work in at least one wave. Because the intratemporal condition holds only for interior hours, the structural estimation uses person-waves with positive hours before and after the (pseudo) health shock. This conditions on a post-shock outcome so any changes in preferences are described among those who remain attached to work. Tables 4, 5, and 6 report the group- and period-specific estimates; even though pure controls make up a small share of full

controls, we conduct a separate estimation on both pure and other controls as well as on full controls. Because other controls account for the main share of control observations, the full control estimates are close to those for other controls. Table 7 reports the change in each parameter within the treated and control groups and the resulting difference-in-differences, and Figure 6 plots these changes for the treated group and for each of the three control groups.

Table 4: Steps 1 and 3 estimates of the share parameter α and time discounting δ; Entropy-balanced sample.

	$\hat{\alpha}$					$\hat{\delta}$				
	(1)	(2)	(3)	(4)		(5)	(6)	(7)	(8)	
g :	C_{pre}	C_{post}	T_{pre}	T_{post}	N	C_{pre}	C_{post}	T_{pre}	T_{post}	N
Panel A: Men										
All controls	0.161***	0.137***	0.195***	0.151***	121,485	0.995**	0.999	0.969***	0.978**	31,323
	(0.002)	(0.001)	(0.012)	(0.010)		(0.002)	(0.002)	(0.010)	(0.009)	
Pure controls	0.162***	0.139***	.	.	20,457	0.984***	0.980***	.	.	4,750
	(0.004)	(0.004)				(0.004)	(0.006)			
Other controls	0.161***	0.137***	.	.	102,602	0.996*	0.988	.	.	27,059
	(0.002)	(0.001)				(0.002)	(0.018)			
Panel B: Women										
All controls	0.153***	0.132***	0.198***	0.150***	117,764	0.987***	0.998	0.981	0.976*	31,288
	(0.001)	(0.001)	(0.009)	(0.008)		(0.003)	(0.003)	(0.012)	(0.013)	
Pure controls	0.159***	0.129***	.	.	14,354	0.976***	0.991*	.	.	3,593
	(0.003)	(0.002)				(0.005)	(0.005)			
Other controls	0.153***	0.132***	.	.	105,793	0.987***	0.995	.	.	28,330
	(0.001)	(0.001)				(0.003)	(0.009)			

Notes: Columns (1)–(4) report estimates of the share parameter α, estimated from the intratemporal condition, and columns (5)–(8) estimates of time discounting δ, estimated from the system of Euler equations. Columns labelled C_{pre} and C_{post} refer to control estimates in the pre- and post-shock-period samples; T_{pre} and T_{post} refer to the treated group. Pure and other controls rows report control-subgroup estimates. N is the number of observations used to estimate α and δ, respectively. Standard errors in parentheses. For $\hat{\alpha}$, stars denote significance against zero; for $\hat{\delta}$, stars denote a Wald test of H_0: $\delta = 1$. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The clearest difference between treated and comparison individuals is in the consumption–leisure margin. The nondurable consumption share falls between the pre- and post-shock periods in every group, but it falls more among the treated. Among women, $\hat{\alpha}$ declines from 0.198 to 0.150 in the treated group (4.8 pp), compared with a decline from 0.153 to 0.132 in the full control group (2.1 pp). Results with other controls are identical. Pure controls, however, show a larger decline for women - from 0.159 to 0.129 in the

Table 5: Step 2 estimates of the relative risk aversion parameter σ; Entropy-balanced sample.

$\hat{\sigma}$ $g:$	(1) C_{pre}	(2) C_{post}	(3) T_{pre}	(4) T_{post}	N
Panel A: Men					
All controls	**1.001** (0.0061)	**1.000** (0.0037)	**1.006** (0.016)	**1.024** (0.017)	66,596
Pure controls	**1.004** (0.0091)	**0.996** (0.0073)	.	.	9,998
Other controls	**1.001** (0.0069)	**1.000** (0.0041)	.	.	57,384
Panel B: Women					
All controls	**1.001** (0.0044)	**1.006** (0.0048)	*1.025* (0.015)	**1.028** (0.022)	68,505
Pure controls	**0.984** (0.012)	**1.015** (0.012)	.	.	7,725
Other controls	**1.002** (0.0046)	**1.005** (0.0052)	.	.	62,028

Notes: Columns labelled C_{pre} and C_{post} refer to control estimates in the pre- and post-shock-period samples. Columns labelled T_{pre} and T_{post} refer to treated group. Pure and other rows report control-subgroups estimates. N is the number of observations. Standard errors in parentheses. Coefficients in bold are not significantly different from one at the 10% level; coefficients in italic are different from one at the 10% level but not at the 5% level.

pure control group (3.0 pp), and a similar decline for men - from 0.162 to 0.139 in the pure control group (2.3 pp). The difference-in-differences among women in treated and full/other control groups of -0.027 is significant at the 5% level, implying that women who experience a health shock shift weight away from nondurable consumption and towards leisure. The difference-in-differences between treated and pure control groups of -0.018 is not statistically different from zero at any conventional level. For men the difference-in-differences is similar in size across all types of control groups (-0.020) but not statistically significant.

Estimates of σ from Step 2 of the estimation are tightly clustered around one in every group and period, ranging from 0.98 to 1.03 (Table 5). With one exception, none of the estimated coefficients is significantly different from one at the 10% level. Only the pre-shock estimate for treated women (1.025) differs from one at the 10% level but

not at the 5%, based on a Wald test of the null hypothesis that the coefficient on wage growth equals zero. Two implications follow. First, preferences are close to logarithmic, so consumption and leisure are separable in the within-period utility function. Second, there is no evidence that the shock changes the curvature of utility, meaning that the change in behaviour documented above mainly operates through the composition of utility, captured by α, rather than through σ.

The Step 3 estimates of the discount factor show a clear life-cycle pattern among controls (see Table 4). For both men and women, $\hat{\delta}$ rises towards one after the (pseudo) shock period: from 0.995 to 0.999 for men and from 0.987 to 0.998 for women in full control group. This corresponds to annual discount rates falling from 0.5% to 0.1% for men, and from 1.3% to 0.2% for women. Treated individuals discount the future more heavily throughout. Treated men have discount factors of 0.969 before and 0.978 after the shock, implying annual discount rates of about 3.2% and 2.2%. Their pre-shock gap relative to controls is large (2.6 pp difference in discount factor) and is not driven by the shock itself: treated men become slightly more patient after the shock, in line with the controls, so the difference-in-differences is close to zero. For women the picture is different. Before the shock, treated women discount only slightly more than controls (0.6 pp difference). After the shock, however, they do not follow the controls' increase in patience. For treated women, the discount factor goes down to 0.976 (a 2.5% discount rate compared to 1.9 % discounting before the shock). The resulting difference-in-differences of -0.016 implies that, relative to full controls, treated women's annual discount rate rises by about 1.6 percentage points. While these findings are economically sizeable, there is no strong evidence of their statistical significance. The difference-in-differences estimates

for both men and women are not statistically distinguishable from zero at this stage.

Table 6: Step 3 estimates of the elasticity of the marginal utility of consumption with respect to health ϕ; Entropy-balanced sample.

$\hat{\phi}$ $g:$	(1) C_{pre}	(2) C_{post}	(3) T_{pre}	(4) T_{post}	N
Panel A: Men					
All controls	0.811^{***}	0.125	-0.017	0.313^{*}	31,323
	(0.188)	(0.239)	(0.145)	(0.167)	
Pure controls	0.543^{**}	0.890^{***}	.	.	4,750
	(0.264)	(0.297)			
Other controls	0.721^{***}	-0.443	.	.	27,059
	(0.197)	(0.334)			
Panel B: Women					
All controls	0.829^{***}	0.903^{**}	-0.046	-0.078	31,288
	(0.199)	(0.387)	(0.169)	(0.483)	
Pure controls	-0.467^{*}	0.170	.	.	3,593
	(0.267)	(0.207)			
Other controls	0.847^{***}	0.722^{**}	.	.	28,330
	(0.181)	(0.329)			

Notes: Columns labelled C_{pre} and C_{post} refer to control estimates in the pre- and post-shock-period samples. Columns labelled T_{pre} and T_{post} refer to treated group. Pure and other rows report control-subgroups estimates. N is the number of observations. Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Step 3 estimation delivers suggestive estimates of the elasticity with respect to health $\hat{\phi}$, which is positive for all controls of both genders before the shock (0.81 for men and 0.83 for women, see Table 6). These estimates are consistent with health and consumption being complements, which has been shown by Finkelstein et al. (2013) - the marginal utility of consumption is higher in good health. For treated individuals, $\hat{\phi}$ is close to zero before the shock, implying that health changes have no effect on the marginal utility of consumption. For men, the difference-in-differences between treated and full controls is large and positive (1.017, significant at the 1% level). But this estimate is mainly driven by the control group, whose $\hat{\phi}$ falls from 0.81 to an insignificant 0.13, rather than by the treated group, whose $\hat{\phi}$ increases to 0.31 (significant at 10%) from zero. Using pure controls alone yields a difference-in-differences close to zero (-0.017). For

women, none of the changes is significant. We therefore do not interpret the health-shock effect on ϕ as robust. This reflects the sensitivity of $\hat{\phi}$ to the choice of comparison group. Figure 6 provides a visual representation of the difference-in-differences estimations of all three parameters of interest - α, δ, ϕ.

Table 7: Difference-in-differences comparison of the share parameter α, time discounting δ, and elasticity with respect to health ϕ; Entropy-balanced sample.

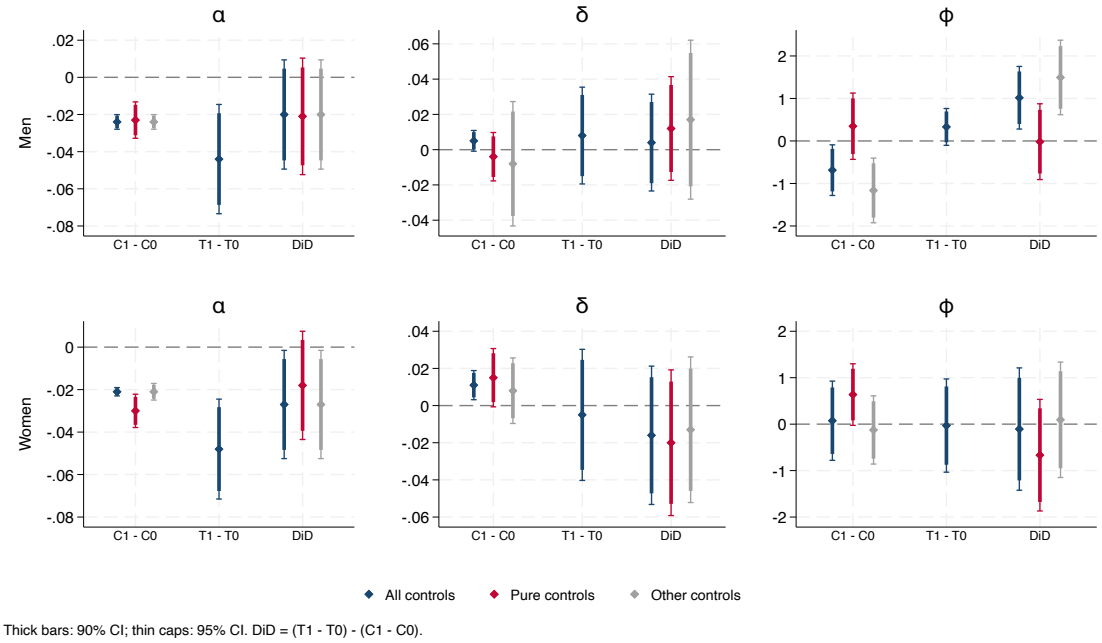
	$\hat{\alpha}$			$\hat{\delta}$			$\hat{\phi}$		
	(1) $C_{post} - C_{pre}$	(2) $T_{post} - T_{pre}$	(3) DiD	(4) $C_{post} - C_{pre}$	(5) $T_{post} - T_{pre}$	(6) DiD	(7) $C_{post} - C_{pre}$	(8) $T_{post} - T_{pre}$	(9) DiD
Panel A: Men									
All controls	−0.024***	−0.044***	−0.020	0.005	0.008	0.004	−0.686**	0.331	1.017***
	(0.002)	(0.015)	(0.015)	(0.003)	(0.014)	(0.014)	(0.304)	(0.221)	(0.376)
Pure controls	−0.023***	.	−0.021	−0.004	.	0.012	0.347	.	−0.017
	(0.005)		(0.016)	(0.007)		(0.015)	(0.398)		(0.455)
Other controls	−0.024***	.	−0.020	−0.008	.	0.017	−1.164***	.	1.495***
	(0.002)		(0.015)	(0.018)		(0.023)	(0.388)		(0.447)
Panel B: Women									
All controls	−0.021***	−0.048***	−0.027**	0.011***	−0.005	−0.016	0.074	−0.031	−0.106
	(0.001)	(0.012)	(0.013)	(0.004)	(0.018)	(0.019)	(0.435)	(0.512)	(0.672)
Pure controls	−0.030***	.	−0.018	0.015*	.	−0.020	0.637*	.	−0.668
	(0.004)		(0.013)	(0.008)		(0.020)	(0.338)		(0.613)
Other controls	−0.021***	.	−0.027**	0.008	.	−0.013	−0.125	.	0.094
	(0.002)		(0.013)	(0.009)		(0.020)	(0.375)		(0.634)

Notes: For each parameter, the first column reports the change in the control group, $C_{post} - C_{pre}$, the second the change in the treated group, $T_{post} - T_{pre}$, and the third the difference-in-differences, $DiD = (T_{post} - T_{pre}) - (C_{post} - C_{pre})$. Pure and other rows use the corresponding control subgroups; the treated-group change does not depend on the control subgroup and is reported once. Standard errors in parentheses. We do not account for the positive correlation between pre- and post-shock estimates of δ for the same individuals and standard errors are therefore conservative.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

6 Discussion and Next Steps

This paper asks whether cancer, stroke and heart attack change the preferences that influence consumption and labour supply decisions. Taken together, our results point to the consumption-leisure margin rather than to pure time preferences as the channel through which preferences may respond to an acute health shock. In the event study, men reduce employment and hours only temporarily, while women start to work fewer hours one year before the shock is reported, which may be a sign of early symptoms of

Figure 6: Difference-in-differences comparison of the share parameter α, time discounting δ, and elasticity with respect to health ϕ; Entropy-balanced sample. Corresponding results from Table 7.



Notes: Correspondence between the Figure and Table 7 notations is as follows.
$C1 \equiv C_{post}$; $C0 \equiv C_{pre}$; $T1 \equiv T_{post}$; $T0 \equiv T_{pre}$.

cancer - the most common health shock among treated women in the sample. In the structural estimates, women who stay in work put less weight on consumption relative to leisure after the shock: the decline in this parameter is 2.7 percentage points larger than among women in the control group, whereas for men the difference is of a similar size (2.0 pp) but imprecisely estimated. We find no significant change in the curvature of utility or in the discount factor, although the point estimates imply that, relative to the control group, the discount rate of treated women rises by about 1.6 percentage points.

These results should be read with some caution, and we see five main limitations of the current analysis. First, α is pinned down by the ratio of the value of leisure to consumption (wl/c) and anything that moves this ratio around the shock is attributed

to preferences. In our estimation, the time endowment M does not depend on health, so hours spent on treatment, recovery or coping with fatigue are counted as leisure in the current setup. Holding α fixed, the change we estimate for women would require that about 15% of the leisure measured after the shock is in fact time absorbed by illness (about 11% for men). With a time endowment of 5,840 hours per year (16 hours per day) and the average post-shock hours of treated individuals in our sample, this corresponds to about 630 hours per year for women (roughly 12 hours per week) and about 415 hours per year for men (roughly 8 hours per week), which is somewhat larger than the time cost of bad health estimated by French (2005). Without modelling the time endowment explicitly as a function of health, a change in preferences and a time cost of this size cannot be distinguished in the intratemporal condition.

Second, nondurable consumption is measured by household expenditure on food and groceries, which is the only expenditure item available in both surveys. Since food is a necessity and time spent at home has been shown to substitute for food expenditure (Aguiar and Hurst, 2005), a health shock that changes household resources or time allocation may change food expenditure without any change in overall consumption or in an individual's preferences. Additionally, in our data the food and grocery bill is measured at the household level, while hours and wages are at the individual level, so the labour supply response of household members without a health shock may affect the ratio that identifies α.

Third, the structural estimation uses individuals who remain in work after the shock, because the intratemporal condition holds only for interior hours. This conditions on a post-shock outcome. If the individuals whose preferences change most stop working,

our estimates understate the average change in preferences, though they describe selected individuals who remain attached to the labour market.

Fourth, treated individuals are interviewed a few years earlier than individuals in the control group. Differences in interest rates, prices and survey design between these periods may affect the levels of the estimated parameters, in particular the discount factor, which is identified from consumption growth net of the interest rate. This is why we focus on the difference-in-differences rather than on the comparison of levels between the two groups.

Fifth, the discount factor is estimated from log-linearised Euler equations, which identify δ only together with the variance of consumption growth (Carroll, 2001). If consumption risk increases or borrowing constraints become tighter after a health shock, expected consumption growth rises, which may lead to an increase in the estimated discount factor and hide a fall in patience. Thus, at this stage, given the small number of treated individuals in the sample used for the structural estimation (143 men and 209 women), the absence of a significant change in δ should be read as an absence of evidence rather than as evidence of no change.

The next step in this project is the simulation exercise, main goal of which is to address the first of the current limitations and to quantify how much of the observed change in behaviour can be explained by the change in preferences. Under perfect foresight and interior hours, the closed-form solutions derived in Appendix E give the optimal paths of consumption, labour and leisure as functions of the preference parameters, wage, survival probabilities and time endowment. They will be used in three ways. First, to

simulate the optimal paths when a health shock affects only the constraints, i.e. the wage, survival and the time endowment, holding preferences fixed at their pre-shock levels, and compare them with the paths when the shock changes α and δ as estimated in Section 5. Comparing both sets of paths with the event-study estimates allows to see how much of the observed change in hours can be explained by each mechanism separately. Second, to ask how large a time cost of illness would have to be to generate the observed change in the ratio between leisure and consumption with fixed preferences, and whether such a time cost is consistent with the observed profile of hours. Third, to generate larger panels with the same structure as our data and re-estimate the model with our three-step procedure, and to show how well the model and the estimation procedure recover the true parameters, in particular the discount factor. Additional extensions in our pipeline are to model the time endowment as a function of health, using information on health limitations available in the data, and to estimate α and δ by event time rather than only before and after the shock, to see whether individuals return to their pre-shock preferences over time. Finally, if preferences do change after a health shock, life-cycle models that keep preferences fixed attribute all post-shock behaviour to constraints and may overstate the welfare loss from reduced work and the employment gains that can be achieved by policies that aim to bring people back to work. Quantifying these effects is the natural next step once the simulation exercise is complete.

References

Aguiar, M. and Hurst, E. (2005). Consumption versus expenditure. *Journal of Political Economy*, 113(5):919–948.

Baji, P. and Bíró, A. (2018). Adaptation or recovery after health shocks? Evidence using subjective and objective health measures. *Health Economics*, 27(5):850–864.

Banks, J., Bassoli, E., and Mammi, I. (2020). Changing attitudes to risk at older ages: The role of health and other life events. *Journal of Economic Psychology*, 79:102208.

Becker, G. S. and Mulligan, C. B. (1997). The endogenous determination of time preference. *The Quarterly Journal of Economics*, 112(3):729–758.

Blundell, R., Borella, M., Commault, J., and De Nardi, M. (2024). Old age risks, consumption, and insurance. *American Economic Review*, 114(2):575–613.

Blundell, R., Britton, J., Costa Dias, M., and French, E. (2023). The impact of health on labor supply near retirement. *Journal of Human Resources*, 58(1):282–334.

Callaway, B. and Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

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

Carroll, C. D. (2001). Death to the log-linearized consumption Euler equation! (and very poor health to the second-order approximation). *Advances in Macroeconomics*, 1(1).

Caumette, E., Di Meglio, A., Vaz-Luis, I., Charles, C., Havas, J., Ruiz de Azua, G., Martin, E., Vanlemmens, L., Delaloge, S., Everhard, S., Martin, A.-L., Dhaini Merimeche,

A., Rigal, O., Coutant, C., Fournier, M., Jouannaud, C., Soulié, P., Cottu, P.-H., Tredan, O., Menvielle, G., and Dumas, A. (2023). Change in the value of work after breast cancer: Evidence from a prospective cohort. *Journal of Cancer Survivorship*, 17(3):694–705.

Chao, L.-W., Szrek, H., Sousa Pereira, N., and Pauly, M. V. (2009). Time preference and its relationship with age, health, and survival probability. *Judgment and Decision Making*, 4(1):1–19.

Chuang, Y. and Schechter, L. (2015). Stability of experimental and survey measures of risk, time, and social preferences: A review and some new results. *Journal of Development Economics*, 117:151–170.

De Nardi, M., French, E., and Jones, J. B. (2010). Why do the elderly save? The role of medical expenses. *Journal of Political Economy*, 118(1):39–75.

Decker, S. and Schmitz, H. (2016). Health shocks and risk aversion. *Journal of Health Economics*, 50:156–170.

Disney, R., Emmerson, C., and Wakefield, M. (2006). Ill health and retirement in Britain: A panel data-based analysis. *Journal of Health Economics*, 25(4):621–649.

Dobkin, C., Finkelstein, A., Kluender, R., and Notowidigdo, M. J. (2018). The economic consequences of hospital admissions. *American Economic Review*, 108(2):308–352.

Finkelstein, A., Luttmer, E. F. P., and Notowidigdo, M. J. (2013). What good is wealth without health? The effect of health on the marginal utility of consumption. *Journal of the European Economic Association*, 11(S1):221–258.

French, E. (2005). The effects of health, wealth, and wages on labour supply and retirement behaviour. *The Review of Economic Studies*, 72(2):395–427.

García-Gómez, P., van Kippersluis, H., O'Donnell, O., and van Doorslaer, E. (2013). Long-term and spillover effects of health shocks on employment and income. *Journal of Human Resources*, 48(4):873–909.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Hainmueller, J. (2012). Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis*, 20(1):25–46.

Heckman, J. (1974). Life cycle consumption and labor supply: An explanation of the relationship between income and consumption over the life cycle. *The American Economic Review*, 64(1):188–194.

Jones, A. M., Rice, N., and Zantomio, F. (2020). Acute health shocks and labour market outcomes: Evidence from the post crash era. *Economics & Human Biology*, 36:100811.

Kools, L. and Knoef, M. (2019). Health and consumption preferences; estimating the health state dependence of utility using equivalence scales. *European Economic Review*, 113:46–62.

Lenhart, O. (2019). The effects of health shocks on labor market outcomes: Evidence from UK panel data. *The European Journal of Health Economics*, 20(1):83–98.

Low, H. and Pistaferri, L. (2015). Disability insurance and the dynamics of the incentive

insurance trade-off. *American Economic Review*, 105(10):2986–3029.

MaCurdy, T. E. (1981). An empirical model of labor supply in a life-cycle setting. *Journal of Political Economy*, 89(6):1059–1085.

Sawatzky, R., Sajobi, T. T., Russell, L., Awosoga, O. A., Ademola, A., Böhnke, J. R., Lawal, O., Brobbey, A., Lix, L. M., Anota, A., Sebille, V., Sprangers, M. A. G., and Verdam, M. G. E. (2024). Response shift results of quantitative research using patient-reported outcome measures: A descriptive systematic review. *Quality of Life Research*, 33(2):293–315.

Stigler, G. J. and Becker, G. S. (1977). De gustibus non est disputandum. *American Economic Review*, 67(2):76–90.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175 – 199.

Trevisan, E. and Zantomio, F. (2016). The impact of acute health shocks on the labour supply of older workers: Evidence from sixteen European countries. *Labour Economics*, 43:171–185.

Wooldridge, J. M. (2021). Two-way fixed effects, the two-way Mundlak regression, and difference-in-differences estimators. Working Paper 3906345, SSRN.

Xu, B., Shao, Y.-F., Xi, H.-L., Chang, S.-Y., and Ge, Z.-Q. (2024). Time preference shifts in medical decision-making after serious illness. *Social Science & Medicine*, 360:117321.

Yaari, M. E. (1965). Uncertain lifetime, life insurance, and the theory of the consumer. *The Review of Economic Studies*, 32(2):137–150.

A Entropy balancing on all individuals of working age (25 - 65), including never employed

This appendix drops the restriction to individuals ever observed in employment and balances the full sample of individuals of working age with complete demographic information on the same covariates and moments as in the baseline (Table 2). The sample contains 553 treated and 23,949 control men, and 710 treated and 27,879 control women. Table 8 reports balance. The treated in this sample are considerably older than in the baseline (63.1 against 48.4 years for men, and 59.9 against 44.2 for women), because the baseline restriction disproportionately removes older, non-working treated individuals. Before reweighting, the treated are older than the controls (by 18.7 years for men and 15.0 for women), more likely to be single, and less educated. Standardized differences on these covariates reach 1.08 for men and 0.81 for women. After reweighting, all targeted covariates are balanced, with standardized differences no larger than 0.004 in absolute value. The effective sample size is 8,494 for men (35% of the nominal control sample) and 13,505 for women (48%).

Balance on non-targeted outcomes is weaker than in the baseline, particularly for men. Treated men are 12 pp less likely to be employed at $t = -1$ than weighted controls, and their monthly earnings are lower by £174, with both differences significant at the 1% level. For women, the gaps are smaller: employment differs by 4 pp, earnings by £42, and weekly hours among those in work are higher among the treated by 2.2 hours. As in the baseline, work experience remains higher among the treated (by approximately 1.4 and 2.1 years).

Table 8: Entropy-Balanced Sample at $t = -1$ on all individuals of working age, including never employed: Pre-Treatment Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Treated	Control Mean		Difference	Standardized Diff.		
Variable	Mean	Unweighted	EB	(EB)	Unweighted	EB	N
Panel A: Men							
Weekly hours worked	41.669	40.739	40.838	0.831	0.066	0.059	15,454
Real monthly earnings	476.385	1,112.556	649.962	-173.578***	-0.541	-0.148	24,071
Working status	0.275	0.626	0.395	-0.120***	-0.752	-0.257	24,491
Age	63.134	44.386	63.117	0.016	1.082	0.001	24,502
Marital status: single	0.168	0.079	0.168	0.000	0.272	0.000	24,502
White	0.976	0.818	0.976	0.000	0.540	0.001	24,502
Self-employed	0.212	0.190	0.266	-0.054	0.054	-0.135	15,747
Year of interview	2002	2009	2009	-6.277***	-1.564	-1.417	24,502
Work experience (years)	9.564	8.141	8.204	1.360***	0.221	0.211	19,032
Higher education	0.127	0.297	0.127	0.000	-0.427	0.000	24,502
No education	0.389	0.132	0.389	0.000	0.612	0.000	24,502
Number of individuals	553	23,949	23,949				
Effective sample size			8,494				
Panel B: Women							
Weekly hours worked	32.118	30.844	29.934	2.184**	0.088	0.150	15,244
Real monthly earnings	344.037	620.151	386.121	-42.083	-0.350	-0.053	28,263
Working status	0.310	0.506	0.349	-0.039**	-0.406	-0.081	28,579
Age	59.859	44.850	59.846	0.013	0.805	0.001	28,589
Marital status: single	0.331	0.170	0.331	0.000	0.378	0.000	28,589
White	0.986	0.827	0.985	0.001	0.568	0.004	28,589
Self-employed	0.109	0.090	0.105	0.003	0.063	0.012	15,467
Year of interview	2002	2009	2009	-6.908***	-1.843	-1.646	28,589
Work experience (years)	9.957	7.549	7.908	2.050***	0.385	0.327	20,017
Higher education	0.163	0.308	0.163	0.000	-0.345	0.000	28,589
No education	0.413	0.167	0.413	0.000	0.563	0.000	28,589
Number of individuals	710	27,879	27,879				
Effective sample size			13,505				

Notes: The control group is reweighted by entropy balancing (Hainmueller, 2012), separately by sex, to match the treated group on the first and second moments of the variables shown in bold, measured at $t = -1$; balance on these moments holds by construction. The remaining variables are not targeted and are reported as checks. Column (2) reports unweighted control means and column (3) entropy-balanced control means in the same sample. Column (4) reports the coefficient on the treatment indicator from a weighted regression, with robust standard errors. Columns (5) and (6) report standardized differences, $\frac{\text{mean}(T) - \text{mean}(C)}{\sqrt{[\text{Var}(T) + \text{Var}(C)]/2}}$, computed without and with the entropy-balancing weights. The effective sample size is $\frac{(\sum_i w_i)^2}{\sum_i w_i^2}$ over the control group. Column (7) reports the number of observations in the corresponding regression.
Significance level: *** p<0.01, ** p<0.05, * p<0.1

B Data visualisation

Table 9: A schematic representation of the data for an event study model with $K = 4$ and $L = 5$.

i	year	a	$Y_{i,a}$	$D_{i,a}$	E_i	$D_{i,a}^{-K}$	...	$D_{i,a}^{-1}$	$D_{i,a}^{0}$	$D_{i,a}^{1}$	...	$D_{i,a}^{L}$
1	2001	30	$y_{1,30}$	0	∞	0	...	0	0	0	...	0
1	2002	31	$y_{1,31}$	0	∞	0	...	0	0	0	...	0
1	2003	32	$y_{1,32}$	0	∞	0	...	0	0	0	...	0
⋮	⋮	⋮	⋮	⋮	⋮	⋮	...	⋮	⋮	⋮	...	⋮
1	2007	36	$y_{1,36}$	0	∞	0	...	0	0	0	...	0
1	2008	37	$y_{1,37}$	0	∞	0	...	0	0	0	...	0
2	2002	21	$y_{2,21}$	0	26	1	...	0	0	0	...	0
2	2003	22	$y_{2,22}$	0	26	1	...	0	0	0	...	0
2	2004	23	$y_{2,23}$	0	26	0	...	0	0	0	...	0
2	2005	24	$y_{2,24}$	0	26	0	...	0	0	0	...	0
2	2006	25	$y_{2,25}$	0	26	0	...	1	0	0	...	0
2	2007	26	$y_{2,26}$	1	26	0	...	0	1	0	...	0
2	2008	27	$y_{2,27}$	1	26	0	...	0	0	1	...	0
⋮	⋮	⋮	⋮	⋮	⋮	⋮	...	⋮	⋮	⋮	...	⋮
$N-1$	2003	42	$y_{N-1,42}$	0	45	0	...	0	0	0	...	0
$N-1$	2004	43	$y_{N-1,43}$	0	45	0	...	0	0	0	...	0
$N-1$	2005	44	$y_{N-1,44}$	0	45	0	...	1	0	0	...	0
$N-1$	2006	45	$y_{N-1,45}$	1	45	0	...	0	1	0	...	0
N	2004	43	$y_{N,43}$	0	44	0	...	1	0	0	...	0
N	2005	44	$y_{N,44}$	1	44	0	...	0	1	0	...	0
N	2006	45	$y_{N,45}$	1	44	0	...	0	0	1	...	0
⋮	⋮	⋮	⋮	⋮	⋮	⋮	...	⋮	⋮	⋮	...	⋮
N	2010	49	$y_{N,49}$	1	44	0	...	0	0	0	...	1
N	2011	50	$y_{N,50}$	1	44	0	...	0	0	0	...	1
N	2012	51	$y_{N,51}$	1	44	0	...	0	0	0	...	1

Notes: the data is an unbalanced panel, so some individuals may have as little as two observations.

C Raw means

Figures 7 - 12 show unconditional raw means of the outcomes of interest for treated and entropy-balanced control groups by time relative to the (pseudo) health shock, with 95% confidence intervals. Since treated individuals are interviewed earlier than controls, treated and control means at the same relative time partly reflect changes over time in prices and survey design, as well as differences in age. This is most visible for individual consumption of nondurable goods (Figure 11), which rises among the treated and falls among the controls before the shock. The event-study estimates in Section 4 account for time and individual fixed effects, and show no significant differences before the shock in almost all outcomes, including consumption.

Figure 7: Raw means for treated and EB weighted controls for $\theta = \ln\left(\frac{wl}{c}\right)$

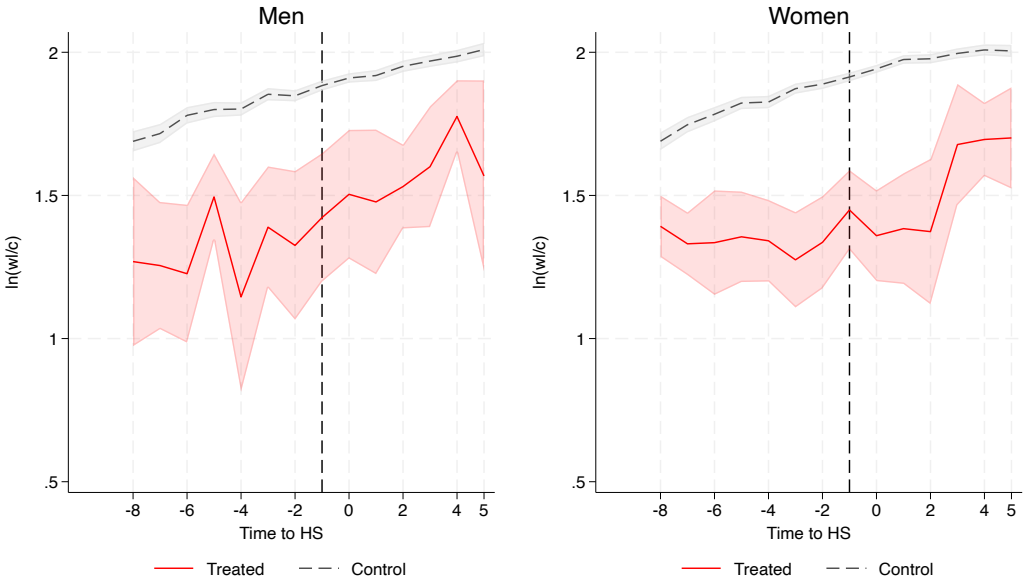


Figure 8: Raw means for treated and EB weighted controls for the subjective satisfaction with health

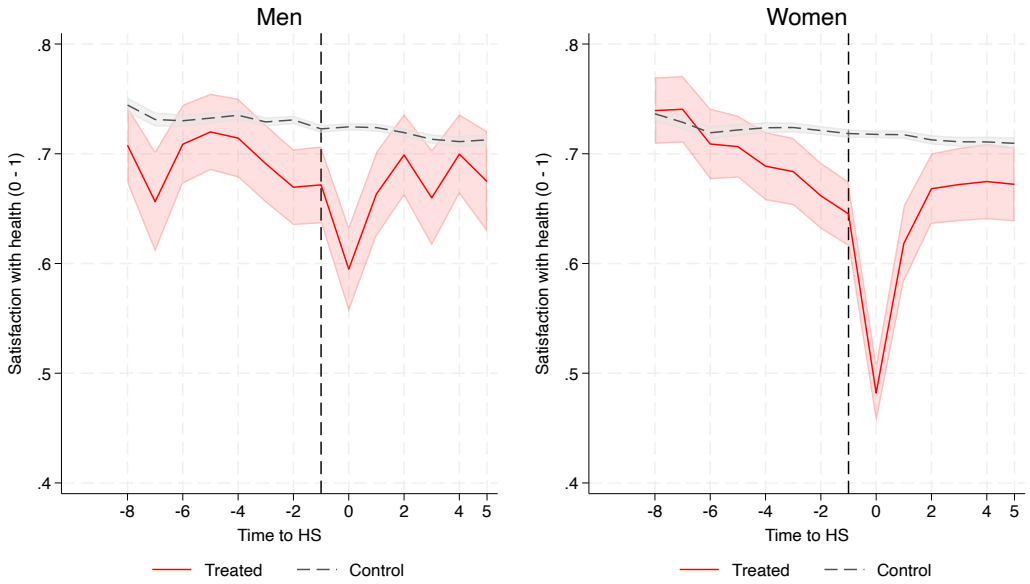


Figure 9: Raw means for treated and EB weighted controls for whether working

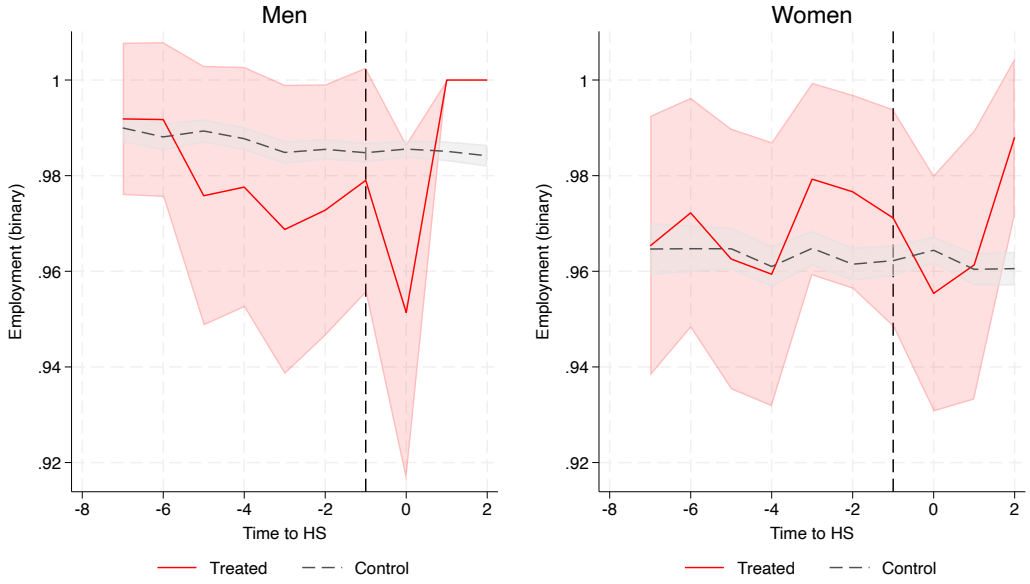


Figure 10: Raw means for treated and EB weighted controls for the number of working hours per year

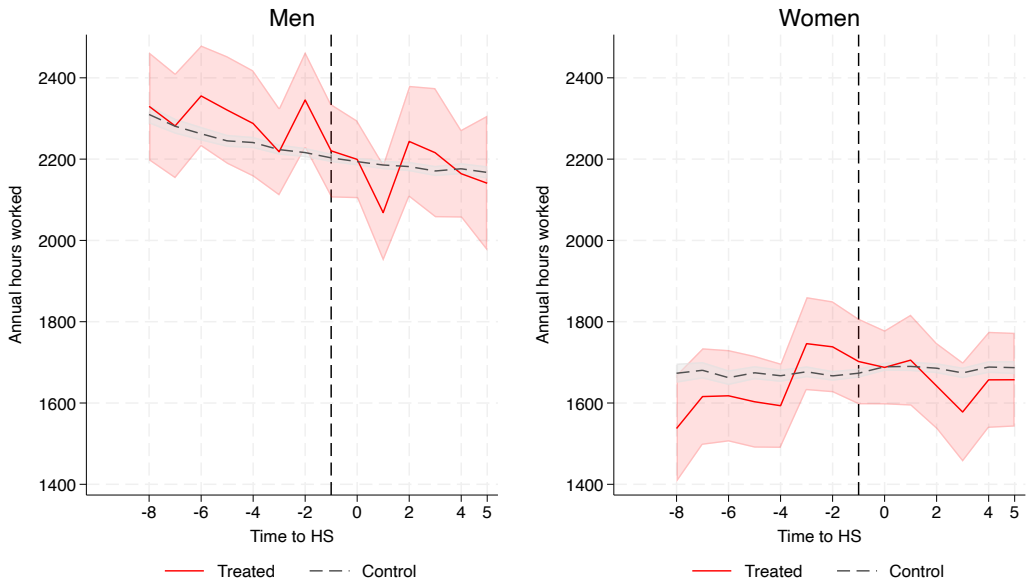


Figure 11: Raw means for treated and EB weighted controls for the annual individual consumption of nondurable goods

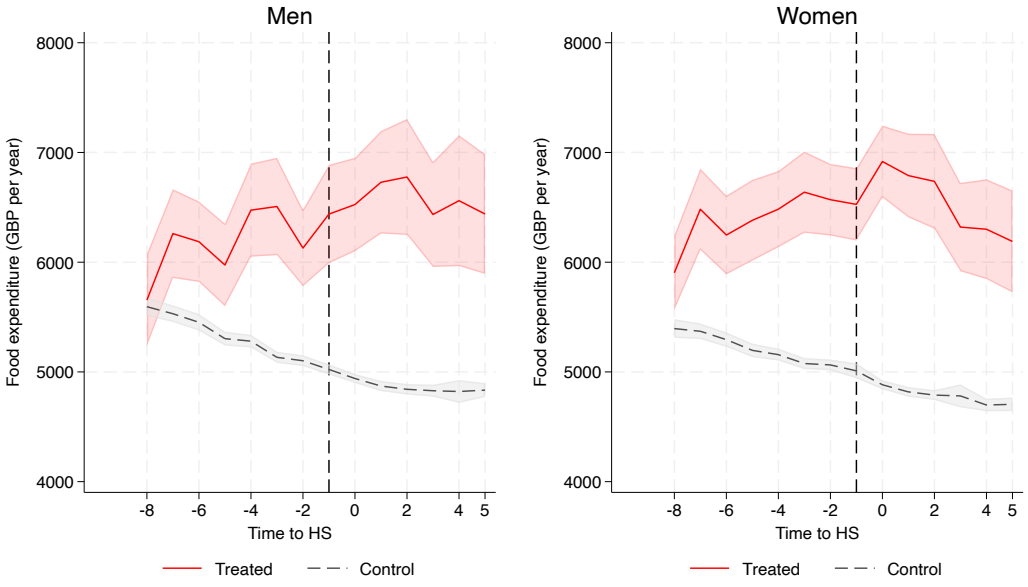
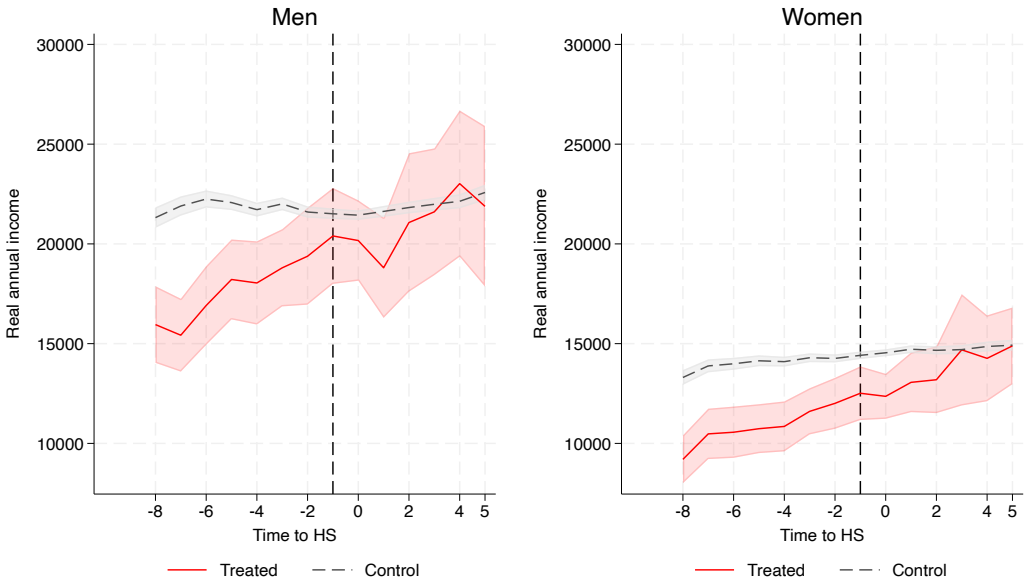


Figure 12: Raw means for treated and EB weighted controls for the annual labour income in levels



D Event-study on all individuals of working age, including never employed

Table 10 re-estimates the event study on the full sample of individuals of working age, including those never observed in employment throughout the survey. Relative to the baseline (Table 3), the sample increases from 14,008 to 18,231 men and from 13,741 to 19,353 women. Overall, the results are very similar to the baseline.

The decline in self-reported satisfaction with health at onset of a health shock is almost unchanged (0.116 against 0.118 for men, and the same for women), and it returns to the pre-shock levels within one to two years. The mild pre-trend in women's health is weaker in the full sample. The effects on labour supply are also robust. Men's probability

Table 10: Event-study estimates for all men and women of working age, including never-employed; Entropy-balanced sample; Outcomes in annual levels.

Dep. variable:	Health satisf.		Employment (binary)		Labour hours		Food expenditure (£)		Labour income (£)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	men	women	men	women	men	women	men	women	men	women
μ_{-4}	-0.0004	0.022	0.0008	0.031***	22.02	19.18	29.74	242.4*	-358.9	-727.5
	(0.015)	(0.017)	(0.012)	(0.0092)	(42.16)	(44.13)	(161.1)	(141.7)	(588.7)	(502.8)
μ_{-3}	-0.014	0.0080	-0.025*	0.029	-59.68	91.17*	166.8	201.8	-131.9	-560.8
	(0.024)	(0.020)	(0.014)	(0.018)	(55.35)	(51.66)	(234.6)	(180.9)	(938.7)	(587.2)
μ_{-2}	-0.013	0.0011	-0.017	0.0081	48.14	93.74*	-231.3	-184.0	-370.4	-145.8
	(0.024)	(0.016)	(0.015)	(0.019)	(49.11)	(48.83)	(210.6)	(164.8)	(1,213)	(411.7)
μ_0	-0.116***	-0.199***	-0.058**	-0.036*	-29.68	-8.358	248.9	288.8**	821.3	-505.8
	(0.039)	(0.032)	(0.023)	(0.020)	(45.81)	(41.84)	(186.3)	(142.5)	(825.7)	(434.3)
μ_1	-0.025	-0.041*	0.0045	-0.023	-131.6**	27.76	276.8	110.6	-352.1	-425.1
	(0.025)	(0.022)	(0.010)	(0.021)	(64.62)	(59.42)	(199.1)	(151.7)	(902.2)	(629.6)
μ_2	0.0041	0.020	0.0012	0.016	15.68	6.289	352.7	77.90	-936.6	-850.1
	(0.026)	(0.021)	(0.017)	(0.012)	(60.21)	(66.47)	(261.5)	(198.8)	(1,020)	(689.3)
μ_3	0.021	0.015	0.0003	0.019	56.10	86.43	245.2	170.8	1,846	1,148
	(0.027)	(0.021)	(0.012)	(0.015)	(69.22)	(62.87)	(268.0)	(215.7)	(1,320)	(1,184)
μ_4	0.039	0.028	-0.021	-0.015	-60.23	-12.71	277.4	31.78	298.2	-450.6
	(0.025)	(0.022)	(0.021)	(0.021)	(91.15)	(73.08)	(263.3)	(202.8)	(1,637)	(779.4)
μ_5	0.028	0.033*	-0.0030	-0.019	129.2**	-12.89	594.1***	-144.2	1,466	54.61
	(0.017)	(0.017)	(0.0073)	(0.013)	(55.25)	(74.36)	(203.8)	(153.8)	(1,055)	(778.5)
Single	-0.010	0.011	-0.0062	0.0061	13.16	-59.52	2,743***	873.5***	29.80	-502.4
	(0.013)	(0.012)	(0.0080)	(0.012)	(37.38)	(45.17)	(469.9)	(174.3)	(1,181)	(469.8)
High education	-0.058**	-0.043*	0.011	0.070***	15.39	246.6***	185.1	545.8***	2,428	2,418**
	(0.029)	(0.026)	(0.012)	(0.022)	(75.27)	(75.01)	(422.1)	(157.6)	(1,663)	(1,229)
No education	-0.0036	0.025	-0.0039	-0.027	-90.38	-116.7	-25.54	-128.8	742.4	-1,129
	(0.014)	(0.045)	(0.0055)	(0.022)	(55.25)	(129.5)	(190.6)	(146.4)	(1,103)	(1,031)
Self employed	-0.0024	0.013	0.023***	0.040***	68.13	7.795	-224.8	-28.98	-5,012***	-3,816***
	(0.013)	(0.012)	(0.0066)	(0.0086)	(41.46)	(73.51)	(194.4)	(96.01)	(1,326)	(771.8)
Observations	114,732	122,910	142,959	143,774	142,959	143,774	138,709	138,972	142,695	143,636
R-squared	0.536	0.504	0.359	0.301	0.664	0.695	0.634	0.569	0.676	0.729
Num. of ind.	15,378	16,852	16,950	17,812	16,950	17,812	16,770	17,575	16,945	17,812

a. Standard errors clustered by individual in parentheses. *** p<0.01, ** p<0.05, * p<0.1

b. Acute health shock happens between periods $e = -1$ and $e = 0$, i.e. μ_0 measures the ATT in the first time period a health shock had been reported.

c. Interaction-weighted estimates Sun and Abraham (2021); treated cohorts are defined by age at the health shock, never-treated individuals form the control group. All specifications include individual, year and age fixed effects.

of working falls by 5.8 percentage points at onset (6.3 in the baseline) and returns to pre-shock levels in the following year, and their annual hours fall by 132 hours one year after the shock (130 in the baseline). For women, the decline in employment at onset is of similar magnitude (3.6 against 4.1 pp), and it is followed by a recovery two to three years after the shock. The positive pre-shock coefficients for women's hours spent on work persist in the full sample, although they are somewhat smaller (91 and 94 hours at μ_{-3} and μ_{-2}, against 118 and 117), so the anticipatory dip in the year before the shock is not driven by sample selection or an employment restriction.

The results for income and consumption are also qualitatively unchanged. Labour income in levels shows no significant response for both men and women in either sample. Women's nondurable consumption increases at the onset of a shock by £289 (£220 in the baseline), while there is a significant increase for men five or more years after the shock - £594 against £684 in the baseline.

Figure 13: Estimates and 95% CI of μ_e for the subjective satisfaction with health (see Table 10, columns (1) - (2))

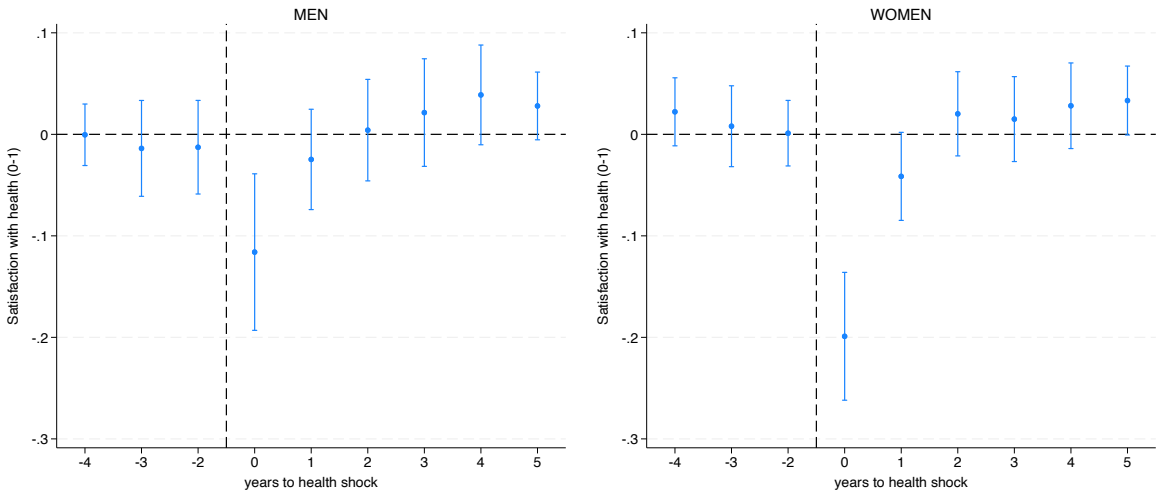


Figure 14: Estimates and 95% CI of μ_e for whether working (see Table 10, columns (3) - (4))

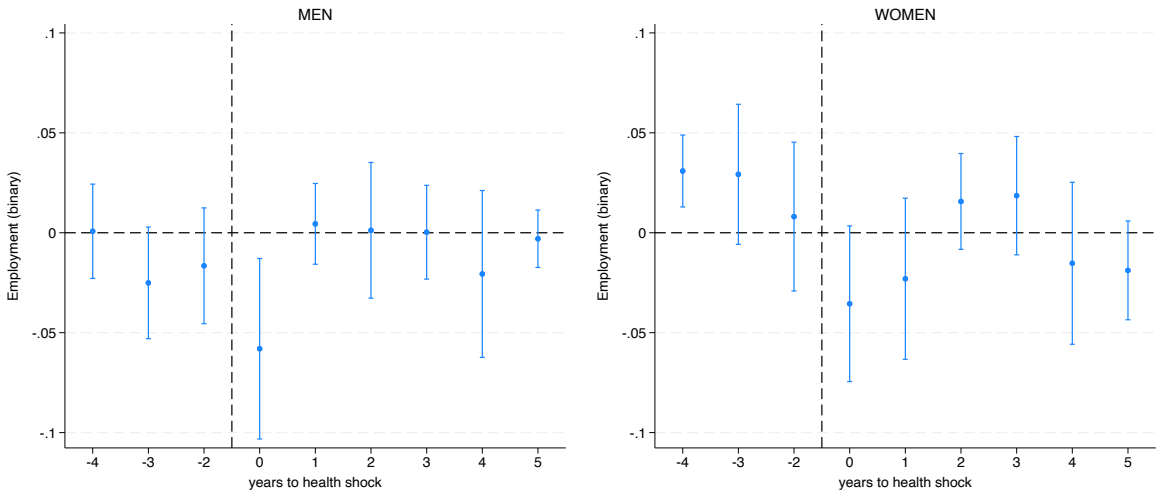


Figure 15: Estimates and 95% CI of μ_e for the number of working hours per year (see Table 10, columns (5) - (6))

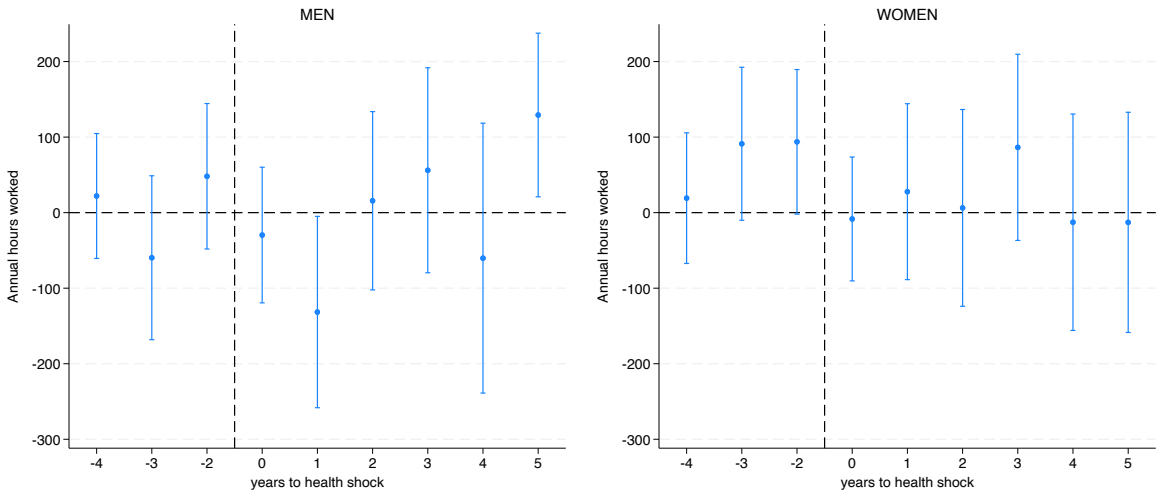


Figure 16: Estimates and 95% CI of μ_e for the annual consumption of nondurable goods (see Table 10, columns (7) - (8))

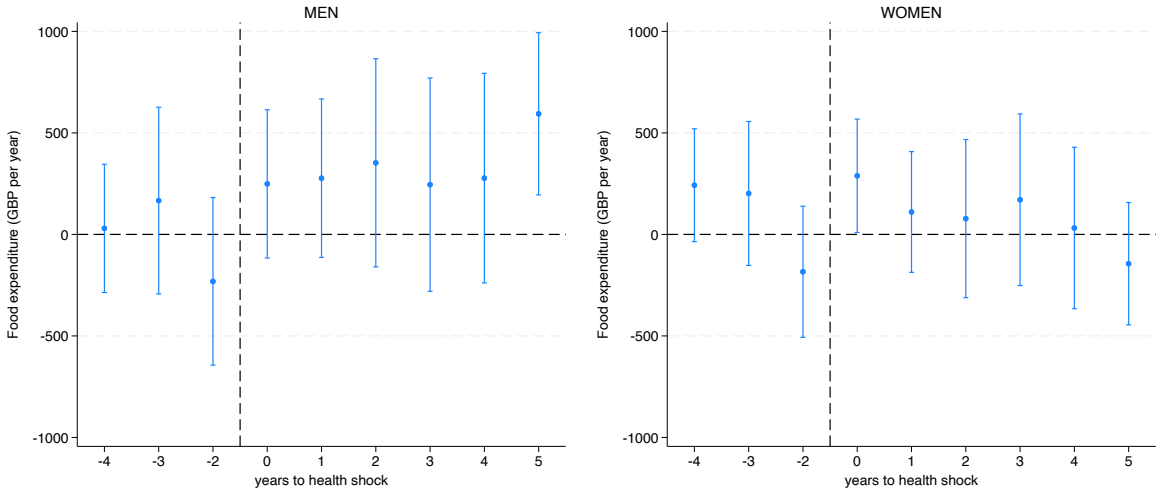
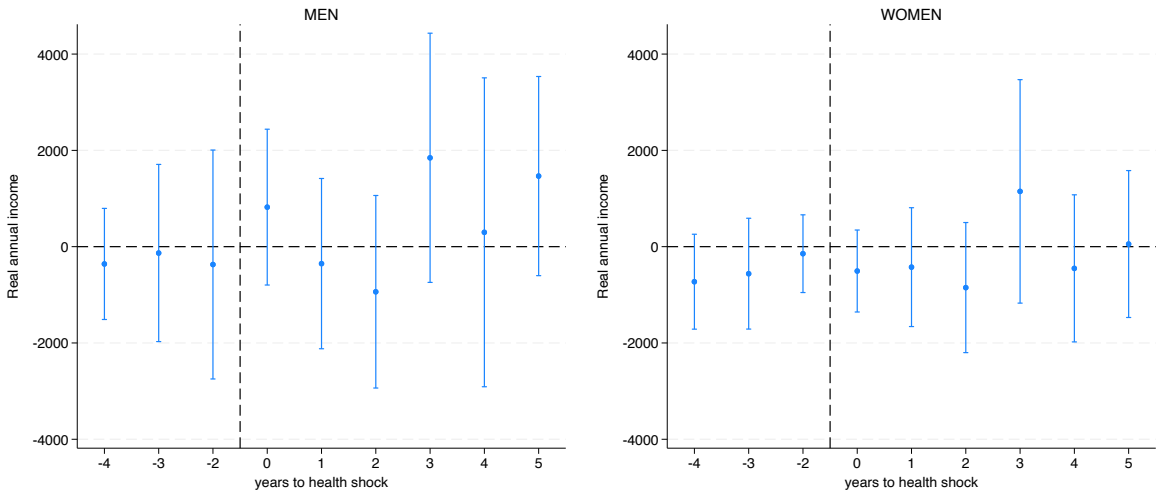


Figure 17: Estimates and 95% CI of μ_e for the annual labour income in levels (see Table 10, columns (9) - (10))



E Closed form solutions

Under perfect foresight, interior hours $0 < h_t < M_t$, and a slack no-borrowing constraint $a_{t+1} > 0$, the optimality conditions admit closed-form solutions for the optimal paths of consumption and hours. These solutions are derived from using the intratemporal condition (IT), and Euler equation (EE) for consumption. As the main timeline of interest in this project is working age, the closed form solutions only describe the working phase $t < R$. Telescoping the Euler equation we get a consumption at any time period t as a function of c_{t_0}:

$$c_t = c_{t_0} \left\{ \delta^{t-t_0} \prod_{k=t_0+1}^{t} (1+r_k) \prod_{k=t_0}^{t-1} p_k \left(\frac{w_t}{w_{t_0}} \right)^{-(1-\alpha)(1-\sigma)} \left(\frac{H_t}{H_{t_0}} \right)^{\phi} \right\}^{\frac{1}{\sigma}}. \tag{15}$$

From the budget constraint (3) and the terminal condition $a_{T+1} = 0$, iterating forward and substituting $w_t h_t = w_t(M_t - l_t) = w_t M_t - \frac{1-\alpha}{\alpha} c_t$ from (IT) gives the in-

tertemporal budget constraint:

$$\frac{1}{\alpha}\sum_{t=t_0}^{R}\frac{c_t}{\prod_{k=t_0}^{t}(1+r_k)} = a_{t_0} + \sum_{t=t_0}^{R}\frac{w_t M_t + v_t}{\prod_{k=t_0}^{t}(1+r_k)}, \tag{16}$$

which, in combination with (15), gives us a level of optimal consumption at the beginning of a life cycle as well as at any other point in time:

$$c_{t_0} = \frac{\alpha(1+r_{t_0})\left[a_{t_0} + \sum_{t=t_0}^{R}\frac{w_t M_t + v_t}{\prod_{k=t_0}^{t}(1+r_k)}\right]}{\sum_{t=t_0}^{R}\left\{\delta^{t-t_0}\prod_{k=t_0}^{t-1}p_k\left[\prod_{k=t_0+1}^{t}(1+r_k)\right]^{1-\sigma}\left(\frac{w_t}{w_{t_0}}\right)^{-(1-\alpha)(1-\sigma)}\left(\frac{H_t}{H_{t_0}}\right)^{\phi}\right\}^{\frac{1}{\sigma}}}; \tag{17}$$

$$c_t = c_{t_0}\left\{\delta^{t-t_0}\prod_{k=t_0+1}^{t}(1+r_k)\prod_{k=t_0}^{t-1}p_k\left(\frac{w_t}{w_{t_0}}\right)^{-(1-\alpha)(1-\sigma)}\left(\frac{H_t}{H_{t_0}}\right)^{\phi}\right\}^{\frac{1}{\sigma}}; \tag{18}$$

with leisure and hours worked expressed as

$$l_t = \frac{1-\alpha}{\alpha}\frac{c_t}{w_t}; \quad h_t = M_t - l_t. \tag{19}$$