

WHAT GOOD IS WEALTH WITHOUT HEALTH? THE EFFECT OF HEALTH ON THE MARGINAL UTILITY OF CONSUMPTION

Amy Finkelstein

Massachusetts Institute of Technology

Erzo F. P. Luttmer

Dartmouth College

Matthew J. Notowidigdo

University of Chicago,
Booth School of Business

Abstract

We estimate how the marginal utility of consumption varies with health. To do so, we develop a simple model in which the impact of health on the marginal utility of consumption can be estimated from data on permanent income, health, and utility proxies. We estimate the model using the Health and Retirement Study's panel data on the elderly and near-elderly, and proxy for utility with measures of subjective well-being. Across a wide range of alternative specifications and assumptions, we find that the marginal utility of consumption declines as health deteriorates, and we are able to clearly reject the null of no state dependence. Our point estimates indicate that a one-standard-deviation increase in the number of chronic diseases is associated with a 10%–25% decline in the marginal utility of consumption relative to this marginal utility when the individual has no chronic diseases. We present some simple, illustrative calibration results that suggest that state dependence of the magnitude we estimate can have a substantial effect on important economic problems such as the optimal level of health insurance benefits and the optimal level of life-cycle savings. (JEL: D12, I1).

The editor in charge of this paper was Stefano DellaVigna.

Acknowledgments: We are grateful to Robert Barro, Raj Chetty, David Cutler, Larry Katz, Jean-Paul L'Huillier, Erik Hurst, Olivia Mitchell, Jim Poterba, Ivan Werning, Justin Wolfers, many anonymous referees, and seminar participants at the MIT Public Finance lunch, the Kennedy School faculty lunch, Wharton, the Chicago Fed, the Harvard/MIT/BU Health Seminar, the NBER Summer Institute, Toronto, Columbia, Dartmouth, the Harris School, the AEA meetings, Berkeley, Urbana-Champaign, London School of Economics, and Pompeu Fabra for helpful comments, and to James Wang for excellent research assistance. This research was supported by the US Social Security Administration through grant #10-P-98363-1-05 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium and by the National Institute of Aging through grant #T32-AG000186. The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, any agency of the Federal Government, or the NBER. All errors are our own. Finkelstein and Luttmer are Research Associates and Notowidigdo is a Research Fellow at the NBER. Luttmer is a Research Fellow at IZA.

E-mail: afink@mit.edu (Finkelstein); erzo.fp.luttmer@dartmouth.edu (Luttmer); noto@chicagobooth.edu (Notowidigdo)


1. Introduction

It has long been recognized that a dependence of the shape of the utility function on health status has implications for a range of important economic behaviors (e.g., Zeckhauser 1970, Arrow 1974). Yet it is standard practice in applied work to assume that the shape of the utility function does not vary with health. State independence is routinely assumed by papers that estimate the demand for (or value of) health-related insurance products such as acute health insurance (e.g., Feldstein 1973, Feldman and Dowd 1991), long-term care insurance (e.g., Brown and Finkelstein 2008), annuities (e.g., Mitchell et al. 1999, Davidoff, Brown, and Diamond 2005), or disability insurance (e.g., Golosov and Tsyvinski 2006). State independence is also assumed in modeling the demand for health care (e.g., Hall and Jones 2007), and in calibrations of individuals' optimal life-cycle savings (e.g., Engen, Gale, and Uccello 1999, Scholz, Seshadri, and Khitatrakun 2006), although there are some important exceptions (e.g., French 2005, De Nardi, French, and Jones 2010). Yet even a moderate amount of state dependence can have a substantial effect on the conclusions of such calculations. Moreover, not only the magnitude but also the sign of any potential state dependence is a priori ambiguous. On the one hand, the marginal utility of consumption could decline with deteriorating health, as many consumption goods—such as travel—are complements to good health. On the other hand, the marginal utility of consumption could increase with deteriorating health, as other consumption goods—such as prepared meals or assistance with self-care—are substitutes for good health.


Despite the potential importance, there has been relatively little empirical work on how the marginal utility of consumption varies with health. This presumably reflects the considerable empirical challenges involved in constructing credible estimates. We begin our paper with a discussion of possible empirical approaches for estimating state dependence and their limitations.

In this paper, we adopt an approach in which we examine how the difference in individual utility between healthy and sick states of the world varies with consumption. If the difference in utility increases with consumption (as in Figure 1A), we infer that the marginal utility of consumption declines as health deteriorates, a phenomenon we refer to as negative state dependence. By contrast, if the difference in utility declines with consumption (as in Figure 1B), we conclude that marginal utility increases as health deteriorates (positive state dependence). Moreover, the magnitude of the change in the difference in utility across health states by consumption level allows us to quantify the magnitude of any state-dependent utility.

There are two key practical challenges to implementing this conceptually straightforward approach. First, data with broad-based consumption measures are notoriously scarce. We therefore develop a model of optimizing behavior that allows us to infer, under conditions which appear to be empirically valid, how marginal utility of *consumption* varies with health status from estimates of how marginal utility of *permanent income* (which is more widely available) varies with health status. Second, the approach requires a valid proxy for utility. In our context, a valid proxy is one



Panel A: Marginal utility declines with sickness



Panel B: Marginal utility increases with sickness

FIGURE 1. State-dependent utility functions.

whose signal-to-noise ratio is high enough to detect any true state dependence, and that is not contaminated by measurement error that varies with health differentially by permanent income. Our baseline utility proxy is a measure of subjective well-being (SWB), specifically whether or not the individual agrees with the statement “Much of the time during the past week I was happy.”

Economists, rightly, tend to be skeptical about the use of SWB measures. However, as we discuss in detail in what follows, most of the well-known issues would tend to decrease the precision of our estimates without biasing the coefficients. We also thoroughly investigate the concern that our estimates could be biased due to systematic errors in our assumed mapping from latent von Neumann–Morgenstern utility to measured SWB; we conclude that such errors are unlikely to be driving our estimates and provide some intuition for why this potential type of bias should not be a concern in practice in our particular setting.

We implement our approach using the Health and Retirement Study’s (HRS) panel data of a representative sample of the elderly and near-elderly in the United States. We estimate the effect of chronic disease on the marginal utility of nonmedical

consumption, evaluated at a constant level of nonmedical consumption. To minimize the effect of health on nonmedical consumption—and hence estimated marginal utility—through changes in labor income or because of medical expenses, we restrict our sample to individuals without labor income and with medical insurance; we also discuss the adjustment needed to our estimate to account for any change in nonmedical consumption.

Across a wide range of alternative specifications and assumptions, we find statistically significant evidence that the marginal utility of consumption declines as health deteriorates. The point estimates of our baseline specifications indicate that, relative to marginal utility of consumption when the individual has no chronic diseases, a one-standard-deviation increase in an individual's number of chronic diseases is associated with a 10%–25% decline in marginal utility. The point estimates of alternative specifications generally also fall in this range, but not in every case.¹ In short, our paper provides clear evidence on the sign of state dependence but leaves considerable uncertainty on its exact magnitude.

While our main estimates are based on regression specifications that we derive from a model of optimal consumption and health expenditure choices, the key empirical finding can be illustrated by a simple figure that is the empirical analogue of Figure 1. We take all individuals in our baseline sample who transition from above-median health status (“good health”) to below-median health status (“poor health”), and group them into terciles according to their permanent income, which we use as a proxy for consumption for this illustration. Figure 2 plots the percentage reporting to be happy by health status and permanent-income terciles. Note that, by construction, the mean happiness of those in good health and the mean happiness of those in poor health are calculated using the same set of individuals within each income tercile, so changes in mean happiness by permanent-income tercile are based on within-individual comparisons. The figure illustrates that mean happiness drops more after a deterioration in health for those with higher permanent incomes than those in lower income groups. In other words, there is negative state dependence—the gradient of happiness with respect to permanent income becomes lower when health deteriorates.

To illustrate the potential implications of state dependence, we examine the impact of our estimates for the optimal levels of health insurance benefits and life-cycle savings. The results from simple calibration exercises suggest that, relative to a state-independent benchmark, accounting for our estimate of state dependence lowers the optimal share of medical expenditures reimbursed by health insurance by about 20 to 45 percentage points and lowers the optimal fraction of earnings saved for retirement by about 3 to 5 percentage points. Of course, considerable caution should be exercised in using the results of our extremely stylized calibrations. Nonetheless, at a basic level,

1. In particular, estimates for state dependence are considerably smaller in magnitude for two alternative measures of subjective well-being, though these estimates are not statistically significantly different from our baseline estimates. In addition, the magnitude and significance of the estimates for state dependence are sensitive to the inclusion of other interactions with permanent income, though they are not statistically significantly different from our baseline estimates.




FIGURE 2. The impact of health on happiness by income tercile. Sample consists of all 12,047 individual \times wave observations in our baseline HRS sample (age ≥ 50 , not in labor force, with health insurance) who transition from good health (less than the median number of chronic diseases of two) to poor health (two or more chronic diseases). Happy is a dummy variable for the respondent agreeing to the statement “Much of the time during the past week I was happy.” Error bars denote 95% confidence intervals.

they underscore the likely substantive importance of the state dependence that we detect.

The rest of the paper proceeds as follows. Section 2 describes the main options for estimating state-dependent utility. Section 3 presents our approach. Section 4 presents our main results. Section 5 discusses potential threats to our identifying assumptions. Section 6 presents some simple numerical calculations that help gauge the magnitude of our estimate of state dependence. The last section concludes.

2. Approaches to Estimating the State Dependence of the Utility Function

There are two broad classes of possible approaches to estimating health-state dependence.² One class of approaches is based on individuals’ revealed demand for moving resources across health states. For example, health insurance demand—with its associated state-dependent payoff stream—would seem to be a natural way

2. We offer a more detailed discussion of these possible approaches in Finkelstein, Luttmer, and Notowidigdo (2009).

to infer state dependence. In practice, however, inferences about state dependence based on health insurance demand would be sensitive to assumptions about other parameters of the utility function that affect insurance demand, such as risk aversion. This may pose challenges given the wide range of estimates of risk aversion in the literature (see, e.g., Cohen and Einav 2007). Another practical difficulty is that moral hazard issues presumably preclude the offering of health insurance policies that pay out *more* than medical expenditures; thus if marginal utility increases as health deteriorates, this could be difficult to uncover from estimates based on health insurance demand.

A related approach would be to try to infer state dependence from information on how the time profile of consumption varies across otherwise identical individuals with different health trajectories; optimizing individuals will adjust their consumption path to increase consumption in periods when marginal utility is high and decrease it in periods when marginal utility is low. Conceptually, the most attractive way to implement this approach would be to compare the consumption profiles of otherwise identical individuals who differ in their expected future health. But in practice, we know of no dataset with information on consumption and in which expected future health differs for currently similar individuals. Alternatively, one could examine how consumption responds to unexpected health shocks. However, a key limitation to such an approach is that, since the life-cycle budget constraint must be satisfied (resources must be consumed or left to the next generation), any inference about even the existence and sign of state dependence (let alone the magnitude) will be sensitive to what is assumed about bequest motives. At a practical level, there is a marked lack of consensus on how to model bequest motives, and the choice of model can have a profound effect on the inference about state dependence drawn from the consumption trajectory. For example, if we assume that individuals receive no utility from bequests (e.g., all observed bequests are accidental as in, for instance, Abel 1985), then an individual with state-dependent utility who experiences a permanent health shock will not change her level of consumption because the only thing on which she can spend her remaining lifetime income is consumption. In contrast, if intentional bequests are allowed, an unexpected permanent health shock can affect consumption even when utility is not state dependent. For example, in the strategic bequest model of Bernheim, Shleifer, and Summers (1985), if utility from seeing your children increases when your health declines, you may want to consume less (and thus bequeath more) when you receive an unexpected negative health shock.³

The second broad class of approaches is to estimate how the within-individual utility change associated with a health shock varies across individuals of different

3. Despite these obstacles, Lillard and Weiss (1997) pursue just such an approach. They build a structural model of consumption in which they make the key assumption that marginal utility of consumption depends on health but the marginal utility of bequests does not. Inferring consumption from panel data on asset changes and income flows in the Retirement History Survey, they compare consumption paths across individuals who vary in their predicted probability of entering poor health (based on different demographic characteristics such as education and wealth). They estimate that the marginal utility of consumption rises as health decreases.

consumption levels or resources. This is the approach we adopt, and the next section describes in detail how we operationalize it. It of course also has its limitations, which we discuss in detail in what follows. However, given the formidable obstacles to the revealed demand approach, we believe that it offers a valuable and complementary avenue for obtaining estimates of state-dependent utility.

In a series of papers, Viscusi and co-authors pursue a very different version of this general type of approach. They survey individuals regarding how much money they would require to compensate them for hypothetical exposure to specific health risks, and examine how these self-reported compensating differentials vary with income. This approach has yielded a wide range of estimates, from those indicating no state dependence (Evans and Viscusi 1991), to estimates suggesting that marginal utility in the diseased state is just 8% of the marginal utility in the healthy state (Sloan et al. 1998), as well as intermediate estimates (Viscusi and Evans 1990).⁴

3. Empirical Approach

Figure 1 illustrated the intuition behind our empirical approach: if sickness causes a larger decline in utility for individuals with higher consumption than for individuals with lower consumption, the utility curve for good health must be steeper than the one for poor health, which means that the marginal utility of consumption falls in poor health. Conversely, if the drop in utility is smaller at higher levels of consumption, the marginal utility of consumption increases in poor health. If we could observe information on health, consumption, and a proxy for utility, we could simply regress the utility proxy on consumption, health, and the interaction of consumption and health, and the coefficient on the interaction term would give an estimate of state-dependent utility. In practice, however, we know of no panel dataset with sufficient sample size that contains information on *consumption*, health, and utility proxies. We therefore use the Health and Retirement Study (HRS), a nationally representative panel of the elderly and near-elderly, which contains data on *permanent income*, health, and utility proxies.

We outline a model of optimizing consumption behavior that yields conditions under which we can infer how marginal utility of consumption varies with health status from empirical evidence on the relationship between permanent income, health, and utility proxies. The key requirement is that wealth in the sick state is pre-determined, meaning that health shocks do not lead to changes in wealth. This requirement guides our selection of a baseline sample limited to individuals who are not in the labor force

4. The range of estimates may reflect the different populations and diseases studied, as well as different assumptions regarding the coefficient of relative risk aversion (i.e., how the marginal utility of income changes with income), which determines how the empirical estimate of the relationship between income and the size of the compensating differential is translated into a relationship between income and the marginal utility of health.

(so that health shocks do not have a first-order effect on income).⁵ We investigate potential threats to the validity of this assumption in Section 5.2.

3.1. Theoretical Framework and Estimating Equation

We build a parsimonious two-period model in which optimizing individuals take account of how marginal utility of consumption depends on health in choosing both their time-path of consumption and the allocation of their consumption between nonhealth consumption and health services. For expositional ease, we assume that health is binary. All individuals are healthy ($S = 0$) in the first period and have probability p of falling ill ($S = 1$) in period 2. We adopt the Epstein–Zin (1989) and Weil (1990) formulation of preferences in order to allow the intertemporal elasticity of substitution (IES) and the level of risk aversion to vary independently. Therefore, lifetime utility aggregates first-period nonhealth consumption C_1 and expected second-period utility according to

$$U = \left(\frac{1}{1-\gamma} \right) \left(C_1^{1-\theta} + \frac{1}{1+\delta} ((1-\gamma)E_1[U_2])^{(1-\theta)/(1-\gamma)} \right)^{(1-\gamma)/(1-\theta)}, \quad (1)$$

where γ denotes the coefficient of relative risk aversion, $1/\theta$ is the elasticity of intertemporal substitution, and δ denotes the discount rate. We assume that $\gamma \geq 0$ and $\theta \geq 0$. $E_1[U_2]$ is the expectation of utility in period 2 from the perspective of period 1. Second-period utility is stochastic from the perspective of period 1 because of uncertainty about one's future health state.⁶ Second-period utility is given by

$$U_2 = (1 + \varphi_1 S) \frac{1}{1-\gamma} C_2^{1-\gamma} + S \Psi(H), \quad (2)$$

where C_2 denotes second-period nonhealth consumption, H denotes consumption of health services, and S is the indicator of sickness. The second-period utility function is a standard CRRA utility function for healthy individuals. Sickness affects second-period utility in two ways. First, sickness multiplies the marginal utility of second-period consumption by a factor of $(1 + \varphi_1)$. The objective of our empirical analysis is to recover an unbiased estimate of φ_1 .⁷ Second, only sick individuals derive utility

5. We further limit the sample to individuals with health insurance. We do this to minimize the size of the adjustment that we have to make to our estimate to account for changes in nonmedical consumption caused by out-of-pocket medical expenses. Because this adjustment also depends on the coefficient of relative risk aversion, our estimate is less sensitive to assumptions about the coefficient of relative risk aversion for more fully insured individuals.

6. In addition, we will allow the rate of return on savings to be stochastic in our calibration exercise in Section 6. However, we treat the rate of return on savings as nonstochastic here to obtain a closed-form solution.

7. Sickness may also have a direct effect on the level of utility. We abstract from this in the theoretical model to simplify the exposition and to obtain closed-form solutions when $\gamma \neq \theta$. Consistent with this assumption, in Section 6 we simulate the effects of our estimated state dependence on the marginal utility on savings or optimal insurance by normalizing the intercept of the second-period utility function such that health does not affect the level of utility. Of course, in the empirical analysis, we include a control for any direct effects of health on subjective well-being to ensure that our estimates are solely identified by the interaction of health and permanent income.

from health services, H . We choose as our baseline functional form for $\Psi(H)$ a power function with exponent $1-\gamma$, so $\Psi(H) = \varphi_2 (1-\gamma)^{-1} H^{1-\gamma}$.⁸

We assume that all individuals have health insurance that covers a fraction b of second-period health expenditures, with $0 \leq b \leq 1$; this is financed by a tax rate τ on permanent income, Y . In period 1, before the individual knows her future health status, the individual chooses how much of permanent income Y to consume in the first period and how much to save for the second period. In the second period, health is realized, and the individual chooses how to allocate her savings between nonhealth consumption and health services. The resulting budget constraint is

$$Y(1 - \tau) = C_1 + \frac{1}{1+r}(C_2 + (1-b)H), \quad (3)$$

where r denotes the real interest rate.

We solve the model backwards. If the person is healthy in the second period ($S = 0$), nonhealth consumption is simply equal to second-period wealth $W = (1+r)(Y(1-\tau) - C_1)$, hence

$$U_{2,S=0} = \frac{1}{1-\gamma} W^{1-\gamma}. \quad (4)$$

If the person is sick in the second period, she chooses nonhealth consumption and health services to solve

$$\begin{aligned} \max_{C_2, H} U_2(C_2, H) &= \max_{C_2, H} (1 + \varphi_1) \frac{1}{1-\gamma} C_2^{1-\gamma} + \varphi_2 \frac{1}{1-\gamma} H^{1-\gamma} \\ \text{subj. to: } &C_2 + (1-b)H = W. \end{aligned} \quad (5)$$

Conditional on being sick, the resulting optimal consumption and health services are given by

$$C_2 = W / \left(1 + \varphi_2^{1/\gamma} (1 + \varphi_1)^{-1/\gamma} (1-b)^{1-1/\gamma} \right) \quad (6)$$

and

$$H = \varphi_2^{1/\gamma} (1 + \varphi_1)^{-1/\gamma} (1-b)^{-1/\gamma} W / \left(1 + \varphi_2^{1/\gamma} (1 + \varphi_1)^{-1/\gamma} (1-b)^{1-1/\gamma} \right). \quad (7)$$

Substituting (6) and (7) into the second-period utility function yields second-period utility as a function of second-period wealth for sick individuals ($S = 1$):

$$U_{2,S=1} = \frac{1}{1-\gamma} (1 + \varphi_1) \left(1 + (1 + \varphi_1)^{-1/\gamma} (1-b)^{1-1/\gamma} \varphi_2^{1/\gamma} \right)^\gamma W^{1-\gamma}. \quad (8)$$

8. Again, this is done for expositional ease. In Online Appendix A, we show that a more general second-period utility function that is a CES aggregate of nonhealth consumption and health services leads to identical estimating equations as the ones we derive here. For our simulations in Section 6, we use this more general CES formulation, which breaks the relationship between the elasticity of demand for health services and risk aversion.

We calculate expected second-period utility as the weighted average of equations (4) and (8), and use the intertemporal budget constraint and the lifetime utility function to solve for C_1 and W as functions of permanent income Y . We find that optimal second-period wealth is proportional to permanent income, $W = wY$. In Online Appendix A, we derive the expression for the proportionality factor w , which turns out to be a complicated function of preference parameters ($\varphi_1, \varphi_2, \gamma, \theta, \delta$), prices (r, b, τ), and the probability of sickness (p).

Substituting $W = wY$ into equations (4) and (8) yields indirect utility $v(Y, S)$ in the second period for the healthy state and the sick state, respectively:

$$v(Y, 0) = \frac{1}{1 - \gamma} (wY)^{1-\gamma} \quad (9)$$

and

$$v(Y, 1) = \frac{1}{1 - \gamma} (1 + \varphi_1) \left(1 + \varphi_2^{1/\gamma} (1 + \varphi_1)^{-1/\gamma} (1 - b)^{1-1/\gamma} \right)^\gamma (wY)^{1-\gamma}. \quad (10)$$

These indirect utility functions suggest a nonlinear regression of the following form:

$$v = \beta_1 S \times Y^{\beta_2} + \beta_3 Y^{\beta_2} + \eta, \quad (11)$$

which yields the parameter estimates

$$\beta_1 = (1 + \varphi_1) \left(1 + \varphi_2^{1/\gamma} (1 + \varphi_1)^{-1/\gamma} (1 - b)^{1-1/\gamma} \right)^\gamma \frac{w^{1-\gamma}}{1 - \gamma} - \frac{w^{1-\gamma}}{1 - \gamma},$$

$$\beta_2 = 1 - \gamma, \quad \text{and} \quad \beta_3 = \frac{w^{1-\gamma}}{1 - \gamma}. \quad (12)$$

Interpreting the Parameter Estimates. Taking the ratio of the incremental income gradient of utility in the sick state relative to the healthy state (β_1) to the income gradient of utility in the healthy state (β_3) yields

$$\beta_1/\beta_3 = (1 + \varphi_1) \left(1 + \varphi_2^{1/\gamma} (1 + \varphi_1)^{-1/\gamma} (1 - b)^{1-1/\gamma} \right)^\gamma - 1. \quad (13)$$

Dividing equation (6) by equation (7), substituting in to equation (13) and re-arranging yields an expression that we will use to infer state dependence:⁹

$$(\varphi_1 + 1) = \frac{\beta_1/\beta_3 + 1}{(1 + m(1 - b))^\gamma}, \quad (14)$$

where $m (=H/C_2)$ is the ratio of health services to nonhealth consumption, b is the fraction of health expenditures covered by insurance, and γ is the coefficient of relative risk aversion.

9. This step implicitly calibrates the parameter φ_2 based on the observed consumption of health services, m . In the calibration exercises of Section 6, we take care to keep φ_2 constant as we vary other parameters that affect the consumption of health services. In Section 4, we discuss how we calibrate m and b for our sample.

Equation (14) shows the relationship between the state-dependence parameter in our model (φ_1) and our estimated coefficients β_1 and β_3 . When insurance is full insurance (i.e., $b = 1$), equation (14) indicates that the ratio of our estimated coefficients β_1/β_3 directly gives us our estimate of state dependence (φ_1). Moreover, under full insurance, our estimate of the coefficient β_1 , the coefficient on the interaction term between permanent income and sickness in equation (11), provides a direct test of whether (and in what direction) the marginal utility of consumption is affected by health. This motivates our sample selection of insured individuals.

When insurance is less than full (i.e., $0 \leq b < 1$), equation (14) shows that the true amount of state dependence (φ_1) is lower than our coefficient ratio β_1/β_3 , since the term $(1 + m(1 - b))^\gamma$ is weakly greater than 1. The estimate of β_1 now provides a direct test only of negative state dependence; specifically, $\beta_1 < 0$ implies that marginal utility declines as health deteriorates (i.e., that $\varphi_1 < 0$). To quantify the direct effect of sickness on the marginal utility of consumption (i.e., when consumption is held fixed), we must correct our estimate of the effect of sickness on the marginal utility of wealth (β_1/β_3) for the term that accounts for the direct effect of sickness on the level of consumption ($1 + m(1 - b)$). Intuitively, with less than full insurance, more wealth allows you to divert more second-period wealth to health consumption, so that richer people increase their utility more (in absolute terms) than poorer people in the sick state than in the healthy state, thus biasing the coefficient ratio β_1/β_3 toward positive state dependence. Given that the elasticity of the marginal utility of consumption with respect to the level of consumption equals $-\gamma$, the reduction in consumption increases the marginal utility by a factor of $(1 + m(1 - b))^\gamma$.¹⁰

Translation to Empirical Estimation. We estimate equation (11) by running a fixed-effects regression of the following form:

$$UtilityProxy_{it} = g \left(\beta_1 S_{it} \times \bar{Y}_i^{\beta_2} + \beta_4 S_{it} + X_{it} \Gamma_1 + \alpha_i \right), \quad (15)$$

where i indexes individuals and t indexes time periods. Since all individuals in our sample are older (mean age 72) and out of the labor force, we think of the time periods in the empirical analysis as repeated observations on an individual in period 2 of the theoretical model. The repeated observations serve to isolate the effect of health on the utility proxy from time-invariant omitted variables that are correlated with both health and the utility proxy. In the interest of parsimony, we did not model such omitted variables in the theoretical model, but we do control for them in the empirical analysis. The explanatory variables consist of a measure of sickness (S), a measure of permanent income (\bar{Y}), and demographic covariates (X). The individual fixed effects (α_i) absorb

10. Note that whether insurance is full or less than full, the coefficient ratio β_1/β_3 can be interpreted as the effect of sickness on the marginal utility of second-period wealth (i.e., holding second-period wealth fixed). Since the price of consumption is unity and individuals have strictly positive consumption levels, the marginal utility of second-period wealth equals the marginal utility of consumption. However, the effect of health on the marginal utility of wealth is not equal to the effect of health on the marginal utility of consumption because in the former case wealth is held constant while in the latter case consumption is held constant.

any direct effect of permanent income and any other time-invariant characteristics on utility. Finally, $\beta_1, \beta_2, \beta_4$, and the vector Γ_1 are coefficients to be estimated, and $g(\cdot)$ is a monotonically increasing mapping from latent von Neumann–Morgenstern utility to the utility proxy.¹¹

The effect of permanent income on utility in the healthy state (i.e., β_3 from equation (11)), which we use to interpret the magnitude of state dependence, is absorbed by the individual fixed effects in equation (15). We recover an estimate of β_3 by running an auxiliary regression of the estimated fixed effects from equation (15) on permanent income and demographic controls:

$$\hat{\alpha}_i = \beta_3 \bar{Y}_i^{\beta_2} + X_{it} \Gamma_2 + \xi_{it}. \quad (16)$$

3.2. Identifying Assumptions

For ease of exposition, we focus our discussion here on the case of full insurance ($b = 1$). We discuss the assumptions needed to interpret β_1 as a test of state dependence, and the additional assumptions needed to identify the magnitude of state dependence (β_1/β_3). As the discussion in the previous section pointed out, with less than full insurance ($0 \leq b < 1$) the same identifying assumptions allow us to interpret β_1 as a test of negative state dependence and the coefficient ratio β_1/β_3 as an upper bound on the magnitude of state dependence.

Two identifying assumptions are required to interpret the estimate of β_1 as a test of state dependence. First, the difference between our imposed mapping $g(\cdot)$ and the true mapping from latent von Neumann–Morgenstern utility to the utility proxy cannot vary differentially with health by permanent income. In other words, there may be error in our assumed mapping $g(\cdot)$ from latent utility to the utility proxy, but this error cannot vary systematically with health by permanent income. In particular, misspecification of the mapping $g(\cdot)$ can cause incorrect inference of the true magnitude—or even the true sign—of state dependence; this issue resembles the dissimilarity that can arise between the magnitude and sign of a marginal effect and the magnitude and sign of an interaction term in a nonlinear model (Ai and Norton 2003). With this concern in mind, we assess in great detail the sensitivity of our estimates to different assumptions and estimates of $g(\cdot)$ in Section 5.1.

Second, we assume that conditional on sickness S , control variables X , and fixed effects α , there are no omitted determinants of utility that vary with health differentially by permanent income (i.e., are correlated with $S \times \bar{Y}$). This assumption is considerably more palatable in the panel than it would be in a cross-section, where there may well be person-specific characteristics (such as optimism/pessimism) that are correlated with utility, health, and permanent income. In the panel, in the same spirit as typical identification for difference-in-differences estimation, it is only a problem for our

11. For example, in our OLS specification, $g(v)$ equals $v + \eta_{it}$, with η_{it} denoting the error term in utility. In the probit specification, $g(v)$ equals $1\{v + \eta_{it} > 0\}$, with $1\{\cdot\}$ denoting the indicator function and η_{it} denoting a standard normal error term. In Section 5.1, we explore more general specifications of $g(\cdot)$.

analysis if health changes within individuals vary across individuals of different permanent income in ways that are correlated with utility. In Section 5.2 we discuss and investigate potential threats to the validity of this identifying assumption in greater depth, including (but not limited to) the possibility that health may affect wealth, thereby violating the assumption of predetermined wealth.

Identification of β_3 requires the additional identifying assumption that, conditional on S , X , and $S \times \bar{Y}$, there are no omitted determinants of utility that are correlated with \bar{Y} . In other words, the only reason why people with higher permanent incomes have higher levels of utility is because of their higher levels of consumption, and not because of any other determinants of utility that are correlated with permanent income. We emphasize, however, that this strong assumption is needed only to provide a way of scaling (interpreting) our key parameter β_1 ; it is not fundamental to our identification of the existence of state-dependent utility.

3.3. Data and Baseline Specification

We use all cohorts in the first seven waves of the HRS, which span the years 1992 through 2004. We limit the sample to individuals (and their spouses) aged 50 and older who are not in the labor force and who have health insurance. Our baseline sample consists of 11,514 unique individuals and 45,447 person-years. Online Appendix B provides more detail on our sample and variable definitions. Table 1 presents descriptive statistics. Our sample is 63% female and 87% white, with an average age of 72. About three-fifths of person-years are married.

To operationalize equations (15) and (16) we must choose measures of the explanatory variables and a functional form. We describe our baseline choices here and later show robustness to a wide range of alternative choices. We measure an individual's permanent income (\bar{Y}_i) by the average across all waves of their total annual household income (adjusted for household composition) plus 5% of the household's current financial wealth to account for the fact that elderly households may be spending down their accumulated financial savings.¹² Average permanent income in our sample is about \$29,200. Our baseline measure of health status (S_{it}) is the number of chronic diseases that individual i in wave t has ever been told by a doctor that she has had ($NUM_DISEASE_{it}$); we code each disease as an absorbing state. Following standard practice in the HRS (see, e.g., Smith 1999), we consider the following seven diseases that are asked consistently over time: hypertension, diabetes, cancer, heart disease, chronic lung disease, stroke, and arthritis.¹³ We selected this measure as our baseline because of its widespread prior use and because it is not subjective. The latter is important because our outcome variable is subjective and having subjective variables on both sides of the equation can lead to spurious results due to correlated measurement

12. Results are unaffected if we instead assume a 10% or 0% "drawdown" rate of financial wealth.

13. We exclude psychiatric illnesses since they may affect the reporting of subjective well-being. They are often excluded from measures of chronic disease (see, e.g., Smith 1999); our results are not sensitive to including them.

TABLE 1. Descriptive statistics.

Sample: Age \geq 50 & Not in labor force & Has health insurance							
	Obs.	Mean	Std. dev.	5th percentile	Median	95th percentile	Std. dev. (within-indiv.)
Demographics							
<i>NUM_WAVE</i>	45447	4.52	1.50	2	4	7	0
\bar{Y} (Permanent income, \$)	45447	29224	33297	6236	20667	77285	0
<i>FEMALE</i>	45447	0.63	0.48	0	1	1	0
<i>NON_WHITE</i>	45447	0.13	0.33	0	0	1	0
<i>SINGLE</i>	45447	0.40	0.49	0	0	1	0.21
<i>AGE</i>	45447	72.39	9.00	57	73	87	3.28
<i>HOUSEHOLD_SIZE</i>	45447	1.99	1.00	1	2	4	0.56
Health Measure							
<i>NUM_DISEASE</i>	45447	1.95	1.30	0	2	4	0.63
Utility Proxy							
<i>HAPPY</i>	45447	0.87	0.34	0	1	1	0.28

Notes: *NUM_WAVE* is the number of waves that the respondent was interviewed. Permanent income is constructed by taking the average across all waves of total household income plus a 5% annual draw down of current financial wealth. The average is then adjusted using an OECD-style adjustment (divide by 1.0 if single, and divide by 1.7 if married and living with spouse). Household size includes all residents of household (including spouse). *NUM_DISEASE* is the number of diseases reported by the respondent. This variable is calculated as the sum of the answers to the Yes/No questions “Has a doctor ever told you have d?”, where the set of diseases is the following: {hypertension, diabetes, cancer, heart disease, chronic lung disease, stroke, arthritis}. *HAPPY* is a dummy variable for Yes/No response to the following statement: “Much of the time the past week I was happy.”

error in the subjective variables. On average, a person in our sample has 1.95 diseases; the within-person standard deviation in number of diseases (which is our key right-hand side variation) is 0.625.

Our utility proxy is the response to the question: “Much of the time during the past week I was happy. (Would you say yes or no?)” We code this as an indicator variable *HAPPY* in which an affirmative answer is given the value 1 and consider this to be an index of latent utility.¹⁴ On average, 87% of person-years respond in the affirmative. The within-person standard deviation in the response (which is our key left-hand side variation) is 0.28. As is standard in the happiness literature, we interpret the happiness question as a proxy for von Neumann–Morgenstern (flow) utility.¹⁵ We recognize that this proxy is not perfect. For example, Benjamin et al. (2012) show that, while predicted

14. The happiness question is asked as part of eight questions from the CESD depression scale. Within the set of CESD questions, the happiness questions follows a question on restless sleep except in 1994, when it follows a question about feeling sad. The CESD module comes after questions about height and weight in the 1992–1996 waves and after word-recall questions in the 1998–2004 waves. In all waves, questions on health conditions are asked somewhere in the survey before the CESD module, and income questions come somewhere after the CESD module.

15. There is no strong evidence establishing whether happiness is closer to flow utility or lifetime utility (i.e., the PDV of current and future flow utilities). The latter interpretation seems inconsistent with the empirical finding that happiness increases with age for older people (unless one believes growing older means fewer future years with negative flow utility). Kimball and Willis (2006) propose that happiness is a combination of “elation”, which comes from news about innovations to lifetime utility, and “baseline mood”, which is a component of current flow utility. If news about the health shocks has been processed by

happiness has been shown to be a powerful predictor of hypothetical choices, there are also systematic differences between hypothetical choices and predicted happiness.¹⁶

The covariates (X_{it}) in equation (15) are wave fixed effects, age, age squared, household size, and an indicator for whether the individual is currently single; for equation (16) we use these same covariates and also add controls for race and gender, which are absorbed in (15) by the individual fixed effects; we do this so as not to confound the relationship between permanent income and SWB in (16) with demographics that are correlated with permanent income.

For our baseline specification, we assume a linear mapping $g(\cdot)$ and a coefficient of relative risk aversion of 1 ($\beta_2 \rightarrow 0$ and thus $\gamma \rightarrow 1$). We relax both assumptions in Section 5.1.

4. Main Results

4.1. Baseline Results

Our baseline estimates of equation (15) and the auxiliary regression equation (16) are shown in Table 2, column (1), Panel A. As expected, our utility proxy is increasing in permanent income ($\beta_3 > 0$) and decreasing in number of diseases ($\beta_4 < 0$). Because \bar{Y} is demeaned, the coefficient β_4 on S_{it} of -0.011 (s.e. = 0.003) indicates that, for someone of average permanent income, an increase of one chronic disease is associated with a statistically significant 1.1 percentage point decline in the probability the individual is happy; this is identified within-person using variation in individuals' health state over time. The coefficient on log permanent income (β_3) of 0.048 indicates that a 10% increase in permanent income is associated with a 0.48 percentage point increase in the probability an individual reports that he is happy most of the time in the past week (off of a mean of 87%). Of course, this cross-sectional comparison of happiness across individuals of different permanent incomes may conflate the causal effect of permanent income with the effects of other characteristics of high-permanent-income individuals that are themselves determinants of subjective well-being. For this reason, we do not rely on such cross-sectional variation to identify the existence of state-dependent utility (i.e., β_1).

The key coefficient of interest is β_1 , which we estimate to be -0.009 (s.e. = 0.004). The negative sign indicates that the marginal utility of permanent income declines as health worsens. It is statistically significant at the 5% level, which implies that we reject the null of state-independent utility. Note that even if insurance is partial, under

the time our respondents are interviewed, there should be no systematic differences in changes in elation by income, and our estimates would be unbiased.

16. In particular, Benjamin et al. (2012) find that individuals are more sensitive to money in hypothetical choices than in predicted self-rated happiness. Because we estimate state dependence as a relative change in the income-happiness gradient, the reduced sensitivity to money in happiness ratings will not affect our results as long as this reduction is proportional. If the reduced sensitivity is not proportional, our estimates will be biased but the sign remains correct.

TABLE 2. Estimated magnitude of state-dependent utility.

Dependent Variable: <i>HAPPY</i>			
Sample restrictions:		(1) Age ≥ 50 , NILF, Has Health Insurance	(2) Age ≥ 65 , NILF, Has Health Insurance
Panel A: Estimates			
$NUM_DISEASE_{it} \times \log(\bar{Y}_i): \beta_1$		-0.009 (0.004) [0.018]	-0.008 (0.004) [0.048]
$\log(\bar{Y}_i): \beta_3$		0.048 (0.003) [0.000]	0.038 (0.003) [0.000]
$NUM_DISEASE_{it}: \beta_4$		-0.011 (0.003) [0.001]	-0.013 (0.004) [0.000]
R^2		0.474	0.470
N		45447	37829
Number of individuals		11514	10108
Panel B: Implied State Dependence Assuming Full Insurance ($b = 1$)			
<i>% change in marginal utility for a 1 unit increase in</i> $NUM_DISEASE_{it} (\beta_1/\beta_3)$		-17.9%	-21.0%
(bootstrapped standard error)		(5.8%)	(8.2%)
(bootstrapped 95% CI)		(-4.2%, -26.9%)	(-0.5%, -32.5%)
[bootstrapped p-value]		[0.018]	[0.048]
<i>% change in marginal utility for a 1 standard deviation</i> <i>increase in</i> $NUM_DISEASE_{it} (\sigma \beta_1/\beta_3)$		-11.2%	-13.4%
(bootstrapped standard error)		(3.6%)	(5.2%)
(bootstrapped 95% CI)		(-2.7%, -16.8%)	(-0.3%, -20.8%)
[bootstrapped p-value]		[0.018]	[0.048]
Marginal utility with 0 diseases (13.1%, 12.7% of samples)		0.048	0.038
Marginal utility with 1 disease (26.7%, 26.8% of samples)		0.039	0.030
Marginal utility with 2 diseases (28.8%, 29.1% of samples)		0.030	0.022
Marginal utility with 3 diseases (18.8%, 19.4% of samples)		0.021	0.014
Marginal utility with 4 diseases (9.1%, 8.9% of samples)		0.012	0.005
(3.5%, 3.2% of samples have 5+ diseases)			

Notes: **Panel A** reports coefficients β_1 and β_4 from estimating equation (15) and coefficient β_3 from estimating equation (16). Both equations are estimated using OLS. The dependent variable in equation (15) is a dummy variable that equals 1 if the respondent agrees with the following statement: "Much of the time during the past week I was happy." The mean of the dependent variable is 0.87. The dependent variable in equation (16) is the individual fixed effect estimated in equation (15). $NUM_DISEASE_{it}$ is a composite health measure which is the total number of reported diseases. $\log(\bar{Y}_i)$ is the average across all waves of the individual's total annual household income, adjusted for household composition and a 5 percent annual draw down of current financial wealth. We demean $\log(\bar{Y}_i)$ before interacting with $NUM_DISEASE_{it}$. In addition to the covariates shown in the table, the estimates of equation (15) include individual fixed effects, wave fixed effects, and controls for Age, Age², Household size, and a dummy for whether the individual is single; estimates of equation (16) include wave fixed effects and a non-white dummy, a female dummy, Age, Age², Household size, and a dummy for being currently single. Standard errors for β_1 , β_3 , and β_4 are in brackets. Online Appendix B contains more detail on data and variable definitions; summary statistics are in Table 1.

Panel B reports several ways of interpreting the estimates in Panel A. In Panel B we report bootstrapped standard errors, 95% confidence intervals, and p -values for β_1/β_3 and $\sigma \beta_1/\beta_3$ based on 10,000 bootstrap iterations, resampling individuals with replacement. The within-person standard deviation change in $NUM_DISEASE_{it}$ (σ) is 0.625 in column (1) and 0.634 in column (2).

our identifying assumptions our negative estimate of β_1 provides evidence of negative state dependence.

Panel B of Table 2 reports several simple ways to interpret the magnitude of our estimate of β_1 under the assumption that individuals are fully insured (i.e., $b = 1$),

so that state dependence is given by the ratio β_1/β_3 . Recall that with less than full insurance the implied state dependence is more negative than this estimated parameter ratio. We find, for example, that a one-standard-deviation increase in the number of chronic diseases is associated with an 11.2% decline in marginal utility for a previously healthy individual (i.e., $\sigma\beta_1/\beta_3 = -11.2\%$), with a 95% confidence interval ranging from 2.7%–16.8%.¹⁷ We also find that marginal utility falls from 0.048 for a healthy (disease-free) individual, to 0.039 with one disease, 0.030 with two diseases, 0.021 with three diseases, and 0.012 with four diseases (only 3.5% of our sample has more than four diseases); this is the regression-based counterpart to the stylized picture in Figure 1A of utility curves “fanning out”.

An important potential concern with our baseline estimates in column 1 of Table 2 is that our sample selection criteria of individuals who are not in the labor force may bias us towards finding negative state dependence since those most likely to exit the labor force due to poor health may be precisely those for whom the marginal benefit of work (i.e., the marginal utility of consumption) falls more with poor health. This potential sample selection issue is much less of a concern when looking at a sample of individuals who are unlikely to work regardless of health. Therefore, in column 2 we report results in which we further limit our baseline sample to individuals who are 65 and older (as opposed to 50 and older); this sample is also insured and not in the labor force. While only 53.2% of individuals 50 and older who are insured are not in labor force, that number rises to 79.8% for 65 and older, suggesting that sample selection is less of a concern. The results in column 2 are reassuring. Our estimate of the decline in marginal utility of consumption associated with a one-standard-deviation increase in the number of chronic diseases rises from 11.2% with full insurance in our baseline specification (column 1) to 13.4% with full insurance when the baseline sample is further restricted to individuals 65+ (column 2).

4.2. Implied State Dependence under Partial Insurance

Table 3 reports the implied state dependence from our estimates when we do not assume full insurance. As seen in equation (14), the adjustment factor needed to translate the coefficient ratio β_1/β_3 into the implied state dependence φ_1 depends on out-of-pocket health expenditure as a share of nonhealth consumption $m(1 - b)$ and the coefficient of relative risk aversion γ . We use MEPS data to calculate the out-of-pocket shares (see Online Appendix B for details).¹⁸ We report the implied state dependence for

17. Because of our two-step estimation procedure, we bootstrap the sample at the individual level (10,000 replications) to calculate standard errors, p -values, and confidence intervals of our estimates of β_1/β_3 and $\sigma\beta_1/\beta_3$ in this table and in all subsequent tables.

18. Since H is the incremental out-of-pocket health spending associated with becoming sick, we approximate it using data from the 1996 MEPS limited to those who meet our sample selection criteria. Our approximation of H is based on the difference in mean out-of-pocket health spending for those whose medical spending is above the median and those whose medical spending is below the median. We scale this difference in mean out-of-pocket health spending by mean annual consumption, which is \$41,648 using the consumption data in the HRS CAMS survey (described in more detail in Online Appendix

TABLE 3. Estimated magnitude of state-dependent utility.

		Dependent Variable: <i>HAPPY</i>								
		(1a)	(1b)	(1c)	(2a)	(2b)	(2c)	(3a)	(3b)	(3c)
Sample restrictions:		Age ≥ 50, NILF, Has Health Insurance [Baseline Sample]		Age ≥ 65, NILF, Has Health Insurance		Age ≥ 65, NILF, Has Health Insurance + (Medicaid, VA, or Medicare HMO)				
Coefficient of relative risk aversion, γ :		$\gamma = 1$	$\gamma = 3$	$\gamma = 5$	$\gamma = 1$	$\gamma = 3$	$\gamma = 5$	$\gamma = 1$	$\gamma = 3$	$\gamma = 5$
Panel A: Implied State Dependence Assuming Full Insurance ($b = 1$)										
% change in marginal utility for a 1 standard deviation increase in $NUM_DISEASE_{it}$ ($\sigma\beta_1/\beta_3$)		-11.2% [0.018]	-11.4% [0.012]	-11.9% [0.050]	-13.4% [0.048]	-14.2% [0.020]	-14.4% [0.105]	-12.0% [0.106]	-13.5% [0.048]	-13.8% [0.120]
Panel B: Implied State Dependence Assuming Partial Insurance ($b < 1$)										
Out-of-pocket health expenditure share, $m(1-b)$		0.023								
% change in marginal utility for a 1 standard deviation increase in $NUM_DISEASE_{it}$ ($(\sigma\beta_1/\beta_3 + 1)/(1 + m(1 - b))\gamma - 1$)		-13.2% [0.008]	-17.1% [0.003]	-21.2% [0.014]	-15.5% [0.022]	-20.2% [0.008]	-24.1% [0.028]	-13.6% [0.064]	-18.1% [0.026]	-21.3% [0.049]
R^2		0.474	0.474	0.474	0.470	0.470	0.470	0.595	0.595	0.595
N		45447		37829		10108		10537		3056
Number of Individuals										

Notes: Panel A reports the implied magnitude of state dependence under full insurance, and Panel B reports the implied magnitude of state dependence under partial insurance. The within-person standard deviation change in $NUM_DISEASE_{it}$ (σ) is 0.625, 0.634, and 0.559 for the three samples, respectively. The degree of insurance coverage (b) is 85.1%, 84.9%, and 91.7% for the three samples, respectively. See Table 2 for more details. We report bootstrapped p -values based on 10,000 bootstrap iterations, resampling individuals with replacement. See text for details.

coefficients of relative risk aversion of 1, 3, and 5,¹⁹ and use the functional form of our estimating equations (15) and (16) that is consistent with the assumed coefficient of risk aversion (i.e., we set β_2 equal $1 - \gamma$ for $\gamma > 1$).

Columns (1a)–(1c) of Table 3 show the sensitivity of the results in our baseline sample to the assumed coefficient of relative risk aversion. Panel A shows the results assuming full insurance, so that the parameter ratio $\sigma \beta_1 / \beta_3$ gives the implied change in the marginal utility of consumption associated with a one-standard-deviation increase in the number of chronic diseases. The estimates are not sensitive to the assumed degree of risk aversion; they range from -11.2% to -11.9% as we vary the curvature on permanent income in our estimating equation to be consistent with relative risk aversion between 1 and 5. In Panel B, we allow for partial insurance and we calculate the implied state dependence according to equation (14) when out-of-pocket health expenditure as a share of nonhealth consumption equals the observed value of 2.3% for the baseline sample. For a coefficient of relative risk aversion of 3, the estimates imply that a standard deviation increase the number of chronic diseases reduces the marginal utility of consumption by 17.1% . The implied state dependence under partial insurance, however, is much more sensitive to assumptions about risk aversion; it ranges from -13.2% for $\gamma = 1$ to -21.2% for $\gamma = 5$. Columns (2a)–(2c) show that results remain similar if we restrict the baseline sample to those aged 65 or more.

To reduce the importance of the adjustment factor in our estimated state dependence for partial insurance, columns (3a)–(3c) of Table 3 shows results for a subsample with a higher degree of insurance and a lower out-of-pocket share. The sample in columns (3a)–(3c) is the same as the sample in columns (2a)–(2c), except that we further require the observations to have their health insurance, which is Medicare for 97% of the sample, supplemented by some other insurance (Medicaid, VA, or Medicare HMO). We calculate the out-of-pocket share for this sample at 1.9% compared to 2.5% for the full sample in columns (2a)–(2c) and 2.3% for the baseline sample in columns 1a–1c. The estimates in panel B remain quite similar to our baseline estimates, with our estimate of state dependence ranging from -13.6% for $\gamma = 1$ to -21.3% for $\gamma = 5$.

Taken together, the estimates of Table 3 show clear evidence of significant negative state dependence. The exact magnitude of the negative state dependence is sensitive to assumptions about the degree of insurance and the coefficient of relative risk aversion. However, under a wide variety of plausible values for these assumptions, the point estimate of state dependence ranges between about -10% and -25% .

In the interest of parsimony, we conduct further analyses assuming that $\gamma = 1$ unless indicated otherwise. To facilitate comparability across specifications, our discussion of magnitude focuses on the implied percentage change in marginal utility for a healthy person associated with a one-standard-deviation decline in health (i.e.,

B). Consumption in the CAMS is calculated on a household basis, so we convert consumption to an individual-level measure using the OECD adjustment for household composition (see Online Appendix B for details).

19. There is, of course, little consensus on the coefficient of relative risk aversion. We use a range of values corresponding to commonly used values in calibration models of insurance and savings.

$\sigma\beta_1/\beta_3$), as this provides a scale-free way of comparing different estimates. Since the results in Panel A of Table 3 indicated that, under the assumption of full insurance, the implied magnitude of state dependence is not sensitive to the assumed level of risk aversion, interested readers can adjust our reported magnitude ($\sigma\beta_1/\beta_3$) to account for partial insurance using their preferred estimate of risk aversion and the degree of insurance.

4.3. *Additional Analyses*

We performed a large number of additional sensitivity analyses on our baseline results from column (1) of Table 2. A full set of results from these additional sensitivity analyses are presented in Online Appendix F; we briefly summarize the main ones here.

We find that the sign of our results (and in general the estimated magnitude) is not sensitive to a series of natural alternative specifications including excluding the demographic controls (X_{it}), or restricting the sample to individuals who are always single (for whom there is no concern about potential correlations in health changes within a couple confounding the estimate of own health on own marginal utility). We also find no evidence of habituation or adaptation effects; the decline in marginal utility after a negative health shock does not appear to diminish over time, but we cannot rule out sizable adaptation effects either. The literature on the existence and size of adaptation to health shocks is mixed (see, e.g., Lucas 2007, Oswald and Powdthavee 2008). To the extent there is habituation, our estimates already reflect habituation that occurs on average during the years we observe individuals after a change in health status. Any further habituation causes our estimates to overstate the degree of state dependence in the long run.

Another set of sensitivity analyses investigates whether the impact of disease on the marginal utility of consumption varies across our seven individual diseases. Not surprisingly, looking disease-by-disease reduces our statistical power. Indeed, we estimate statistically significant state dependence only for blood pressure and lung disease. Nonetheless, with the exception of heart disease and arthritis, the point estimates on the interaction terms are uniformly negative; moreover, we are unable to reject at the 10% level the hypothesis that all seven interaction terms are equal (p -value = 0.14). A priori, one might expect a stronger response in marginal utility to symptomatic diseases than to asymptomatic diseases, but we cannot reject that the effect of an additional disease is the same for symptomatic and asymptomatic diseases (p -value = 0.60).²⁰ Finally, we examined whether the magnitude of the drop in marginal utility from an additional disease depends on the number of diseases that the individual already has. We find no evidence of such nonlinearities and cannot reject the hypothesis that the effect of an additional disease on the marginal utility is the same for each number of pre-existing diseases (p -value = 0.36).

20. We classify lung disease, stroke, arthritis, and cancer as symptomatic diseases and high blood pressure, heart disease, and diabetes as asymptomatic diseases.

A final set of sensitivity analyses investigated the sensitivity of our results to alternative measures of our key variables. Specifically, we investigated how results are affected by alternative measures of permanent income (such as education or wealth), alternative standard measures of health (such as limitations to activities of daily living), and alternative utility proxies derived from the CESD depression scale. Our finding of negative state dependence persists across all the alternative measures, but the estimated magnitudes vary. Overall, the estimated magnitudes tend to lie centered around (and are relatively similar to) our baseline estimate.

5. Threats to Identifying Assumptions

We discuss and explore potential threats to each of the two key identifying assumptions discussed in Section 3.2 that allow us to interpret the estimate of β_1 as a test of state dependence. Our reading of the results is that they greatly alleviate concerns that violations of our identifying assumptions are responsible for our finding of negative state-dependent utility.

5.1. *Differential Errors in the Mapping from Utility to Subjective Well-Being*

The first identifying assumption is that there are no errors in our specification of the mapping $g(\cdot)$ from latent von Neumann–Morgenstern utility to the utility proxy that vary differentially with health by permanent income. There is growing evidence that measures of self-reported well-being are meaningful measures of utility: for example, people who rate themselves as happy are more likely to be rated happy by others, self-reports of happiness correlate in the expected direction with objective life circumstances, and happier people are less likely to commit suicide.²¹ Of course, there is doubtless considerable measurement error in SWB measures as utility proxies; for example, it is well documented that answers to these questions can be sensitive to wording, framing, or question order, among other factors (see, e.g., Bertrand and Mullainathan 2001).

However, random (even non-mean-zero) measurement error in SWB does not bias our estimate of state dependence. What is required is that a given *change* in true underlying von Neumann–Morgenstern utility associated with a given *change* in health must map into the same *change* in the latent variable corresponding to our proxy for utility at different levels of permanent income. To see how the shape of the mapping

21. See Frey and Stutzer (2002), Di Tella and MacCulloch (2006), or Krueger et al. (2009) for more extensive discussions of the validity and usefulness of happiness data in economic research. Additionally, in our context, an important advantage of using proxies for utility to infer state-dependent utility is that they do not require individuals to accurately forecast how their utility function will change once they become ill. Loewenstein, O'Donoghue, and Rabin (2003) and Conlin, O'Donoghue, and Vogelsang (2007) show that individuals suffer from "projection bias" by putting too much weight on their current preferences when forecasting their future preferences. This would imply that healthy people underestimate the effect of sickness on their marginal utility and, as a result, do not sufficiently adjust their demand for insurance or their consumption behavior.

$g(\cdot)$ from true utility to the utility proxy may cause this assumption to fail, suppose that the true mapping from utility to the utility proxy is quadratic, but that we use a linear specification for $g(\cdot)$ in our estimation. Suppose that the onset of an illness causes true utility for high-income individuals to drop from 5 to 4 and causes true utility for low-income individuals to drop from 2 to 1; in this example the utility function is state-independent. If the true mapping is quadratic, the onset of disease would cause the utility proxy to drop by 9 (from 25 to 16) for high-income individuals but only by 3 (from 4 to 1) for low-income individuals. If we were to use a linear specification of $g(\cdot)$ in our estimation, we would wrongly infer negative state dependence.

Misspecification of $g(\cdot)$ led to an incorrect inference in this example because the high-income individual had higher utility levels than the low-income individual so that different ranges of the mapping function were relevant for each of them. However, if we were to compare the effect of the onset of a disease within bins of high-income and low-income individuals with same mean levels of utility, *any* monotonically increasing mapping $g(\cdot)$ from true utility to the utility proxy will produce correct inference about the sign of state dependence. This insensitivity arises because within such bins, high- and low-income individuals are on average at the same range of the $g(\cdot)$ function. While misspecification in $g(\cdot)$ will lead to errors that vary by utility, by construction utility does not vary with permanent income within the bin, and these errors in the estimated effect of the onset of a disease are therefore not correlated with permanent income.

This observation forms the basis for our nonparametric test of state dependence. Although high-income individuals have higher utility levels on average than low-income individuals in our data set, there is considerable overlap in the levels of the utility proxy by income because of variation across individuals in other (non-income-related) determinants of utility. This considerable overlap allows us to implement our nonparametric test.

For each bin based on mean values of the utility proxy and mean values of the number of diseases,²² we run an OLS regression of the utility proxy on number of diseases, permanent income, and an interaction of the number of diseases with permanent income.²³ Since the number of observations in these narrowly defined bins is low, the estimates of the coefficients on the interaction term are imprecise. Still, under the null hypothesis of no state dependence, the t -statistics of these coefficients follow a t -distribution with the number of degrees of freedom determined by the number of observations in each bin. The CDF of each t -statistic is a draw from the uniform (0,1) distribution under the null hypothesis of no state dependence, and therefore the mean of the CDFs should be 0.5. If the mean of the CDFs is significantly lower than 0.5, we can reject the null of no state dependence in favor of the alternative of negative state dependence.

22. To ensure that we are comparing individuals with similar levels of health, we further define the bins by mean values (over time) of the number of diseases.

23. We also include a control for permanent income; however, since permanent income does not vary over time and the mean level of the utility proxy has virtually no variation within bins by construction, the coefficient on permanent income is not typically identified.

TABLE 4. Nonparametric tests of state dependence.

	(1)	(2)	(3)	(4)	(5)
Bin width for mean happiness	0.1	0.05	0.2	0.1	0.1
Bin width for mean number of diseases	0.5	0.5	0.5	0.2	1.0
Number of cells	140	280	70	350	70
Number of usable cells	85	118	56	155	47
Number of individuals in usable cells	3250	3227	3256	3147	3254
Average number of individuals per used cell	38	27	58	20	69
<i>Results for interaction term, NUM_DISEASE_{it} × log(\bar{Y}_i)</i>					
Mean of CDFs of <i>t</i> -statistics (<i>F</i>)	0.402	0.420	0.383	0.420	0.393
<i>p</i> -value of test that <i>F</i> = 0.5	[0.003]	[0.003]	[0.004]	[0.001]	[0.022]
Fraction of estimated interaction terms < 0	0.624	0.576	0.661	0.600	0.660
<i>Results for direct effect of NUM_DISEASE_{it}</i>					
Mean of CDFs of <i>t</i> -statistics (<i>F</i>)	0.335	0.363	0.340	0.383	0.313
<i>p</i> -value of test that <i>F</i> = 0.5	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Fraction of estimated direct terms < 0	0.718	0.669	0.696	0.606	0.723

Notes: This table reports results of nonparametric tests as described in Section 5.1, using the sample in column (1) of Table 2 (i.e., Age ≥ 50, NILF, and Has Health Insurance). The estimation sample includes only individuals with variation in the variable *HAPPY*; i.e., we drop all individuals with mean values of *HAPPY* of either 0 or 1 during the sample period. Individual mean values of *HAPPY* and mean values of *NUM_DISEASE_{it}* are used to group individuals into bins. Bins for mean happiness and for mean number of diseases are defined to be inclusive below and exclusive above. The maximum number of diseases is 7, so the number of cells is 7/(*hd*), where *h* is bin width for happiness and *d* is bin width for number of diseases. The number of usable cells is the total number of cells with more than one individual. For each bin, we run a regression of *HAPPY* on log of permanent income, *NUM_DISEASE_{it}*, and their interaction. The cumulative distribution of the *t*-statistics on the interaction term is computed, and we report the *p*-value from a two-sided test of whether the average of these *t*-statistics is equal to 0.5 to test for state dependence in brackets. We repeat the same procedure for the *t*-statistics for number of diseases.

Table 4 shows the results of the nonparametric test. Column 1 shows our baseline case in which we have sorted individuals into ten categories based on their mean level of happiness and into 14 categories based on their mean number of diseases. Of the resulting 140 bins, 85 have sufficient observations that we can estimate a *t*-statistic on the interaction term between number of diseases and permanent income. A majority (62%) of these *t*-statistics is negative, which provides a first indication of negative state dependence. Moreover, the mean of the CDFs of the *t*-statistics is 0.402, which is significantly different from the expected value of 0.5 under the null hypothesis of no state dependence (*p* = 0.003). Columns (2)–(5) show that we obtain similar results for alternative choices for the bin sizes. Hence, the nonparametric test indicates that our finding of negative health state dependence is not driven by differential errors in the mapping from von Neumann–Morgenstern utility to the utility proxy.

This nonparametric test allows us to flexibly estimate the sign of health state dependence but does not provide information on the magnitude of state dependence. We therefore also examine the robustness of our quantitative estimates to alternative

parametric assumptions or estimates of $g(\cdot)$. Table 5 summarizes the results from three different classes of approaches to investigating the sensitivity of our quantitative estimates to potential misspecifications of $g(\cdot)$.

As a first step, we present results from a fixed effects probit model.²⁴ In our baseline specification, the linear probability model implicitly assumes that the probability that an individual responds “Yes” to the happiness question is a cardinal measure of true utility. In the probit specification, by contrast, the implicit assumption is that the latent variable in the probit model is a cardinal measure of true utility. Our sample size is reduced because we cannot include individuals who never change their response to the happiness question, and this will mechanically increase the magnitude of the estimate. Hence, to assess the sensitivity of our result to the specification of the mapping function of utility to subjective well-being, the estimate from the probit specification (column (2)) should be compared to column (3), where we report our baseline (linear) specification on the sample used for the fixed-effects probit. Both estimates are statistically significant and similar in magnitude (−28.7% versus −30.7%).

Our second approach is to assess the plausibility of the assumed mapping by estimating the level of risk aversion (curvature) implied by our mapping function. In our baseline specification, we imposed a coefficient of relative risk aversion of 1. We can instead estimate the implied curvature by estimating equation (16) by nonlinear least squares; the coefficient of relative risk aversion γ is given by one minus the exponent on permanent income (i.e., $\gamma = 1 - \beta_2$).²⁵ We conduct this exercise for both the linear and probit mappings. Column (4) shows that the linear mapping function implies a coefficient of relative risk aversion of 2.1, and column (5) shows that the fixed effects probit specification implies a coefficient of relative risk aversion of 3.5. Hence, both the linear and the probit mappings imply latent utility measures that contain plausible levels of curvature.²⁶ Moreover, when we estimate the risk aversion parameter in columns (4) and (5), our estimates of negative state dependence remain statistically significant and are very close in magnitude to the linear and probit specifications in columns (1) and (2), which imposed a coefficient of relative risk aversion of 1.

Of course, these findings do not rule out the possibility that alternative mappings exist which are also consistent with reasonable levels of risk aversion but which would yield different estimates of the magnitude of state dependence. Our third approach

24. Because of the large number of fixed effects in the nonlinear model, we follow the algorithm in Heckman and MaCurdy (1980) to maximize the log-likelihood function. We report marginal effects evaluated at the mean, but we use the original probit coefficients to calculate the magnitude of state dependence (i.e., $\sigma\beta_1/\beta_3$).

25. We estimate the coefficient of relative risk aversion using the cross-sectional variation in subjective well-being and permanent income because we do not have variation in permanent income within individuals. We then use our estimate of β_2 from the cross-sectional equation (16) in our panel equation (15). We iterate until β_2 converges.

26. While there is hardly a consensus in the literature on the magnitude of the coefficient of relative risk aversion (see e.g. Cohen and Einav 2007 for a review of the range of existing estimates), our estimates are well within the range of the most commonly used values. They are somewhat higher than those by Layard, Mayraz, and Nickell (2008), who estimate the coefficient of relative risk aversion based on the curvature of the relationship between subjective well-being and income in six different surveys. They posit a linear mapping and find estimates of the coefficient of relative risk aversion ranging from 1.19 to 1.30.

TABLE 5. Robustness to alternative mappings from cardinal utility to the utility proxy.

		Sample restrictions: Age ≥ 50 , NILF, Has Health Insurance							
Coefficient of relative risk aversion (γ):		(1) $\gamma = 1$	(2) $\gamma = 1$	(3) $\gamma = 1$	(4) γ is estimated	(5) $\gamma = 1$	(6) $\gamma = 1$	(7) $\gamma = 3$	(8) $\gamma = 5$
Specification:		OLS	FE Probit	OLS (drop all 0s and all 1s)	OLS w/ NLLS	FE Probit w/ NLLS	Semiparametric Fixed Effects Probit		
$NUM_DISEASE_{it} \times \bar{Y}_t^{1-\gamma} / (1-\gamma); \beta_1$		-0.009 (0.004) [0.018]	-0.027 (0.010) [0.008]	-0.027 (0.011) [0.011]	-0.225 (0.071) [0.002]	-22.10 (11.04) [0.045]	-0.013 (0.008) [0.123]	-15.17 (4.86) [0.002]	-582.9 (174.6) [0.001]
$\bar{Y}_t^{1-\gamma} / (1-\gamma); \beta_3$		0.048 (0.003) [0.000]	0.059 (0.005) [0.000]	0.056 (0.005) [0.000]	1.253 (0.032) [0.000]	57.94 (2.56) [0.000]	0.045 (0.020) [0.025]	34.86 (10.70) [0.001]	1470.9 (397.6) [0.001]
$NUM_DISEASE_{it}; \beta_4$		-0.011 (0.003) [0.001]	-0.038 (0.010) [0.000]	-0.036 (0.010) [0.000]	-0.010 (0.003) [0.001]	-0.111 (0.027) [0.000]	-0.148 (0.053) [0.005]	-0.132 (0.027) [0.000]	-0.143 (0.019) [0.002]
$\gamma: 1 - \beta_2$					2.100 (0.148) [0.000]	3.464 (0.467) [0.000]			
R^2		0.474		0.167	0.474				
N		45447	14040	14040	45447	14040	14040	14040	14040
Number of individuals		11514	3260	3260	11514	3260	3260	3260	3260
Within-person std dev of $NUM_DISEASE_{it}$		0.625	0.638	0.638	0.625	0.638	0.638	0.638	0.638
(σ)									
% change in marginal utility for a 1 std. dev. change in $NUM_DISEASE_{it}$		-11.2% [0.018]	-28.7% [0.012]	-30.7% [0.016]	-11.2% [0.007]	-24.3% [0.002]	-19.2% [0.113]	-27.7% [0.000]	-25.3% [0.004]
$(\sigma \beta_1 / \beta_3)$									

Notes: Column (1) reports the results from the baseline specification in Table 2; see notes to Table 2 (Panel A) for more details. Subsequent columns report results using different specifications and estimation methods. Column (2) reports marginal effects (at the mean) from a fixed effects probit specification. Column (3) reports the results from the baseline specification with sample used in column (2). Column (4) iteratively estimates the coefficient of relative risk aversion (CRRA) using a cross-sectional, non-linear, least squares (NLLS) regression using the estimated fixed effects from the baseline specification. The estimated CRRA coefficient is plugged back into the baseline specification to generate a new set of estimated fixed effects; this procedure is iterated until convergence. Column (5) does the same iterative estimation procedure but uses a fixed effects probit specification instead of the baseline (linear probability model) specification. Columns (6) through (8) report results from a semiparametric approach; see online Appendix D for more details. Columns (6) through (8) drop the Age, Age², Household Size, and Single control variables, which are included in all other columns. Standard errors are in parentheses and are adjusted to allow for an arbitrary variance-covariance matrix for each individual over time; p -values are in brackets. In columns (1) through (5), the p -value for $\sigma \beta_1 / \beta_3$ is bootstrapped base on 10,000 iterations, resampling individuals with replacement. Because of computational complexity, in columns (6) through (8) the p -value for $\sigma \beta_1 / \beta_3$ is based on an asymptotic t -test using the t -statistic derived from the point estimate of $\sigma \beta_1 / \beta_3$ divided by the bootstrapped standard error of this ratio, based on 100 bootstrap iterations. In columns (6) through (8), all standard errors are bootstrapped standard errors based on 100 bootstrap iterations, resampling individuals with replacement.

uses semiparametric estimation to further explore this possibility. Here, we once again impose a value for the coefficient of relative risk aversion, but we now flexibly estimate the monotone mapping $g(\cdot)$ such that the latent utility variable exhibits the assumed degree of risk aversion. Online Appendix D provides the technical details of this procedure. Then, using this estimate of the mapping $g(\cdot)$, we estimate the sign and magnitude of health state dependence.

Columns (6)–(8) show the estimates for state dependence when we select the mapping $g(\cdot)$ such that the coefficient of relative risk aversion in latent von Neumann–Morgenstern utility equals 1, 3, and 5, respectively. As is in the case with the standard fixed-effects probit regression, individuals who never change their response to the happiness question must be excluded from the sample. The estimates in columns (6)–(8) should therefore be compared to the estimate of the standard fixed effects probit regression in column 2. We find that our estimate of state dependence ranges between -19% and -28% for these three mappings, which is roughly similar to the estimate of -29% in the standard fixed-effects probit regression.

Overall, we believe the evidence greatly alleviates concerns that our results are primarily due to errors in the specification of the mapping from von Neumann–Morgenstern utility to the utility proxy that are correlated with the interaction between permanent income and health, although we recognize that it is impossible to completely rule out this concern. We suspect that the estimates of state dependence turn out to be relatively insensitive to the specification of mapping $g(\cdot)$ because, in practice, there is a large amount of overlap in utility levels for individuals of different levels of permanent income; this presumably reflects the importance of other individual-specific factors uncorrelated with income in determining utility.

5.2. *Omitted Determinants of Utility that Vary with Health by Permanent Income*

We also explore possible threats to our second identifying assumption that, conditional on sickness S , control variables X , and fixed effects α , there are no omitted determinants of utility that vary with health differentially by permanent income. We consider three possible threats: (i) differential trends over time in utility by permanent income, (ii) differential disease reporting by permanent income, and (iii) failure of the assumption of predetermined wealth. Results summarized but not presented here are discussed and presented in Online Appendix F.

5.2.1. Differential Trends over Time in Utility by Permanent Income. If the consumption path of the poor increases more (or declines less) than that of the rich, this tendency could show up in our estimates as negative state dependence. Since the number of diseases increases over time, it could look like the rich have a greater drop in utility with the onset of a disease simply due to different trends in underlying utility. Reassuringly, the evidence suggests that, if anything, the consumption path of the poor declines (in percentage terms) relative to that of the rich over time, which

would bias us against our finding negative state dependence.²⁷ A related issue is that our estimates of the differential effect of health changes by permanent income may in part capture differential effects of other time-varying covariates by permanent income. Reassuringly, when we allow the effect of permanent income in our baseline specification to vary not only with number of diseases but also with marital status and with household size, our estimate of state dependence remains similar in magnitude, although statistically less precise. If we further include an interaction with age, the point estimate of state dependence becomes unstable, though it is not statistically different from our baseline estimate.

5.2.2. Differential Reporting of Diseases by Permanent Income. If, conditional on reporting a disease, the severity of the disease varies by permanent income, this would violate our identifying assumption and bias our inferences. For example, if, conditional on reporting a disease, severity is greater for the rich than the poor, we would estimate a larger decline in utility for those with higher permanent income, thus biasing us toward finding negative state dependence; the converse would bias us in the opposite direction. Note, however, that no bias arises from differential rates of disease occurrence by permanent income because this rate merely affects the frequency with which we observe the sick state.

Our reading of the evidence suggests that any reporting differences by socio-economic status would likely bias against our finding of negative state dependence. Banks et al. (2006) find in the National Health and Nutrition Examination Survey (NHANES) that, conditional on having a biologically measured health condition, individuals of lower education are, if anything, less likely to self-report it. Similarly, in our HRS data we find that, conditional on reporting that a doctor has told them they have a particular disease, individuals of lower permanent income are more likely to report conditions that indicate a more severe form of the disease. Both of these findings suggest that the threshold for reporting a disease is higher for the poor, since they are less likely to report the disease and, conditional on reporting that they have a disease, they are more likely to have a severe form of the disease. Under the (reasonable) assumption that this underreporting by the poor exacerbates the difference in health status among the poor between those who report that they have a disease and those who do not, this would bias against our finding of negative state-dependent utility.

5.2.3. What if Second-Period Wealth Is not Predetermined? We discuss how plausible violations of the assumption of predetermined wealth would result in a bias against our finding of negative state-dependent utility. We also present complementary empirical evidence that is consistent with the assumption of predetermined second-period wealth in our baseline sample.

27. Online Appendix B provides a description of the (limited) consumption data that we use for this analysis.

Relaxing the Assumption of Predetermined Second-Period Wealth. Second-period wealth would not be predetermined if health has a mechanical effect on the resources available in the second period. For example, health may affect labor income, longevity, and household production. In addition, consumption would not be predetermined if optimizing individuals are able to reallocate resources over time in response to their state-dependent preferences; this could occur, for example, if health shocks are anticipated or transitory.

When wealth is predetermined, health can affect marginal utility of consumption through two distinct channels: its direct effect on the slope of the utility function (i.e., the state-dependent parameter of interest), and its indirect effect through the reduction in consumption necessary to finance out-of-pocket medical expenditures. These two effects were analyzed in Section 3.1. When wealth is not predetermined, health can also affect marginal utility of consumption through a third channel: the effect of state-dependent income on second-period wealth, which in turn affects second-period consumption. Online Appendix E formally extends the baseline model to allow for health to directly affect second-period wealth. Here, we limit ourselves to an informal discussion designed to convey the intuition behind the bias that may result.

Figure 3 illustrates the natural case in which second-period wealth falls with declining health, due for example to reduced income or home production. Because our regression specification does not correct for the fall in wealth (we proxy wealth by permanent income, which does not change), the estimated utility curve runs through points A and B. The nature of any bias this creates for our estimate of state-dependent utility depends on two parameters: the curvature of the utility function and the proportionality (or not) of the drop in wealth associated with the health decline. Curvature is captured by the parameter γ , the coefficient of relative risk aversion.

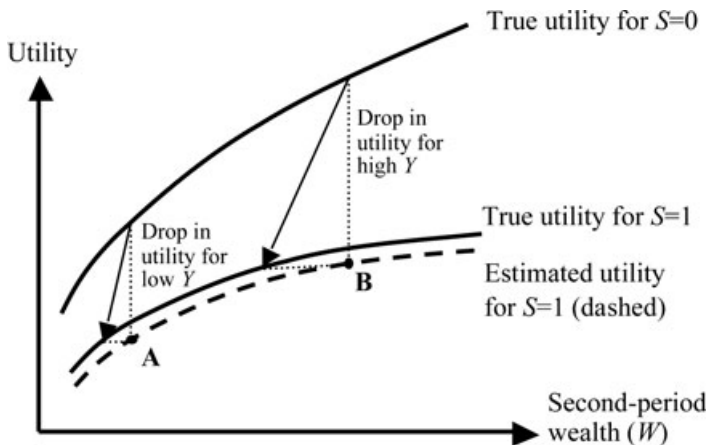


FIGURE 3. Estimates of state dependence when second-period wealth varies with health.

Consider first the case in which the drop in second-period wealth is proportional to second-period wealth in the healthy state. If $\gamma = 1$, the slope of the utility curve is inversely proportional to second-period wealth. Thus, the bias in the estimate of utility (=absolute drop in consumption \times slope of utility curve) is the same independent of income or consumption, as illustrated in Figure 3. Since the estimate of the slope of the utility curve in poor health is unbiased, the estimate of φ_1 remains unbiased. While there is some support in the literature for a coefficient of relative risk aversion of one (e.g., Metrick 1995; Chetty 2006), many papers estimate a substantially higher level of relative risk aversion (e.g., Gertner 1993; Cohen and Einav 2007). If $\gamma > 1$, then marginal utility falls more than proportionally with wealth, resulting in a smaller drop in utility for rich individuals than for poor individuals. This will cause the slope of the estimated utility curve for poor health to be biased upward, leading to a positive bias in our estimate of φ_1 and biasing against our finding of negative state-dependent utility.

Other cases follow the same basic intuition. For example, the drop in wealth may be less than proportional to wealth in the healthy state if there are absolute wealth declines associated with poor health (e.g., the reduction in the value of home production may be roughly constant across households). As a result of these absolute declines, the drop in utility associated with poor health is smaller at higher consumption levels, creating even more positive bias in our estimate of φ_1 , assuming $\gamma \geq 1$.²⁸

Of course, if deterioration in health causes an increase in wealth, the sign of the bias is exactly the opposite of what has just been discussed, and therefore biases in favor of our estimate of negative state-dependent utility. This could occur when health deteriorations reduce an individual's life expectancy; absent full annuitization, a reduction in life expectancy increases effective wealth (i.e., wealth per remaining year of life). We show below that any such bias is likely to be small in practice.

Finally, the same basic intuition applies when health changes are anticipated or transitory, or when there is an outside future consumption good with a state-independent marginal utility, such as bequests as in Lillard and Weiss (1997).²⁹ In such cases, individuals would adjust their savings (re-allocate wealth) to equate marginal utility across current and future periods (and/or the outside good). For example, for negative state-dependent utility (i.e., $\varphi_1 < 0$), people who know they will become sick will save less than they otherwise would have for future consumption. As a result of such reallocation, our estimate of state-dependent utility would be biased toward zero (for $\gamma > 1$) and would be unbiased for $\gamma = 1$, whether or not the true state dependence were positive or negative.

28. The opposite would be true for a coefficient of relative risk aversion of less than 1 (for which there is little support in the literature) or potentially for the case in which the drop in consumption is disproportionately larger for richer individuals (which seems empirically less relevant).

29. Our health conditions were chosen to be chronic ones, so as to minimize the chance they are transitory. Indeed, they are coded as absorbing events. The extent to which they are anticipated is less clear. Smith (2005) argues with respect to the same chronic disease measures that while people may to some extent be able to anticipate the onset of a chronic disease, "much of the actual realization and especially the timing may be unanticipated."

TABLE 6. Income and consumption response to disease.

Sample restrictions: Age \geq 50, NILF, Has Health Insurance						
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	Current income		Total consumption		Non-durable consumption	
Number of Diseases per Person _{it}	0.017 (0.008) [0.038]	0.016 (0.008) [0.053]	0.033 (0.033) [0.308]	0.035 (0.033) [0.287]	0.035 (0.035) [0.326]	0.037 (0.035) [0.301]
Number of Diseases per Person _{it} \times $\log(\bar{Y}_i)$		-0.003 (0.012) [0.774]		0.036 (0.034) [0.292]		0.051 (0.037) [0.172]
Within-Household std. dev. in Number of Diseases per Person _{it}	0.560	0.560	0.427	0.427	0.427	0.427
R^2	0.784	0.784	0.775	0.775	0.779	0.780
N	36721	36721	5014	5014	5014	5014

Notes: Table reports results from an OLS regression of the dependent variable on the covariates shown in the table, household fixed effects, wave fixed effects, and controls for a quadratic in average household age, household size, and a dummy for whether the household is single. The Number of Diseases per Person is the total number of diseases in the household divided by the number of respondents in the household. The dependent variable in columns (1) and (2) is the current household income. The dependent variables in columns (3) through (6) are household consumption measures. All consumption measures include out-of-pocket medical expenditures. See Online Appendix B for more information on the consumption measures used in this table. All dependent variables are in logs. Standard errors, adjusted to allow for an arbitrary variance-covariance matrix for each household over time, are in parentheses and p-values are in brackets.

Is Wealth Predetermined in our Sample? Some Suggestive Evidence. We now examine how current income changes as health deteriorates. In addition, for about 10% of our person-year observations, we are also able to examine whether *consumption* changes with adverse health events.³⁰ We find no evidence of an economically significant change in either income or consumption following an adverse health event, but recognize that our estimates have sizable standard errors.

Table 6 shows the results. Since both income and consumption are household measures, our health measure becomes the number of chronic diseases per respondent in the household—that is, the average number of diseases per person³¹—and we include household fixed effects instead of individual fixed effects. Column 1 indicates that a one-standard-deviation increase in the number of household diseases is associated with marginally statistically significant but economically insignificant 0.95% increase

30. Online Appendix B provides more detail on the small Consumption and Activities Mail Survey (CAMS) module of the HRS, which allows us to construct a broad-based measure of total consumption as well as nondurable consumption. These consumption measures include out-of-pocket medical expenditures, so they can be considered proxies for second-period wealth.

31. This is an admittedly ad-hoc way of parameterizing family health. An alternative approach would be to perform our analysis just for single individuals, or to allow the effects to differ for single and married individuals and also across diseases. The results from these exercises (which we report in Online Appendix F) show similar or even smaller changes in income and consumption associated with health shocks for single individuals, but are very imprecise.

in current household income. The results in columns 3 and 5 respectively suggest that a one-standard-deviation increase in the number of household diseases is associated with a statistically and economically insignificant 1.41% increase in total consumption and 1.49% increase in nondurable consumption. We calculate that the resultant downward bias in our estimate of φ_1 if we use the point estimate of the 1.5% increase in consumption following a one-standard-deviation decline in health is too small to explain our finding of negative state dependence.³²

Estimates of state dependence will also be biased if a decline in health leads to a proportional consumption change that differs by level of permanent income. For example, if compared to the case where a health shock leads to the same proportional change in consumption for everyone, a health shock leads the poor to consume relatively more than the rich, then we will underestimate the utility decline due to the health shock for the poor and overestimate this decline for the rich, thereby biasing toward our finding of negative state dependence. However, in columns (2), (4), and (6) of Table 6, we find that neither the income nor the consumption response differs significantly by level of permanent income. Moreover, the point estimate on the interaction term between the number of diseases and permanent income is generally positive, indicating that, if anything, a health shock leads to a relative consumption increase for the rich, which would generate a positive bias in the estimate of state dependence.

The most plausible reason why consumption might increase following an adverse health event—and thus biasing toward our finding of negative state dependence—is that the onset of disease reduces life expectancy, resulting in an increase in the effective resources available for consumption for individuals who are not fully annuitized. To gauge this potential bias, Table 7 compares the results in the full sample (column (1)) to results in a sample limited to those with more than 50% of their permanent income annuitized (column (2)) or more than 75% annuitized (column (3)). The implied magnitude of state dependence, shown in the bottom row, remains remarkably stable, though the decrease in sample size causes the estimate in column (3) to not be statistically significant at conventional levels.

32. As we derive in Online Appendix E, the bias can be computed given an assumption about the curvature of utility and the proportionality of the effect of health on second-period wealth. We assume that any increase in wealth following a negative health event is proportional to wealth. Equation (E.9) in Online Appendix E shows that our estimate of β_1/β_3 will be biased by approximately $(1 - \gamma)n$, where γ is the coefficient of relative risk aversion and n is the proportional change in second-period wealth from a health shock. For log utility ($\gamma = 1$) there is no bias. If we assume $\gamma = 3$, which is a standard assumption in models of insurance demand (e.g., Mitchell et al. 1999; Brown and Finkelstein 2008) or savings (e.g., Scholz, Seshadri, and Khitakrakun 2006), then the point estimate of a 1.5% increase in nondurable consumption would produce an estimate of state-dependent utility of -3.0% , which cannot fully explain our point estimate of an 11.2% decline in the marginal utility of consumption for a one-standard-deviation increase in the number of diseases.

TABLE 7. Alternative sample restrictions.

Baseline sample restrictions: Age \geq 50, NILF, Has Health Insurance			
	(1)	(2)	(3)
	Baseline	X = Fraction of permanent income from social security, SSI, SSDI, pensions, and annuities	
		X > 50%	X > 75%
$NUM_DISEASE_{it} \times \log(\bar{Y}_i): \beta_1$	-0.009 (0.004) [0.018]	-0.013 (0.006) [0.033]	-0.014 (0.009) [0.131]
$\log(\bar{Y}_i): \beta_3$	0.048 (0.003) [0.000]	0.074 (0.005) [0.000]	0.083 (0.007) [0.000]
$NUM_DISEASE_{it}: \beta_4$	-0.011 (0.003) [0.001]	-0.012 (0.005) [0.014]	-0.015 (0.007) [0.031]
R^2	0.474	0.459	0.465
N	45447	26506	15940
Number of individuals	11514	6437	3920
Within-person standard deviation of $NUM_DISEASE_{it}$ (σ)	0.625	0.647	0.652
% change in marginal utility for a 1 std. dev. change in $NUM_DISEASE_{it}$ ($\sigma\beta_1/\beta_3$)	-11.2% [0.018]	-11.8% [0.034]	-10.9% [0.133]

Notes: Column (1) reports the results from the baseline specification in Table 2; see notes to Table 2 (Panel A) for more details. Columns (2) and (3) report results limiting the baseline sample based on the fraction of permanent income coming from social security, SSI, SSDI, employer pensions, and annuities. All employer pension income must come from Defined Benefit pensions or else the individual is dropped from the samples in both columns (2) and (3). Standard errors are in parentheses and are adjusted to allow for an arbitrary variance-covariance matrix for each individual over time; p -values are in brackets; the p -value for $\sigma\beta_1/\beta_3$ is bootstrapped based on 10,000 iterations, resampling individuals with replacement.

6. Interpreting Our Estimates Of State-Dependent Utility: Some Simple Numerical Calibrations

Our finding of negative state dependence implies a lower optimal level of health insurance and fraction of earnings saved for retirement than what would be optimal with a state-independent utility function. To provide a rough sense of the effects of our estimated degree of state dependence, we simulate its impact on the optimal level of health insurance and the optimal savings rate using the model presented in Section 3.1, generalized to allow the elasticity of demand for health services to differ from the coefficient of relative risk aversion (see footnote 8 and Online Appendix A). Naturally, this exercise requires a number of modeling assumptions that may influence the quantitative estimates. At a broad level, however, it provides some guidance as to the potential substantive implications of our estimated state dependence.

TABLE 8. Optimal savings rate, s^* .

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	$\gamma = 1$			$\gamma = 3$			$\gamma = 5$		
	$\theta = 0.2$	$\theta = 0.5$	$\theta = 3$	$\theta = 0.2$	$\theta = 0.5$	$\theta = 3$	$\theta = 0.2$	$\theta = 0.5$	$\theta = 3$
$\varphi_1 = 0.0$	25.9%	25.9%	25.9%	25.9%	25.9%	25.9%	25.9%	25.9%	25.9%
$\varphi_1 = -0.2$	18.8%	22.3%	25.1%	20.8%	23.0%	25.2%	21.7%	23.3%	25.2%
$\varphi_1 = -0.4$	14.9%	19.2%	24.4%	18.6%	21.0%	24.5%	20.1%	21.8%	24.5%

Notes: This table presents the results of the life-cycle savings simulation. Each column reports results from alternative values of the coefficient of relative risk aversion (γ) and the inverse of the intertemporal elasticity of substitution (θ). In each column we report optimal savings rates (s^*) for three values of the state dependence parameter (φ_1). The following parameters are used in this simulation: ratio of health consumption to non-health consumption (m) = 0.236, health insurance benefit level (b) = 0.851, the elasticity of substitution between health consumption and non-health consumption (ε) = -0.20, and the probability of entering the sick state (p) = 0.5. The annual discount rate (δ) is set to 0.027 when $\gamma = 3$ and $\theta = 3$. For other columns, δ is chosen so that optimal savings when $\varphi_1 = 0$ is the same in all simulations. The annual rate of return on savings (r) is drawn with replacement from the historical distribution of real S&P 500 returns. See Section 6 and Online Appendix C for more details on the simulation.

The optimal savings rate s^* is a direct outcome of the generalized model (see equation (A.24)), and we define the optimal level of health insurance b^* as the level of b that maximizes lifetime utility (from equation (1)) if individuals treat b and τ as given and the tax rate τ is set to satisfy the government budget constraint. Our empirical estimates report the change in marginal utility in response to a one-standard-deviation increase in the number of diseases. Because our model is based on a binary health variable, we approximate a movement from the healthy to the unhealthy state as a movement from having one standard deviation fewer than the average number of diseases to one standard deviation more than the average number of diseases (a total change of two standard deviations in the number of diseases). Therefore, the appropriate φ_1 for the simulations is roughly twice the size of the φ_1 estimates presented in Table 3. To capture our range of estimates, we report the simulations for $\varphi_1 = -0.2$ and $\varphi_1 = -0.4$. We set the elasticity of demand for health services equal to -0.20 in order to match the empirical estimates from the RAND Health Insurance Experiment (Manning et al. 1987). Rather than assuming a fixed real interest rate, we allow the real interest rate to have the same stochastic properties as the historical S&P 500 real returns. As a result, the model no longer has closed-form solutions and we therefore solve it numerically. The full details of the simulations are given in online Appendix C.

Table 8 presents the results of the life-cycle savings calibration. Each column reports results from alternative values of the coefficient of relative risk aversion ($\gamma = 1, 3, 5$) and the inverse of the intertemporal elasticity of substitution ($\theta = 0.2, 0.5, 3$).³³ In each column, we report optimal savings rates for three values of the state dependence

33. We use the same range of values for the coefficient of relative risk aversion (γ) as in Section 4.2. We adopt Gruber's (2006) estimate of an intertemporal elasticity of substitution of 2 (so $\theta = 0.5$) as our central estimate. To reflect the uncertainty in the value of the IES, we also report calibrations for an IES of 1/3 and an IES of 5.

parameter ($\varphi_1 = 0.0, -0.2, -0.4$). In all columns, the results confirm that negative state dependence reduces the optimal level of savings. The magnitude of the savings response is primarily governed by the intertemporal elasticity of substitution, since a higher IES implies that individuals are more responsive in reallocating consumption between the periods in response to a differential in marginal utilities.³⁴ We find that for plausible ranges of negative state dependence (-0.2 to -0.4) and for a IES of 2, saving rates are about 10%–25% lower than what they would be absent state dependence.

In the second calibration exercise, the social planner chooses the optimal health insurance benefit level (b^*) using the same model used in the individual life-cycle savings calibration. The calibration measures how the optimal health insurance benefit level varies with the state-dependent utility parameter (φ_1). Therefore, in this simulation b and τ are no longer exogenous parameters (as they were previously), but rather b is chosen optimally by the social planner to maximize ex-ante expected utility of the consumer, taking into account the endogenous behavioral response to any change in policy. As in the previous simulation, we require that the government budget must balance in expectation, which determines τ for a given choice of b . Finally, we emphasize that this health insurance calibration follows our model, in which health is a binary variable. Thus, the optimal insurance levels are best thought of as the optimal level of insurance for health shocks that move someone from one standard deviation above median health to one standard deviation below median health. For larger health shocks, optimal health insurance benefit levels would be higher. Hence, this calibration is illustrative only for medium-sized health shocks.

Table 9 presents the results of the optimal insurance calibration. We find that for plausible ranges of negative state dependence the optimal level of health insurance is markedly lower than it would be absent state dependence. For our central estimates of risk aversion and the IES ($\gamma = 3$ and $1/\theta = 2$), negative state dependence in the range of -0.2 to -0.4 reduces optimal level of health insurance by 20 to 45 percentage points. The size of the reduction depends primarily on the coefficient of relative risk aversion because this parameter is also equal to minus the elasticity of the marginal utility of second-period consumption with respect to second-period consumption. A back-of-the-envelope calculation would suggest that if poor health reduces the marginal utility by φ_1 , optimal second-period consumption in the sick state falls by a factor of φ_1/γ . If health expenditures are a fraction m of second-period consumption, second-period consumption can be reduced by a factor of φ_1/γ by reducing the fraction of health expenditures reimbursed by $\varphi_1/(\gamma m)$.³⁵ This back-of-envelope calculation

34. A back-of-the envelope calculation would suggest that the relative change in the ratio of second-period consumption to first-period consumption is equal to the IES ($1/\theta$) times the relative drop in the expected marginal utility in the second-period caused by state dependence. This expected drop in marginal utility is the probability of sickness (p) times our state dependence parameter (φ_1). So, the back-of-the-envelope calculation would suggest that $\Delta \ln(C_2/C_1) = \Delta \ln(s^*/(1-s^*)) \approx (1/\theta) \times p \times \varphi_1$. Applying this calculation to $\theta = 0.5$ and $p = 0.5$ yields $s^* = 22.4\%$ for $\varphi_1 = -0.2$ and $s^* = 19.1\%$ for $\varphi_1 = -0.4$. These back-of-the envelope calculations are reasonably close to our calibrations, which further take into account risk aversion, medical expenditure decisions, and the stochastic nature of interest rates.

35. This back-of-the envelope formula also explains why b^* can be negative for low values of γ . If the optimal reduction in second-period consumption (φ_1/γ) exceeds the medical expenditures that the

TABLE 9. Optimal level of insurance, b^* .

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	
	$\gamma = 1$			$\gamma = 3$			$\gamma = 5$			
	$\theta = 0.2$	$\theta = 0.5$	$\theta = 3$	$\theta = 0.2$	$\theta = 0.5$	$\theta = 3$	$\theta = 0.2$	$\theta = 0.5$	$\theta = 3$	
$\varphi_1 = 0.0$	0.0	7.8%	13.1%	17.7%	27.2%	36.3%	42.8%	37.1%	46.7%	52.8%
$\varphi_1 = -0.2$	-21.6%	-31.3%	-45.8%	14.9%	17.7%	20.3%	32.1%	36.3%	39.5%	
$\varphi_1 = -0.4$	-75.8%	-106.1%	-177.7%	-6.8%	-7.8%	-9.0%	20.2%	22.3%	24.2%	

Notes: This table presents the results of the optimal insurance simulation. Each column reports results from alternative values of the coefficient of relative risk aversion (γ) and the inverse of the intertemporal elasticity of substitution (θ). In each column we report optimal level of insurance (b^*) for three values of the state dependence parameter (φ_1). See notes to Table 8 for the parameter values used in this simulation. See Section 6 and Online Appendix C for more details on the simulation.

suggests that for $\gamma = 3$, $m = 0.236$, and φ_1 in the range of -0.2 to -0.4 , the level of optimal health insurance for a medium-sized health shock is 30 to 55 percentage points lower than what it would be absent state dependence. The results from the calibrations differ from this back-of-the envelope calculation because the calibration further takes into account that insurance distorts the tradeoff between health services and nonmedical consumption and that individuals can partially self-insure through precautionary savings.

7. Conclusion

If the marginal utility of consumption varies with health, a number of well-studied economic problems, including the value of insurance and the optimal profile of life-cycle savings, will be affected. Yet the sign of any such state dependence is a priori ambiguous, and there are relatively few empirical estimates of state dependence.

Our approach is to estimate how within-person adverse health events affect a proxy for utility, and to compare this effect across individuals with different levels of permanent income. We implement this approach using seven waves of panel data on older individuals from the Health and Retirement Study and using a measure of subjective well-being as our primary proxy for utility. Across a wide range of alternative specifications, we find robust evidence that a deterioration in health is associated with a statistically significant decline in the marginal utility of consumption. We estimate that a one-standard-deviation increase in the number of chronic diseases of an individual is associated with a 10%–25% decline in marginal utility of consumption relative to marginal utility of consumption when the individual has no chronic diseases. The results from two highly stylized calibration exercises suggest that this magnitude of state dependence can have a substantial effect on important economic behaviors. For

individual would choose absent insurance, then this consumption reduction can only be achieved by taxing medical expenditures—that is, by having a negative b^* .

example, these exercises suggest that, relative to the standard practice in the applied literature of assuming a state-independent utility function, the level of state dependence we estimate lowers the optimal share of medical expenditures reimbursed by health insurance by about 20 to 45 percentage points and lowers the optimal fraction of earnings saved for retirement by about 3 to 5 percentage points.

Our findings also raise several important questions for future work. We estimate the average effect on marginal utility from the onset of common chronic diseases in a population of older individuals. While the average effect is the relevant one for many economic questions (such as the optimal level of savings), it would nonetheless be interesting to explore whether different chronic diseases have the same effect on marginal utility; unfortunately, we lack the statistical power to do so. Likewise, the data do not permit us to estimate the effect of acute diseases on marginal utility, nor do they permit analysis of state dependence in a prime-age population. In a similar vein, our analysis has focused on the possibility that marginal utility varies with health while leaving unexplored the possibility of other types of state dependence, such as how marginal utility of consumption is affected by leisure relative to labor. We hope that our paper serves as a point of departure for further work on these important topics.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Appendix A: Derivations and Extensions of the Theoretical Model.

Appendix B: Data Appendix.

Appendix C: Details of the Calibrations.

Appendix D: Semiparametric Estimator of the Mapping $g(\cdot)$.

Appendix E: Estimates of State Dependence when Second-Period Wealth Varies with Health.

Appendix F: Robustness Checks and Additional Results.

References

- Abel, Andrew (1985). "Precautionary Savings and Accidental Bequests." *American Economic Review*, 77(4), 771–791.
- Ai, Chunrong and Edward C. Norton (2003). "Interaction Terms in Logit and Probit Models." *Economics Letters*, 80, 123–129.
- Arrow, Kenneth (1974). "Optimal Insurance and Generalized Deductibles." *Scandinavian Actuarial Journal*, 1974, 1–42. Reprinted in *Collected Papers of Kenneth J. Arrow*, Vol. 3. Harvard University Press, 1984, pp. 212–260.
- Banks, James, Michael Marmot, Zoe Oldfield, and James Smith (2006). "Disease and Disadvantage in the United States and England." *Journal of the American Medical Association*, 295, 2037–2045.
- Benjamin, Daniel J., Ori Heffetz, Miles S. Kimball, and Alex Rees-Jones (2012). "What Do You Think Would Make You Happier? What Do You Think You Would Choose?" *American Economic Review*, 102(5), 2083–2110.

- Bernheim, B. Douglas, Andrei Shleifer, and Lawrence Summers (1985). "The Strategic Bequest Motive." *Journal of Political Economy*, 93, 1045–1076.
- Bertrand, Marianne and Sendhil Mullainathan (2001). "Do People Mean What They Say? Implications for Subjective Survey Data." *American Economic Review*, 91(2), 67–72.
- Brown, Jeffrey R. 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.
- Chetty, Raj (2006). "A New Method of Estimating Risk Aversion." *American Economic Review*, 96(5), 1821–1834.
- Cohen, Alma and Liran Einav (2007). "Estimating Risk Preferences from Deductible Choice." *American Economic Review*, 97(3), 745–788.
- Conlin, Michael, Ted O'Donoghue, and Timothy J. Vogelsang (2007). "Projection Bias in Catalog Orders." *American Economic Review*, 97(4), 1217–1249.
- Davidoff, Tom, Jeffrey R. Brown, and Peter A. Diamond (2005). "Annuities and Individual Welfare." *American Economic Review*, 95(5), 1573–1590.
- 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, 37–75.
- Di Tella, Rafael and Robert MacCulloch (2006). "Some Uses of Happiness Data in Economics." *Journal of Economic Perspectives*, 20, 25–46.
- Engen, Eric M., William G. Gale, and Cori E. Uccello (1999). "The Adequacy of Household Savings." *Brookings Papers on Economic Activity*, 2, 65–187.
- Epstein, Larry G. and Stanley E. Zin (1989). "Substitution, Risk Aversion, and the Temporal Behavior or Consumption and Asset Returns: A Theoretical Framework." *Econometrica*, 57, 937–969.
- Evans, William N. and W. Kip Viscusi (1991). "Estimation of State-Dependent Utility Functions Using Survey Data." *Review of Economics and Statistics*, 73, 94–104.
- Feldman, Roger and Bryan Dowd (1991). "A New Estimate of the Welfare Loss of Excess Health Insurance." *American Economic Review*, 81(1), 297–301.
- Feldstein, Martin. 1973. "The Welfare Loss of Excess Health Insurance." *Journal of Political Economy*, 81, 251–280.
- 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), 116–121.
- French, Eric (2005). "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour." *Review of Economic Studies*, 72, 395–427.
- Frey, Bruno S. and Alois Stutzer (2002). "What Can Economists Learn from Happiness Research?" *Journal of Economic Literature*, 40, 402–435.
- Gertner, Robert (1993). "Game Shows and Economic Behavior: Risk-Taking on 'Card Sharks'." *Quarterly Journal of Economics*, 108, 507–521.
- Golosov, Mikhail and Aleh Tsyvinski (2006). "Designing Optimal Disability Insurance: A Case for Asset Testing." *Journal of Political Economy*, 114, 257–269.
- Gruber, Jonathan (2006). "A Tax-Based Estimate of the Elasticity of Intertemporal Substitution." NBER Working Paper No. 11945.
- Hall, Robert E. and Charles I. Jones (2007). "The Value of Life and the Rise in Health Spending." *Quarterly Journal of Economics*, 122, 39–72.
- Heckman, James J. and Thomas E. Macurdy (1980). "A Life Cycle Model of Female Labour Supply." *The Review of Economic Studies*, 47, 47–74.
- Kimball, Miles and Robert Willis (2006). "Utility and Happiness." Working paper, University of Michigan.
- Krueger, Alan B., Daniel Kahneman, David Schkade, Norbert Schwarz, and Arthur A. Stone (2009). "National Time Accounting: The Currency of Life." In *Measuring the Subjective Well-Being of Nations: National Accounts of Time Use and Well-Being*, edited by Alan B. Krueger. University of Chicago Press, pp. 9–86.
- Layard, Richard, Guy Mayraz, and Stephen J. Nickell (2008). "The Marginal Utility of Income." *Journal of Public Economics*, 92, 1846–1857.

- Lillard, Lee A. and Yoram Weiss (1997). "Uncertain Health and Survival: Effects of End-of-Life Consumption." *Journal of Business and Economic Statistics*, 15, 254–268.
- Loewenstein, George, Ted O'Donoghue, and Matthew Rabin (2003). "Projection Bias in Predicting Future Utility." *Quarterly Journal of Economics*, 118, 1209–1248.
- Lucas, Richard E. (2007). "Long-Term Disability Is Associated With Lasting Changes in Subjective Well-Being: Evidence From Two Nationally Representative Longitudinal Studies." *Journal of Personality and Social Psychology*, 92: 717–730.
- Manning, Willard G., Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, Arleen Leibowitz, and M. Susan Marquis (1987). "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment." *American Economic Review*, 77(3), 251–277.
- Metrick, Andrew (1995). "A Natural Experiment in 'Jeopardy!'" *American Economic Review*, 85(1), 240–253.
- Mitchell, Olivia S., James M. Poterba, Mark Warshawsky, and Jeffrey R. Brown (1999). "New Evidence on the Money's Worth of Individual Annuities." *American Economic Review*, 89(5), 1299–1318.
- Oswald, Andrew J. and Nattavudh Powdthavee (2008). "Does Happiness Adapt? A Longitudinal Study of Disability with Implications for Economists and Judges." *Journal of Public Economics*, 92, 1061–1077.
- Scholz, John Karl, Ananth Seshadri, and Surachai Khittrakun (2006). "Are Americans Saving 'Optimally' for Retirement?" *Journal of Political Economy*, 114, 607–643.
- Sloan, Frank A., W. Kip Viscusi, Harrell W. Chesson, Christopher J. Conover, and Kathryn Whetten-Goldstein (1998). "Alternative Approaches to Valuing Intangible Health Losses: The Evidence for Multiple Sclerosis." *Journal of Health Economics*, 17, 475–497.
- Smith, James (1999). "Healthy Bodies and Thick Wallets: The Dual Relation between Health and Economic Status." *Journal of Economic Perspectives*, 13(2), 145–166.
- Smith, James (2005). "Consequences and Predictors of New Health Events." In *Analyses in the Economics of Aging*, edited by David Wise. University of Chicago Press, pp. 213–237
- Viscusi, W. Kip and William N. Evans (1990). "Utility Functions That Depend on Health Status: Estimates and Economic Implications." *American Economic Review*, 80(3), 353–374.
- Weil, Philippe (1990). "Unexpected Utility in Macroeconomics." *Quarterly Journal of Economics*, 105, 29–42.
- Zeckhauser, Richard (1970). "Medical Insurance: A Case Study of the Tradeoff Between Risk Spreading and Appropriate Incentives." *Journal of Economic Theory*, 2, 10–26.