

Income Gradients in Health-State Dependence*

Philippe De Donder[†],
Mathieu Lefebvre[‡],
Marie-Louise Leroux[§],
Quitterie Roquebert[¶]

June 25, 2026

Abstract

This paper investigates how marginal utility varies with health status (e.g. health-state dependence) while allowing this relationship to differ across income levels. Building on the existing literature, we develop a framework that quantifies the income adjustments necessary to maintain individuals' well-being when they become disabled. Using SHARE data, we empirically estimate how health affects the marginal utility of consumption across the income distribution for older adults in Europe. Our results show that health-state dependence is negative among low-income individuals, indicating that their marginal utility of consumption declines when their health worsens. In contrast, at the very top of the income distribution, health-state dependence is positive, implying that marginal utility of consumption rises as health deteriorates.

Keywords: Health-state dependence, Marginal utility of consumption, Income heterogeneity, SHARE survey

JEL Codes: D12, I1, J14

*This work of the Interdisciplinary Thematic Institute MAKERs, as part of the ITI 2021-2028 program of the University of Strasbourg, CNRS and INSERM, was supported by IdEx Unistra (ANR-10-IDEX-0002). Financial support from the ANR (Programme d'Investissement d'Avenir ANR-17-EURE-0010) and from Fonds de Recherche du Québec-Société et Culture (FRQSC, grant number: 2024-SE3-328700) is also gratefully acknowledged. The authors thank Ana Montes Vinas, Pilar Garcia-Gomez, Eugenio Peluso, Igor Francetic, Vincenzo Galasso, Jean-Baptiste Delhay and participants at the 3rd Croatian Health Economics Workshop, le Colloque retraite et vieillissement 2025 de la Caisse des Dépôts, the 2025 (LIS)²ER-SHARE Luxembourg Joint Workshop, the 2026 Journées du BETA, the AFéPOP 2026 Annual Conference, the 2026 Journées Louis-André Gérard-Varet for their useful comments.

[†]Toulouse School of Economics, CNRS, University Toulouse Capitole, Toulouse, France. Email: philippe.dedonder@tse-fr.eu

[‡]University of Strasbourg, BETA, CNRS, Strasbourg, France. Email: mathieu.lefebvre@unistra.fr

[§]Département des Sciences Economiques, ESG-UQAM, Montréal, Canada; CESifo, Munich, Germany; CORE, Université catholique de Louvain, Louvain-la-Neuve, Belgium. Email: leroux.marie-louise@uqam.ca

[¶]University of Strasbourg, BETA, CNRS, Strasbourg, France. Email: roquebert@unistra.fr

1 Introduction

Health is undeniably a key determinant of individuals' well-being. It influences welfare not only directly, since good health is inherently preferable to poor health, but also indirectly, by enhancing or degrading the enjoyment of various life experiences. This indirect influence ties closely to the concept of marginal utility of consumption. Assessing how marginal utility is affected by health is critical for accurately modeling individuals' lifetime decisions, such as savings (Viscusi and Evans, 1990; Finkelstein et al., 2013; Ameriks et al., 2020; Achou et al., 2026), consumption allocation (Blundell et al., 2024), insurance behavior (De Donder and Leroux, 2021), and in turn, for designing efficient public policies.

The relationship between marginal utility and health has thus been studied in the economic literature, with the aim of identifying if marginal utility increases (i.e. positive health-state dependence) or decreases (i.e. negative health-state dependence) when health deteriorates. In the first case, consumption is mostly a substitute to good health. This could be expected for instance for consumption goods such as home modifications, mail delivery services or long-term care services, which may be more valued under bad health. In the second case, consumption is mostly a complement to good health as could be expected for travel or sports goods.

However the relationship between marginal utility and health remains a topic of debate in the economics literature. Some studies, such as Lillard and Weiss (1997), Kools and Knoef (2019), Ameriks et al. (2020), Bassoli (2023), Amo-Agyei (2025) and Achou et al. (2026), find positive health-state dependence, indicating that marginal utility increases under bad health. Conversely, other researchers, including Finkelstein et al. (2013), Hong et al. (2013), Koijen et al. (2016), and Blundell et al. (2024), report negative health-state dependence following a negative health shock, suggesting that marginal utility decreases with declining health. In contrast, De Nardi et al. (2010) analyze the saving decisions of older individuals and find no significant relationship between marginal utility and health. Overall, the evidence is mixed, and it remains unclear whether marginal utility rises or falls with deteriorating health. We provide a brief review of this literature in the next section.

To the best of our knowledge, this literature measures an average effect by implicitly assuming that the sign of health-state dependence is homogeneous across the population under study. As highlighted by

Amo-Agyei (2025), the relationship between marginal utility and health may also depend on individual characteristics (as well as differences in cultural and institutional contexts). Building on this insight, our paper allows for heterogeneity in the health-state dependence of utility, focusing on how it varies with income. We conjecture that depending on income levels, marginal utility may increase or decrease with health. For instance, one could think that richer individuals consume relatively more goods which are complements to good health (travel, restaurants, etc.) while poorer people consume relatively more substitutes (such as streaming entertainment services, processed food or medication). This would correspond to a situation of positive health-state dependence for poor people, and negative health-state dependence for rich people. Or, to the opposite, it may be the case that richer individuals are better able to adjust their consumption patterns to their health status and to purchase costly at-home services (restaurant delivery, comfort services at home). Poorer individuals, by contrast, face tighter constraints and have fewer possibilities to adapt their consumption. This suggests negative (resp. positive) health-state dependence for poor (resp. rich) people.

The objective of this paper is to explore this relationship and to assess whether the sign of the health-state dependence can differ across income levels. For this purpose, we develop a model inspired by Finkelstein et al. (2013) and Kools and Knoef (2019), where individuals live two periods and face a positive probability that their health deteriorates in the second period. Agents, who only differ in income, decide how to allocate their income across these two periods before knowing their future health status. Following the literature on household living standards and exposure to specific health risks (Viscusi and Evans, 1990; Evans and Viscusi, 1991), we calculate the necessary income adjustments (increases or decreases) to maintain individuals' material well-being (i.e., utility from consumption) when their health deteriorates. We first detail the predictions of this theoretical model and then estimate it using data from the Survey of Health, Aging, and Retirement in Europe (SHARE). By linking these estimates to the predictions of our theoretical model, we gain a deeper understanding of the relationships between permanent income, consumption, the health equivalence scale, and the sign of the health dependence.

Our paper contributes to the literature in three main ways. First, rather than focusing on average effects, we show that health-state dependence of marginal utility is strongly heterogeneous across the income distribution: health deterioration lowers the marginal utility of consumption for low-income individuals but raises it at the top of the income distribution. Second, by emphasizing income heterogeneity, we offer a novel

explanation for the lack of consensus in the literature (see the next section) regarding the sign and magnitude of health-state dependence. Finally, we demonstrate that these findings are robust across alternative measures of health, including functional limitations, self-reported health, and chronic conditions, as well as a wide range of sensitivity checks.

The paper is structured as follows. The next section details the existing literature. Section 3 introduces the model used throughout the paper and analytically examines the relationship between health-state dependence, the equivalence scale, and permanent income. Section 4 presents the empirical model while Section 5 describes the data used in the analysis. Section 6 discusses the estimation results, while the final section provides concluding remarks.

2 Literature review

We now briefly review the literature that quantifies the magnitude and sign of health-state dependence. Existing studies differ along several key dimensions, including the object of interest (marginal utility of income versus consumption), the measurement of health deterioration, and whether health changes are treated as temporary shocks or permanent states.

A first strand of the literature focuses on how health status affects the marginal utility of permanent income. Finkelstein et al. (2013) provide one of the seminal contributions in this area. Using U.S. panel data, they proxy individual utility with measures of subjective well-being and approximate consumption through permanent income.¹ Health is measured by the number of chronic conditions. By interacting income with the health status in a well-being equation, they infer how the marginal utility of income varies across health states. Their results point to negative health-state dependence, meaning that income yields lower utility when individuals are in poor health. As the authors illustrate, these findings could have important implications for the optimal design of public health insurance benefits.

Related approaches are adopted by Ameriks et al. (2020) and Achou et al. (2026), who focus on persistently deteriorated health states and, in particular, compare marginal utility across states of dependency and autonomy. Using detailed survey data on older Americans, Ameriks et al. (2020) exploit subjective well-

¹Well-being is measured by whether the respondent reports having been happy for most of the time during the previous week.

being measures and stated preferences to compare marginal utility across these states, and find evidence of positive health-state dependence, meaning that marginal utility increases under disability. Achou et al. (2026) extend this analysis using Canadian data by explicitly contrasting dependent and autonomous states and emphasizing the role of long-term care (LTC, hereafter) settings, distinguishing between care provided at home and care in LTC facilities. Their results also point to positive health-state dependence, which is especially pronounced when LTC is provided at home.

A complementary contribution is Kools and Knoef (2019), who propose an innovative indirect approach based on financial satisfaction rather than life satisfaction. Using European survey data, they estimate how financial satisfaction depends on income and health, and derive a health equivalence scale defined as the income adjustment required in poor health to reach the same level of financial satisfaction as in good health. This framework allows them to infer the sign of health-state dependence without directly observing consumption. Their findings indicate positive health-state dependence, implying that the marginal utility of consumption increases as health, proxied by limitations in activities of daily living (ADLs), deteriorates. Importantly, their results are robust to alternative health measures (chronic diseases, functional limitations and cognitive dysfunction) and to different estimation specifications.

A second strand of the literature studies health-state dependence within structural or life-cycle frameworks. Lillard and Weiss (1997) study changes in households' assets in response to income and health shocks using U.S. panel data. By modeling consumption and saving behaviors under health and survival uncertainty, they conclude that poorer health increases marginal utility of consumption, as households place greater value on resources when facing higher mortality risks and medical needs. Their analysis emphasizes the dynamic nature of health shocks and their interaction with precautionary saving.

More recent work builds on similar ideas using richer data and modern econometric techniques. Blundell et al. (2024) analyze health shocks within a life-cycle consumption framework, allowing health to affect both needs and preferences. They show that health deterioration can raise the marginal utility of consumption through increased needs, even if overall well-being declines. Amo-Agyei (2025) provides further evidence along these lines, using longitudinal data to examine how consumption responses differ before and after health shocks, with results again pointing toward positive health-state dependence.

Finally, several studies (e.g. Hong et al., 2013; Koijen et al., 2016; Viscusi, 2019) focus on the implications

of health shocks for portfolio choice, insurance demand, and consumption smoothing. While not always explicitly framed in terms of marginal utility, their findings are consistent with (negative) health state-dependent preferences with health shocks altering both risk attitudes and the valuation of consumption.

Overall, this literature highlights substantial heterogeneity in both methods and conclusions. Differences in the utility argument (income versus consumption), the nature of health variations (temporary shocks versus permanent dependency), and the measurement of health (different physical and cognitive impairments measures) and utility (life satisfaction, happiness or financial strain) contribute to the mixed evidence on the sign of health-state dependence. These differences underscore the importance of clearly defining both health states and the economic margin of interest when assessing how health affects preferences. In this respect, we are the first to provide empirical evidence that the marginal utility of consumption varies systematically with individual income.

3 Theoretical framework

Our theoretical framework builds on the models developed by Finkelstein et al. (2013) and Kools and Knoef (2019), which we extend to allow marginal utility under poor health to vary with permanent income. The use of permanent income instead of consumption is motivated by the empirical application that will follow. Indeed data with broad-based consumption measures are scarce and thus we develop a behavioral model that, under empirically plausible conditions, infers how marginal utility of consumption depends on health from estimates of how marginal utility of permanent income varies with health.²

In this model, individuals live for two periods with certainty. In the first period, they are in good health. However, in the second period, their health status is uncertain, and they may become sick with a probability $0 \leq p \leq 1$. The only ex ante source of variation across individuals is the amount of income Y they are endowed with at the beginning of period 1.

In each period, an individual's utility depends both on her health status S and on consumption C as

²There are two main approaches to estimating health-state dependence. The first relies on revealed demand, such as health insurance choices or consumption responses to health shocks, but inference is highly sensitive to assumptions about risk aversion, moral hazard, and bequest motives, making identification difficult. The second approach estimates how the utility impact of health shocks varies across individuals with different income levels, which avoids many of these complications. This is the approach we adopt here.

follows

$$U(C, S) = (1 + \phi_1 S) \ln(C) - \phi_0 S \quad (1)$$

where S is a dichotomous variable taking value 0 if the individual is in good health and 1 if in bad health. Parameter $\phi_0 > 0$ is a scaling parameter directly decreasing the utility in bad health while $\phi_1 \geq -1$ is the health-state dependence parameter that we will estimate in Section 6. Log utility is chosen as a parsimonious benchmark that avoids introducing an additional, weakly identified curvature parameter which is difficult to identify in our data.³ Since the term $(1 + \phi_1 S)$ multiplies marginal utility of consumption, if $-1 \leq \phi_1 < 0$, marginal utility is reduced under bad health (i.e., negative health-state dependence), while if $\phi_1 > 0$, marginal utility of consumption is higher under bad health (i.e., positive health-state dependence). If $\phi_1 = 0$, preferences are state independent.

At the beginning of period 1, individuals allocate their permanent income Y between consumption at each period by maximizing their expected utility subject to their inter-temporal budget constraint:

$$\begin{aligned} \max_{C_1, C_2} EU(C_1, C_2) &= \ln C_1 + \frac{-\phi_0 p + (1 + \phi_1 p) \ln C_2}{1 + \delta} \\ \text{s.to } Y &= C_1 + \frac{C_2}{1 + r} \end{aligned}$$

with C_t , the consumption in period $t = \{1, 2\}$, δ the discount rate and r the real interest rate. Note that the second term in the above expected utility function corresponds to the expected utility in period 2.⁴

Expected utility is then maximized for the following values of per-period consumptions:

$$\begin{aligned} C_1^* &= Y - \frac{C_2^*}{1 + r}, \\ C_2^* &= \alpha(\phi_1)Y, \end{aligned}$$

³Whether ϕ_0 varies with income or not plays no role in either the theoretical or empirical analysis.

⁴Our framework explicitly abstracts from health insurance. This choice reflects the European institutional setting, in which most medical expenditures are covered by public insurance. However, in the empirical sections, we discuss further and test the validity of this assumption.

where

$$\alpha(\phi_1) = \frac{(1+r)(1+p\phi_1)}{(2+\delta+p\phi_1)} \quad (2)$$

denotes the fraction of income devoted to second-period consumption and depends on the health-state dependence parameter. In line with intuition, $\alpha(\phi_1)$ is increasing in ϕ_1 as a higher marginal utility of consumption in case of illness induces the agent to consume more in the second period (when she can be sick).

We then obtain second-period indirect utilities in the healthy and the unhealthy states respectively as:

$$\begin{aligned} V(Y, S = 0) &= \ln(\alpha(\phi_1)Y), \\ V(Y, S = 1) &= -\phi_0 + (1 + \phi_1) \ln(\alpha(\phi_1)Y). \end{aligned}$$

We derive a health equivalence scale from these indirect utility functions, similar to how one would derive a household equivalence scale. We equate material well-being (utility from consumption) across different health states, rather than overall well-being. In our context, this means we do not account for the direct utility loss from being sick (i.e., $-\phi_0$). Instead, we focus on how health affects utility derived from consumption. The health equivalence scale, denoted by μ , is therefore defined as the income adjustment that equalizes material well-being across health states, satisfying the following condition:

$$V(Y, S = 0) = V(\mu Y, S = 1) + \phi_0 \quad (3)$$

Solving for the above equality, we obtain that⁵

$$\mu = (\alpha(\phi_1)Y)^{\frac{-\phi_1}{1+\phi_1}} \quad (4)$$

⁵This formulation assumes that the proportion of income allocated to second-period consumption remains unaffected by the construction of the health equivalence scale. This simplifying assumption can be justified by two arguments. First, adjusting income by a factor μ has a second-order effect on $\alpha(\cdot)$, which is negligible and can be disregarded when μ is close to one. Second, extending the model's timeline, we can assume that the income variation μ is introduced after the individual has already decided what fraction of their permanent income Y to allocate to second-period consumption. In this scenario, the fraction α is determined by Y and remains constant as μ is introduced. This also facilitates the empirical strategy below, since it allows to estimate μ directly from individual permanent income and consumption data.

The empirical part of the paper will be dedicated to the estimation of individuals' values of this health equivalence scale μ . Since our interest is in the sign of ϕ_1 (i.e., whether preferences exhibit negative or positive health-state dependence) and how ϕ_1 varies with income, we can first provide two analytical results before moving to the empirical estimations.

Result 1. *A value $\mu > 1$ can be observed if and only if $-1 < \phi_1 < 0$, implying negative health-state dependence. To the contrary, $0 < \mu < 1$ can be observed if and only if $\phi_1 > 0$, implying positive health-state dependence. Preferences are health-state independent ($\phi_1 = 0$) when $\mu = 1$.*

Proof. Direct consequence of equation (4), noting that $C_2 = \alpha(\phi_1)Y > 1$ as Y is measured in € in the data used below. □

Result 1 states that positive (resp., negative) health-state dependency is observed when $\mu < 1$ (resp., $\mu > 1$). Note that this is intuitive: with positive health-state dependence, marginal utility is higher when sick ($\phi_1 > 0$) and this must be compensated by giving less income to the individual ($\mu < 1$) under the bad health state so as to keep the same material well-being as under good health. The opposite holds true with negative health-state dependence.

The next result goes one step forward and looks at how ϕ_1 varies with the empirically estimated value of μ .

Result 2. *ϕ_1 decreases with μ if $\phi_1 > 0$. The sign of the variation is ambiguous if $-1 < \phi_1 \leq 0$.*

Proof. It is actually easier to look at $d\mu/d\phi_1$. Take a log-transformation of eq. (4) to obtain

$$\ln \mu = \frac{-\phi_1}{1 + \phi_1} [\ln \alpha(\phi_1) + \ln Y] \quad (5)$$

and differentiate this expression with respect to ϕ_1 :

$$\begin{aligned} \frac{1}{\mu} \frac{d\mu}{d\phi_1} &= -\frac{\ln \alpha(\phi_1)Y}{[1 + \phi_1]^2} - \frac{\phi_1}{1 + \phi_1} \frac{d\alpha(\phi_1)/d\phi_1}{\alpha(\phi_1)} \\ &= \frac{1}{1 + \phi_1} \left[\frac{\ln \mu}{\phi_1} - \frac{d\alpha(\phi_1)}{d\phi_1} \frac{\phi_1}{\alpha(\phi_1)} \right] \end{aligned} \quad (6)$$

where the first term in the second line is obtained by replacing for eq. (5). If $\phi_1 > 0$, then $0 < \mu < 1$ as shown in Result 1. The first term is then negative as is the second one, since $d\alpha(\phi_1)/d\phi_1 > 0$. In that case, $d\mu/d\phi_1 < 0$ so that the higher the estimated μ , the smaller the value of ϕ_1 . To the opposite, as shown in Result 1, if $-1 < \phi_1 \leq 0$, $\mu \geq 1$, the first term inside brackets is positive while the second term is negative. In that situation, the sign of $d\mu/d\phi_1$ is analytically ambiguous. \square

The effect of ϕ_1 on μ is twofold, with each component corresponding to a term in eq. (6). The first effect, which holds α constant with respect to ϕ_1 , is negative: a higher marginal utility of consumption when sick (i.e., a higher value of ϕ_1) implies a lower value of μ . With negative health dependence, an increase in $\phi_1 < 0$ makes health dependence less negative, reducing the need to raise income in the sick state; accordingly, $\mu > 1$ decreases and moves closer to one. With positive health dependence, an increase in ϕ_1 strengthens the incentive to reduce income when sick, so that $\mu < 1$ must decrease further. The second effect relates to how ϕ_1 affects C_2 . A larger ϕ_1 boosts C_2 by increasing $\alpha(\phi_1)$, which enhances second-period utility in both sick and healthy states. Simple computations reveal that utility is more responsive to consumption when sick than when healthy if and only if $\phi_1 > 0$.⁶ In this scenario, μ must be further decreased to rebalance the utilities between the sick and healthy states. Thus, the second effect of ϕ_1 on μ reinforces the first, resulting in a clear negative impact. However, when $\phi_1 < 0$, the second effect requires an increase in μ to equalize utilities, making the overall impact of ϕ_1 on μ ambiguous. In the next section, we will empirically determine how ϕ_1 varies with μ .

To tackle the main point of this paper, namely how health-state dependence varies with income, we now look at how ϕ_1 should vary with Y . Fully differentiating equation (5) with respect to Y , we obtain after some rearrangements that

$$\frac{d\phi_1}{dY} = \frac{1 + \phi_1}{Y} \frac{\frac{\phi_1}{1 + \phi_1} + \frac{d\mu}{dY}}{\left[\frac{\ln Y}{\phi_1} - \frac{d\alpha(\phi_1)}{d\phi_1} \frac{\phi_1}{\alpha(\phi_1)} \right]}$$

Given that ϕ_1 can be either positive or negative and that $d\mu/dY$ is derived from our data, the sign of the expression is inherently ambiguous. In the following section, we will empirically determine this sign.

⁶To see this, let us rewrite the utility under good health as a function of consumption as $\ln C_2$ and under bad health as $(1 + \phi_1) \ln C_2$. Differentiating both utilities with respect to C_2 yields $1/C_2$ and $(1 + \phi_1)/C_2$ respectively.

4 Empirical approach

If utility changes are observable, one can exploit within-individual responses to a health shock across income levels to identify how such shocks affect marginal utility of consumption. This identification, however, requires a proxy of utility that is sufficiently informative to capture genuine state dependence and that is not contaminated by health-related measurement error correlated with permanent income, as is often the case for subjective well-being or self-reported happiness (Finkelstein et al., 2013).

Accordingly, following Kools and Knoef (2019), we move away from cross-income comparisons of overall well-being and instead examine within-individual changes in financial satisfaction, proxied by respondents' ability to make ends meet (see Section 5.2). This design addresses potential bias from unobserved income-related characteristics that may confound the link between health shocks and well-being. The intuition is straightforward: if an individual initially reports being able to make ends meet but reports the opposite after a health shock despite unchanged financial resources, this indicates a shift in the implicit benchmark for material sufficiency.

We now explain how we estimate the health equivalence scale μ . For each individual i at time t , we observe a self-reported measure of financial satisfaction, z_{it} , recorded on an ordinal scale from 1 to K , where higher values indicate greater satisfaction. We assume that financial satisfaction depends on income and individual time t -specific benchmark levels and that individual i reports financial satisfaction level k if his or her permanent income Y_i is at or above a certain benchmark $B_{k-1,it}$, but below $B_{k,it}$ with

$$\ln B_{k,it} = \alpha_k + \rho \ln Y_i + \beta S_{it} + \gamma S_{it} \ln Y_i + X_{it} \eta + v_i + \varepsilon_{it} \quad \forall k \in [1, K-1], B_{0,it} = -\infty \text{ and } B_{K,it} = \infty$$

where S_{it} and Y_i are health status and permanent income, respectively. X_{it} is a vector of time constant and time varying variables of individual i in period t , v_i an individual specific effect, and ε_{it} errors assumed to be distributed as standard normal with mean zero and variance one. Time periods in the empirical analysis can be thought of as repeated observations of an individual in period 2 of the theoretical model. We follow the usual approach in the literature and assume a loglinear specification. We deviate from Kools and Knoef (2019) by introducing the interaction between S and $\ln Y$ since our theoretical framework stresses the interaction

between income and health in explaining financial satisfaction.

If z_{it} denotes the answer to the question on financial satisfaction, the probability of observing a given level of financial satisfaction is given by

$$\begin{aligned}
\text{Prob}(z_{it} = k) &= \text{Prob}(B_{k-1,it} \leq Y_i < B_{kit}) \\
&= \text{Prob}(\ln B_{k-1,it} \leq \ln Y_i < \ln B_{kit}) \\
&= \text{Prob}(\alpha_{k-1} + \rho \ln Y_i + \beta S_{it} + \gamma S_{it} \ln Y_i + X_{it}\eta + v_i + \varepsilon_{it} \leq \ln Y_i < \alpha_k + \rho \ln Y_i + \beta S_{it} \\
&\quad + \gamma S_{it} \ln Y_i + X_{it}\eta + v_i + \varepsilon_{it}) \\
&= \text{Prob}(-\alpha_k + (1 - \rho) \ln Y_i - \beta S_{it} - \gamma S_{it} \ln Y_i - X_{it}\eta - v_i \leq \varepsilon_{it} < -\alpha_{k-1} + (1 - \rho) \ln Y_i \\
&\quad - \beta S_{it} - \gamma S_{it} \ln Y_i - X_{it}\eta - v_i) \\
&= \Phi(-\alpha_{k-1} + (1 - \rho) \ln Y_i - \beta S_{it} - \gamma S_{it} \ln Y_i - X_{it}\eta - v_i) \\
&\quad - \Phi(-\alpha_k + (1 - \rho) \ln Y_i - \beta S_{it} - \gamma S_{it} \ln Y_i - X_{it}\eta - v_i), k = 1, \dots, K
\end{aligned}$$

From which, we can estimate a random effects ordered probit model:

$$z_{it}^* = \theta \ln Y_i + \beta S_{it} + \gamma S_{it} \ln Y_i + \eta X_{it} + v_i + \varepsilon_{it} \quad (7)$$

where $\theta = -(1 - \rho)$ and z_{it} is the observed ordinal response linked to an unobserved latent continuous variable defined as follows:

$$z_{it} = \begin{cases} 1 & \text{if } -\alpha_1 < z_{it}^*, \\ 2 & \text{if } -\alpha_2 < z_{it}^* \leq -\alpha_1, \\ \vdots & \\ K, & \text{if } z_{it}^* \leq -\alpha_{K-1}. \end{cases}$$

To estimate the effects of within individual changes in the variables of interest, we follow Mundlak (1978) and employ a correlated random effect. The individual specific effect is expressed as a linear function of the average time-varying explanatory variables over time, plus a random individual specific effect that is assumed

to be independent of the explanatory variables:

$$v_i = \zeta_0 \bar{S}_i + \bar{X}_i \zeta_1 + \xi_i$$

where \bar{S}_i and \bar{X}_i are the individual means of S_{it} and X_{it} respectively and ξ is i.i.d. normal distributed with mean zero and variance σ_ξ^2 . Equation (7) can then be rewritten as

$$z_{it}^* = \theta \ln Y_i + \beta S_{it} + \gamma S_{it} \ln Y_i + \eta X_{it} + \zeta_0 \bar{S}_i + \zeta_1 \bar{X}_i + \xi_i + \varepsilon_{it} \quad (8)$$

Recall that μ_i is the proportionality factor that equalizes financial satisfaction z_{it} across health states. Specifically, μ is defined such that financial satisfaction under bad health ($S_{it} = 1$) with income $\mu_i Y_i$ equals financial satisfaction under good health ($S_{it} = 0$) with income Y_i . Under this definition and after some rearrangements, μ_i is given by

$$\mu_i = \exp\left(\frac{-\gamma \ln Y_i - \beta}{\theta + \gamma}\right). \quad (9)$$

This equivalence factor can be directly predicted from our data and is individual-specific, as Y is measured at the individual level.

However, we are primarily interested in the sign and magnitude of the health-state dependence parameter ϕ_1 . Result 1 implies that the sign of ϕ_1 can be directly inferred from the observed value of μ (and whether it is above or below unity). We will further look at how ϕ_1 varies with income and at how μ varies with ϕ_1 (e.g. Result 2).

5 Data and sample

We use data from SHARE. This international, interdisciplinary and longitudinal survey provides approximately 380,000 in-depth interviews with 140,000 people aged 50 or older from 28 European countries and Israël. Implemented every two years since 2004, the questions relate to the health status (objective and subjective) of the respondents and their potential partner, and to the economic and social situation of the household, including relations with their children and close relatives and their overall well-being.

5.1 Sample selection

We use waves 2 to 8, excluding waves 3 and 7 that are dedicated to collecting retrospective information, so that our sample runs from 2004 until 2020. For all households, we keep the household respondent who answers the questions related to well-being that we are going to use in our empirical estimations. Because the health status and the decision to work may be simultaneously determined, we focus on respondents aged 65 and over who are no longer in the labor market. This allows us to get rid of any endogeneity between health and labor income, and to treat wealth in the sick state as predetermined. We focus on households of one or two people, keeping only singles and couples in which the partner is also 65+, and we drop individuals once they enter a nursing home. Finally, in order to measure changes in the health status, we restrict the sample to individuals observed in at least two waves and meeting previous criteria in each wave, with non-missing responses to all three health-measure questions (detailed below). This selection leaves us with a sample of 48,922 observations on 17,931 individuals,⁷ covering 15 countries: Austria, Belgium, Switzerland, Germany, Denmark, Spain, France, Greece, Italy, the Netherlands, Sweden, Czech Republic, Poland, Estonia, and Slovenia.⁸

Following the classification in the literature (see for instance, Ucceddu et al. 2019; Lefebvre et al. 2025; Srakar 2018), these countries can be broadly grouped into four categories: Scandinavian or Northern Europe (Sweden, Denmark, Netherlands); Continental or Western Europe (Austria, Belgium, France, Germany, Switzerland); Mediterranean or Southern Europe (Greece, Italy, Spain); and Eastern Europe (Czech Republic, Estonia, Poland, Slovenia). This regional classification closely mirrors the LTC regime typology of Costa-Font and Vilaplana-Prieto (2025), who includes fewer countries but reach similar groupings. Both classifications account not only for institutional differences but also for broader cultural and socio-economic environments. As explained by Kools and Knoef (2019) and Costa-Font and Vilaplana-Prieto (2025), in Northern European countries the government is primarily responsible for the universal provision of care, whereas in Southern Europe families bear most of the responsibility for providing care. In Western and Eastern Europe, the responsibility is shared between the state and families, although the demand for LTC is higher in the former

⁷Of the 17,931 individuals in the sample, 50% are observed in two waves, 31% in three waves, 16% in four waves, and 3% in all five waves. On average, each individual is observed for three waves.

⁸We drop Israël because their surveys were conducted at different points in time and differ slightly from those of the other countries.

than in the latter. These differences are relevant when assessing potential heterogeneity in our results—an issue we examine later in the paper.

5.2 Variables of interest

Following Kools and Knoef (2019), we proxy material well-being with the question on financial satisfaction:⁹ “Thinking of your household’s total monthly income, would you say that your household is able to make ends meet?”. Respondents can answer on a scale taking values 1 (with great difficulty), 2 (with some difficulty), 3 (fairly easily) or 4 (easily). Answers to this question will constitute our dependent variable. We measure an individual’s permanent income as the average over all survey waves of their total annual household income plus 5% of the household’s current wealth (including financial and non-financial wealth minus debts) to account for the fact that elderly households may be spending down their accumulated financial savings (Finkelstein et al., 2013; Kools and Knoef, 2019).¹⁰ Note that, by construction, permanent household income accounts for changes in household composition, such as the death of a partner or divorce, and is adjusted, as all amounts, using the OECD equivalence scale.

The SHARE survey provides a broad set of health indicators, encompassing self-reported health status, functional limitations, major chronic conditions, health care utilization, and measures of physical performance. For comparability, we consider three alternative measures of health status available in the SHARE data. For each measure, we construct an indicator variable, *Unhealthy*, which captures poor health. The first measure is limitations in ADLs.¹¹ The prevalence of functional limitations increases with age, and using ADL limitations as a proxy for health has the advantage of focusing on a dimension of health that is particularly relevant for mediating the relationship between health and consumption.¹² An individual is classified as being in poor health if she reports at least one ADL limitation.¹³ The second health measure is the self-reported health status using a dedicated question in SHARE. On a scale from 1 (excellent) to 5

⁹Kools and Knoef (2019) build on the approach of Pradhan and Ravallion (2000), who derive household-size-dependent poverty lines using survey questions on financial satisfaction.

¹⁰Results are unaffected if we instead assume a 10% of wealth or even 0%.

¹¹Six different limitations are usually considered: 1) dressing, including putting on shoes and socks; 2) walking across a room; 3) bathing or showering; 4) eating, such as cutting up your food; 5) getting in or out of bed; 6) using the toilet, including getting up or down.

¹²Loss of autonomy entails substantial functional limitations (most importantly reduced mobility) that directly constrain opportunities to consume.

¹³In Appendix A.1, we show that our results are robust to assuming that poor health corresponds to having at least two ADLs.

(poor), respondents are asked to evaluate their health. An individual is considered to be in poor health if she reports a self-assessed health status of 4 (“fair”) or 5 (“poor”). The third measure is the number of chronic conditions an individual reports.¹⁴ A respondent is considered in bad health if she has been diagnosed with at least two chronic diseases, corresponding to the median value in our sample.

These three measures can be regarded as complementary. The ADL and chronic condition questions are sufficiently precise and objective to ensure that, despite being self-reported, the answers closely reflect respondents’ actual health status and should not suffer from reporting behaviors (Layes et al., 2012).¹⁵ Self-reported health is also a valuable measure as it is likely to reflect the way individuals feel globally and has been shown to be correlated to major health outcomes and events, such as long-term consumption (Roquebert et al., 2021) and mortality (Jylhä, 2009). Overall, each measure captures situations in which individuals may experience different changes in their consumption patterns following a change in their health status. Appendix A.2 provides details on the distribution of observations and the number of healthy periods in our sample.

Finally, in all estimations, we control for age, gender, if the respondent lives with a partner, number of children, as well as for the household’s net wealth. Table 1 summarizes the main characteristics of our sample by region. Women represent 65% of respondents, the average age is 76.5, and just under half live with a partner. Depending on the health measure used, the share of individuals classified as unhealthy varies substantially, confirming that these indicators capture different dimensions of health. As presented above, we observe on average individuals during three waves of the the survey and they are unhealthy, still on average, during 1.5 to 2 waves, depending on the health measure. The average financial satisfaction score is around 2.9, indicating a generally positive level of satisfaction.

¹⁴Chronic diseases are: a heart attack including myocardial infarction or coronary thrombosis or any other heart problem including congestive heart failure; high blood pressure or hypertension; high blood cholesterol; a stroke or cerebral vascular disease; diabetes or high blood sugar; chronic lung disease such as chronic bronchitis or emphysema; cancer or malignant tumour, including leukaemia or lymphoma, but excluding minor skin cancers; stomach or duodenal ulcer, peptic ulcer; Parkinson disease; cataracts; hip fracture; other fractures; Alzheimer’s disease, dementia, organic brain syndrome, senility or any other serious memory impairment; other affective or emotional disorders, including anxiety, nervous or psychiatric problems; rheumatoid arthritis; osteoarthritis, or other rheumatism.

¹⁵This feature is particularly valuable given that our outcome of interest (the ability to make ends meet) is inherently subjective. When both the dependent variable and the explanatory health variables are subjective, correlated reporting errors could arise and may bias the estimates.

Table 1: Descriptive statistics of the sample

	North	Central	East	South	Total
N	9,261 (18.9%)	18,931 (38.7%)	12,683 (25.9%)	8,047 (16.4%)	48,922 (100.0%)
<i>Demographics</i>					
Age (mean, s.d.)	76.585 (6.878)	76.624 (6.810)	76.205 (6.484)	76.949 (6.610)	76.562 (6.711)
Female	61.1%	61.7%	73.3%	63.3%	64.9%
Has a partner in hh	50.0%	43.7%	40.9%	52.0%	45.5%
Number of children	2.259 (1.321)	2.090 (1.451)	1.880 (1.112)	2.076 (1.448)	2.065 (1.351)
Net wealth	261,063 (348,311)	308,417 (531,381)	73,111 (101,990)	191,545 (273,152)	219,226 (395,303)
Permanent income	20,943 (10,647)	26,791 (20,689)	8,587 (5,579)	11,803 (6,571)	17,623 (15,406)
<i>Utility proxy</i>					
Financial satisfaction (mean, s.d.)	3.368 (0.805)	3.171 (0.867)	2.559 (0.928)	2.372 (0.967)	2.918 (0.968)
<i>Unhealthy</i>					
1 or more ADL limitation	10.4%	16.5%	17.7%	15.3%	15.4%
Self-perceived	32.1%	39.0%	64.7%	54.1%	46.9%
Chronic conditions	55.1%	58.8%	67.2%	66.2%	61.5%

6 Results

Table 2 presents the estimation results of equation (8) for the three health measures (ADL, self-reported health, chronic diseases). In Panel A, all three estimations present similar results with a negative effect of a health shock on our utility proxy, i.e. the ability to make ends meet ($\beta < 0$) but a positive interaction term between income and health ($\gamma > 0$). The positive sign indicates that the marginal utility of permanent income as defined in equation (8) increases as health worsens. It is statistically significant at the 1% level for two health measures (self-reported health and chronic disease) and at the 5% level for ADL measure. We can thus reject the null hypothesis of state-independent utility.

Table 2: Estimation results

	(1)	(2)	(3)
	ADL	Self-reported	Chronic
A. Estimates			
$\ln(Y):\theta$	1.151*** (0.036)	1.024*** (0.040)	0.927*** (0.068)
Unhealthy: β	-1.072** (0.464)	-1.804*** (0.377)	-1.633*** (0.356)
$\ln(Y) \times$ Unhealthy: γ	0.109** (0.049)	0.168*** (0.040)	0.160*** (0.037)
Age	-0.019** (0.009)	-0.005 (0.009)	-0.025*** (0.009)
Female	-0.249*** (0.035)	-0.236*** (0.035)	-0.224*** (0.035)
Has a partner in hh	0.398*** (0.060)	0.401*** (0.060)	0.399*** (0.060)
Number of children	-0.051*** (0.011)	-0.055*** (0.011)	-0.053*** (0.011)
Net wealth (in 10,000s)	0.003*** (0.001)	0.003*** (0.001)	0.003*** (0.001)
Cut-off point 1	10.344 (0.433)	8.650 (0.457)	6.934 (0.704)
Cut-off point 2	13.135 (0.436)	11.445 (0.459)	9.723 (0.705)
Cut-off point 3	15.591 (0.438)	13.905 (0.461)	12.180 (0.706)
σ^2	2.606 (0.066)	2.558 (0.065)	2.628 (0.066)
Observations	48922	48922	48922
B. Health state dependence:			
average estimated μ ($\bar{\mu}$)	1.021	1.181	1.105
95% CI	[1.021, 1.023]	[1.180, 1.183]	[1.104, 1.107]

Data: SHARE, wave 2, 4, 5, 6, 8.

Sample: 48,922 observations corresponding to 17,931 individuals aged 65+, living alone or with a 65+ partner.

Notes: In the above table, reported figures are the regression coefficients of the random effects ordered probit model (equation 8). Standard errors clustered at the individual level are reported in parentheses. All specifications contain country and wave fixed effects as well as Mundlak terms. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

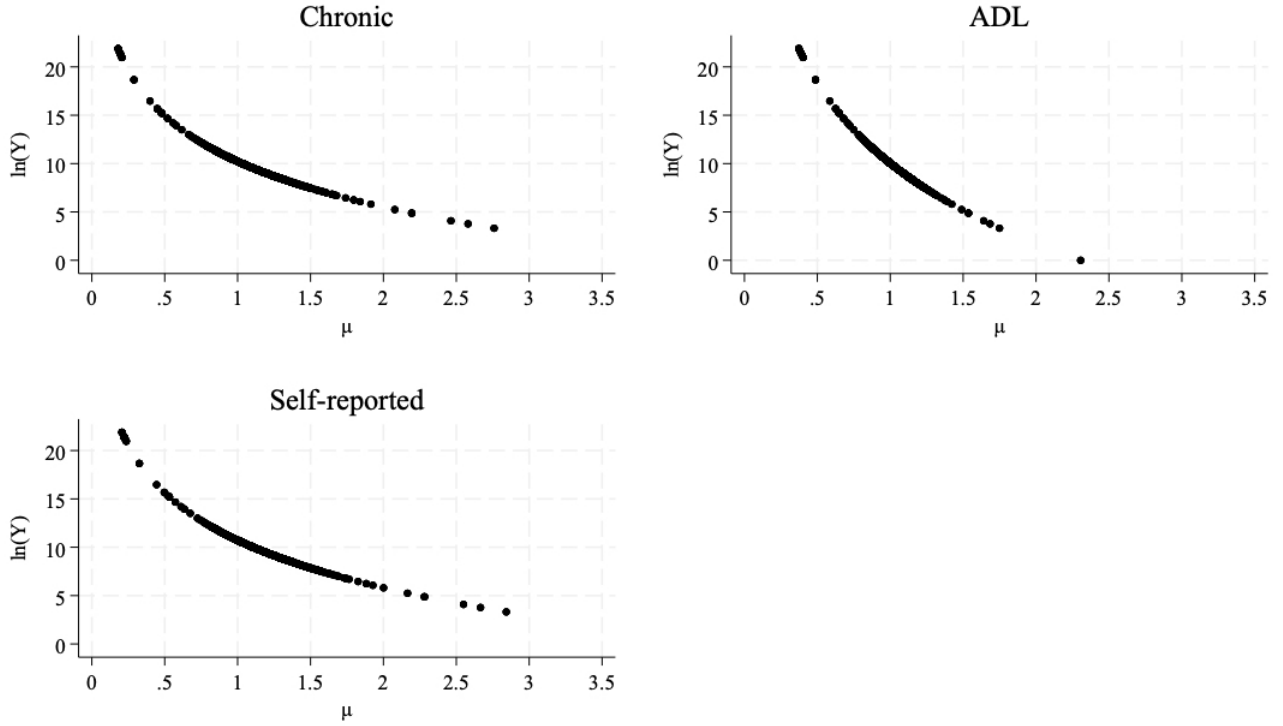
Using equation (9), we compute the individual-specific health equivalence scale parameter μ as a function of each individual's permanent income level. Table 2 (Panel B) reports the average estimated value of μ across individuals, denoted by $\bar{\mu}$. Depending on the measure of health retained, $\bar{\mu}$ ranges between 1.021 and 1.181. This implies that individuals on average need between 2.1% and 18.1% more income after a health shock to reach the same material well-being as before. This is in accordance with results obtained by Kools and Knoef (2019) who find an average value for μ of 1.133 in their baseline specification, using the same SHARE data.

Figure 1 illustrates how the health equivalence parameter μ varies with permanent income for the three health measures considered. We find a clear negative relationship between μ and income such as μ exceeds one at lower income levels but falls below one as income rises. The income level at which $\mu = 1$ also varies substantially across health measures.¹⁶ When health is proxied by limitations in ADLs, the average income associated with $\mu = 1$ is approximately €18,000. Using self-assessed health yields a substantially higher threshold, around €45,000, while the measure based on the number of chronic conditions produces an intermediate estimate of roughly €27,000. These differences underscore the sensitivity of the implied income threshold to the choice of health indicator.

Moreover, when health is proxied by limitations in ADLs, the estimated values of μ_i are less dispersed than those obtained using self-reported health or the presence of chronic conditions. These findings complement the results of Kools and Knoef (2019), who estimate an average value of μ for the population as a whole but do not investigate heterogeneity in health-state dependence across the income distribution.

¹⁶To estimate the income threshold corresponding to $\mu = 1$, we use a non-parametric procedure based on local variation in the data. Specifically, we identify observations for which μ lies within a narrow interval around one and compute the average income among these observations.

Figure 1: $\hat{\mu}$ as a function of permanent income



Data: SHARE, wave 2, 4, 5, 6, 8.

Note: Estimated income levels corresponding to $\mu = 1$ are approximately €18,000 (ADL limitations), €27,000 (number of chronic conditions), and €45,000 (self-assessed health).

Sample: 48,922 observations corresponding to 17,931 individuals aged 65+, living alone or with a 65+ partner.

At this stage, it is important to discuss the role of health insurance in shaping individual preferences. In our baseline setting, we have implicitly assumed that medical expenses are fully covered by either public or private insurance systems. This is indeed quite reasonable for the vast majority of the European countries under study. Yet, it may still be the case that individuals face incomplete insurance and incur out-of-pocket medical costs. In this case, the observed positive health-state dependence we estimate could be driven by higher medical spending rather than by a genuine increase in the marginal utility of non-medical consumption. To address this concern, we follow Kools and Knoef (2019) and re-estimate the regressions reported in Table 2 by including out-of-pocket medical expenditures as an additional control. The corresponding results are

presented in Appendix A.3. Although the inclusion of this variable substantially reduces the sample size, the estimates remain qualitatively unchanged and our conclusions are unaffected.¹⁷ The estimated average equivalence scale factor lies between 1.038 and 1.117.

Using Result 1, we can further infer the sign of the health-state-dependence parameter ϕ_1 directly from Figure 1. Independently from the measure of health considered, the different figures suggest negative health-state dependence for lower-income individuals (e.g. $\phi_1 < 0$ when $\mu > 1$) and positive health-state dependence for higher-income individuals (e.g. $\phi_1 > 0$ when $\mu < 1$). This means that consumption and health are more likely to be complements in the lower part of the income distribution, although they are more likely to be substitutes at the top of the income distribution.

However, obtaining precise individual-specific values of ϕ_1 is more challenging, as we now explain. As shown in Eq. 2, α is itself a function of ϕ_1 . Substituting the expression $\alpha(\phi_1)$ into Eq. 4 yields:

$$\mu = \left(\frac{(1+r)(1+p\phi_1)}{2+\delta+p\phi_1} Y \right)^{-\frac{\phi_1}{1+\phi_1}} \quad (10)$$

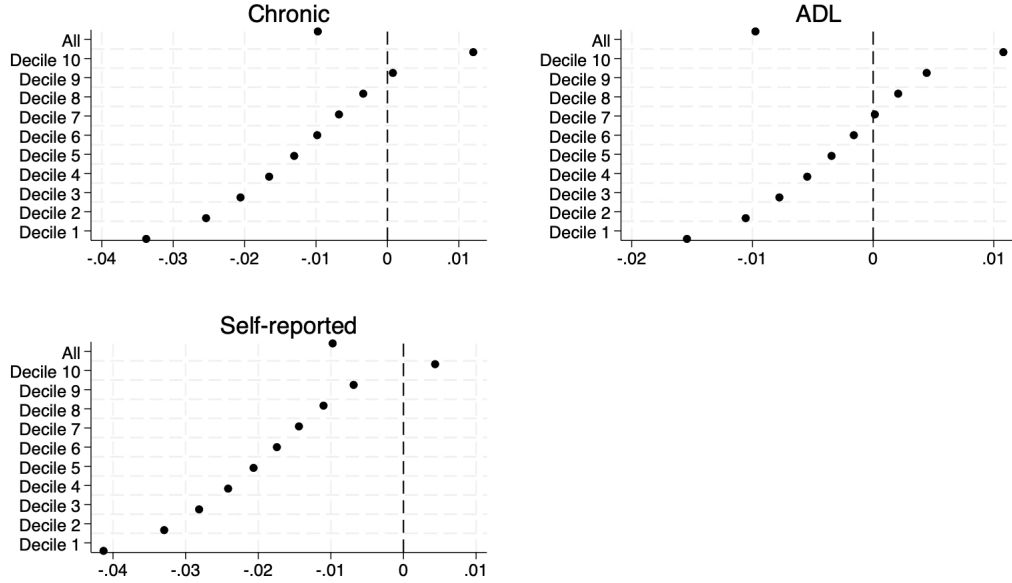
Consequently, μ is a non-linear function of ϕ_1 , which hinders the identification of ϕ_1 as a function of Y_i .

In order to circumvent this issue, we propose two illustrative numerical attempts. The first exercise estimates ϕ_1 using non-linear least squares applied to Eq. (10). To this end, we set $r = 0.03$ and $\delta = 0.98$, while p is calibrated using our data as the average probability of transitioning into an unhealthy state.¹⁸ Figure 2 displays the estimates of ϕ_1 by income decile, for the three indicators of health status. It shows similar patterns across the different measures of health with ϕ_1 increasing with income and remaining negative up to the top deciles of the income distribution. For instance, in the case of chronic conditions, the average value of $\hat{\phi}_1$ is -0.01 , and the individual estimated values of ϕ_1 become positive only at the very top of the income distribution, with values close to zero in the ninth decile and strictly positive for the highest decile. For ADL limitations, the average value of ϕ_1 is very close to zero, with individual estimated values of ϕ_1 becoming positive at lower deciles (from the eighth decile).

¹⁷Out-of-pocket medical expenditure data in SHARE are subject to substantial non-response and changes in questionnaire design. As a result, wave 4 must be excluded from the analysis, leading to a significant loss of observations.

¹⁸From the SHARE data, we estimate the probability of having at least one limitation to be 0.17, the probability of self-reporting poor health to be 0.47, and the probability of having at least one chronic condition to be 0.61. Using the average probability by decile of permanent income yields very similar results.

Figure 2: Estimation of $\hat{\phi}_1$ by income decile



Data: SHARE, wave 2, 4, 5, 6, 8.

Sample: 48,922 observations corresponding to 17,931 individuals aged 65+, living alone or with a 65+ partner.

As a second exercise, we estimate individual values of ϕ_1 by making the strong assumption that the propensity to consume, i.e. the parameter α , is independent of ϕ_1 . This allows us to isolate ϕ_1 by log-linearizing Eq. (4) and rearranging terms, as follows:

$$\hat{\phi}_1 = -\frac{A}{1+A}$$


$$\text{with } A = \frac{\text{Log}(\hat{\mu})}{\text{Log}(\alpha Y)}.$$

Recognizing that second-period consumption satisfies $C_2^* = \alpha Y$ (see Section 3) and using our estimates of $\hat{\mu}$, we can recover the individual-specific parameter $\hat{\phi}_1$ from the equation above.¹⁹ Note that second-period consumption C_2 is approximated using reported total household expenditures observed in the data, with each observation (wave in the data) corresponding to a new realization of C_2 . We keep the last observation of

¹⁹An alternative approach is to treat α as a random parameter and compute $\hat{\phi}_1$ for different values of $0 \leq \alpha \leq 1$. The corresponding results are reported in Figure A.4 in Appendix A.4.

C_2 for each individual.²⁰ As shown in Figure 3, this second exercise confirms the presence of negative (resp. positive) health-state dependence for low-income (resp. high-income) individuals.

Figure 3: $\hat{\phi}_1$ as a function of Y - using second-period consumption



Data: *SHARE*, wave 2, 4, 5, 6, 8.

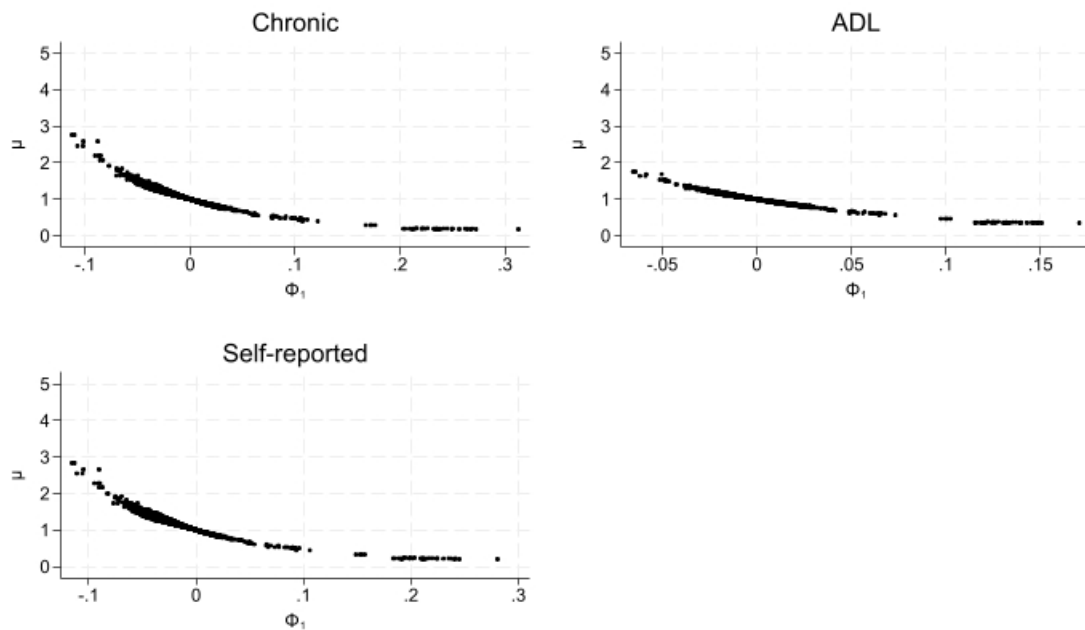
Sample: 48,922 observations corresponding to 17,931 individuals aged 65+, living alone or with a 65+ partner.

Lastly, we want to test Result 2 (Section 3) which predicts a negative relationship between μ and ϕ_1 for positive values of ϕ_1 , whereas the sign of this relationship is indeterminate when $-1 < \phi_1 \leq 0$. Our empirical estimations allow us to look at these variations. Figure 4 shows a negative relationship between μ and ϕ_1 across the entire range of ϕ_1 , with a steeper slope for negative values of ϕ_1 . This means that the smaller (or more negative) the health-state dependence parameter, the higher the income required to adjust for poor health. Coming back to the intuition for this result, this implies that the first-order effect (i.e. a higher

²⁰As a robustness check, we also construct C_2 analogously to permanent income Y , by averaging consumption across all periods in the data. Results are presented in Figure A.5 in Appendix A.4 and are very similar to those obtained in Figure 3.

marginal utility when sick lowers the value of μ) always dominates the second-order effect of ϕ_1 through C_2 .

Figure 4: $\hat{\phi}_1$ and $\hat{\mu}$ - using second-period consumption to estimate ϕ_1



Data: SHARE, wave 2, 4, 5, 6, 8.

Sample: 48,922 observations corresponding to 17,931 individuals aged 65+, living alone or with a 65+ partner.

7 Heterogeneity and sensitivity analysis

We next conduct a series of additional estimations to explore potential heterogeneity and assess the robustness of our results. Since the results for the three health measures are extremely similar (as shown in Table 2), we focus when necessary on being unhealthy according to the number of chronic conditions.²¹

Group of countries. We first examine heterogeneity across groups of European countries. As discussed in Section 5, these countries differ markedly in their health care systems, and out-of-pocket medical expenditures are therefore likely to be higher where public provision of health care is more limited. In Table 3, we stratify the analysis by country group and estimate the average health equivalence factor, $\bar{\mu}$, separately for each group.

We find that the health dependence results reported in Panel A of Table 2 remain statistically significant only for Central and Eastern European countries: for these two groups, the coefficients θ , β , and γ are all jointly significant. The average health equivalence scale parameter $\bar{\mu}$ is larger in Central and Southern Europe than in Northern and Eastern Europe. In Northern countries, $\bar{\mu}$ is close to one, indicating little or no health-state dependence. Finally, Figure 5 shows that the dispersion of $\bar{\mu}$ is substantially greater in Central Europe than in any other European region.

Marital status, gender and age. Singles and couples may interpret the “making ends meet” question differently. In particular, individuals living in a couple may benefit from income insurance provided by the partner, which can attenuate the impact of individual health shocks on perceived financial strain. Moreover, correlated health shocks within couples may blur the identification of the effect of one’s own health on one’s own marginal utility. Gender differences in income levels and health trajectories may also lead to heterogeneous health-state dependence between men and women. Finally, among retired people, age can affect the health dependence parameter, due to differences in consumption expectations.

Accordingly, in Table 4, we estimate the health dependence scale parameter separately for singles and couples, as well as for men and women. We find that the health equivalence factor is higher for singles than for couples (in which case, it is close to 0), and higher for women than men (also closer to 0). In Figure 6,

²¹We also examined heterogeneity by home ownership status. These results are not reported because home ownership is highly prevalent in the sample (70% overall) and displays little variation in several countries, including Spain (94%), Italy (85%), France (77%), and Belgium (75%). The resulting small number of renters limits meaningful comparisons.

for all four groups, we find that, as previously observed (see Figure 1), $\hat{\mu}$ is decreasing with income, with lower-income (resp. higher-income) individuals exhibiting negative (resp. positive) health-state dependence.

Table 3: Estimation results by group of countries

	North	Central	South	East
A. Estimates:				
ln(Y): θ	1.551*** (0.272)	0.592*** (0.104)	1.644*** (0.197)	0.846*** (0.179)
Unhealthy: β	-1.116 (1.531)	-1.962*** (0.627)	-0.381 (1.051)	-2.050** (0.915)
ln(Y)*Unhealthy: γ	0.114 (0.156)	0.183*** (0.063)	0.014 (0.113)	0.226** (0.102)
N	9,261	18,931	8,047	12,683
B. Health state dependence:				
$\bar{\mu}$	0.999	1.198	1.167	1.032
95 % CI	[0.998,0.999]	[1.195,1.201]	[1.167,1.168]	[1.031,1.033]

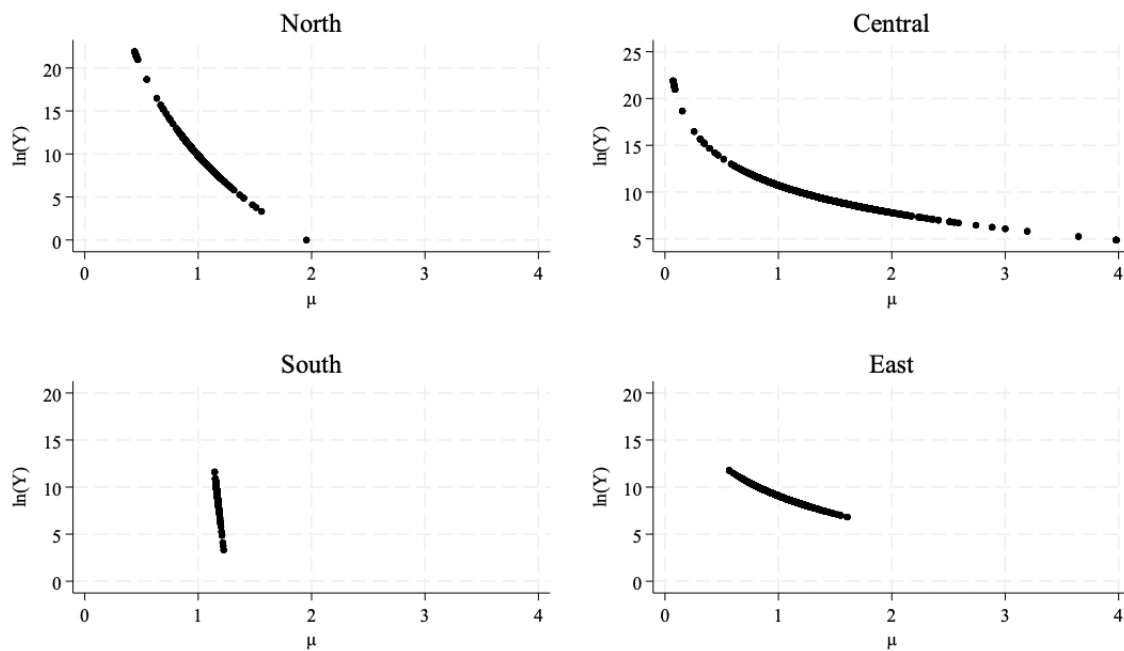
Notes: Standard errors clustered at the individual level are reported in parentheses. All specifications contain controls for age, gender, the number of children, the presence of a partner and the net wealth. The health measure used is the number of chronic diseases. Country and wave fixed effects are included as well as Mundlak terms. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Estimation results by marital status, gender and age

	By marital status		By gender		By age	
	Single	Couple	Men	Women	Age \leq 75	Age $>$ 75
A. Estimates:						
ln(Y): θ	0.794*** (0.091)	1.121*** (0.108)	0.797*** (0.110)	1.026*** (0.088)	1.050*** (0.104)	0.806*** (0.093)
Unhealthy: β	-1.868*** (0.467)	-1.186** (0.581)	-1.768*** (0.612)	-1.340*** (0.447)	-2.876*** (0.561)	-1.368*** (0.477)
ln(Y)*Unhealthy: γ	0.183*** (0.049)	0.118** (0.060)	0.175*** (0.063)	0.128*** (0.047)	0.289*** (0.059)	0.132*** (0.050)
N	26664	22258	17190	31732	22176	26746
B. Health state dependence:						
$\bar{\mu}$	1.168	1.035	1.073	1.116	1.091	1.129
95 % CI	[1.167,1.171]	[1.034,1.036]	[1.070,1.075]	[1.115,1.117]	[1.088,1.092]	[1.128,1.131]

Notes: Standard errors clustered at the individual level are reported in parentheses. All specifications contain controls for age, gender, the number of children, the presence of a partner and the net wealth. The health measure used is the number of chronic diseases. Country and wave fixed effects are included as well as Mundlak terms. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Figure 5: $\hat{\mu}$ as a function of Y - by group of countries

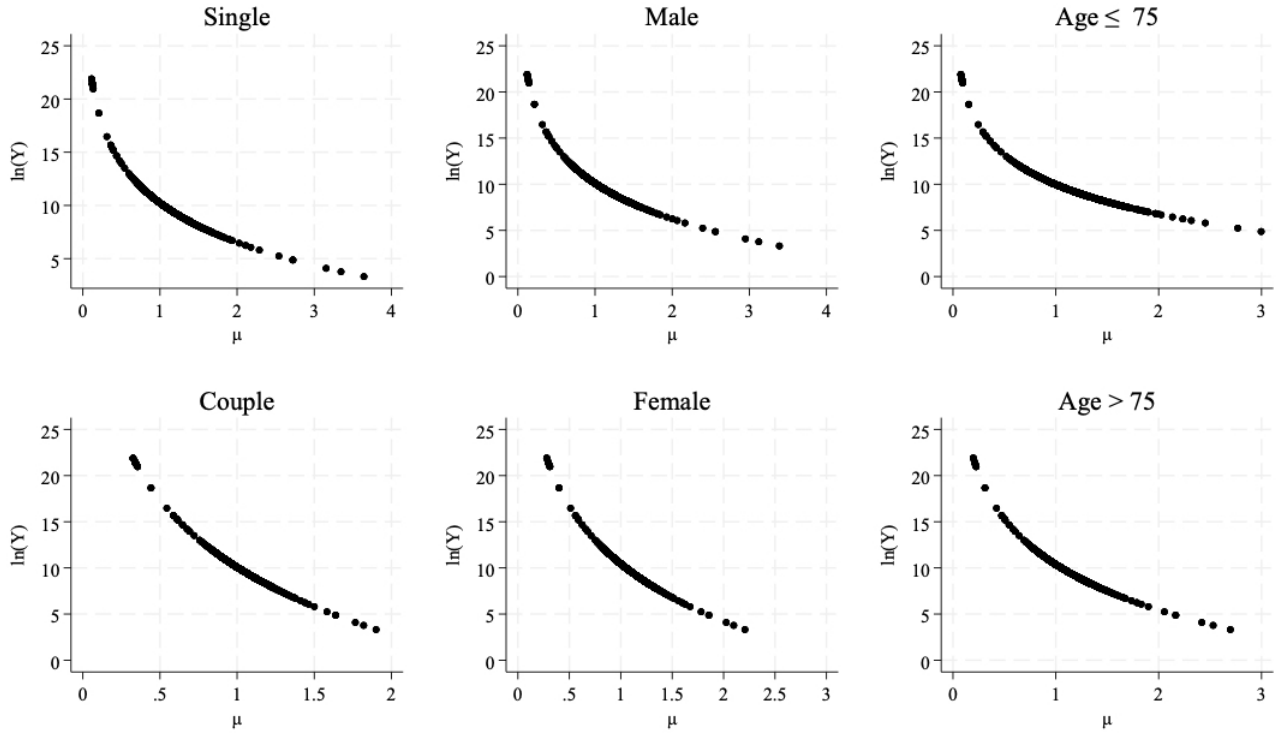


Data: SHARE, wave 2, 4, 5, 6, 8.

Sample: 48,922 observations corresponding to 17,931 individuals aged 65+, living alone or with a 65+ partner.

Note: Estimated income levels corresponding to $\mu = 1$ are approximately €18,450 in Northern countries, €45,000 in Central countries, and €9,400 in Eastern countries.

Figure 6: $\hat{\mu}$ as a function of Y - by marital status, gender and age



Data: SHARE, wave 2, 4, 5, 6, 8.

Sample: 48,922 observations corresponding to 17,931 individuals aged 65+, living alone or with a 65+ partner.

Note: Estimated income levels corresponding to $\mu = 1$ are approximately €28,000 for singles, €24,000 for couples, €24,000 for male, 34,000 for female, 21,000 for people aged less or equal than 75, and 32,000 for people aged more than 75.

Robustness checks. Finally, we conduct three robustness checks using chronic diseases as the health measure. First, we assess the robustness of our results to alternative specifications. We look at how the results are dependent on the exclusion of covariates (Table 5, first column). It shows that the estimated health state dependence parameter is very similar to Table 2. All our estimations are obtained using a random effects ordered probit model, but a linear specification is another option. The second column of Table 5 shows the

results for the linear specification and here also it does not change much the qualitative results. Finally, we include country-by-wave fixed effects rather than country effects and wave effects separately (Table 5, third column). It allows us to account for possible country-specific shocks. Results are stable.

Second, health shocks may affect both actual and perceived life expectancy, thereby influencing consumption behavior and reported material satisfaction. Since marginal utility of income is likely to decline as individuals approach death, such mechanisms could bias our estimates of health-state dependence downward. To assess robustness, we therefore re-estimate the model excluding individuals who died during the observation period (Table 5, fourth column). We find an estimated average health equivalence factor $\bar{\hat{\mu}}$ equal to 1.100, close to our baseline estimation of 1.105 in Table 2.

Third, as noted by Kools and Knoef (2019), financial satisfaction is reported on a four-point scale, which introduces potential bias due to its finite range. Our baseline results indicate that on average, individuals report lower financial satisfaction following a health shock. However, for respondents already at the lowest category prior to the shock, further declines cannot be expressed in the survey measure. This censoring at the lower bound may therefore bias downward the estimated health-state dependence parameter and affect its variation with income. As a robustness check, we repeat the main regression, while excluding those in the lowest income. In addition, we also exclude individuals in the highest income quartile to assess whether the estimated income gradient is driven by observations at the top of the income distribution, where financial satisfaction may be less sensitive to health shocks due to higher levels of financial buffering, differential reporting behavior, or ceiling effects in subjective measures. Results from these restricted samples are reported in the last two columns of Table 5. Reassuringly, the estimated effects remain very similar in magnitude and significance, indicating that our findings are not driven by either censoring at the bottom or leverage from the top of the income distribution.

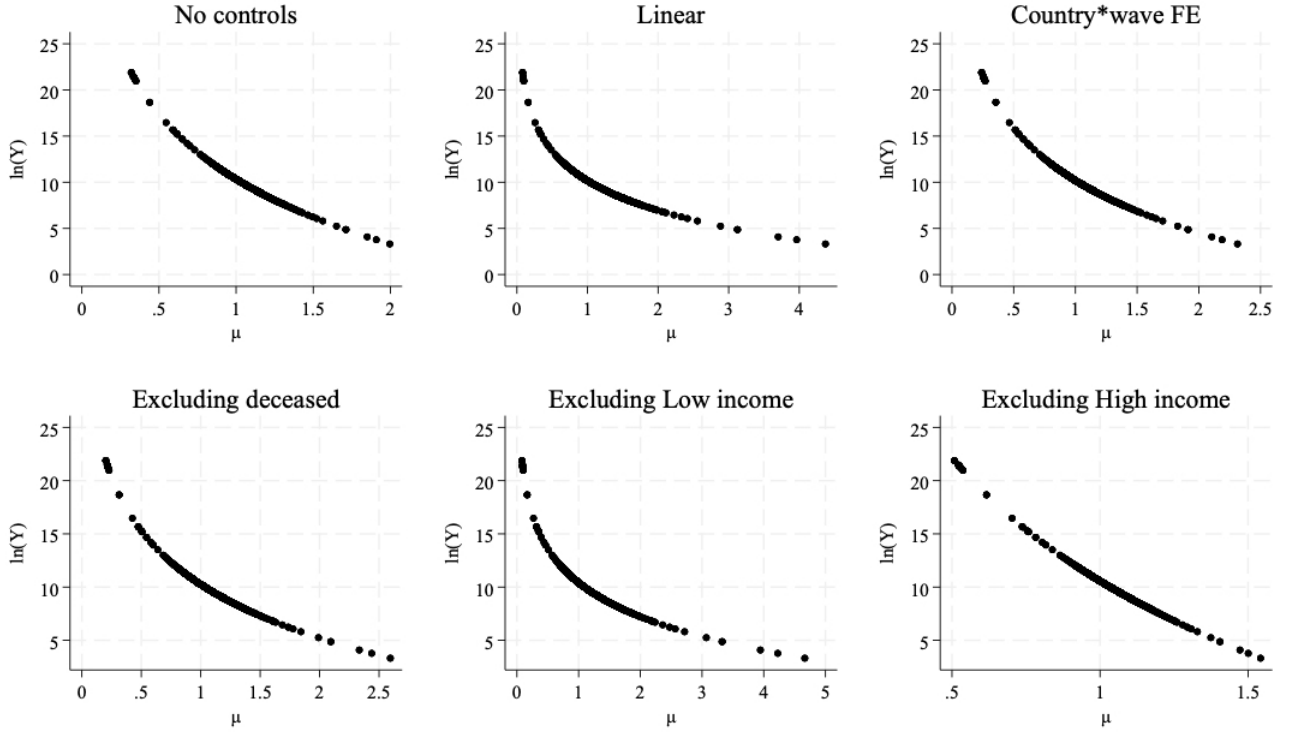
Finally, Figure 7 which reports the variations of $\hat{\mu}$ by income under these alternative specifications shows results similar to those observed in our baseline estimations. Low-income individuals exhibit negative health-state dependence (e.g. $\hat{\mu} > 1$) while high-income individuals exhibit positive dependence (e.g. $\hat{\mu} < 1$).

Table 5: Robustness analysis

	Specification			Exclude individuals		
	No controls	Linear	Country× year FE	Deceased	With low Y (Q1)	With high Y (Q4)
A. Estimates:						
ln(Y): θ	1.328*** (0.068)	0.225*** (0.017)	1.067*** (0.068)	0.949*** (0.074)	0.664*** (0.092)	1.557*** (0.110)
Unhealthy: β	-1.499*** (0.361)	-0.627*** (0.094)	-1.512*** (0.360)	-1.552*** (0.385)	-1.918*** (0.531)	-1.045* (0.553)
ln(Y)*Unhealthy: γ	0.145*** (0.038)	0.062*** (0.010)	0.148*** (0.038)	0.151*** (0.040)	0.184*** (0.054)	0.099* (0.059)
N	48922	48922	48922	42500	36690	36693
B. Health state dependence:						
$\bar{\mu}$	1.080	1.149	1.082	1.100	1.135	1.081
95 % CI	[1.079,1.081]	[1.147,1.151]	[1.081,1.083]	[1.099,1.101]	[1.133,1.136]	[1.080,1.081]

Notes: Standard errors clustered at the individual level are reported in parentheses. All specifications, except the column "No controls", contain controls for age, gender, the number of children, the presence of a partner and the net wealth. The health measure used is the number of chronic diseases. Country and wave fixed effects are included as well as Mundlak terms. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Figure 7: $\hat{\mu}$ as a function of Y - by different specifications



Data: SHARE, wave 2, 4, 5, 6, 8.

Sample: 48,922 observations corresponding to 17,931 individuals aged 65+, living alone or with a 65+ partner.

8 Conclusion

This paper examines the relationship between health status and marginal utility of consumption, with a particular focus on how this relationship varies across income levels. While the existing literature typically provides an average estimate of the health-state dependence parameter for the population as a whole, we explicitly account for heterogeneity in health-state dependence by income. We develop a model in which individuals differ only in income and allocate resources over two periods, facing uncertainty about a health shock in the second period. Guided by the model's predictions, we then use SHARE data to estimate the empirical relationship between health and the marginal utility of consumption across the income distribution.

Our empirical results provide clear evidence that health-state dependence varies systematically along the income distribution. For low-income individuals, health deterioration is associated with a decline in the marginal utility of consumption, indicating negative health-state dependence. As income rises, this effect weakens and eventually reverses: at the top of the income distribution, marginal utility increases when health worsens, implying positive health-state dependence. These findings are robust across alternative measures of health—including ADL limitations, self-reported health, and chronic conditions—and to a wide range of sensitivity analyses accounting for institutional differences, demographic heterogeneity, and potential confounding factors.

This pattern may reflect differences in individuals’ ability to adjust the composition of consumption in response to health deterioration. Higher-income individuals are better able to substitute toward goods and services that are compatible with poorer health, such as at-home comfort services, home-delivered meals, or in-home care, whereas lower-income individuals face tighter constraints in adapting their consumption bundle. This interpretation is consistent with Amo-Agyei (2025), who shows that the marginal utility of consumption increases under poor health, and that this is mainly driven by an increase in marginal utility for luxury goods (including tourism and leisure, non-medical services, and alcohol and tobacco) rather than for necessity goods (such as basic food, housing-related expenses, and utilities).

By documenting this income gradient, our results help reconcile the mixed evidence in the literature on health-state dependence (see the Introduction and the Literature Review). Studies that focus on different segments of the income distribution, or that consider averages across population groups, may therefore reach opposing conclusions about whether marginal utility increases or decreases following health deterioration. Our findings suggest that income heterogeneity in health-state dependence, rather than methodological inconsistency, is a primary source of divergence in the existing literature.

These results have important policy implications. In particular, they question the welfare optimality of uniform income transfers or LTC compensation schemes following health deterioration at older ages. Because low-income individuals experience a sharp decline in the marginal utility of consumption when health worsens, they require proportionally larger income compensation to maintain a level of material well-being comparable to that enjoyed in good health. By contrast, among higher-income individuals, health deterioration may raise the marginal value of consumption—potentially because consumption can be redirected toward health-related

goods and services—implying that smaller compensatory transfers may be sufficient to offset welfare losses. This heterogeneity suggests that redistributive policies and long-term care benefits should be tailored to both income and health status. At the same time, this paper remains agnostic about the precise mechanisms underlying these patterns and the planner’s objective so that further research would be needed to better understand the channels at work and to determine the precise forms these optimal transfers should take.

More broadly, our findings underscore the importance of accounting for state-dependent preferences in the evaluation of social insurance programs, particularly in aging societies where health shocks and long-term care needs are increasingly prevalent. Future research could extend this framework by explicitly modeling long-term care choices, informal care provision, or bequest motives, as well as by exploring how health-state dependence interacts with wealth accumulation and insurance coverage over the life cycle.

Bibliography

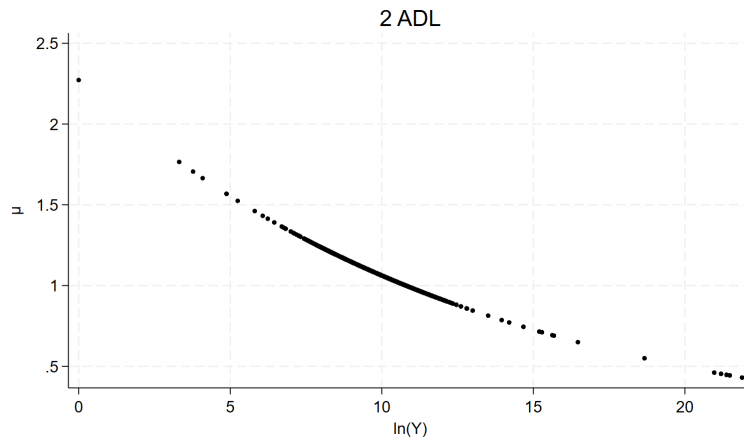
- Achou, B., De Donder, P., Glenzer, F., Lee, M., and Leroux, M.-L. (2026). At home versus in a nursing home: Long-term care settings and marginal utility. *Journal of Economic Dynamics and Control*, 183:105254.
- Ameriks, J., Briggs, J., Caplin, A., Shapiro, M. D., and Tonetti, C. (2020). Long-term-care utility and late-in-life saving. *Journal of Political Economy*, 128(6):2375–2451.
- Amo-Agyei, S. (2025). The effect of health on the marginal utility of consumption: Evidence based on household consumption expenditure. *Journal of Economic Behavior & Organization*, 229:106809.
- Bassoli, E. (2023). An empirical analysis of health-dependent utility on share and elsa data. *Italian Economic Journal: A Continuation of Rivista Italiana degli Economisti and Giornale degli Economisti*, 9(3):1217–1261.
- Blundell, R., Borella, M., Commault, J., and De Nardi, M. (2024). Old age risks, consumption, and insurance. *American Economic Review*, 114(2):575–613.
- Costa-Font, J. and Vilaplana-Prieto, C. (2025). The hidden value of adult informal care in europe. *Health Economics*, 34(4):791–812.
- De Donder, P. and Leroux, M.-L. (2021). Long term care insurance with state-dependent preferences. *Health Economics*, 30(12):3074–3086.
- 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.
- Evans, W. N. and Viscusi, W. K. (1991). Estimation of State-Dependent Utility Functions Using Survey Data. *The Review of Economics and Statistics*, 73(1):94–104.
- 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*, 11:221–258.
- Hong, J. H., Pijoan-Mas, J., and Rios-Rull, J.-V. (2013). Health heterogeneity and preferences. *working paper*.
- Jylhä, M. (2009). What is self-rated health and why does it predict mortality? towards a unified conceptual model. *Social science & medicine*, 69(3):307–316.

- Koijen, R., Van Nieuwerburgh, S., and Yogo, M. (2016). Health and mortality delta: Assessing the welfare cost of household insurance choice. *The Journal of Finance*, 71(2):957–1010.
- Kools, L. and Knoef, M. (2019). Health and consumption preferences; estimating the health state dependence of utility using equivalence scales. *European Economic Review*, 113(C):46–62.
- Layes, A., Asada, Y., and Kephart, G. (2012). Whiners and deniers—what does self-rated health measure? *Social science & medicine*, 75(1):1–9.
- Lefebvre, M., Pestieau, P., and Schoenmaeckers, J. (2025). Grandchild care and eldercare. a quid pro quo arrangement. *Economic Modelling*, 146:107006.
- Lillard, L. A. and Weiss, Y. (1997). Uncertain health and survival: Effects on end-of-life consumption. *Journal of Business & Economic Statistics*, 15(2):254–268.
- Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica*, 46(1):69–85.
- Pradhan, M. and Ravallion, M. (2000). Measuring poverty using qualitative perceptions of consumption adequacy. *The Review of Economics and Statistics*, 82(3):462–471.
- Roquebert, Q., Sicsic, J., Rapp, T., and Consortium, S.-T. (2021). Health measures and long-term care use in the european frail population. *The European Journal of Health Economics*, 22(3):405–423.
- Srakar, Andrej; Nagode, M. (2018). Why are older informal carers in better health? solving a causality problem. *Public Sector Economics*, 42(3):303–323.
- Uccheddu, D., Gauthier, A. H., Steverink, N., and Emery, T. (2019). Gender and socioeconomic inequalities in health at older ages across different european welfare clusters: Evidence from share data, 2004–2015. *European Sociological Review*, 35(3):346–362.
- Viscusi, W. K. (2019). Utility functions for mild and severe health risks. *Journal of Risk and Uncertainty*, 58(2):143–166.
- Viscusi, W. K. and Evans, W. N. (1990). Utility Functions That Depend on Health Status: Estimates and Economic Implications. *American Economic Review*, 80(3):353–374.

A Appendix

A.1 Results with at least two ADL limitations

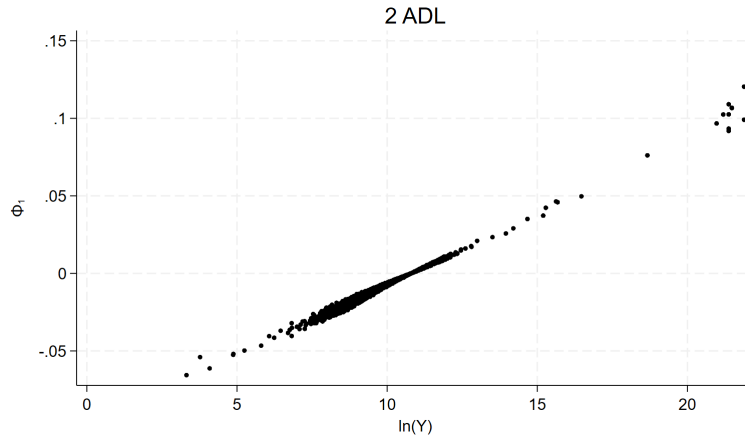
Figure A.1: $\hat{\mu}$ as a function of permanent income (at least 2 ADL limitations)



Data: SHARE, wave 2, 4, 5, 6, 8.

Sample: 48,922 observations corresponding to 17,931 individuals aged 65+, living alone or with a 65+ partner.


Figure A.2: $\hat{\phi}_1$ as a function of Y - using second-period consumption (at least 2 ADL limitations)



Data: SHARE, wave 2, 4, 5, 6, 8.

Sample: 48,922 observations corresponding to 17,931 individuals aged 65+, living alone or with a 65+ partner.

Figure A.3: Estimation of $\hat{\phi}_1$ by income decile (at least 2 ADL limitations)



Data: SHARE, wave 2, 4, 5, 6, 8.

Sample: 48,922 observations corresponding to 17,931 individuals aged 65+, living alone or with a 65+ partner.

A.2 Number of observations and healthy periods

A.3 Taking into account medical costs

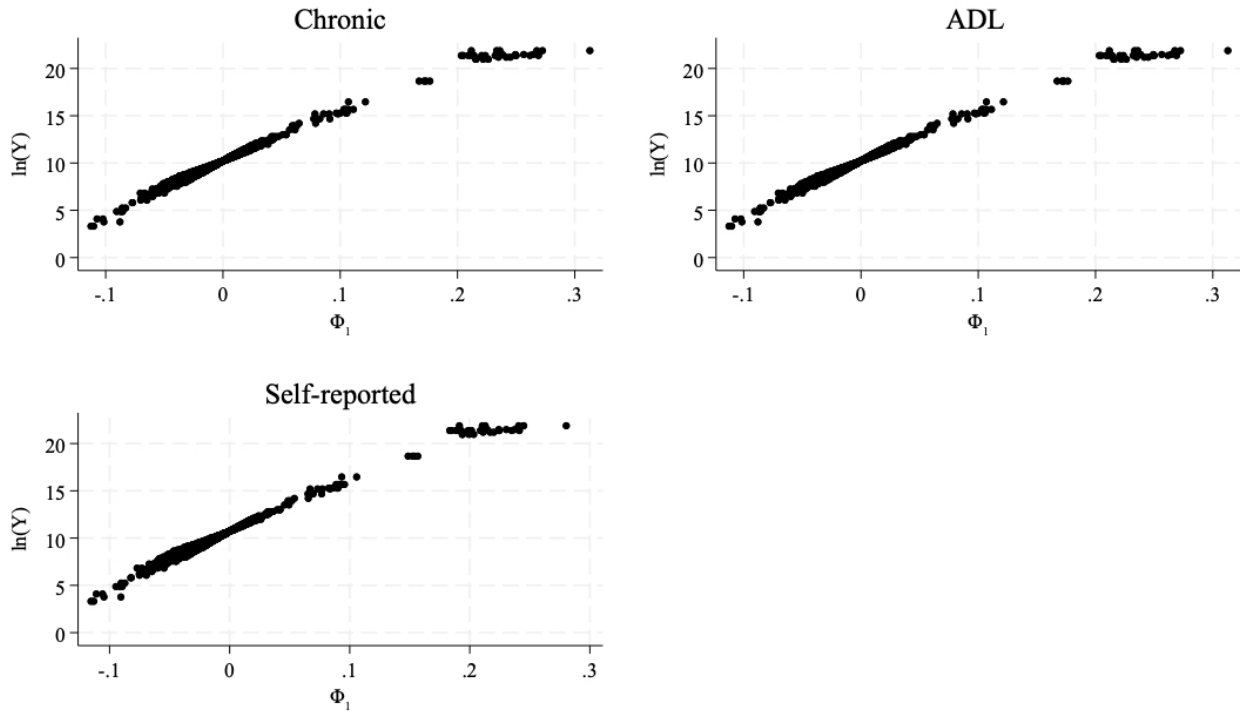
Table A.1: Estimation results with OOP expenditures

	Baseline			With OOP expenditures		
	(1) ADL	(2) Self-reported	(3) Chronic	(4) ADL	(5) Self-reported	(6) Chronic
A. Estimates:						
ln(Y): θ	1.293*** (0.043)	1.181*** (0.048)	0.938*** (0.087)	1.298*** (0.043)	1.185*** (0.048)	0.942*** (0.087)
Unhealthy: β	-1.178* (0.608)	-1.445*** (0.488)	-2.439*** (0.464)	-1.236** (0.609)	-1.477*** (0.489)	-2.473*** (0.464)
ln(Y)*Unhealthy: γ	0.117* (0.064)	0.136*** (0.051)	0.244*** (0.048)	0.124* (0.064)	0.139*** (0.051)	0.248*** (0.048)
Age	-0.017* (0.010)	-0.004 (0.010)	-0.024** (0.010)	-0.017* (0.010)	-0.005 (0.010)	-0.025** (0.010)
Female	-0.268*** (0.039)	-0.254*** (0.039)	-0.239*** (0.039)	-0.266*** (0.039)	-0.252*** (0.039)	-0.236*** (0.039)
Number of children	-0.051*** (0.013)	-0.056*** (0.013)	-0.053*** (0.013)	-0.051*** (0.013)	-0.056*** (0.013)	-0.053*** (0.013)
Has a partner in hh	0.349*** (0.091)	0.365*** (0.091)	0.359*** (0.091)	0.347*** (0.091)	0.363*** (0.091)	0.356*** (0.091)
Net wealth	0.004*** (0.001)	0.004*** (0.001)	0.004*** (0.001)	0.004*** (0.001)	0.004*** (0.001)	0.004*** (0.001)
OOP expenditures (in 1000s)	-	-	-	-0.014 (0.020)	-0.013 (0.19)	-0.014 (0.020)
<i>N</i>	30656	30656	30656	30656	30656	30656
B. Health state dependence:						
$\bar{\mu}$	1.041	1.119	1.103	1.038	1.117	1.101
95 % CI	[1.040, 1.042]	[1.117, 1.120]	[1.101, 1.105]	[1.036, 1.038]	[1.115, 1.118]	[1.098, 1.103]

Notes: Standard errors clustered at the individual level are reported in parentheses. All specifications contain country and wave fixed effects as well as Mundlak terms. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

A.4 Estimating the health-state dependence parameter $\hat{\phi}_1$

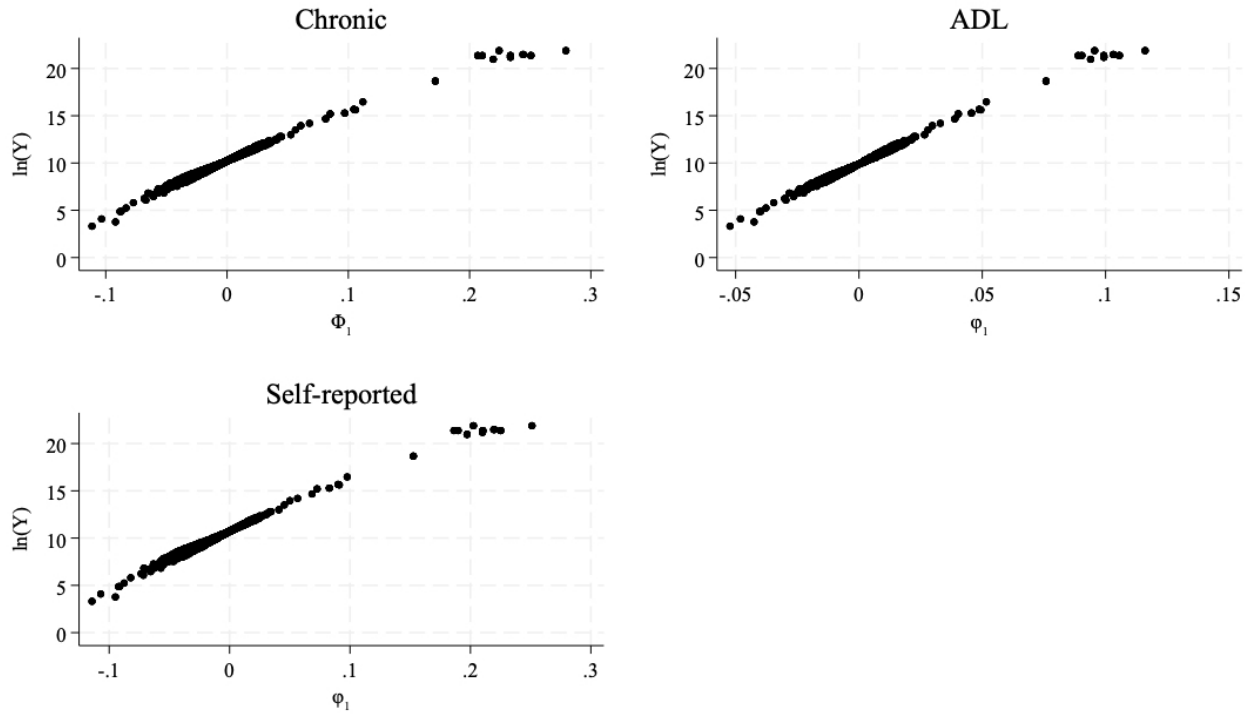
Figure A.4: $\hat{\phi}_1$ as a function of Y - for given (exogenous) values of α



Data: SHARE, wave 2, 4, 5, 6, 8.

Sample: 48,922 observations corresponding to 17,931 individuals aged 65+, living alone or with a 65+ partner.

Figure A.5: $\hat{\phi}_1$ as a function of Y - using permanent consumption



Data: SHARE, wave 2, 4, 5, 6, 8.

Sample: 48,922 observations corresponding to 17,931 individuals aged 65+, living alone or with a 65+ partner.

Number of healthy periods, with: ADL definition self-reported definition chronic condition definition	Number of observations for the individual			
	2	3	4	5
0 (always disabled)	7% <u>34%</u> 45%	5% <u>29%</u> 39%	3% <u>24%</u> 34%	1% <u>12%</u> 29%
1	17% <u>27%</u> 30%	8% <u>19%</u> 25%	5% <u>15%</u> 22%	2% <u>10%</u> 18%
2	(always healthy) 76% <u>38%</u> 25%	16% <u>20%</u> 20%	8% <u>13%</u> 16%	5% <u>13%</u> 16%
3		(always healthy) 70% <u>32%</u> 16%	18% <u>18%</u> 14%	9% <u>12%</u> 14%
4			(always healthy) 67% <u>30%</u> 14%	21% <u>18%</u> 11%
5				(always healthy) 61% <u>35%</u> 12%

Notes: Figures provided for the ADL definition are neither underlined nor in bold;
Figures provided for the self-reported health definition are underlined;
Figures provided with the **chronic condition definition** are in bold.