

Old Age Risks, Consumption, and Insurance[†]

By RICHARD BLUNDELL, MARGHERITA BORELLA,
JEANNE COMMAULT, AND MARIACRISTINA DE NARDI*

In the United States, after age 65, households face income and health risks, and a large fraction of these risks are transitory. While consumption significantly responds to transitory income shocks, out-of-pocket medical expenses do not. In contrast, both consumption and out-of-pocket medical expenses respond to transitory health shocks. Thus, most US elderly keep their out-of-pocket medical expenses close to a satiation point that varies with health. Consumption responds to health shocks mostly because adverse health shocks reduce the marginal utility of consumption. The effect of health on marginal utility changes the optimal transfers due to health shocks. (JEL D12, E21, G22, G51, I10, J14, J26)

What risks do households face, and to what extent are they insured by the government, themselves, their family, and their community? Previous work mainly studies the effects of income shocks on consumption for people of working age. Because people are living longer, studying older people is becoming more important, and because health shocks are prevalent at older ages, broadening the analysis to both income and health shocks is becoming essential.

For working-age people, the response of consumption to income shocks is typically used to measure the degree of insurance against these shocks. But when health shocks are another important source of risk, because they might affect both resources and ability to derive utility from consumption, the interpretation of consumption fluctuations as lack of insurance is no longer straightforward. Hence, to better understand the extent to which people are exposed to risks, it is necessary

*Blundell: University College London, IFS, TSE, and CEPR (email: r.blundell@ucl.ac.uk); Borella: University of Torino and CeRP-Collegio Carlo Alberto (email: margherita.borella@unito.it); Commault: Sciences Po (email: jeanne.commault@sciencespo.fr); De Nardi: University of Minnesota, Federal Reserve Bank of Minneapolis, CEPR, and NBER (email: denar001@umn.edu). Mikhail Golosov was the coeditor for this article. Blundell gratefully acknowledges financial support from the ESRC Centre for the Microeconomic Analysis of Public Policy (CPP) the IFS (ES/M010147/1). Borella gratefully acknowledges financial support from the project funded by Next Generation EU “Age-It—Ageing well in an ageing society” (PE0000015), National Recovery and Resilience Plan (NRRP)-PE8-Mission 4, C2, Intervention 1.3. Commault gratefully acknowledges financial support from the 2019 Banque de France-Sciences Po partnership fund. De Nardi gratefully acknowledges financial support from NSF grant SES-2044748. We are grateful to the coeditor, Marco Bassetto, Eric French, and three anonymous referees for helpful comments and suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the Banque de France, the CEPR, the National Bureau of Economic Research, the European Commission, the European Union, the Federal Reserve Bank of Minneapolis, or the Federal Reserve System.

[†]Go to <https://doi.org/10.1257/aer.20220555> to visit the article page for additional materials and author disclosure statements.

both to measure income and health risks and to disentangle the effects of health shocks on resources and marginal utility of consumption.

We develop a semistructural approach, new identification techniques, and use high-quality data to measure the effects of income and health shocks on consumption among US households over age 65. We also propose a novel methodology to decompose the consumption response to a transitory health shock into its effect on resources and on the marginal utility of consumption. More specifically, we estimate income and health risks and the pass-through of transitory risks to consumption and out-of-pocket medical expenses. We do so by using a flexible specification for the policy functions determining consumption and medical expenses. We also use a rich structural model to derive novel implications that allow us to disentangle the effects of transitory health shocks on resources and medical expenses.

In terms of data, we use the Health and Retirement Study (HRS) (2014) and its Consumption and Activities Mail Survey (CAMS) (HRS 2001–2015a). The HRS is a longitudinal panel study that, starting in 1994, is conducted every other year and is representative of the US population over the age of 50 and their spouses. Since 2001, CAMS collects detailed data on nondurable consumption and out-of-pocket medical expense subcategories. Hence, it allows us to analyze both consumption and out-of-pocket medical expenses and their subcategories. We complement them with the Rand HRS Fat files (HRS 2000–2014) and the Rand HRS tax calculations (HRS 2001–2015b).

Our analysis yields several important and novel findings. First, after age 65, households are subject to significant temporary fluctuations in both income and health. In terms of magnitudes, the variance of the current transitory component of income explains 41 percent of the variance of changes in income, and the variance of the current transitory component of health explains 31 percent of the variance of changes in health (after we detrend income and health from the effect of observed demographic characteristics). The bulk of these shocks cannot be attributed to measurement error for two reasons: the HRS has been documented to be of excellent quality,¹ and we find that these transitory shocks have a significant impact on households' decision variables.

Second, transitory income shocks have sizable and statistically significant effects on nondurable consumption. In contrast, they have statistically insignificant effects on out-of-pocket medical expenses.

More specifically, our estimated average pass-through of transitory income shocks implies that a 10 percent increase in transitory income is associated with a 1.3 percent increase in current nondurable consumption. This magnitude is comparable to the results obtained using working-age households. Among the lower-wealth households (which we define as those in the bottom quintile of the wealth distribution), the effect is twice as large.

Turning to the effects of income shocks on out-of-pocket medical expenses, because these expenses make up for a small fraction of total expenses, our estimated

¹Hurd and Rohwedder (2009) discuss the CAMS data quality and show that spending totals are close to those measured in the Consumer Expenditures Survey (CEX), and the age profiles of wealth changes implied by spending and after-tax income are similar to the wealth change in the HRS data. French, Jones, and McCauley (2017) find that the HRS data are of high quality.

pass-through coefficient implies that the level of out-of-pocket medical expenses fluctuates little with transitory income shocks. The same finding applies for the lower-wealth households. This small response suggests that, for the existing level of insurance, most US elderly are satiated in their consumption of out-of-pocket medical expenses.

Third, transitory health shocks affect both nondurable consumption and out-of-pocket medical expenses. Our estimated average pass-through of transitory health shocks to consumption implies that a 1 standard deviation transitory decrease in one's health index is associated with a 2.4 percent decrease in nondurable consumption. This effect is larger among lower-wealth households, for whom the same decrease in health implies a 5.6 percent decrease in consumption. Our estimated average pass-through of transitory health shocks to out-of-pocket medical expenses implies that a 1 standard deviation decrease in one's health index translates into a 7 percent increase in medical expenses. The corresponding number for lower-wealth households is 21.3 percent. These findings indicate that people's satiation point for medical goods and services varies with their health.

Fourth, in our overall sample, 98.3 percent of the response of consumption to transitory health shocks is due to the fact that health shocks change the marginal utility of consumption, while only 1.7 percent is due to its effect on resources (through a change in medical expenses). For lower-wealth households, these numbers are 94.1 percent and 5.9 percent, respectively. Both effects are significant for both samples. For lower-wealth households, the resource effect is larger because, when they are hit by a negative health shock, their out-of-pocket medical expenses increase more and because, for them, a given decrease in resources leads to a larger decrease in consumption.

To better understand the mechanisms behind our findings, we also examine the responses of various consumption subcategories. Here, we find that the subcategories that respond to a transitory income shock are different from those that respond to a transitory health shock. More specifically, in our overall sample, necessities (food, utilities, car-related expenses) and luxuries (leisure activities, equipment) both respond to an income shock. In contrast, only luxuries respond to a health shock. This is consistent with our finding that the shift in marginal utility plays an important role in the response of consumption to a transitory health shock because if the response of consumption to a health shock were caused by its resources effect, a health shock should affect the same consumption subcategories as an income shock—since a health shock is equivalent to a loss in resources.

We also use these subcategories to estimate a demand system. This methodology evaluates how consumption and medical expenses react to total health and income changes (which thus include both transitory and permanent shocks), while holding spending constant. Our estimated demand system yields two main findings. First, it reveals that a change in health generates a reallocation across consumption goods subcategories. That is, even absent any variation in resources, sick people do not consume the same goods as healthy people. This indicates that people's marginal utility of certain goods changes with their health (else, holding resources constant, a change in health would have no effect on the allocation of consumption across subcategories). It therefore confirms our finding that the marginal utility channel is important to understand the consumption response to health shocks. Second, it

highlights that even the sum of permanent and transitory health shocks affects the marginal utility of consumption.

Finally, after having shown that transitory income shocks only affect resources and that transitory health shocks mostly impact the marginal utility of consumption, we examine their qualitative implications in terms of optimal insurance. While shocks to resources result in fewer resources in some states than in others, shocks to marginal utility generate a mismatch between resources and the ability to take advantage of resources across states. That is, people might end up with a relatively large amount of resources while they are in bad health and have low marginal utility of consumption, or low resources while they are in good health and have high marginal utility of consumption. While a benevolent planner insuring households against income shocks allocates the same level of consumption and medical expenses to all, a planner insuring households against marginal utility shocks allocates more medical expenses to households experiencing an increase in their marginal utility of medical expenses and more consumption to households who do not experience a reduction in their marginal utility of consumption.

In sum, our main contribution is showing that older households face substantial transitory income and health risks, that they react to these risks, and that transitory health shocks have important effects on households' marginal utility of consumption. Our contribution has implications for both the positive and normative literature on households' savings and insurance. That is, positive models should include transitory income and health risks and should imply responses to health shocks that are consistent with our findings. Normative analysis should include health shocks and account for their effects on the marginal utility of consumption.

Our paper relates to the literature studying the impact of a specific one-time (and hence transitory) resource shock on consumption and finding that transitory shocks such as tax rebates, lottery gains, or changes in current assets significantly affect consumption.² It also relates to the literature on consumption insurance, the literature on savings and risks during retirement, and the literature testing whether the utility from consumption depends on health.³ We contribute to these branches of the literature by focusing on the retirement period, by showing that both transitory income and health shocks are large, and that adverse transitory health shocks reduce the marginal utility of consumption in a quantitatively important way. Our findings thus suggest that, even though households over age 65 are covered by Medicare, the income and health risks that they face during old age are large, and there might be scope to rethink the current insurance scheme to take into account the effects of health on the marginal utility of consumption.

²See, for instance, Mian, Rao, and Sufi (2013); Fagereng, Holm, and Natvik (2021); and Cloyne et al. (2019).

³For some important contributions on consumption insurance, see Cochrane (1991); Attanasio and Davis (1996); Blundell, Pistaferri, and Preston (2008); Heathcote, Storesletten, and Violante (2009); Kaplan and Violante (2010); Blundell, Low, and Preston (2013); Farhi and Werning (2013); and Golosov, Troshkin, and Tsyvinski (2016). Important works on savings and risks, including during retirement, include Hubbard, Skinner, and Zeldes (1994, 1995); Palumbo (1999); Brown and Finkelstein (2008); Love, Palumbo, and Smith (2009); De Nardi, French, and Jones (2010, 2016); Blundell et al. (2016); Braun, Kopecky, and Koreshkova (2017); De Nardi, Pashchenko, and Porapakarm (2017); Poterba, Venti, and Wise (2018); and Ameriks et al. (2020). For testing for health-dependence in utility, see, for instance, Viscusi and Evans (1990); Evans and Viscusi (1991); Finkelstein, Luttmer, and Notowidigdo (2009, 2013).

I. The Model

We are able to measure the effects of transitory income and health shocks on consumption and medical expenses by making minimal assumptions. In contrast, disentangling the sources of these effects necessitates a structural model.

To make the logic of our analysis more cohesive, in this section we develop a structural framework for our analysis. In the next section, we examine the implications for the consumption and medical expenses responses to transitory shocks. We then turn to identification and explain which parameters of our analysis can be recovered with fewer assumptions than in the full structural model outlined here.

Our structural model is fairly general and embeds the majority of models used in the structural literature on health risks and savings. We generalize previous work by allowing for two important features. First, we allow for both transitory income and health shocks, and second, we let a household’s utility flexibly depend on consumption, total medical expenses, and health status.

$$(1) \quad \max_{\{c_t, m_t\}_{t=0}^T} E_0 \sum_{t=0}^T \beta^t \left\{ s_t(\{\pi^h\}_t) \left[u(c_t, \tilde{m}(m_t), h_t) \right] \right\}$$

subject to

$$(2) \quad p_{t+1} a_{t+1} = (1 + r_t) p_t a_t + p_t y_t - p_t^m m_t - p_t^c c_t, \quad \forall 0 \leq t \leq T,$$

$$(3) \quad a_T \geq 0,$$

$$(4) \quad \ln(y_t) = \pi_t^y + \varepsilon_t^y, \quad \pi_t^y = \pi_{t-1}^y + \eta_t^y,$$

$$(5) \quad h_t = \pi_t^h + \varepsilon_t^h, \quad \pi_t^h = \pi_{t-1}^h + \eta_t^h,$$

$$\text{cov}(\eta_t^y, \eta_t^h) \neq 0, \quad \text{cov}(\varepsilon_t^y, \varepsilon_t^h) = 0,$$

$$\text{cov}(\eta_t^y, \eta_{t'}^y) = 0, \quad \text{cov}(\varepsilon_t^y, \varepsilon_{t'}^y) = 0, \quad \text{cov}(\eta_t^h, \eta_{t'}^h) = 0, \quad \text{cov}(\varepsilon_t^h, \varepsilon_{t'}^h) = 0.$$

Starting from age 66 (which we renormalize as period 0) and until age T , a household chooses its consumption, c_t , and out-of-pocket medical expenses, m_t , to maximize its expected utility.

Health affects both one’s survival probability and marginal utility of consumption. The term $s_t(\{\pi^h\}_t) = \prod_{l=0}^t \tilde{s}_l(\pi_l^h)$ denotes the cumulative survival probability of a household at age t , conditional on being alive at age 66. It is a function of the history of the permanent health component $\{\pi^h\}_t$. The rationale for excluding transitory shocks from it is that people recover fully from transitory shocks after at most two years. Hence, they should not be affecting their survival probability.

The within-period utility function, u , is a function of consumption c_t , total medical expenses $\tilde{m}(m_t)$ (which relate to out-of-pocket medical expenses m_t through the function $\tilde{m}(\cdot)$), and health h_t during that period. The utility function can be nonseparable in its arguments, and the expected value of future utility is taken with respect to uncertain income and health. Within-period utility is twice differentiable in its arguments and strictly increasing and concave in its first argument: $u_c(\cdot, \tilde{m}, h)$

> 0 and $u_{cc}(\cdot, \tilde{m}, h) < 0$. Utility is time additive, and β is the discount factor. We drop demographics from our model's exposition to simplify notation, but our framework allows for utility to be influenced by demographic characteristics. Our empirical strategy accounts for demographics.

The timing is the following. At the beginning of each period, income and health shocks are realized, income is received, and households optimally choose consumption, medical expenses, and savings. At the end of the period, mortality risk is realized.

Our maximization problem is subject to four constraints. Equation (2) is the budget constraint. The household can use an asset, a_t , to store its wealth from one period to the next at a possibly stochastic rate of return r_t . During each period, the household receives stochastic income $p_t y_t$, and spends $p_t^m m_t$ and $p_t^c c_t$ on out-of-pocket medical expenses and on consumption, where p_t is the price index for output, p_t^m the one for medical expenses, and p_t^c the one for consumption. The terminal condition on wealth states that households cannot hold negative assets during their last period, when they die with certainty.

Equations (4) and (5) govern the evolution of the log income and health (net of the effect of demographics, which we purge in our empirical strategy), which are the sum of a permanent component π that evolves as a random walk, and of a transitory component ε that is an MA(0) process. The shocks are not required to be drawn from normal distributions. They are centered around zero. Hence, many households might receive small positive shocks, while a few might be hit by large negative shocks. In addition, different households might draw shocks from different distributions, and the same household might draw shocks from different distributions over time.⁴ A positive health shock is not necessarily an absolute increase in health. Rather, it is a health increase relative to what health would have been absent the shock (for instance, health deteriorates less than demographics would predict since these processes model the evolution of health and income net of demographics—to simplify notation, we do not make explicit the effect of demographics in this section).

For each household, we allow the permanent health and income shocks η^y and η^h to be correlated within a period. We can also let the contemporaneous transitory income and health shocks ε^y and ε^h be correlated with each other. Yet, because we estimate that correlation to be small and not statistically significant (see Table 3), we set it to zero in our main analysis. Online Appendix E relaxes this assumption and shows that our results are unaffected by it. We also impose that none of the shocks are serially correlated.

Our transitory-permanent specifications are consistent with the observed autocovariances of log income growth and health growth. In contrast, several alternative statistical models are ruled out: the observed autocovariances reject that the transitory components are more persistent than an MA(0) in biennial data and suggest that the permanent component evolves as a random walk rather than an AR(1). We detail this in Section IVB, Table 1.⁵

Our model assumes that medical expenses during retirement generate utility during the current period and are endogenously chosen by the household but that

⁴This point is detailed in Commault (2022) in Section I.A, footnote 5.

⁵We also conduct robustness checks in which our income process is an AR(1). Online Appendix J shows that the results are very similar.

medical expenses do not affect one's future health, which evolves exogenously. The papers on health and medical expenses that we mention earlier in this section make either more restrictive or similar assumptions.

Importantly, our modeling of health is also consistent with much empirical evidence showing that the effects of medical expenses on health and mortality are small for US retirees. Two reasons can explain this finding. First, the medical expenses that we are considering are supplementing Medicaid, Medicare, and insurance-provided medical goods and services, which cover most life-threatening conditions. Second, the stock of health carried by an older person is in large part determined by health investments that were made in the past, including those made by the person's parents during their childhood, and even before birth. Hence, the effects of additional health investments for people aged 66 and older are not as large as in earlier stages of life.

In terms of empirical evidence, many papers find that medical expenses have small effects on health and mortality. To start with, in the RAND Health Insurance Experiment, a random set of individuals were given co-payment-free health insurance over a three- to five-year period, while a control group faced standard co-payments. Brook et al. (1983) found that even though the group with free health care utilized medical services much more intensively than the control group, the additional treatments had only minor effects on subsequent health outcomes. Moreover, some empirical studies show that even programs such as Medicare, which sometimes help pay for critical treatments, do not significantly increase life expectancy (Fisher et al. 2003; Finkelstein and McKnight 2008). Card, Dobkin, and Maestas (2009) find that Medicare caused a small reduction in mortality among 65-year-olds admitted through emergency rooms for what they refer to as "nondeferrable" conditions. Using a different method that compares uninsured individuals between age 50 and 61 with matched uninsured individuals, Black et al. (2017) show that the uninsured consume fewer health care services but that their health (while alive) does not deteriorate relative to that of the insured and that their mortality is similar. Khwaja (2010) estimates a structural model in which medical expenses both provide utility and improve health and finds that 80 percent of medical utilization is mitigative, in the sense of just increasing current utility, while the remaining 20 percent is curative, in the sense that it does improve one's health. Blau and Gilleskie (2008) reach similar conclusions. Given that the existing evidence indicates that the effect of additional medical spending on subsequent health and life expectancy is small, and that we study older people, we focus on the utility effects of medical expenses.

II. The Consumption and Medical Expenses Responses to Shocks

The goal of this section is to formalize the intuition of the channels through which permanent and transitory shocks affect decisions and to use our structural model to decompose the consumption and medical expenses responses to transitory shocks into quantities that we can estimate.

A. The Transmission Channels

To start, note that the policy functions (and their partial derivatives) are informative about the total effects of income and health shocks on consumption and medical

expenses but say little about the channels at play. Indeed, the problem described by equations (1)–(5) implies the following policy functions for c_t and m_t :

$$(6) \quad c_t = c_t(a_t, \pi_t^y, \pi_t^h, \varepsilon_t^y, \varepsilon_t^h),$$

$$(7) \quad m_t = m_t(a_t, \pi_t^y, \pi_t^h, \varepsilon_t^y, \varepsilon_t^h).$$

In these expressions, the partial derivatives with respect to ε_t^y and ε_t^h capture a combination of channels. To make them explicit, we start from the Euler equation. It relates the marginal utility of current consumption (which depends on current consumption, current medical expenses, and current health) to the expected marginal utility of future consumption (which depends on future consumption, future medical expenses, and future health), weighted by the future survival probability

$$(8) \quad u_c(c_t, \tilde{m}(m_t), h_t) \geq E_t[u_c(c_{t+1}, \tilde{m}(m_{t+1}), h_{t+1})\tilde{s}_{t+1}(\pi_{t+1}^h)R_{t+1}],$$

where $R_{t+1} \equiv \beta(1 + r_{t+1})$ is a factor capturing all intertemporal substitution motives other than the survival probability.

We substitute c_{t+1} and m_{t+1} in equation (8) using the policy functions (6) and (7) and use $a_{t+1} = [(1 + r_t)p_t a_t + p_t y_t - p_t^m m_t - p_t^c c_t]/p_{t+1}$, $\pi_{t+1}^y = \pi_t^y + \eta_{t+1}^y$, and $\pi_{t+1}^h = \pi_t^h + \eta_{t+1}^h$

$$(9) \quad \underbrace{u_c(c_t, \tilde{m}(m_t), h_t)}_{(i)} \geq E_t[u_c(c_{t+1}, \underbrace{((1 + r_t)p_t a_t + p_t y_t - p_t^m m_t - p_t^c c_t)}_{(ii)})/p_{t+1},$$

$$\underbrace{\pi_t^y + \eta_{t+1}^y, \varepsilon_{t+1}^y}_{(iii)}, \underbrace{\pi_t^h + \eta_{t+1}^h, \varepsilon_{t+1}^h}_{(iii)}),$$

$$\tilde{m}(m_{t+1}, \underbrace{((1 + r_t)p_t a_t + p_t y_t - p_t^m m_t - p_t^c c_t)}_{(ii)})/p_{t+1},$$

$$\underbrace{\pi_t^y + \eta_{t+1}^y, \varepsilon_{t+1}^y, \pi_t^h + \eta_{t+1}^h, \varepsilon_{t+1}^h}_{(iii)}),$$

$$\underbrace{\pi_t^h + \eta_{t+1}^h + \varepsilon_{t+1}^h}_{(iii)}\tilde{s}_{t+1}(\underbrace{\pi_t^h + \eta_{t+1}^h}_{(iv)})R_{t+1}.]$$

The resulting expression is an optimality condition relating current consumption c_t to m_t and to the state variables at t ($a_t, h_t = \pi_t^h + \varepsilon_t^h, y_t = \pi_t^y + \varepsilon_t^y$) in a way that makes explicit the following effects at play:

- (i) Out-of-pocket medical expenses and health change the marginal utility of current consumption (in equation (9), m_t and h_t affect $u_c(c_t, \tilde{m}(m_t), h_t)$);

- (ii) Assets, income, and medical expenses determine the resources that will remain after consumption at the current period, which affects consumption at the next period—thus, the value of current consumption that equalizes the current and expected future marginal utilities—(in the expression above, a_t , y_t , and m_t enter the value of assets at the beginning of $t + 1$, which is $(1 + r_t)p_t a_t + p_t y_t - p_t^m m_t$);
- (iii) Independently of the resources passed on to next period, the current permanent components of income and health influence the value of income and health at the next period, thus, consumption at the next period (in the expression above, π_t^y appears in the expression of π_{t+1}^y , which is $\pi_t^y + \eta_{t+1}^y$, and π_t^h appears in the expression of π_{t+1}^h , which is $\pi_t^h + \eta_{t+1}^h$); and
- (iv) Independently of the resources passed on to next period and of the distribution of income and health at the next period, the current permanent component of health determines the next survival probability (in the expression above, π_t^h appears in the survival probability $\tilde{s}_{t+1}(\pi_t^h + \eta_{t+1}^h)$).

While, for simplicity of exposition, we abstract from specifying borrowing constraints in our model, and hence, the Euler equation always holds at equality, we write it here as an inequality for generality. In the case in which there are binding borrowing constraints, the Euler equation holds as an inequality because the Lagrange multiplier on resources is positive. The value of this multiplier depends on whether future resources (our second channel) are below a certain threshold. Hence, currently binding borrowing constraints play a similar role to that of future resources and manifest through one of the four main channels in (10).

Because this optimality condition implicitly defines consumption (and therefore log consumption) as a function of these four channels, it is convenient to write it as

$$(10) \quad \ln(c_t) = f^{c,t} \left(\underbrace{m_t, h_t}_{\substack{\text{affect } c_t \text{ through} \\ \text{marginal utility} \\ u_c(\cdot, \tilde{m}(m_t), h_t)}}, \underbrace{(1 + r_t)p_t a_t + p_t y_t - p_t^m m_t}_{= R_t \text{ affect } c_t \text{ through} \\ \text{the budget constraint}}, \underbrace{\pi_t^y, \pi_t^h}_{\substack{\text{affect } c_t \text{ through} \\ \text{the distribution of} \\ y_{t+1} \text{ and } h_{t+1} \text{ (holding} \\ \tilde{s}_{t+1}(\pi_{t+1}^h) \text{ constant)}}}, \underbrace{\pi_t^h}_{\substack{\text{affects } c_t \\ \text{through survival} \\ \text{probability} \\ \tilde{s}_{t+1}(\pi_{t+1}^h)}} \right).$$

A similar expression holds for log medical expenses, with $f^{m,t}$ the function determining $\ln(m_t)$. In these expressions, the partial derivatives of $f^{c,t}$ and $f^{m,t}$ with respect to each argument correspond to the effect of each channel holding the rest constant (e.g., the partial derivative of $f^{c,t}$ with respect to the first argument corresponds to the effect of a change in medical expenses through the marginal utility channel only, holding the budget constraint, the distributions of future income and health, and the next survival probability constant).

B. Decomposing the Response to Transitory Shocks

Because transitory shocks have no effects on the future distribution of income and health, nor on people's survival probability, they only influence consumption and medical spending through two channels: the marginal utility channel and the resources channel. We now turn to studying the case of transitory shocks and the factors determining the magnitude of these two channels.

When we take the derivative of equation (10) with respect to transitory income and health shocks, only the first three derivatives of the function $f^{c,t}$ appear. We denote them as $f_m^{c,t}, f_h^{c,t}$ (which both correspond to the marginal utility channel) and $f_R^{c,t}$ (which correspond to the resources channel).⁶ Hence, the partial derivatives with respect to transitory shocks are

$$(11) \quad \frac{d\ln(c_t)}{d\varepsilon_t^y} = \underbrace{f_m^{c,t} \frac{dm_t}{d\varepsilon_t^y}}_{\text{Marginal utility}} + \underbrace{f_R^{c,t} \left(p_t \frac{dy_t}{d\varepsilon_t^y} - p_t^m \frac{dm_t}{d\varepsilon_t^y} \right)}_{\text{Resources}},$$

$$(12) \quad \frac{d\ln(c_t)}{d\varepsilon_t^h} = \underbrace{f_m^{c,t} \frac{dm_t}{d\varepsilon_t^h} + f_h^{c,t} \frac{dh_t}{d\varepsilon_t^h}}_{\text{Marginal utility}} - \underbrace{f_R^{c,t} p_t^m \frac{dm_t}{d\varepsilon_t^h}}_{\text{Resources}},$$

where we have used the lack of correlation between the transitory shocks to set $(dh_t)/(d\varepsilon_t^y) = 0$ (and hence, the term containing $f_h^{c,t}$ drops out) and also $(dy_t)/(d\varepsilon_t^h) = 0$ (and hence, available resources to consume and save only change in response to a health shock because of the change in medical expenses).

Because our income and health processes imply $(dy_t)/(d\varepsilon_t^y) = [d\ln(y_t)]/(d\varepsilon_t^y) \times y_t = 1 \times y_t$ and $(dh_t)/(d\varepsilon_t^h) = 1$, and we have that $(dm_t)/(d\varepsilon_t^h) = [d\ln(m_t)]/(d\varepsilon_t^h) m_t$, we can simplify (11) and (12) as

$$(13) \quad \frac{d\ln(c_t)}{d\varepsilon_t^y} = \underbrace{f_m^{c,t} \frac{dm_t}{d\varepsilon_t^y}}_{\text{Marginal utility}} + \underbrace{f_R^{c,t} \left(p_t y_t - p_t^m \frac{dm_t}{d\varepsilon_t^y} \right)}_{\text{Resources}},$$

$$(14) \quad \frac{d\ln(c_t)}{d\varepsilon_t^h} = \underbrace{f_m^{c,t} \frac{dm_t}{d\varepsilon_t^h} + f_h^{c,t}}_{\text{Marginal utility}} - \underbrace{f_R^{c,t} p_t^m \frac{d\ln(m_t)}{d\varepsilon_t^h} m_t}_{\text{Resources}}.$$

Finally, we assume that, after age 65, people do not adjust their level of out-of-pocket medical expenses when experiencing transitory income changes (that is, $(dm_t)/(d\varepsilon_t^y) \approx 0$). Two points are important here. To start, not only the response of out-of-pocket medical expenses is statistically insignificant, but, because out-of-pocket medical expenses are low, they fluctuate little with income, given our estimated pass-through coefficient. In addition, in what follows, we discuss the consequences of relaxing this assumption.

⁶To ease notation, we denote $f_m^{c,t}(m_t, h_t, (1+r_t)p_t a_t + p_t y_t - p_t^m m_t, \pi_t^y, \pi_t^h, \pi_t^h)$ as $f_m^{c,t}$, and similarly for $f_h^{c,t}$ and $f_R^{c,t}$.

We obtain

$$(15) \quad \frac{d\ln(c_t)}{d\varepsilon_t^y} = \underbrace{\widehat{f_R^{c,t}}}_{\text{Resources}} p_t y_t,$$

$$(16) \quad \frac{d\ln(c_t)}{d\varepsilon_t^h} = \underbrace{f_m^{c,t} \frac{dm_t}{d\varepsilon_t^h}}_{\text{Marginal utility}} + f_h^{c,t} - \underbrace{\widehat{f_R^{c,t}} p_t^m \frac{d\ln(m_t)}{d\varepsilon_t^h}}_{\text{Resources}} m_t.$$

Hence, the elasticity of consumption to a transitory income shock only depends on the strength of the resource channel. It is the product of the change in resources caused by the shock $p_t(dy_t)/(d\varepsilon_t^y) = p_t y_t$ (by construction a 1-unit transitory income shock corresponds to 100 percent income change) and the multiplier $f_R^{c,t}$ —which measures how much the pass-through of a shock to consumption increases when the shock raises resources by 1 unit. In contrast, the elasticity of consumption to a transitory health shock depends on both the marginal utility channel and the resources channel. The latter is, again, the product of the shock-induced resources change $(p_t^m(dm_t))/(d\varepsilon_t^h) = p_t^m [d\ln(m_t)]/(d\varepsilon_t^h)m_t$ and the multiplier $f_R^{c,t} = 0$.

Online Appendix D relaxes the assumption that income shocks do not affect medical expenses. It shows that if medical expenses do respond to transitory income shocks, ignoring this effect leads to underestimating the share of the consumption response that is due to a shift in marginal utility. Hence, our estimate of the effects of health on the marginal utility of consumption, which is an important contribution of our paper, is conservative in this regard.

It is worth noting that an alternative condition that yields a similar expression, but does not require imposing that income shocks do not affect medical expenses to obtain it, is assuming that $f_m^{c,t} = 0$. That is, that, conditional on health, the marginal utility of consumption is unaffected by medical expenses. Indeed, the marginal utility channel of the transitory income pass-through $f_m^{c,t}(dm_t)/(d\varepsilon_t^y)$ is zero either when $(dm_t)/(d\varepsilon_t^y) = 0$ (which we estimate to be the case) or when $f_m^{c,t} = 0$. Note that under the alternative case of separability, the interpretation of the marginal utility channel is different. Instead of capturing both the effects of health on marginal utility and the (potentially counterbalancing) effect of the response of medical expenses on marginal utility $f_m^{c,t}(dm_t)/(d\varepsilon_t^h) + f_h^{c,t}$, it only relates to the effect of health on marginal utility $f_h^{c,t}$ because $f_m^{c,t} = 0$.

For completeness, online Appendix L provides the exact mapping between the marginal utility and resources channels expressed in terms of partial effects on consumption $f_m^{c,t}$, $f_h^{c,t}$, and $f_R^{c,t}$ (that we present here) and expressed in terms of partial effects on the marginal utility of consumption (u_{cc}, u_{cm}, u_{ch}) .

III. Identification and Implementation

We now turn to discussing how we identify the partial derivatives with respect to transitory shocks and, within them, the resources and marginal utility channels that compose them, and how we implement this identification strategy.

A. Identification

Variance of the Income and Health Shocks.—Transitory shocks are not directly observed in our data, which report income and health. Assuming a transitory-permanent specification allows us to identify the variances and covariances of the transitory and permanent components of income and health by using moment conditions. For transitory shocks, as Meghir and Pistaferri (2004) and Blundell, Pistaferri, and Preston (2008)—henceforth, BPP—we use equations (4) and (5) to derive $\Delta \ln(y)$ and Δh and obtain moment conditions that we can estimate:

$$(17) \quad \text{cov}(\Delta \ln(y_t), -\Delta \ln(y_{t+1})) = \text{var}(\varepsilon_t^y),$$

$$(18) \quad \text{cov}(\Delta h_t, -\Delta h_{t+1}) = \text{var}(\varepsilon_t^h),$$

$$(19) \quad \text{cov}(\Delta h_t, -\Delta \ln(y_{t+1})) = \text{cov}(\varepsilon_t^y, \varepsilon_t^h),$$

$$(20) \quad \text{cov}(\Delta \ln(y_{t+1}), -\Delta h_t) = \text{cov}(\varepsilon_t^y, \varepsilon_t^h).$$

Intuitively, these moments identify the variances of transitory shocks because future growth (at $t + 1$) filters out the permanent component of current growth (at t): a current transitory shock generates current positive growth and future negative growth, while a current permanent shock generates current positive growth and no growth afterward. This identification only requires that income and health evolve as in equations (4) and (5).

We also estimate the variances and covariances of *permanent shocks* for the purpose of comparing them with those of transitory shocks. For this, as Meghir and Pistaferri (2004) and BPP, we use that $\text{cov}(\Delta \ln(y_t), \Delta \ln(y_{t-1}) + \Delta \ln(y_t) + \Delta \ln(y_{t+1})) = \text{var}(\eta_t^y)$ and $\text{cov}(\Delta h_t, \Delta h_{t-1} + \Delta h_t + \Delta h_{t+1}) = \text{var}(\eta_t^h)$.

Pass-Through Coefficients.—To identify the partial derivatives with respect to transitory shocks, $[d \ln(c_t)] / (d\varepsilon)$ —which we refer to as the pass-through coefficients—we linearize log consumption. It makes sense to do so because the pass-through captures the total effect of a shock.

Because the consumption policy function states that log consumption is determined by a household's assets, permanent and transitory income, and permanent and transitory health, we linearize it around the point where all these variables are at their average sample values (we denote with $|_0$ any variable taken at this approximation point):

$$(21) \quad \ln(c_t) \approx \ln(c_t)|_0 + (a_t - E[a_t]) \frac{d \ln(c_t)}{da_t} \Big|_0 + (\pi_t^y - E[\pi_t^y]) \frac{d \ln(c_t)}{d\pi^y} \Big|_0 \\ + (\pi_t^h - E[\pi_t^h]) \frac{d \ln(c_t)}{d\pi_t^h} \Big|_0 + (\varepsilon_t^y - E[\varepsilon_t^y]) \frac{d \ln(c_t)}{d\varepsilon_t^y} \Big|_0 \\ + (\varepsilon_t^h - E[\varepsilon_t^h]) \frac{d \ln(c_t)}{d\varepsilon_t^h} \Big|_0.$$

We then take the covariance of both sides of (21) with ε_t^y (or ε_t^h) and divide both sides by $\text{var}(\varepsilon_t^y)$ (or $\text{var}(\varepsilon_t^h)$). The linearization implies that the ratios $[\text{cov}(\ln(c_t), \varepsilon_t^y)] / [\text{var}(\varepsilon_t^y)]$ and $[\text{cov}(\ln(c_t), \varepsilon_t^h)] / [\text{var}(\varepsilon_t^h)]$ coincide with the pass-through coefficients of the transitory income and health shocks at the approximation point $[d\ln(c_t)] / (d\varepsilon_t^y)|_0$ and $[d\ln(c_t)] / (d\varepsilon_t^h)|_0$.⁷ Now, the covariance between log consumption and transitory shocks is the same as the covariance between log consumption *growth* and transitory shocks: this is because shocks at t are true shocks and thus, orthogonal to variables at $t - 1$. We denote these key ratios $[\text{cov}(\Delta \ln(c_t), \varepsilon_t^y)] / [\text{var}(\varepsilon_t^y)]$ and $[\text{cov}(\Delta \ln(c_t), \varepsilon_t^h)] / [\text{var}(\varepsilon_t^h)]$ as ϕ_c^y and ϕ_c^h . As a result, we have

$$(22) \quad \phi_c^y = \frac{\text{cov}(\Delta \ln(c_t), \varepsilon_t^y)}{\text{var}(\varepsilon_t^y)} = \frac{\text{cov}(\ln(c_t), \varepsilon_t^y)}{\text{var}(\varepsilon_t^y)} \approx \left. \frac{d\ln(c_t)}{d\varepsilon_t^y} \right|_0,$$

$$(23) \quad \phi_c^h = \frac{\text{cov}(\Delta \ln(c_t), \varepsilon_t^h)}{\text{var}(\varepsilon_t^h)} = \frac{\text{cov}(\ln(c_t), \varepsilon_t^h)}{\text{var}(\varepsilon_t^h)} \approx \left. \frac{d\ln(c_t)}{d\varepsilon_t^h} \right|_0.$$

Similarly, the ratios of the covariance between growth in medical expenses and the transitory shocks over the variance of the shocks, denoted ϕ_m^y and ϕ_m^h , coincide with the pass-through of transitory shocks to medical expenses at the approximation point.

As we already identify the variances of our transitory shocks from moments (17)–(18), we now only need to identify some covariances to obtain these ratios. To do so, we use the same insight as in the identification of the variances: we filter out the effect of permanent shocks on current log consumption growth by taking its covariance with future income growth:

$$(24) \quad \text{cov}(\Delta \ln(c_t), -\Delta \ln(y_{t+1})) = \text{cov}(\Delta \ln(c_t), \varepsilon_t^y) = \phi_c^y \text{var}(\varepsilon_t^y),$$

$$(25) \quad \text{cov}(\Delta \ln(c_t), -\Delta \ln(h_{t+1})) = \text{cov}(\Delta \ln(c_t), \varepsilon_t^h) = \phi_c^h \text{var}(\varepsilon_t^h).$$

Similarly,

$$(26) \quad \text{cov}(\Delta \ln(m_t), -\Delta \ln(y_{t+1})) = \text{cov}(\Delta \ln(m_t), \varepsilon_t^y) = \phi_m^y \text{var}(\varepsilon_t^y),$$

$$(27) \quad \text{cov}(\Delta \ln(m_t), -\Delta \ln(h_{t+1})) = \text{cov}(\Delta \ln(m_t), \varepsilon_t^h) = \phi_m^h \text{var}(\varepsilon_t^h).$$

This identification strategy is the same as in Commault's (2022) robust version of the BPP estimator: because it only uses the covariances between current and future growth, it identifies the pass-through of transitory shocks without imposing a particular specification on log consumption.

The identification of the ratios ϕ_m^y and ϕ_m^h thus only requires that the law of motion for income and for health is well specified (equations (4) and (5)) and that

⁷Commault (2022) shows that the same equality between the ratios $[\text{cov}(\ln(c_t), \varepsilon_t)] / [\text{var}(\varepsilon_t)]$ and the pass-through coefficients $[d\ln(c_t)] / (d\varepsilon_t)$ holds in the case of a second-order approximation when the shocks have zero skewness—which we find to be the case empirically.

consumption is orthogonal to future shocks (at the next period). This orthogonality means that the shocks are unanticipated and thus, uncorrelated with the variables in the previous period. Online Appendix J discusses how people, anticipating changes in their income or health at the next period, induce a downward bias in our pass-through estimates.

To interpret these ratios as approximations of the pass-through coefficients $[d\ln(c_t)]/(d\varepsilon_t^y)|_0$ and $[d\ln(c_t)]/(d\varepsilon_t^h)|_0$, we additionally need log consumption to be approximately linear in transitory shocks ε_t^y and ε_t^h —but we do not need log consumption to be linear in its other determinants.

In contrast, if one wants to identify the pass-through of *permanent shocks*, one needs stronger assumptions about consumption. For example, BPP impose that log consumption is a random walk. Note that in equation (21), although we take a first-order approximation around the state variables, we do not impose random walk behavior. In particular, log consumption growth and log medical expenses growth depends on assets growth, which in general depends on past assets (except in the absence of uncertainty). It is also worth noting that, even if we were willing to make strong assumptions about the specification of log consumption to identify the pass-through of *permanent shocks*, we could not decompose their values into different channels, as we do for the pass-through of transitory shocks. This is because permanent shocks influence consumption through all channels and not just by changing the current resources and current marginal utility of consumption.

Thus, our focus on transitory shocks provides two important advantages: (i) it makes it possible first to *identify* the pass-through coefficient with a relatively small set of assumptions, while the identification of the pass-through of permanent shocks would require much stronger hypotheses; (ii) it makes it possible to *interpret* the results since the transitory shocks only affect consumption and medical expenses through two channels into which we can decompose the pass-through coefficients.

Decomposition of the Pass-Through Coefficients.—Applying the decomposition derived in (15) and (16) to the pass-through coefficients at the approximation point yields

$$(28) \quad \phi_c^y \approx \frac{d\ln(c_t)}{d\varepsilon_t^y} \Big|_0 = \overbrace{f_R^{c,t}|_0}^{\text{Multiplier}} \underbrace{p_t y_t|_0}_{E[y_t]}$$

$$(29) \quad \phi_c^h \approx \frac{d\ln(c_t)}{d\varepsilon_t^h} \Big|_0 = \overbrace{f_m^{c,t}|_0 \frac{dm_t}{d\varepsilon_t^h} \Big|_0}^{\text{Contribution of marginal utility}} + f_h^{c,t}|_0 - \overbrace{f_R^{c,t}|_0 p_t^m \frac{d\ln(m_t)}{d\varepsilon_t^h} \Big|_0}_{\approx \phi_m^h} \underbrace{m_t|_0}_{\approx E[m_t]}$$

Indeed, at the approximation point, y_t equals its average sample $E[y_t]$, and m_t approximately equals its average value $E[m_t]$.⁸ Hence, in these expressions we have only two unobserved components that need to be identified, $f_R^{c,t}|_0$ and

⁸More precisely, $m_t|_0 = m_t(E_t[a_t], E_t[\pi_t^y], E_t[\pi_t^h], E_t[\varepsilon_t^y], E_t[\varepsilon_t^h]) \approx E_t[m_t]$. Note that we could choose an approximation point $|_0$ at which we have both $y|_0 = E[y_t]$ and $m_t|_0 = E[m_t]$.

$(f_m^{c,t}|_0(dm_t)/(d\varepsilon_t^h)|_0 + f_h^{c,t}|_0)$, and two expressions to identify them. Expression (28) makes it possible to recover the value of $f_R^{c,t}|_0$, which measures by how much a transitory reduction in resources affects the pass-through of this transitory shock to consumption. This value can in turn be plugged in (29), to obtain $(f_m^{c,t}|_0(dm_t)/(d\varepsilon_t^h)|_0 + f_h^{c,t}|_0)$, which measures the contribution of the shift in the marginal utility of consumption to the pass-through of transitory health shocks. More precisely, this term is the sum of the effect of a change in health on the marginal utility of consumption $f_h^{c,t}|_0$, plus the effect of the endogenous adjustment in medical expenses caused by the change in health on the marginal utility of consumption $f_m^{c,t}|_0(dm_t)/(d\varepsilon_t^h)|_0$,

$$(30) \quad f_R^{c,t}|_0 = \frac{\phi_c^y}{p_t E[y_t]},$$

$$(31) \quad \left(f_m^{c,t}|_0 \frac{dm_t}{d\varepsilon_t^h} \Big|_0 + f_h^{c,t}|_0 \right) = \phi_c^h - \phi_h^m p_t^m E[m_t] \frac{\phi_c^y}{p_t E[y_t]}.$$

The intuition for our identification is as follows. In our model, a transitory shock only affects consumption through two channels: the resources that can be devoted to current and future consumption, and the ability to derive utility from current consumption. Holding the ability to derive utility constant, the effect of resources on consumption is the same whether the change in resources comes from a change in medical expenses or from a change in income. That is, the multiplier $f_R^{c,t}$ on the change in resources is the same in (28) and (29). Intuitively, having to pay a US\$1,000 hospital bill is equivalent to earning US\$1,000 less in net income for nonmedical consumption if the hours spent at the hospital do not change your ability to enjoy it. As a result, when income and health shocks are uncorrelated and income does not affect medical expenses (so an income shock only affects consumption through its impact on resources), we can measure this multiplier $f_R^{c,t}|_0$ from the pass-through of a transitory income shock to consumption. Knowing both this multiplier and the effect of a transitory health shock on medical expenses (which we have estimated previously as ϕ_m^h), we can predict the pass-through of a transitory health shock to consumption that would take place if marginal utility were unaffected and only the resource channel was at play. We can then recover the contribution of the shift in marginal utility as the difference between the pass-through that we measure and the pass-through that would take place if only the resources channel was at play.

B. Implementation

We construct “detrended” health and income variables, that is, net of observed demographic characteristics (see online Appendix C for details). We then use equations (17)–(20), (24)–(27), and (30)–(31) and the Generalized Method of Moments (GMM) to jointly estimate the variances and covariances of our income and health processes, the pass-through coefficients, the multiplier $f_R^{c,t}|_0$, and the marginal utility contribution $(f_m^{c,t}|_0(dm_t)/(d\varepsilon_t^h)|_0)$. We also estimate the average change in resources caused by a transitory health shock. This is not subject to any additional identification problem, as it is given by the product of the pass-through of

health to medical expenses and of the average sample value of medical expenses $\phi_h^m p_t^m E[m_t]$.⁹

When estimating, we pool observations for all years. Our identification strategy requires that the transitory and permanent shocks are not serially correlated. In estimation we allow for the errors in the moment conditions, which may come from measurement error, to be serially correlated within households. To accommodate this, we cluster at the individual level, which allows for general serial correlation of the residuals. Because of this clustering, our GMM weighting matrix is robust to heteroskedasticity.

IV. Key Facts about Our Variables of Interest

We use the Health and Retirement Study data, a longitudinal survey representative of the US population over the age of 50 and their spouses. It contains rich information on health, income, demographics, and many other variables. We combine information from the HRS core interviews and its Consumption and Activities Mail Survey, a supplementary study collecting data on household spending that is administered to a subset of HRS respondents.

Both surveys are biennial. The CAMS is conducted on the years in between the HRS surveys, but the information lines up well because income questions refer to the past year, while consumption questions refer to current consumption.¹⁰ Our merged sample covers the years 2001 to 2013 and drops Medicaid recipients, who make up 9.6 percent of our observations. Online Appendix A describes our sample selection in detail.

The rest of this section starts by detailing the construction of our variables of interest, continues by describing their first moments and percentiles, and concludes by discussing their variances.

A. Variables Construction

Consumption includes food at home and away from home, utilities, car-related expenses, leisure, and equipment. Medical expenses include out-of-pocket costs for drugs, medical services, and medical supplies. Each category is deflated by the corresponding item-specific price index of the Bureau of Labor Statistics (Bureau of Labor Statistics 1991–2015).

Our health index is constructed as follows. We attribute a numerical value from 5 to 1 to the answers to the following survey question: “Would you say your health is excellent, very good, good, fair, or poor?” Then, we predict its value by regressing the resulting variable on dummies for reporting difficulties in activities of daily living (ADLs) or for being diagnosed with certain health conditions. This procedure eliminates both changes in self-reported health that are not caused by any change in

⁹We estimate the variances of the permanent shocks separately because they require a sample of households observed for four consecutive periods, which reduces our sample size.

¹⁰The health questions refer to current health, so the overlap between health and consumption is only partial. We discuss the consequences of this feature of the data in online Appendix J.

objective health measures and changes in objective measures that do not translate into changes in self-reported health (Blundell et al. 2023 follow a similar strategy).

We construct our health measure so that a higher health index corresponds to better health. A one-unit change in our health index has the same interpretation as a one-unit change in self-reported health (although by taking the predicted value, we limit ourselves to the changes driven by our regressors). That is, a one-unit change in the index corresponds to a change from one level of response to the next.

By treating the possible self-reported health statuses as incremental numbers, we assume that changes are homogeneous, so that, for instance, the change from “excellent” to “very good” corresponds to the same quantitative decrease in health than the change from “good” to “fair.” As a result, a 0.1 increase in our health index corresponds to a health improvement of one-tenth of the health difference between “good” and “very good” (or any other two consecutive levels). In the case of households composed of a head and spouse, the health index is the average of their predicted values. Hence, a one-unit change captures any combination of $1 - x$ change in the health of one spouse and x change in the health of the other.

Net worth is the sum of all assets less all liabilities. We deflate it with the Consumer Price Index (CPI) for total consumption (Bureau of Labor Statistics 1991–2015). We take family size into account by dividing the wealth of couples by the square root of two. We define as “lower-wealth households” those with equivalized wealth below US\$75,000 and as “higher-wealth households” the rest. This breakdown splits the bottom quintile of households in terms of equivalized wealth from the other four wealth quintiles. Because we are focusing on older people who are wealthier than the general population, even people with positive net worth are part of the bottom quintile.

Our measure of income is net income. This is because we want to measure the response of consumption to income shocks after engaging in self-insurance (through both labor supply and savings) and receiving government insurance. More specifically, net income includes earnings (wages, salaries, bonuses), capital income (business or farm income, self-employment, rents, dividend and interest income, and other asset income), private pensions (income from employer pension or annuity), benefits (social security retirement income, income from transfer programs and workers’ compensations), and other income (alimony, other income, lump sums from insurance, pension, and inheritance) of both the household’s head and spouse, if present, net of taxes and transfers. We deflate it with the CPI for total consumption (Bureau of Labor Statistics 1991–2015). We report more details about our variables’ construction and some descriptives about their distributions in online Appendix B.

B. Autocovariances of Income and Health Growth and Their Cross-Covariances with Consumption Growth

For brevity, in the remainder of the paper, we refer to our main income measure—that is, the natural logarithm of detrended net income—simply as “income” and to our detrended health index simply as “health index,” or “health.” Similarly, we refer to the natural logarithm of detrended real nondurable consumption expenses as “consumption” and to the natural logarithm of detrended real out-of-pocket medical expenses as “medical expenses.”

TABLE 1—COVARIANCE OF CURRENT INCOME, HEALTH, AND CONSUMPTION GROWTH WITH CURRENT AND FUTURE INCOME AND HEALTH GROWTH

	$\Delta \ln(y_t)$	$\Delta \ln(y_{t+1})$	$\Delta \ln(y_{t+2})$	$\Delta \ln(y_{t+3})$
$\text{cov}(\Delta \ln(y_t), \cdot)$	0.213 (0.007)	-0.087 (0.005)	-0.008 (0.005)	-0.002 (0.006)
$\text{cov}(\Delta \ln(c_t), \cdot)$	0.017 (0.003)	-0.011 (0.003)	-0.001 (0.004)	0.005 (0.006)
Observations	4,999	4,999	3,094	1,915
	Δh_t	Δh_{t+1}	Δh_{t+2}	Δh_{t+3}
$\text{cov}(\Delta h_t, \cdot)$	0.064 (0.002)	-0.020 (0.001)	-0.003 (0.002)	0.002 (0.002)
$\text{cov}(\Delta \ln(c_t), \cdot)$	0.005 (0.002)	-0.003 (0.002)	0.004 (0.002)	-0.004 (0.003)
Observations	4,999	4,999	3,045	1,882

Note: Standard errors in parentheses.

To support our assumptions on the income and health processes, Table 1 presents the autocovariances of income and health growth and the cross-covariances of consumption growth with income growth and health growth.

The first lines of the top and bottom panels of Table 1 report the autocovariances of income growth and health growth. They show that both income and health can be well represented by the sum of a random walk permanent component and of a transitory component that is an MA(0). More specifically, the first two lines of the top panel show that the covariance between income growth at t and $t + 1$ is statistically significant at the 1 percent level, while it is not significant between t and $t + 2$. This is consistent with transitory income being i.i.d. In fact, if transitory income were an MA(k) process with $k > 0$, the covariance between income growth at t and $t + 2$ would be significant. In addition, if the permanent component of income were an AR(1) with a coefficient different from one, rather than a random walk, the covariance between income growth at t and all future periods would be significant, while we fail to find evidence of this. Given the first two lines of the bottom panel, the same reasoning implies that health is also well represented by the sum of a random walk and an MA(0) component.

The second lines in the top and bottom panels of Table 1 report the covariances between consumption growth and current and future income and health growth. These covariances imply that consumption covaries significantly and positively with transitory income and health shocks and that permanent health shocks are partly anticipated, at most two periods ahead. More specifically, the first column indicates that the covariances of consumption growth with contemporaneous income growth and health growth are significant and positive. Under a transitory-permanent specification of income and health, they correspond to $\text{cov}(\Delta \ln(c_t), \eta_t^x + \varepsilon_t^x - \varepsilon_{t-1}^x), x \in y, h$. Thus, a positive value already suggests that permanent and transitory income and health shocks both have a positive impact on contemporaneous consumption.

The second column indicates that the covariances of consumption growth with income growth and health growth at the next period are significant and negative.

Although the covariance between $\Delta \ln(c_t)$ and Δh_{t+1} displays only one star, the p -value of the test that it is zero is 0.051. Hence, this moment is very close to being significant at the 5 percent level. Under a transitory-permanent specification of income and health, these covariances correspond to $\text{cov}(\Delta \ln(c_t), \eta_{t+1}^x + \varepsilon_{t+1}^x - \varepsilon_t^x)$, $x \in y, h$. The fact that these covariances are negative indicates that contemporaneous transitory income and health shocks ε_t^x raise consumption by more than future (and possibly anticipated) shocks $\eta_{t+1} + \varepsilon_{t+1}$.

The third column indicates that the covariance between consumption growth and income growth two periods later is small and not significant, while the covariance between consumption growth and health growth two periods later is significant and positive. Under a transitory-permanent specification of income and health, these covariances are given by $\text{cov}(\Delta \ln(c_t), \eta_{t+2}^x + \varepsilon_{t+2}^x - \varepsilon_{t+1}^x)$, $x \in y, h$. The fact that the covariance with income ($x = y$) is small suggests that income shocks are not anticipated. The fact that the covariance with health ($x = h$) is positive and significant suggests that households partly anticipate the realization of their future permanent health shocks η^h (at most two periods ahead since the covariance after $t + 2$ is small and not significant) and that the effect of this anticipation on current consumption is positive.¹¹ Online Appendix J discusses that the presence of anticipation tends to reduce our estimated pass-through of transitory health shocks to consumption. Intuitively, when people receive advance signals about their future health, the value of our main estimating moment, $\text{cov}(\Delta \ln(c_t), -\Delta h_{t+1})$, is attenuated: consumption does not increase as much with a decrease in future health $-\Delta h_{t+1}$ because such a decrease captures both a positive transitory health shock at t , ε_t^h , and a negative signal about future permanent health at $t + 1$. Since one important message of our paper is that consumption responds to transitory health shocks, this remains true when people partly anticipate future health changes.

To further validate our assumptions, we compute additional moments, which we report in online Appendix C. They show that the covariances between health growth and subsequent income growth are small (between 0.002 and 0.003) and not significant. The same is true of the cross-covariances between income growth and subsequent health growth. This is consistent with our assumption that transitory income and health shocks are uncorrelated (although we relax this assumption in online Appendix E). Moreover, none of the cross-covariances between medical expenses growth and current and future income growth are statistically significant. In addition, the point estimate of the contemporaneous covariance is small (0.007). In contrast, for consumption, the contemporaneous covariance with income growth is significant and equal to 0.017. Furthermore, the cross-covariances of medical expenses growth and current and future health growth suggest that medical expenses respond to transitory health shocks and that these health shocks are partly anticipated (these cross-covariances are similar to the cross-covariances of (nonmedical) consumption growth and health growth, so the same reasoning applies). These results are consistent with our baseline assumption that people adjust their medical expenses

¹¹Note that this covariance is unlikely to be driven by one of the transitory shocks ε in $\Delta h_{t+2} = \eta_{t+2}^h + \varepsilon_{t+2}^h - \varepsilon_{t+1}^h$ since (i) Δh_{t+2} covaries negatively with medical expenses, while it would covary positively if $-\varepsilon_{t+1}^h$, a negative health shock, were the variable causing a significant reaction at t , and (ii) Δh_{t+3} no longer covaries significantly with $\Delta \ln(c_t)$, while it would if ε_{t+2}^h were the variable causing a significant reaction at t . However, our reasoning also holds when people partly anticipate the realization of future transitory shocks.

TABLE 2—VARIANCE OF THE TRANSITORY AND PERMANENT INCOME SHOCKS

	All	Lower wealth	Higher wealth
$\text{var}(\Delta \ln(y_t))$	0.213 (0.007)	0.165 (0.013)	0.225 (0.008)
$\text{var}(\varepsilon_t^y)$	0.087 (0.005)	0.066 (0.009)	0.093 (0.005)
Observations	4,999	970	4,029
$\text{var}(\eta_t^y)$	0.029 (0.006)	0.017 (0.010)	0.031 (0.006)
Observations	3,401	623	2,778

Note: Standard errors in parentheses.

in response to transitory changes in their health but not in response to transitory changes in their income (although we show that we can relax this assumption in online Appendix D).

C. Variances of the Income and Health Shocks

Table 2 highlights that, even at advanced ages, households face substantial income risk. More precisely, the first line of this table reports the variance of the changes in income, $\text{var}(\Delta \ln(y_t))$, across households and periods in our sample. It turns out to be 0.213 and significant. The second line reports the variance of transitory income shocks, $\text{var}(\varepsilon_t^y)$. It has a point estimate of 0.087, and it is significant. This means that current transitory shocks explain 41 percent of the variance of income growth. The third line reports the variance of permanent shocks, $\text{var}(\eta_t^y)$, which has a point estimate of 0.029 and is also significant. Current permanent shocks thus explain about 14 percent of the variance of income growth. Past transitory shocks explain the remainder of this variance.¹²

Lower-wealth households face less income risk than higher-wealth households, in particular, in terms of permanent income risk. This is consistent with a larger fraction of their income coming in from benefits.

We do not assume that our shocks are normal, but we estimate the third and fourth moments of the transitory shocks distribution to better understand its characteristics. The third moment is small and not significant, suggesting the distribution of shocks to income displays little skewness. The fourth moment is large and significant, and its point estimate is more than four times what a normal distribution would imply. This suggests that the distribution of shocks to our income measure has fat tails. Online Appendix C, section “Skewness and kurtosis of the shocks,” presents additional results on this.

Given that we find that older households face substantial income risk, one might wonder what are its sources, especially since previous papers assume that there is no such risk. To explore this question, we compute the standard deviation of each

¹²Indeed, with a transitory-permanent income process, the variance of the changes in income is $\text{var}(\Delta \ln y_t) = \text{var}(\varepsilon_t^y) + \text{var}(\varepsilon_{t-1}^y) + \text{var}(\eta_t^y)$.

TABLE 3—VARIANCE OF TRANSITORY AND PERMANENT HEALTH SHOCKS

	All	Lower wealth	Higher wealth
$\text{var}(\Delta h_t)$	0.064 (0.002)	0.098 (0.006)	0.056 (0.002)
$\text{var}(\varepsilon_t^h)$	0.020 (0.001)	0.033 (0.004)	0.017 (0.001)
$\text{cov}(\varepsilon_t^y, \varepsilon_t^h)$	0.002 (0.001)	0.003 (0.003)	0.002 (0.001)
Observations	4,999	970	4,029
$\text{var}(\eta_t^h)$	0.020 (0.002)	0.026 (0.005)	0.018 (0.002)
$\text{cov}(\eta_t^y, \eta_t^h)$	0.002 (0.002)	-0.003 (0.004)	0.004 (0.002)
Observations	3,401	623	2,778

Note: Standard errors in parentheses.

(detrended) income component in the population. “Benefits” (including social security retirement income) is the income category that displays the smallest variations. The standard deviations of pensions income, earnings, and other income (including inheritances or insurance claims) are relatively similar and twice as large as the standard deviation of benefits. Capital income (including business income) is the category that displays the highest standard deviation. We report these results in detail in online Appendix C, section “Standard deviations of the different components of income.”

Turning to our results for health shocks, Table 3 highlights that households face substantial health risk. More precisely, the first line of this table reports the variance of the changes in health, $\text{var}(\Delta h_t)$, across households and periods in our sample. It turns out to be 0.064 and significant. The second line reports the variance of transitory health shocks, $\text{var}(\varepsilon_t^h)$. It has a point estimate of 0.020 and is significant. This means that current transitory shocks explain one-third of the variance of health growth. The fourth line reports the variance of permanent shocks, $\text{var}(\eta_t^h)$, which has a point estimate of 0.020 and is also significant. Current permanent shocks thus explain another third of the variance of health growth. Past transitory shocks explain the remainder.

A variance of 0.020 implies that, overall in the population and across periods, the shocks are drawn from a distribution with a standard deviation of 0.141. This means that a transitory health shock corresponding to a 1 standard deviation change is a change in health index by 0.141—that is, a change in health corresponding to 14.1 percent of the health difference between two health levels, for example, from “good” to “very good.”

The third line of the table indicates that there is a very small covariance (0.002) between transitory income and health shocks, which is not significant. Because this correlation is tiny and not significant, in most of our analysis, we assume that temporary income and health shocks are uncorrelated. Still, we relax this assumption in online Appendix E. In it, we posit the existence of underlying “pure income” and “pure health” transitory shocks that are uncorrelated but can affect the transitory components of both income and health, resulting in a covariance between the two.

The results from this alternative approach show that the variances of the underlying shocks are almost indistinguishable from the variances of our transitory components, and that, more generally, relaxing this assumption changes our results very little.

The magnitude of the covariance between the permanent shocks is similar to that of the covariance between the transitory shocks (although because variance of the permanent shocks is smaller, in relative terms, the covariance between permanent shocks could be more important), and its estimate is not statistically significant. Note that our identification strategy would be robust to this covariance being nonzero.

Unlike in the case of income shocks, lower-wealth households face higher variances of both transitory and permanent health shocks. The variance of the transitory health shocks is twice as large among lower-wealth households than among higher-wealth households, at 0.033 versus 0.017. Permanent health risk is also larger among lower-wealth households than among higher-wealth households.

Thus, Table 3 and the right-hand-side graph of Figure B.4 in online Appendix B show that lower-wealth households are less healthy than higher-wealth households and experience more health fluctuations, both transitory and permanent. These results are not inconsistent with our assumption that the stock of health carried by an older person is in large part determined by its past life events: people who arrive in old age with lower wealth likely had less means and time (and possibly had parents with less means and time) to build their health stock earlier in life.

Here, too, we do not need to assume that the shocks are normally distributed, but we estimate the third and fourth moments of the distribution of the transitory income shocks because it is interesting. The point estimate of the third moment is zero and not significant, suggesting the distribution is not substantially skewed. The fourth moment is large and significant, and the point estimate is more than five times what a normal distribution would imply, suggesting the distribution has fat tails. Online Appendix C, section “Skewness and kurtosis of the shocks” presents these additional results. It is also worthwhile noticing that, because shocks are centered around zero, some people are subject to positive health shock. However, because these are shocks to detrended health, a positive health shock is best interpreted not as an actual increase in health but as health not deteriorating as fast as its trend would predict, either permanently or temporarily.

To further investigate the determinants of a change in our health index, Table 4 presents the results from a regression of our detrended health index over changes in the reported difficulty to perform instrumental activities of daily living.¹³ As we would expect, the results show that all kinds of difficulties have either a negative and significant impact on our health index or an insignificant one but no positive and significant impact. In terms of magnitudes, the first line, for instance, indicates that if the household head reports a decreased ability to walk for several blocks (from this activity being “not at all difficult” to “very difficult/can’t do”), everything else being equal, its health index decreases by 0.201 (statistically significant at the 1 percent level), which is a bit more than 1 standard deviation of a transitory health

¹³The way we obtain Table 4 is by regressing changes in the health index over changes in the reported difficulties in instrumental activities of daily living. We thus select households for whom difficulties are observed at two consecutive periods. In the absence of spouse, the changes to the spouse variables are set to zero.

TABLE 4—TEMPORARY HEALTH CHANGES AND THE HEALTH INDEX

Change in health index	Coefficient	Standard deviation
Difficult to walk several blocks—head	−0.201	(0.016)
Difficult to walk one block—head	−0.053	(0.022)
Difficult to sit two hours—head	−0.032	(0.016)
Difficult to get up from chair—head	−0.009	(0.015)
Difficult to climb several flights stairs—head	−0.056	(0.011)
Difficult to climb one flight stairs—head	−0.103	(0.017)
Difficult to climb stoop/kneel/crouch—head	−0.049	(0.013)
Difficult to lift/carry 10 lbs—head	−0.066	(0.016)
Difficult to pick up a dime—head	0.007	(0.027)
Difficult to extend arms—head	−0.003	(0.019)
Difficult to push/pull large object—head	−0.119	(0.013)
Difficult to walk several blocks—spouse	−0.145	(0.024)
Difficult to walk one block—spouse	−0.097	(0.034)
Difficult to sit two hours—spouse	−0.054	(0.021)
Difficult to get up from chair—spouse	−0.035	(0.021)
Difficult to climb several flights stairs—spouse	−0.077	(0.016)
Difficult to climb one flight stairs—spouse	0.032	(0.024)
Difficult to climb stoop/kneel/crouch—spouse	−0.036	(0.018)
Difficult to lift/carry 10 lbs—spouse	−0.019	(0.019)
Difficult to pick up a dime—spouse	0.014	(0.043)
Difficult to extend arms—spouse	−0.085	(0.027)
Difficult to push/pull large object—spouse	−0.051	(0.016)
Observations	3,261	

shock. The coefficients of the spouse are of similar magnitude and significance as those of the head.

V. Estimated Pass-Through and Their Decomposition

How do nondurable consumption and out-of-pocket medical expenses respond to temporary shocks in income and health? The first row of Table 5 reports the effects of transitory income and health shocks on nondurable consumption, while the second row displays their effect on out-of-pocket medical expenses.

Transitory income shocks imply significant changes in consumption. In particular, the average pass-through coefficient of income shocks to nondurable consumption is 0.127, significant at the 1 percent level. Hence, a 10 percent transitory decrease in current income leads to a 1.27 percent decrease in nondurable consumption (and vice versa for an increase). This estimate implies that a US\$100 decrease in income reduces nonmedical consumption by US\$6.45 at the average levels of income and consumption.¹⁴

The response of consumption among lower-wealth households is more than twice as large than for all households (0.202) and is statistically significant at the 5 percent level. This suggests that poorer households find it more difficult to self-insure against transitory income shocks. Perhaps surprisingly, even the consumption of

¹⁴One can translate our pass-through coefficient $[d\ln(c_t)]/[d\ln(y_t)]$, which is an elasticity, into the change in the level of consumption that is implied by a change in the level of income $(dc_t)/(dy_t)$. We can do so by using that $(dc_t)/(dy_t) = [d\ln(c_t)]/[d\ln(y_t)] \times (c_t)/(y_t)$. At the average levels of income and consumption in our sample, this corresponds to $(dc_t)/(dy_t) = 0.127 \times (24,279)/(47,825) = 0.0645$.

TABLE 5—PASS-THROUGH ESTIMATES

	Income shock			Health shock		
	Total	Lower wealth	Higher wealth	Total	Lower wealth	Higher wealth
Consumption ϕ_c^ε	0.127 (0.036)	0.202 (0.100)	0.115 (0.038)	0.173 (0.088)	0.306 (0.132)	0.112 (0.114)
Medical expenses ϕ_m^ε	0.132 (0.102)	0.234 (0.288)	0.114 (0.107)	-0.493 (0.232)	-1.171 (0.364)	-0.177 (0.286)
Observations	4,999	970	4,029	4,999	970	4,029

Note: Standard errors in parentheses.

higher-wealth households responds to transitory income shocks. Their pass-through is 0.115, and it is statistically significant. Further disaggregating households in this group into those with low-liquid wealth (or “hand-to-mouth”) and those with high-liquid wealth reveals that the pass-through is largest (0.232) and significant at the 1 percent level for the former group but much smaller (0.070) and only significant at the 10 percent level for the latter group (see online Appendix F).

While transitory income shocks have statistically insignificant effects on out-of-pocket medical expenses, their pass-through coefficient is 0.132, which is not very different from that of consumption. However, because the level of out-of-pocket medical expenses is small, a pass-through coefficient of this magnitude implies only small fluctuations in the level of medical expenses—which are substantially smaller than the consumption fluctuations implied by a pass-through of the same magnitude. More specifically, because average out-of-pocket medical expenses are 10 percent of total consumption, our estimates imply that a US\$100 decrease in income lowers nonmedical consumption by US\$6.45 but reduces medical expenses by only US\$0.66 at average income, consumption, and medical expenses.¹⁵ This is important because it is the change in the level of medical expenses that matters for our identification strategy of the channels decomposition (we set this change to zero in our baseline case, although we relax this assumption in online Appendix D).

Our pass-through estimates of income to medical expenses are also not statistically significant within the groups of lower- and higher-wealth households. Again, the point estimates are not very different from those of consumption, but because the level of out-of-pocket medical expenses is small in both groups (and smaller among lower-wealth households, at US\$2,515 compared with US\$3,024 for all households, as given in online Appendix B, Table B3), the change in the level of medical expenses generated by a change in income is small for these two groups as well.

Transitory health shocks also imply significant changes in consumption. The top row of the right-hand-side panel of Table 5 shows that the point estimate of this pass-through is 0.173, statistically significant at the 5 percent level. This means that a 0.1 transitory decrease in our health index generates a 1.73 percent decrease in consumption, that is, a US\$420 decrease for a household with the average

¹⁵The computation of the change in the level of medical expenses is similar to that of consumption. At the average levels of income and medical expenses in our sample, this corresponds to $(dm_t)/(dy_t) = 0.127 \times (2,515)/(47,825) = 0.0066$.

consumption. A 1 standard deviation decrease in health, that is, a 0.141 decrease, implies a $0.173 \times 0.141 = 2.4\%$ decrease in consumption, that is, a US\$592 decrease for a household with the average consumption level.

The breakdown by wealth shows that, among the lower-wealth households, the pass-through of transitory health shocks to consumption (0.306) is almost twice as large as in our overall sample. In this group, a 1 standard deviation transitory decrease in health, that is, a 0.182 decrease, is associated with a $0.306 \times 0.182 = 5.6\%$ decrease in nondurable consumption. Among the higher-wealth households, the point estimate of the health pass-through is not significant.

Transitory health shocks imply significant changes in medical expenses. The bottom row of the right-hand-side panel of Table 5 shows that the pass-through of transitory health changes to medical expenses is negative, large (-0.493), and statistically significant at the 5 percent level. This means that a 0.1 transitory decrease in our health index generates a 4.93 percent increase in medical expenses. At the average medical expenses level of US\$3,024, this corresponds to a $0.1 \times 0.493 \times \text{US\$}3,024 = \text{US\$}149$ increase in medical expenses. A transitory decrease in health by 1 standard deviation, that is, a 0.141 decrease, is associated with a $0.493 \times 0.141 = 7.0\%$ increase in medical expenses, which corresponds to a US\$210 increase at their average level.

Importantly, we find that the effect of transitory health shocks on medical expenses is heterogeneous by wealth. The average pass-through coefficient is more than twice as large (-1.171) and statistically significant at the 1 percent level among lower-wealth households than in the whole sample. Because their average medical spending is US\$2,515, their medical expenses change by $0.1 \times 1.171 \times 2,515 = \text{US\$}295$ when the health index changes by 0.1, twice as much as in the whole sample. A 1 standard deviation decrease in health, that is, a 0.182 decrease, is associated with a $1.171 \times 0.182 = 21.3\%$ increase in medical expenses, which corresponds to a US\$536 increase at their average level.

Among higher-wealth households, this effect is much lower (-0.177) and not statistically significant, even at the 10 percent level. This finding is consistent with the fact that lower-wealth households, on average, spend only half as much in medical insurance as higher-wealth households, even after removing the effect of demographics.¹⁶ This suggests that, despite Medicare, the medical expenses of people with less private insurance are less insured against transitory health shocks: their out-of-pocket medical expenses increase with a temporary decline in their health.

The observation that medical expenses are little impacted by transitory income shocks but do respond to transitory health shocks suggests that most people tend to be close to satiation in their consumption of medical goods and services but that this satiation point varies with their health. The presence of Medicare is likely important in generating this result because it tends to make the level of extra out-of-pocket expenses required to stay at one's satiation point relatively small. As a result, even lower-wealth households with less private insurance can afford to stay close to their medical consumption satiation point, even if, for them, a health shock implies a significant change in resources.

¹⁶The net expense in medical insurance is on average US\$(2015)1,698 among older lower-wealth households and US\$(2015)2,914 among older higher-wealth households. See online Appendix B, Table B3.

Marital Status.—In online Appendix G, we also break down our sample in two subsamples: that of single households (2,255) and that of couples (2,744). Separately looking at couples and singles is interesting because being in a couple is both a source of risks (the health and resource risks of one's partner) and insurance (pooling risks, economies of scale, and potentially being able to help each other in case of sickness). The point estimates of the pass-through coefficients for income shocks to consumption are 0.143 for singles and 0.113 for couples. Those for health shocks are 0.183 for singles and 0.160 for couples. This is consistent with couples' consumption being less affected by transitory income and health shocks. However, breaking down the sample reduces statistical power. As a result, the differences between the coefficients of couples and singles are not statistically significant. In line with the results in our overall sample, the pass-through of income to medical expenses is small and not significant for both singles and couples. Finally, the pass-through of health shocks to medical expenses is -0.342 for singles and -0.704 for couples, which indicates that couples react to transitory health shocks by spending more in medical goods and services compared with singles. The coefficient for couples is significant, but the estimates for singles and couples are not statistically different. We report these results in online Appendix G.

Robustness.—In our baseline framework, we assume that income shocks are discrete events occurring at the same time every year,¹⁷ that there is no measurement error in income and health, that people do not anticipate future health shocks, and that there is a complete overlap between the consumption and health periods of observation. Relaxing the first three of these assumptions would lead to a modest downward bias in our pass-through estimates, while the effects of the fourth one is ambiguous (see online Appendix J).

Comparison with Existing Estimates.—There is a large literature that relies on natural experiments to measure the effects of transitory income shocks on consumption. It suggests that, among working-age households, the average marginal propensity to consume (MPC) is around 0.25 over the next quarter (see, e.g., the review by Kaplan and Violante 2018). Few of these studies examine the behavior of people in old age, but age seems to be negatively associated with the MPC (sometimes weakly): using lottery wins, Fagereng, Holm, and Natvik (2021) find that, while the average MPC of total—not just nondurable—consumption out of a lottery win is 0.59 over the next year, it drops to 0.44 among people above age 63 and that the difference between the two is significant. Johnson, Parker, and Souleles (2006) and Parker et al. (2013), exploit the 2001 and 2008 tax rebates and break down the sample into different age categories but find no significant differences.

Commault (2022) focuses on the comparison between the pass-through estimates obtained with semistructural methods and the MPCs obtained from natural experiments. She shows that the robust semistructural methods that we use here yield

¹⁷ We do not need to make this assumption about health shocks, which we compute by comparing the stocks of health at two points in time. In contrast, income is a flow that we observe every other year, so we need an assumption about the point in time when a change in the flow occurs to determine the magnitude of the change from the difference in yearly flow.

TABLE 6—DECOMPOSITION

	All	Lower wealth	Higher wealth
Consumption ϕ_c^h	0.173 (0.088)	0.306 (0.132)	0.112 (0.114)
Resources channel	0.003 (0.002)	0.018 (0.011)	0.001 (0.002)
<i>Change in expenses</i> $-\phi_m^h E[mp_m]$	-1,117.91 (528.194)	-2,137.017 (690.707)	-420.397 (679.599)
<i>Multiplier</i> $f_3^c _0 (10^{-6})$	3.093 (0.874)	8.382 (4.120)	2.530 (0.840)
Marginal utility channel $f_1^c \frac{dm_i}{de_i^h} + f_2^c$	0.170 (0.088)	0.288 (0.132)	0.111 (0.114)
Observations	4,999	970	4,029

Note: Standard errors in parentheses.

results that are consistent with those from natural experiments but that one needs to be careful when comparing results from studies in which the data are measured at different frequencies—for instance, quarterly in the natural experiments and biennially in our case. In particular, the pass-through estimated on biennial data turns out to be smaller than the one from annual data. Commault (2022) also shows that the biennial pass-through is 0.125 in the more recent waves of the Panel Study of Income Dynamics (PSID) and among working-age households. Our finding of a pass-through of transitory income shocks to consumption is based on the robust estimator. It yields a pass-through of 0.127, which is thus close to the one estimated from the PSID.

Our findings thus show that temporary changes in income and health affect consumption in old age but do not address our main question: to what extent do consumption fluctuations later in life reflect lack of insurance against fluctuations in one's resources as opposed to fluctuations in one's needs to consume? We now turn to these results.

Decomposing the Impact between Marginal Utility and Resources.—Table 6 reports the results of the decomposition of the pass-through of transitory health shocks to consumption into the two channels that compose it in our structural model. Over our whole sample, the contribution of the resources channel is significant at the 10 percent level, but its point estimate is only 0.003. This means that, if health and medical expenses had no effect on the marginal utility of consumption, the pass-through of a transitory health shock to consumption would be 0.003 instead of 0.173, that is, only 1.7 percent of its actual value.

More precisely, equation (29) shows that the resources channel comprises the product of the change in medical expenses caused by a transitory health shock and a multiplier parameter. This multiplier determines how much larger (or smaller) is the pass-through of a shock when this shock increases (or decreases) resources by an extra dollar. We compute the changes in medical expenses caused by a transitory health shock as the product of the average pass-through of a health shock to medical expenses times their average amount. Our estimates show that both terms

are significant: the change in medical expenses caused by a 1-unit transitory health shock has a point estimate of $-\text{US}\$1,118$, and the multiplier has a point estimate of $3.093e^{-6}$.

The rest of the pass-through is explained by the marginal utility channel, which is significant at the 10 percent level as well. If a transitory health shock had no impact on the budget constraint but still influenced the utility function (through both changes in health and in medical goods and services consumed), the pass-through of a transitory health shock to consumption would be 0.170, very close to its true value of 0.173.

Among lower-wealth households, the overall pass-through coefficient is larger than in the whole sample, with a point estimate of 0.306. The decomposition shows that this comes from both the resources and the marginal utility channels being larger. The contribution of resources is 0.018, significant at the 10 percent level (despite the small sample size), and 6 times as large as in the whole sample. This resources channel is larger for two reasons: first, among lower-wealth households, a one-unit decrease in health raises medical expenses by more (US\$2,137 instead of US\$1,118 in the whole sample).¹⁸ Second, among lower-wealth households, a given dollar decrease in resources is passed on more strongly to consumption, so the multiplier is larger. The result is that, among lower-wealth households, the resources channel explains 5.9 percent of the overall pass-through, versus 1.7 percent in the whole sample. The marginal utility channel contributes the remaining 94.1 percent. If this channel was the only one at play, the pass-through coefficient would be 0.288, significant at the 5 percent level. The converse holds true for higher-wealth households for whom both channels are smaller than in the whole sample. Among higher-wealth households, the overall pass-through coefficient is not statistically significant, although its point estimate is substantial, at 0.112. The contributions of both the resources channel and the marginal utility channel are weaker.

Why is the magnitude of the marginal utility channel different between lower- and higher-wealth households? The contribution of the marginal utility channel can vary across households (e.g., by wealth) even when they have the same marginal utility function, if their consumption levels are different. Indeed, a shift in utility does not affect consumption uniformly along the marginal utility function. For instance, if consumption is close to a satiation point, a multiplicative shift in the ability to derive utility from consumption is not going to affect consumption too much because consumption might remain close to its satiation point even after the value of consuming today has decreased. In contrast, before that satiation point is reached, there can be consumption levels around which a shift in marginal utility will strongly shift one's consumption decision. We graphically illustrate this point with Figure H.1 in online Appendix H. The observation that the contribution of the marginal utility channel is larger among lower-wealth households could thus be due to their consumption being on a portion of the marginal utility where it is more sensitive to shifts in marginal utility.

¹⁸ As shown in the expressions, we use nondeflated medical expenses in this analysis to account for changes in the price of medical expenses over time: if, for instance, our pass-through estimate implies that a given health generates an increase in real medical expenses of one box of pills, the resources effect of that health shock is incorrectly measured as larger if we use deflated values and convert the price of the box of pills in 2015 US dollars, the year when drugs are more expensive.

VI. Estimated Pass-Through Coefficients and Their Decomposition for Finer Spending Categories

We write our model in terms of total consumption and medical expenses because this formulation relates to many previous studies and provides an important benchmark. Given the richness of our data, we are also able to examine more disaggregated consumption and medical expense categories. This provides further insights.

To estimate the pass-through of the shocks to these disaggregated categories, we do not make further assumptions: as for the pass-through to consumption and medical expenses, we only need income and health to be transitory-permanent processes and their future shocks not to be anticipated. To decompose these pass-through coefficients, the underlying model is very similar to that in Section I, except that households now derive utility from N different categories of goods: $c_t^n, n = 1, \dots, N$. This alternative formulation yields very similar decomposition expressions.

Table 7 reports the effects of transitory income and health shocks on consumption and medical expenses at different levels of disaggregation. Its top left-hand side panel displays that in our overall sample, as we have seen in the previous section, nondurable consumption responds significantly to a transitory income shock, and the point estimate of the pass-through is 0.127. The breakdown by necessities and luxuries reveals that the pass-through to the two categories are of similar magnitude (0.109 and 0.110). Going one level of disaggregation further, necessities, food (at home and away from home), utilities, and car-related expenses (car insurance, repairs and gasoline) have similar pass-through coefficients, but only the responses of utilities and car-related expenses remain significant. Within luxuries, the pass-through on leisure activities (spending on trips, hobbies, and sports equipment) is large (0.219) and significant at the 1 percent level, but the pass-through to expenses on equipment (clothing, personal care, house and garden supplies and services) is small and not significant.

The breakdown of the responses by wealth reveals that the categories of disaggregated consumption that respond the most to a transitory income shock are different among lower- and higher-wealth households. For lower-wealth households, spending increases more on necessities, which are food, utilities, and car maintenance. For higher-wealth households, spending increases more on leisure activities. One interpretation is that lower-wealth households are not satiated in their consumption of necessities and thus, adjust it when transitory income fluctuations hit, while higher-wealth households are satiated in their consumption of necessities but not in their consumption of luxuries.

Turning to the effect of an income shock on medical expenses, the bottom left-hand-side panel of Table 7 shows that there are no statistically significant effects of a transitory income shock on any medical expense subcategory. This also holds when we split the response by wealth.

The right-hand-side panel of Table 7 reports the effects of a transitory health shock on consumption (top panel) and medical expenses (bottom panel). As we have seen in the previous section, in our overall sample, nondurable consumption responds significantly to a transitory health shock, and the point estimate of the pass-through is 0.173. The goods categories breakdown reveals that these effects come from the response of luxuries (their pass-through is 0.366 and significant)

TABLE 7—PASS-THROUGH ESTIMATES

	Income shock			Health shock		
	Total	Lower wealth	Higher wealth	Total	Lower wealth	Higher wealth
Consumption ϕ_c^ε	0.127 (0.036)	0.202 (0.100)	0.115 (0.038)	0.173 (0.088)	0.306 (0.132)	0.112 (0.114)
<i>Necessities</i>	0.109 (0.038)	0.314 (0.111)	0.075 (0.040)	0.076 (0.090)	0.344 (0.141)	-0.046 (0.112)
<i>Food</i>	0.090 (0.062)	0.425 (0.183)	0.033 (0.065)	0.045 (0.152)	0.697 (0.266)	-0.259 (0.183)
<i>Utilities</i>	0.099 (0.053)	0.223 (0.128)	0.077 (0.057)	0.044 (0.128)	-0.127 (0.191)	0.125 (0.165)
<i>Car-related</i>	0.098 (0.046)	0.248 (0.125)	0.073 (0.050)	0.285 (0.117)	0.580 (0.184)	0.147 (0.148)
<i>Luxuries</i>	0.110 (0.063)	-0.186 (0.178)	0.160 (0.066)	0.366 (0.150)	0.206 (0.220)	0.438 (0.193)
<i>Leisure activities</i>	0.219 (0.091)	-0.190 (0.333)	0.290 (0.086)	0.426 (0.231)	0.180 (0.372)	0.536 (0.282)
<i>Equipment</i>	0.023 (0.068)	-0.297 (0.175)	0.077 (0.073)	0.401 (0.159)	0.115 (0.228)	0.536 (0.203)
Medical expenses ϕ_m^ε	0.132 (0.102)	0.234 (0.288)	0.114 (0.107)	-0.493 (0.232)	-1.171 (0.364)	-0.177 (0.286)
<i>Drugs</i>	0.063 (0.109)	0.134 (0.285)	0.050 (0.117)	-0.619 (0.248)	-0.936 (0.409)	-0.472 (0.304)
<i>Services and supplies</i>	-0.021 (0.144)	-0.024 (0.403)	-0.022 (0.152)	0.098 (0.343)	-0.173 (0.524)	0.222 (0.433)
Observations	4,994	966	4,028	4,994	966	4,028

Note: Standard errors in parentheses.

rather than of necessities (their pass-through is 0.076 and not significant). Going one level of disaggregation further, among luxuries, both leisure activities and equipment respond. Among necessities, car-related expenses respond.

The breakdown by wealth reveals that, among those with lower wealth, necessities respond more than in the whole sample but it is unclear how luxuries respond since both the pass-through estimate and its standard error are large. Among higher-wealth households, necessities do not respond (the pass-through is small and not significant), while leisure activities and equipment respond strongly and significantly (their point estimates are both equal to 0.536).

The bottom right-hand side of Table 7 reports the response of out-of-pocket medical spending to a transitory health shock. As we saw in the previous section, the pass-through of adverse transitory changes in health to medical expenses is negative and large, at -0.493. Breaking down the effects of a health shock on the components of medical expenses shows that it is the drugs category that significantly responds to transitory health shocks and drives the overall response of medical expenses.

Table K1 in online Appendix K reports the decomposition of the effects coming from marginal utility and resources for our disaggregation by necessities and luxuries. In our overall sample, we find a small and insignificant impact of temporary health shocks on the consumption of necessities. Neither the resources nor the marginal utility channel are significant—although some components of the resources

channel are. Both the resource channel and marginal utility channel are larger among lower-wealth households, but they are either not significant or barely significant.

For the pass-through to luxuries, which is statistically significant and large (0.366) for our overall sample, the resources channel contributes very little to this pass-through, at 0.001. This means that a change in future resources plays no role in the response of the consumption of luxuries. The shift in marginal utility plays a very large role, as it contributes 0.365 out of the 0.366 pass-through coefficient, significant at the 1 percent level. Among lower-wealth households, the overall pass-through to luxuries is quite large but with a large standard deviation, and neither the resources channel nor the marginal utility channels are significant. Among higher-wealth households, the consumption of luxuries responds significantly to temporary health shocks, and the response is also almost entirely driven by the marginal utility channel. Its contribution is 0.434, that is, 99.1 percent of the overall coefficient. This is consistent with a scenario in which higher-wealth households are close to satiation in necessities but not in luxuries; hence, a shift in their ability to derive utility from luxuries has a large impact on their consumption of these goods.

Health and Demand System Estimation.—Previous literature has studied how demand shares change with total expenses and demographics (see, for example, Banks, Blundell, and Lewbel 1997). In this section, we generalize the well-established tool of estimating demand systems to include health. We do so by estimating a demand system that captures how different commodities are affected by both health status and total expenditure.

This approach complements our main results by focusing on the impact of resources and health on within-period allocations. As we condition on total resources and the health status variable, we do not distinguish between the impact of transitory and permanent shocks to health or resources. Because it is a within-period analysis, it ignores any reallocation from current to future consumption when the marginal utility of current consumption decreases. Also, by construction, it does not tell us by how much total consumption changes with health because total resources are kept constant, and we only study the allocation among goods. In estimation we specify a functional form for the budget shares. See online Appendix I for more details.

Table 8 reports the budget and health elasticities for our five-commodities demand system. Its top line shows that, on average, food expenses compose 27 percent of the budget, utilities 23 percent, car maintenance 16 percent, leisure 21 percent, and equipment the remaining 13 percent. The budget elasticities of the items that we group into necessities (food, utilities, and car expenses) are all lower than one, while leisure and equipment, which we group into luxuries, display budget elasticities above one. Hence, our grouping and variable labeling is supported by the data.¹⁹

The estimated elasticities of demand to health status reveal that an improvement in health does have a differential impact on the marginal utility of different

¹⁹We also further disaggregate food expenses into food at home and food away and find that while food at home is a necessity and has a budget elasticity of 0.6, food away from home is a luxury and has an estimated budget elasticity of 1.3. However, because food away from home makes up for only 6.7 percent of the budget for nondurables, and its reaction to health changes is on average 0, we do not disaggregate further to keep the analysis simple.

TABLE 8—BUDGET AND HEALTH ELASTICITIES, FOR DISAGGREGATED CATEGORIES

	Food	Utilities	Car	Leisure	Equipment
<i>Budget shares</i>	0.271 (0.007)	0.232 (0.006)	0.159 (0.005)	0.208 (0.007)	0.131 (0.005)
<i>Budget elasticities</i>	0.777 (0.023)	0.577 (0.027)	0.797 (0.030)	1.864 (0.033)	1.086 (0.035)
<i>Health elasticities</i>					
<i>Whole sample</i>	-0.117 (0.025)	-0.091 (0.029)	0.104 (0.032)	0.324 (0.035)	-0.235 (0.038)
<i>Lower wealth</i>	-0.121 (0.023)	-0.096 (0.027)	0.200 (0.033)	0.488 (0.067)	-0.203 (0.041)
<i>Higher wealth</i>	-0.107 (0.027)	-0.207 (0.031)	0.058 (0.033)	0.345 (0.034)	-0.084 (0.039)

Note: Standard errors in parentheses.

consumption goods since it shifts the budget shares in favor of certain goods and in disfavor of others. This also confirms the nonseparability between the disaggregated consumption goods and health. On average, an increase in health has its biggest positive impact on leisure expenses, and its effect is also positive for car maintenance. The health elasticities of other goods are negative.

In online Appendix G, we report estimates of the demand system when breaking down the sample between the households composed of single individuals and those composed of couples. Singles have in general higher budget elasticities than couples (except for the one on equipment). Their health elasticities are higher, in absolute value, for food and car-related expenses (which are linked to activities that can be easier to be undertaken by one's partner when in a couple).

VII. Going from Positive to Normative Implications

We have so far performed a positive analysis; that is, we have measured the magnitude of the shocks that people face, and how and why they react to these shocks, given the insurance system that is already in place. In this section, we turn to a normative analysis and use our findings to study the social planner problem for the household that we described in Section II.

The social planner optimally allocates consumption and out-of-pocket medical expenses to households, subject to a resource constraint. This formulation is equivalent to letting the planner allocate consumption and total medical expenses (that is, including both out-of-pocket and insured medical expenses), but we find it convenient to solve the problem in terms of what households optimize over in the decentralized economy.

The implications of the presence of transitory income shocks for insurance are straightforward. A social planner using a utilitarian welfare function would completely offset the effect of the shock and allocate the same consumption and medical expenses to people, whether they are hit or not by an income shock.

Now, we want to know how the consumption of someone experiencing a negative transitory health shock changes in a planned economy. Health shocks are different

because we established that negative transitory health shocks come with two effects: an increase in the marginal utility of medical expenses (which can lead to an increase in medical expenses, thus a decrease in the resources available for consumption in a decentralized economy) and a decrease in the marginal utility of consumption. Let us focus on period t and assume that a negative transitory health shock lowers health below its current value, to $h_t^- = h_t - \Delta h_t$. Following the literature on consumption insurance (for instance, Cochrane 1991 and Attanasio and Davis 1996), the optimal allocation coming from a social planner using a utilitarian welfare function implies that the marginal utility of consumption in the two states is the same:

$$(32) \quad u_c(c_t, \tilde{m}_t, h_t) = u_c(c_t + \Delta c_t, \tilde{m}_t + \Delta m_t, h_t - \Delta h_t),$$

where Δc_t and $\Delta \tilde{m}_t$ are the differences in consumption and medical expenses between the bad and normal health states. Taking an approximation of the left-hand side of the above expression around the point where $\Delta c_t = 0$, $\Delta \tilde{m}_t = 0$, and $\Delta h_t = 0$, we can write the optimal difference in consumption across health states

$$(33) \quad u_c(c_t, \tilde{m}_t, h_t) \approx u_c(c_t, \tilde{m}_t, h_t) + u_{cc}^t \Delta c_t + u_{c\tilde{m}}^t \Delta \tilde{m}_t - u_{ch}^t \Delta h_t,$$

$$\Delta c_t \approx \frac{-1}{u_{cc}^t} (u_{c\tilde{m}}^t \Delta \tilde{m}_t - u_{ch}^t \Delta h_t),$$

where $u_{c\tilde{m}}^t = u_{c\tilde{m}}(c_t, \tilde{m}_t, h_t)$, $u_{ch}^t = u_{ch}(c_t, \tilde{m}_t, h_t)$ and $u_{cc}^t = u_{cc}(c_t, \tilde{m}_t, h_t)$ are the partial derivatives of $u(c_t, \tilde{m}_t, h_t)$.

Equation (33) states that the optimal change in consumption resulting from a negative transitory health shock depends on the extent to which the negative health shock decreases the marginal utility of consumption (measured by the product of the size of the health loss $-\Delta h_t$ and the effect of health on the marginal utility of consumption u_{ch}^t) and on the extent to which the change in medical expenses already compensates this loss in marginal utility (measured by the product of the extra medical expenses received in the bad health state $\Delta \tilde{m}_t$ and the effect of medical expenses on the marginal utility of consumption $u_{c\tilde{m}}^t$).

We thus consider two subcases. The first, in which there is separability in the utility of consumption and of medical expenses conditional on health ($u_{c\tilde{m}}^t = 0$ and $u_{ch}^t \neq 0$), we use to set out the main ideas. The second, which is our preferred case, allows for nonseparability in the utility of consumption, health, and medical expenses ($u_{c\tilde{m}}^t \neq 0$ and $u_{ch}^t \neq 0$).

Separability in the Utility of Consumption and of Medical Expenses.—Because $u_{c\tilde{m}}^t = 0$, we are left with

$$(34) \quad \Delta c_t \approx \frac{-1}{u_{cc}^t} u_{ch}^t (-\Delta h_t).$$

We show in online Appendix L that the contribution of marginal utility (MU) to the pass-through of transitory health to consumption when $u_{c\tilde{m}}^t = 0$ is given by

$$(35) \quad MU_t = (u_{ch}^t) \frac{-1}{c_t \vartheta_t},$$

where $\vartheta_t = E_t[(1/p_{t+1})(c_a^{t+1} u_{cc}^{t+1}) \tilde{s}_{t+1} R_{t+1}]$ has the same sign as $u_{cc} < 0$. Table 6 shows that $MU = 0.17 > 0$, which implies $u_{ch}^t > 0$.

Plugging this implication in expression (34), we find that the planner allocates less consumption to those experiencing a negative transitory health shock:

$$(34') \quad \Delta c_t \approx \underbrace{\frac{-1}{u_{cc}^t}}_{>0} \underbrace{u_{ch}^t (-\Delta h_t)}_{<0} < 0.$$

Intuitively, everything else being equal, people going through a negative transitory health shock derive less enjoyment from an extra unit of resources, and it is thus optimal that they consume less.

While this type of insurance might seem counterintuitive, the rationale is the same as insuring unexpected life expectancy by providing more consumption to people who live longer than they expected: although people prefer to have a long life than a short one, they are happy to enter an insurance scheme giving them more total consumption if their life is long rather than short. Here, although people prefer to live their life in good health—with a high ability to enjoy consumption—than in bad health, they are happy to enter an insurance scheme giving them more consumption when their health state is good.

Our finding that a decrease in health raises out-of-pocket medical expenses, while income changes do not affect them, implies that the response of medical expenses to health is driven by a shift in their marginal utility (rather than from a change in resources coming from a reduction in consumption goods and services). Hence, a utilitarian benevolent social planner would allocate more medical expenses to those experiencing a negative transitory health shock.

As a result, a utilitarian benevolent planner allocates less nonmedical consumption and more medical expenses to households in bad health and more consumption but less medical expenses to households in good health. By doing so, it provides insurance against the risk of not having enough resources to pay for medical expenses while in bad health and the risk of not having as much resources as one would like to consume while in good health.

Relaxing Additive Separability in the Utility of Consumption and Medical Expenses.—In this case, an increase in medical expenses can raise the marginal utility of consumption. It is then possible for the planner to use extra medical expenses to partly compensate for the loss in marginal utility of consumption caused by a negative health shock.

Given that we find that most people are at their out-of-pocket medical expenses satiation point—which is hence attainable given available resources—the planner should optimally give households in both the normal and bad health states their satiation level of medical expenses.²⁰ Hence, the planner should allocate to people experiencing a negative transitory health shock an amount of extra medical

²⁰Indeed, giving more would be inefficient, and giving less than this point in either the normal or bad health state in order to give people more consumption in either the normal or bad health state would violate people's revealed preference that they like to remain close to their satiation point of out-of-pocket medical expenses (even lower-wealth people do).

expenses equal to the extra amount we estimate they consume in the decentralized economy: $\Delta \tilde{m}_t \approx (d\tilde{m}_t/d\varepsilon_t^h)(-\Delta h_t)$. Plugging this into the expression of optimal consumption change (33), we have

$$(36) \quad \Delta c_t \approx \frac{-1}{u_{cc}^t} \left[\frac{dm_t}{d\varepsilon_t^h} \tilde{m}'(m_t) u_{c\tilde{m}}^t + u_{ch}^t \right] (-\Delta h_t).$$

Now, the same term $(dm_t)/(d\varepsilon_t^h) \times \tilde{m}'(m_t) u_{c\tilde{m}}^t + u_{ch}^t$ appears in the expression of the contribution of marginal utility to the pass-through of health shocks to consumption that we derive in online Appendix L:

$$(37) \quad MU_t = \left[\frac{dm_t}{d\varepsilon_t^h} \tilde{m}'(m_t) u_{c\tilde{m}}^t + u_{ch}^t \right] \frac{-1}{c_t \vartheta_t},$$

where $\vartheta_t = E_t[1/(p_{t+1})(c_a^{t+1} u_{cc}^{t+1} + \tilde{m}_a^{t+1} u_{c\tilde{m}}^{t+1}) \tilde{s}_{t+1} R_{t+1}]$ takes the same sign as $u_{cc} < 0$ (because the effect of wealth a on medical expenses \tilde{m}_a^{t+1} is zero, as discussed in online Appendix L). Our empirical finding that $MU > 0$ then implies that

$$(38) \quad \left[\frac{dm_t}{d\varepsilon_t^h} \tilde{m}'(m_t) u_{c\tilde{m}}^t + u_{ch}^t \right] > 0.$$

In other words, our finding that the marginal utility channel is positive implies that the increase in medical expenses people get does not fully compensate the decrease in marginal utility that they experience when their health drops. Otherwise, the marginal utility channel would be zero.

Plugging this implication in expression (36), we find again that the planner optimally allocates less consumption to those experiencing a negative transitory health shock:

$$(36') \quad \Delta c_t \approx \frac{-1}{u_{cc}^t} \underbrace{\left[\frac{dm_t}{d\varepsilon_t^h} \tilde{m}'(m_t) u_{c\tilde{m}}^t + u_{ch}^t \right]}_{>0} \underbrace{(-\Delta h_t)}_{<0} < 0.$$

Intuitively, even if the loss in marginal utility of consumption is being partly compensated by the extra medical expenses that people consume, people going through a negative transitory health shock still derive less utility from an extra unit of resources, so the planner optimally allocates less consumption to them.

As in the previous case, although the utilitarian planner allocates less consumption to those experiencing a negative transitory health shock, it also allocates them more medical goods and services. The reason why people in a bad health state receive more medical expenses is still because they have a higher marginal utility of them and an attainable satiation point exists. However, in that case, there is an extra benefit of giving more in medical expenses to people in bad health, which is that it might partly offset the negative effect of bad health on the marginal utility of consumption.

As a result, our finding that optimal consumption is lower in bad health and higher in good health constitutes at the same time a new risk (a mismatch between one's consumption resources and one's marginal utility of them) and possibly a mitigating

mechanism against the risk of increased out-of-pocket medical expenses, and thus, reduced consumption resources while in bad health. We stress the importance of studying this force. Figuring out to what extent this affects the design of optimal health insurance and whether the current health insurance system is optimal requires a fully quantitative normative analysis. We discuss this in the concluding section.

Note that our finding that the marginal utility of consumption decreases in bad health does not only have consequences for the utilitarian planner. Consider a benevolent planner solving a social welfare function that places more weight on people with lower utility levels. In this case, our results can induce the planner to compensate more the people hit with a negative health shock because of their lowered ability to derive utility from consumption and allocate them more consumption because of that.

VIII. Conclusions and Directions for Future Research

We study the effects of income and health changes on people's consumption and medical expenses in old age and provide several interesting and novel findings. First, we show that transitory income and health shocks are prevalent in old age.

Second, we document the response of consumption to transitory income and health shocks. We find that, even during retirement, consumption responds to income shocks, which indicates that people's consumption is not perfectly insured against shocks to one's resources. This result complements previous work finding that consumption responds to income shocks at younger ages. In terms of the response of consumption to health shocks, we show that consumption significantly decreases with negative health shocks (and increases with positive health shocks). Our data analysis and the implications of a rich structural model allow us to show that the consumption response to negative health shocks mainly takes place because health shocks reduce one's marginal utility from consumption rather than reducing resources. An important group for whom a health shock does have an effect on resources is that of lower-wealth households.

Third, we evaluate the response of medical expenses to transitory income and health shocks. Here, we show that out-of-pocket medical expenses do not respond to transitory income shocks in a statistically significant way. This suggests that, given the level of insurance provided by the current system after age 65, most people are close to being satiated in their consumption of out-of-pocket medical expenses—when people receive extra income, they do not increase their out-of-pocket medical expenses, but they increase their consumption. In contrast, they increase with negative health shocks (and decrease with positive health shocks) in a statistically significant way. This indicates that this satiation point changes with health: people hit by a negative health shock increase their out-of-pocket medical spending.

From a normative standpoint, in the presence of shocks that affect people's resources, a benevolent planner using a utilitarian welfare function fully smooths out their effects and gives the same level of consumption to all, whether hit by a shock or not. In contrast, in the presence of shocks that reduce people's marginal utility, the risk that needs insuring is not a decrease in resources but a mismatch between people's level of resources and their ability to enjoy them: people in better health than expected might not have enough resources to enjoy their good

health, while people suffering from worse health than expected might not have the ability to take full advantage of the resources they accumulated. As a result, a benevolent planner insures shocks that affect people's utility by giving less consumption but more medical expenses in the states associated with increased utility of medical expenses and reduced marginal utility of consumption.

Deriving quantitative normative implications, such as the optimal quantitative compensation associated with a transitory health shock, requires estimating and identifying a structural model that allows for permanent and transitory income and health shocks and for health to affect the marginal utility of consumption. More precisely, it requires taking a stand on all functional forms and parameter values of one's model. In contrast, this is avoided with our approach. We see a quantitative normative analysis as an important direction for future research, especially given that we find that bad health reduces one's marginal utility from consumption and this is a force that has largely been ignored in the normative literature.

Our analysis holds under general conditions, both about how health and income evolve during retirement and about how people optimally choose their consumption and medical expenses in the presence of savings and health and income shocks. However, it does assume that health evolution is, at least at the margin, largely predetermined and exogenous during retirement. As we discuss in Section I, this is a commonly made assumption that is also supported by much empirical evidence. Allowing health to depend on one's spending (and potentially effort) requires fully specifying and parameterizing all aspects of the model and, very importantly, taking a stand on the health production function and its identification. Also, we do not model that consumption might be produced by using time and requiring good health (sick people might need more time to do the same things and this is why they consume less), which could be a way to further micro-found our findings. We see these as important directions for future research.

REFERENCES

- Ameriks, John, Joseph Briggs, Andrew Caplin, Matthew D. Shapiro, and Chris Tonetti. 2020. "Long-Term Care Utility and Late in Life Saving." *Journal of Political Economy* 128 (6): 2375–2451.
- Attanasio, Orazio, and Steven J. Davis. 1996. "Relative Wage Movements and the Distribution of Consumption." *Journal of Political Economy* 104 (6): 1227–62.
- Banks, James, Richard Blundell, and Arthur Lewbel. 1997. "Quadratic Engel Curves and Consumer Demand." *Review of Economics and Statistics* 79 (4): 527–39.
- Black, Bernard, José-Antonio Espín-Sánchez, Eric French, and Kate Litvak. 2017. "The Long-Term Effect of Health Insurance on Near-Elderly Health and Mortality." *American Journal of Health Economics* 3 (3): 281–311.
- Blau, David M., and Donna Gilleskie. 2008. "The Role of Retiree Health Insurance in the Employment Behavior of Older Men." *International Economic Review* 49 (2): 475–514.
- Blundell, Richard, Hamish Low, and Ian Preston. 2013. "Decomposing Changes in Income Risk using Consumption Data." *Quantitative Economics* 4 (1): 1–37.
- Blundell, Richard, Margherita Borella, Jeanne Commault, and Mariacristina De Nardi. 2024. "Replication data for: Old Age Risks, Consumption, and Insurance." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E193623V1>.
- Blundell, Richard, Jack Britton, Monica Costa Dias, and Eric French. 2023. "The Impact of Health on Labor Supply near Retirement." *Journal of Human Resources* 58 (1): 282–334.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review* 98 (5): 1887–1921.

- Blundell, Richard, Rowena Crawford, Eric French, and Gemma Tetlow.** 2016. "Comparing Retirement Wealth Trajectories on Both Sides of the Pond." *Fiscal Studies* 37 (1): 105–30.
- Braun, R. Anton, Karen A. Kopecky, and Tatyana Koreshkova.** 2017. "Old, Sick, Alone, and Poor: A Welfare Analysis of Old-Age Social Insurance Programmes." *Review of Economic Studies* 84 (2): 580–612.
- Brook, Robert, John Ware, William Rogers, Emmett Keeler, Allyson Davies, Cathy Donald, George Goldberg, Kathleen Lohr, Patricia Masthay, and Joseph Newhouse.** 1983. "Does Free Care Improve Adults' Health? Results from a Randomized Controlled Trial." *New England Journal of Medicine* 309 (23): 1426–34.
- Brown, Jeff, and Amy Finkelstein.** 2008. "The Interaction of Public and Private Insurance: Medicaid and the Long-Term Care Insurance Market." *American Economic Review* 98 (3): 1083–1102.
- Bureau of Labor Statistics.** 1991–2015. "CPI for All Urban Consumers (CPI-U) 1982-84=100 (Unadjusted) - CUUR0000SA0." United States Department of Labor. <https://data.bls.gov/cgi-bin/surveymost?bls> (accessed September 2023).
- Card, David, Carlos Dobkin, and Nicole Maestas.** 2009. "Does Medicare Save Lives?" *Quarterly Journal of Economics* 124 (2): 597–636.
- Cloyne, James, Kilian Huber, Ethan Iizetzi, and Henrik Kleven.** 2019. "The Effect of House Prices on Household Borrowing: A New Approach." *American Economic Review* 109 (6): 2104–36.
- Cochrane, John H.** 1991. "A Simple Test of Consumption Insurance." *Journal of Political Economy* 99 (5): 957–76.
- Commault, Jeanne.** 2022. "Does Consumption Respond to Transitory Shocks? Reconciling Natural Experiments and Semistructural Methods." *American Economic Journal: Macroeconomics* 14 (2): 96–122.
- De Nardi, Mariacristina, Eric French, and John Bailey Jones.** 2016. "Savings After Retirement: A Survey." *Annual Review of Economics* 8 (1): 177–204.
- De Nardi, Mariacristina, Eric French, and John B. Jones.** 2010. "Why Do the Elderly Save? The Role of Medical Expenses." *Journal of Political Economy* 118 (1): 39–75.
- De Nardi, Mariacristina, Svetlana Pashchenko, and Pompoje Porapakkarm.** 2017. "The Lifetime Costs of Bad Health." NBER Working Paper 23963.
- Evans, William N., and W. Kip Viscusi.** 1991. "Estimation of State-Dependent Utility Functions Using Survey Data." *Review of Economics and Statistics* 73 (1): 94–104.
- Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik.** 2021. "MPC Heterogeneity and Household Balance Sheets." *American Economic Journal: Macroeconomics* 13 (4): 1–54.
- Farhi, Emmanuel, and Iván Werning.** 2013. "Estate Taxation with Altruism Heterogeneity." *American Economic Review* 103 (3): 489–95.
- Finkelstein, Amy, Erzo F.P. Luttmer, and Matthew J. Notowidigdo.** 2009. "Approaches to Estimating the Health State Dependence of the Utility Function." *American Economic Review* 99 (2): 116–21.
- Finkelstein, Amy, Erzo F.P. Luttmer, and Matthew J. Notowidigdo.** 2013. "What Good Is Wealth Without Health? The Effect of Health on The Marginal Utility of Consumption." *Journal of the European Economic Association* 11: 221–58.
- Finkelstein, Amy, and Robin McKnight.** 2008. "What did Medicare do? The Initial Impact of Medicare on Mortality and Out of Pocket Medical Spending." *Journal of Public Economics* 92 (7): 1644–68.
- Fisher, Elliott S., David E. Wennberg, Thérèse A. Stukel, Daniel J. Gottlieb, F. L. Lucas, and Etoile L. Pinder.** 2003. "The Implications of Regional Variations in Medicare Spending. Part 2: Health Outcomes and Satisfaction with Care." *Annals of Internal Medicine* 138 (4): 288–98.
- French, Eric, John Jones, and Jeremy McCauley.** 2017. "The Accuracy of Economic Measurement in the Health and Retirement Study." *Forum for Health Economics and Policy* 20 (2): 1–16.
- Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski.** 2016. "Redistribution and Social Insurance." *American Economic Review* 106 (2): 359–86.
- Health and Retirement Study (HRS).** 2014. RAND longitudinal file public use dataset, version 1992-2014v2. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant NIA U01AG009740). Ann Arbor, MI, (2014). <https://hrsdata.isr.umich.edu/data-products/rand-hrs-archived-data-products> (accessed May 5, 2018).
- Health and Retirement Study (HRS).** 2015. CONSUMPTION AND ACTIVITIES MAIL SURVEY (CAMS) public use dataset. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant NIA U01AG009740). Ann Arbor, MI, (2001-2015). <https://hrsdata.isr.umich.edu/data-products/public-survey-data> (accessed May 5, 2018).
- Health and Retirement Study (HRS).** 2017. RAND HRS FAT FILES, public use dataset. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant NIA U01AG009740). Ann Arbor, MI, (2000-2014). <https://hrsdata.isr.umich.edu/data-products/rand> (accessed May 5, 2018).

- Health and Retirement Study (HRS).** 2018. RAND HRS Tax calculations 2014 v2, public use dataset. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant NIA U01AG009740). Ann Arbor, MI, (2001-2015). <https://hrsdata.isr.umich.edu/data-products/rand-hrs-tax-calculations-2014> (accessed May 5, 2018).
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni Violante.** 2009. "Quantitative Macroeconomics with Heterogeneous Households." *Annual Review of Economics* 1 (1): 319–54.
- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes.** 1994. "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving." *Carnegie Rochester Series on Public Policy* 40: 59–125.
- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes.** 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy* 103 (2): 360–99.
- Hurd, Michael D., and Susann Rohwedder.** 2009. "Methodological Innovations in Collecting Spending Data: The HRS Consumption and Activities Mail Survey." *Fiscal Studies* 30 (3-4): 435–59.
- Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles.** 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review* 96 (5): 1589–1610.
- Kaplan, Greg, and Giovanni L. Violante.** 2010. "How Much Consumption Insurance beyond Self-Insurance?" *American Economic Journal: Macroeconomics* 2 (4): 53–87.
- Kaplan, Greg, and Giovanni L. Violante.** 2018. "Microeconomic Heterogeneity and Macroeconomic Shocks." *Journal of Economic Perspectives* 32 (3): 167–94.
- Khwaja, Ahmed.** 2010. "Estimating Willingness to Pay for Medicare Using a Dynamic Life-Cycle Model of Demand for Health Insurance." *Journal of Econometrics* 156 (1): 130–47.
- Love, David A., Michael G. Palumbo, and Paul A. Smith.** 2009. "The Trajectory of Wealth in Retirement." *Journal of Public Economics* 93 (1): 191–208.
- Meghir, Costas, and Luigi Pistaferri.** 2004. "Income Variance Dynamics and Heterogeneity." *Econometrica* 72 (1): 1–32.
- Mian, Atif, Kamallesh Rao, and Amir Sufi.** 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics* 128 (4): 1687–1726.
- Palumbo, Michael G.** 1999. "Uncertain Medical Expenses and Precautionary Saving near the End of the Life Cycle." *Review of Economic Studies* 66 (2): 395–421.
- Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland.** 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review* 103 (6): 2530–53.
- Poterba, James M., Steven F. Venti, and David A. Wise.** 2018. "Longitudinal Determinants of End-of-Life Wealth Inequality." *Journal of Public Economics* 162: 78–88.
- Viscusi, W. Kip, and William N. Evans.** 1990. "Utility Functions That Depend on Health Status: Estimates and Economic Implications." *American Economic Review* 80 (3): 353–74.