

What Good Is Wealth Without Health? The Effect of Health on the Marginal Utility of Consumption^{*}

Amy Finkelstein
MIT and NBER

Erzo F.P. Luttmer
Harvard and NBER

Matthew J. Notowidigdo
MIT

May 2009

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. We find robust evidence that the marginal utility of consumption declines as health deteriorates. Our central estimate is that a one-standard-deviation increase in the number of chronic diseases is associated with an 11 percent decline in the marginal utility of consumption relative to this marginal utility when the individual has no chronic diseases. The 95 percent confidence interval allows us to reject declines in marginal utility of less than 2 percent or more than 17 percent. Point estimates from a wide range of alternative specifications tend to lie within this confidence interval. 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.

Key Words: State dependence; health; insurance; marginal utility

JEL Classification Codes: D12; I1

^{*} We are grateful to Raj Chetty, David Cutler, Jean-Paul L'Huillier, Erik Hurst, Olivia Mitchell, Jim Poterba, Ivan Werning, Justin Wolfers, 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, and Urbana-Champaign for helpful comments, and to James Wang for excellent research assistance. This research was supported by the U.S. 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. 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.

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. For example, 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 (Brown and Finkelstein 2008), annuities (e.g., Mitchell et al. 1999, Davidoff et al. 2005), or disability insurance (e.g., Golosov and Tsyvinski 2006). It is also assumed in models that explain why preferences for health spending rise with income (Hall and Jones 2007), and it is also standard in calibrations of individuals' optimal life-cycle savings (e.g., Engen, Gale and Uccello 1999, Scholz, Seshadri and Khitatrakun 2006). Yet, as we show below with some simple, stylized numerical examples, 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 its potential importance, there has been relatively little empirical work on how the marginal utility of consumption varies with health. This presumably reflects the considerable challenges involved in constructing credible estimates of how the shape of the utility function varies with health. We begin our paper with a discussion of the different empirical options for estimating state dependence; each has important limitations, which we describe.

We adopt an approach in which we compare 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.

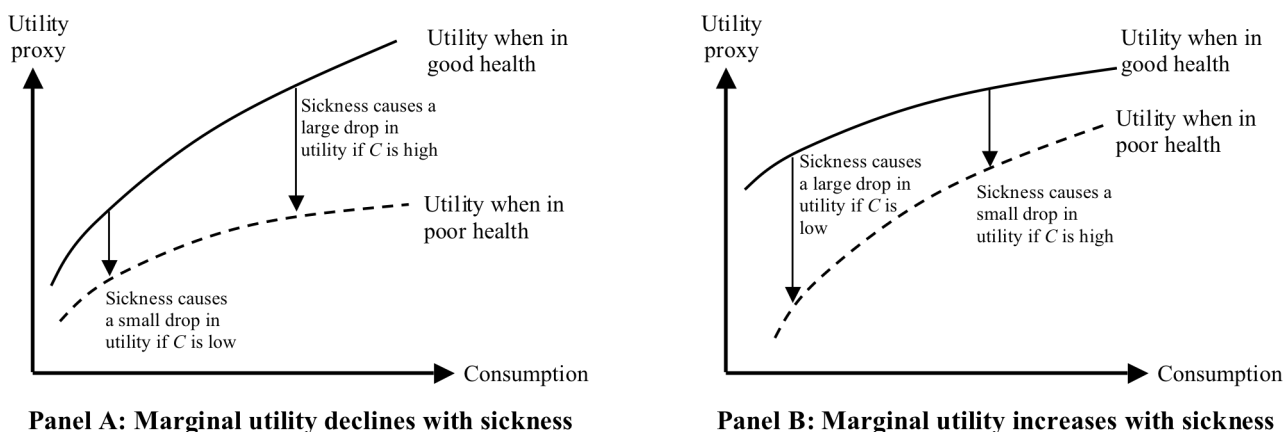


Figure 1: State-dependent utility functions

There are two key practical challenges to implementing this conceptually straightforward approach. First, data with broad-based consumption measures are notoriously scarce, and none exists that contains the other variables needed for the analysis. We therefore develop a simple model of optimizing behavior that allows us, under conditions which appear to be empirically valid, to infer 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 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 subjective data. We discuss some of the concerns with the SWB measures below. However, we emphasize that 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 cardinal 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 from

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 non-medical consumption, evaluated at a constant level of non-medical consumption. So that health does not affect non-medical 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.

Across a wide range of alternative specifications, we find statistically significant evidence that the marginal utility of consumption declines as health deteriorates. Our central estimate is 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 an 11 percent decline in marginal utility. Our 95 percent confidence interval on this central estimate allows us to reject declines in marginal utility associated with a one-standard-deviation decline in health of less than 2 percent or more than 17 percent. The point estimates from a wide range of alternative specifications tend to lie well within this 95 percent confidence interval.

To illustrate the potential implications of these findings, we examine the impact of our central estimate for simple calibration exercises of the optimal levels of health insurance benefits and life-cycle savings. The results suggest that, relative to the standard practice of assuming a state-independent utility function, 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 1 percentage point. Of course, considerable caution should be exercised in using the results of our extremely stylized calibrations. Nonetheless, at a qualitative level, they underscore the likely substantive importance of the state dependence that we detect.

The rest of the paper proceeds as follows. Section two presents some motivating calibrations on the potential importance of even moderate state dependence. Section three describes the main options for estimating state-dependent utility. Section four describes our empirical approach and data. Section five presents the main results and their interpretation. Section six discusses potential threats to the validity of our empirical strategy and demonstrates the robustness of our findings along a number of different dimensions. The last section concludes.

2. The potential importance of state-dependent utility: some simple numerical calibrations

We allow for health-state dependence by assuming a utility function of the form:

$$U(C, S) = \gamma_0 S + (1 + \gamma_1 S)u(C), \quad (1)$$

where C denotes consumption, S denotes sickness (or the state of poor health), and u is a standard CRRA sub-utility function. This is a standard formulation of the utility function (see, e.g., De Nardi et al. 2006 or Palumbo 1999).

For ease of exposition, it is useful to think of S as a binary variable, so that $S=1$ denotes sickness, and $S=0$ denotes good health. The term $\gamma_0 S$ allows for the level of utility to vary with health, but we would not characterize this as state dependence because it leaves the shape of the utility function unchanged. The term $\gamma_1 S$ captures state dependence because it allows the state of health to affect the marginal utility of consumption. For example, a γ_1 of -0.10 would imply that the marginal utility of consumption is 10% lower in sickness ($S=1$) than in health ($S=0$), *when the marginal utilities are evaluated at the same level of consumption*. More generally, we use the term negative state-dependent utility to refer to the case in which sickness reduces the marginal utility of a given level of consumption (i.e., $\gamma_1 < 0$; the utility curves fan out as in Figure 1A). Conversely, positive state-dependent utility refers to the case in which sickness increases the marginal utility of a given level of consumption (i.e., $\gamma_1 > 0$; the utility curves fan in as in Figure 1B). Testing for whether the utility curves fan out or fan in lies at the heart of our empirical test of state-dependent utility.

Our first calibration exercise illustrates the impact of state dependence on the optimal level of health insurance. Following the derivation in Baily (1978) and Chetty (2006a), we find that the optimal level of health insurance under moral hazard is implicitly defined by:

$$\frac{u'_s(C_s) - u'_h(C_h)}{u'_h(C_h)} = \frac{d \log H}{d \log b} = \varepsilon_{H,b},$$

where b is the benefit rate as a fraction of medical expenses, $\varepsilon_{H,b}$ is the elasticity of demand for health care with respect to the benefit level, H is health spending, and $u'_s(C_s)$ and $u'_h(C_h)$ are marginal utility in the sick and healthy states, respectively. Baily and Chetty derive the optimal level of insurance b^* under the assumption that the marginal utility of consumption is the same across health states (i.e., $\gamma_1 = 0$). As we show in Appendix A, allowing for $\gamma_1 \neq 0$ results in the following formula for the optimal level of insurance:

$$b^* = 1 + \frac{C_h}{H} \left(\left(\frac{1 + \gamma_1}{1 + \varepsilon_{H,b}} \right)^{1/\alpha} - 1 \right),$$

where C_h / H is the ratio of consumption (in the healthy state) to health spending and α is the coefficient of relative risk aversion. The optimal level of benefits thus involves a tradeoff between the social cost of moral hazard, which depends on $\varepsilon_{H,b}$, and the gains from increased consumption smoothing, which depends on risk aversion, state dependence, and the ratio of consumption to health expenses.

This expanded formula suggests that even modest amounts of state-dependent utility can have quantitatively important effects on the optimal level of health insurance benefits of the same order of magnitude as modest variations in risk aversion. For this illustrative calibration, we choose estimates of $\varepsilon_{H,b}$ and C_h / H from the literature,¹ and we choose several plausible values of risk aversion (α) and state-dependent utility (γ_1) to produce the following table:

TABLE 1
OPTIMAL HEALTH INSURANCE LEVEL, b^*

α	γ_1		
	-0.10	0.00	0.10
2	59.8%	73.9%	87.2%
3	72.6%	82.3%	91.4%
4	79.2%	86.6%	93.5%

This table shows that, for a given level of risk aversion, allowing health shocks to affect the marginal utility of consumption by just 10% changes the optimal level of health insurance by roughly 10 percentage points. This effect is of the same order of magnitude as the effect of increasing the coefficient of relative risk aversion from 2 to 4.

Our second calibration exercise illustrates the impact of health-state dependence on optimal life-cycle savings. We use a very simple two-period model in which the agent must decide how much of her first-period fixed wage w to save for consumption in the second period, given that she faces a probability p of becoming sick in the second period. Savings yield e^{rT} between periods, where r is the annual rate of return on savings and T is the number of years between

¹ We use $\varepsilon_{H,b} = 0.20$ based on the RAND Health insurance experiment (Manning et al. 1987), and we use $C_h / H = 3$ using data from the MEPS and the CAMS (see Appendix A for more details).

periods. The agent discounts the second period by $e^{\delta T}$ and has CRRA per-period utility (with a coefficient of relative risk aversion α). The full maximization problem of the agent and the formula for the optimal level of savings are given in Appendix A. The following table shows the optimal level of savings for various degrees of risk aversion and state-dependent utility, choosing sensible values for the other parameters ($r = 5\%$ per year, $T = 25$ years, $\delta = 4\%$ per year, $p = 0.5$):

TABLE 2
OPTIMAL LEVEL OF SAVINGS, s^*

α	γ_1		
	-0.10	0.00	0.10
2	24.0%	24.5%	25.0%
3	23.4%	23.7%	24.0%
4	23.1%	23.4%	23.6%

The results suggest that varying risk aversion between 2 and 4 under the assumption of no state dependence (i.e. $\gamma_1=0$) has roughly the same effect on the optimal fraction of earnings saved as holding risk aversion fixed and varying the effect of health shocks on the marginal utility of consumption by 20 percentage points.

3. 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 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

² We offer a more detailed discussion of these possible approaches in Finkelstein et al. (2009).

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 have more for bequests) 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 change varies across individuals of different 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 below. However,

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.

given the formidable obstacles to the revealed demand approach, we believe 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 only 8 percent of the marginal utility in the health state (Sloan et al., 1988), as well as estimates in between these two (Viscusi and Evans 1990).⁴

4. Empirical approach

Figure 1 above 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, it follows that 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 directly and straightforwardly estimate how marginal utility of consumption varies with health. We would simply regress the utility proxy on consumption, health, and the interaction of consumption and health; the coefficient on the interaction term between consumption and health would give an estimate of state-dependent utility. In practice, however, we know of no panel dataset with a sufficient sample size that contains information on *consumption*, health, and utility proxies. Reliable consumption data are notoriously scarce, especially broad-based consumption measures which this approach requires.

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*, in addition to utility

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.

proxies and detailed health measures.⁵ In Section 4.1, we outline a simple model of optimizing consumption behavior that yields conditions under which we can infer how marginal utility of consumption varies with health status from estimates of how marginal utility of permanent income varies with health status. The key requirement, as we explain in more detail below, is that consumption in the sick state is pre-determined, or in other words that health shocks do not lead to changes in consumption.

We select a baseline sample for which this assumption seems a priori plausible: specifically, we restrict our sample to individuals who are not in the labor force – so that deteriorations in health do not have a first-order effect on income – and who have health insurance – to avoid the direct impact of health on consumption through its effect on medical expenditures.⁶ In Section 6.2.3 we extend the baseline model to illustrate that the most likely result of relaxing the assumption that consumption does not vary with health is to bias against our empirical finding of negative state-dependent utility. We also present complementary empirical evidence that is consistent with the validity of our key assumption for our baseline sample and that suggests that any potential biases in our favor (i.e., toward negative state-dependent utility) are, in practice, likely to be small in magnitude.

4.1 Motivating theory and estimating equation

Optimizing individuals allocate lifetime income over consumption in two periods, taking into account that the marginal utility of consumption depends on health status. 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.

Using the notation introduced in Section two, lifetime utility is given by:

$$U(C_1, C_2, S) = u_1(C_1) + \frac{1}{1+\delta} u_2(C_2, S) = \frac{1}{1-\alpha} C_1^{1-\alpha} + \frac{1}{1+\delta} (\gamma_0 S + (1 + \gamma_1 S) \frac{1}{1-\alpha} C_2^{1-\alpha}), \quad (2)$$

where C_1 and C_2 denote first- and second-period consumption respectively, δ denotes the

⁵ As we describe in more detail in Section 6.4 and Appendix B, consumption data are available for a very small subsample of the survey. Not surprisingly, when we use this consumption data to estimate state dependence, the estimates become so imprecise as to preclude any meaningful inference on state dependence.

⁶ While the effect of health on labor income is relevant for the optimal insurance problem (where lost income would be fully replaced if utility is state independent), this is a mechanical effect that is conceptually separable from the effect of the shape of the utility function depending on health. For a similar reason, we also wish to ignore the impact of health on medical expenditures, as does other empirical work on state dependence with respect to health (e.g., Lillard and Weiss, 1997). Ultimately, of course, the optimal level of insurance depends both on our estimate of state dependence, as given by γ_1 , and on the conceptually and practically distinct questions of how health affects net resources (e.g., due to changes in labor income and medical expenses).

discount rate, and α is the coefficient of relative risk aversion. The utility function is a standard additively separable CRRA utility function except for the terms involving S . The objective is to recover an unbiased estimate of γ_1 .

In period 1, before the individual knows her future health status, the individual allocates permanent income (or wealth) Y over first- and second-period consumption. The resulting budget constraint is:

$$Y = C_1 + \frac{1}{1+r} C_2.$$

From the perspective of period 1, health in period 2 is a random variable, and the individual maximizes expected lifetime utility in period 1. Using the budget constraint to eliminate C_1 , we find expected utility as a function of C_2 :

$$E[U] = \frac{1}{1-\alpha} \left(Y - \frac{C_2}{1+r} \right)^{1-\alpha} + \frac{1}{1+\delta} \left(\gamma_0 P + \frac{1}{1-\alpha} (1 + p\gamma_1) C_2^{1-\alpha} \right).$$

Expected utility is maximized for:

$$C_2^* = \frac{\left((1 + p\gamma_1)(1+r) / (1+\delta) \right)^\alpha}{1 + \left((1 + p\gamma_1)(1+r) / (1+\delta) \right)^\alpha / (1+r)} Y \equiv cY,$$

where c is a parameter that expresses the optimal second-period consumption as a fraction of permanent income (Y). Note that second-period consumption is increasing in permanent income and does not depend on the realization of the health shock.

Substituting the optimal level of second-period consumption into the second-period utility function yields indirect utility, $v(Y, S)$, in the second period:

$$v(Y, 0) = \frac{1}{1-\alpha} (cY)^{1-\alpha}, \text{ and } v(Y, 1) = \gamma_0 + \frac{1+\gamma_1}{1-\alpha} (cY)^{1-\alpha}.$$

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

$$v = \beta_1 S \times Y^{\beta_2} + \beta_3 S + \beta_4 Y^{\beta_2} + \varepsilon, \quad (3)$$

which yields the parameter estimates:

$$\beta_1 = \left(\frac{c^{1-\alpha}}{1-\alpha} \right) \gamma_1, \quad \beta_2 = 1-\alpha, \quad \beta_3 = \gamma_0, \quad \text{and } \beta_4 = \frac{c^{1-\alpha}}{1-\alpha}.$$

The estimate of β_1 , the coefficient on the interaction term between permanent income and sickness in equation (3), measures whether (and in what direction) the marginal utility of consumption is affected by health. We reject the hypothesis of state independence if we reject the

null hypothesis $\beta_1 = 0$. Furthermore, $\beta_1 < 0$ indicates that marginal utility declines as health deteriorates, while $\beta_1 > 0$ suggests that marginal utility increases as health deteriorates.

We operationalize equation (3) by running a fixed effects regression of the following form:

$$UtilityProxy_{it} = g(\beta_1 S_{it} \times \bar{Y}_i^{\beta_2} + \beta_3 S_{it} + X_{it} \Psi_1 + \theta_i), \quad (4)$$

where i indexes individuals and t indexes time periods. 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 (θ_i) absorb any direct effect of permanent income and any other time-invariant characteristics on utility. Finally, β_1 , β_2 , β_3 , and the vector Ψ_1 are coefficients to be estimated, and $g(\cdot)$ is a monotonically increasing mapping from latent cardinal utility to the utility proxy.⁷

4.2 Identifying assumptions

The coefficient β_1 denotes the *additional* marginal utility of permanent income when sick relative to marginal utility when healthy. To interpret the estimate of β_1 as a test of state dependence, we need to make two identifying assumptions. Our first identifying assumption is: the difference between our imposed mapping $g(\cdot)$ and the true mapping from latent cardinal utility to the utility proxy does not vary with health differentially by permanent income. In other words, there may be error in our assumed mapping $g(\cdot)$ from latent cardinal utility to the utility proxy, but this error cannot vary systematically with health times permanent income. This assumption allows us to compare changes in latent cardinal utility across individuals of different levels of permanent income and health by comparing observed changes in the utility proxy. 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 6.1.

Our second identifying assumption – needed to interpret the estimate of β_1 as a test of state

⁷ For example, in our OLS specification, $g(v)$ equals $v + \varepsilon_{it}$, with ε_{it} denoting the error term in utility. In the probit specification, $g(v)$ equals $\Phi(v)$, with $\Phi(\cdot)$ denoting the standard normal cumulative density function. In Section 6.1, we explore more general specifications of $g(\cdot)$.

dependence – is the following: conditional on sickness S , control variables X , and fixed effects θ , 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. In a cross section, there might well be person-specific characteristics (such as optimism / pessimism) that are correlated with utility, health, and permanent income (e.g., more optimistic people are happier, work harder, and are healthier). In the panel, however, the individual fixed effects absorb any such person-specific characteristics. In the same spirit as typical identification for difference-in-differences estimation, it is only a problem for our analysis if health changes within individuals vary across individuals of different permanent income in ways that are correlated with utility. In Section 6.2 we discuss potential threats to the validity of this identifying assumption in greater depth. We offer a number of alternative specifications and auxiliary tests, which are generally supportive of this identifying assumption.

To interpret the magnitude of β_1 , we scale the coefficient by the effect of permanent income on utility in the healthy state, i.e., by β_4 from equation (3). Since this effect is absorbed by the individual fixed effects in equation (4), we recover it by running an auxiliary regression of the estimated fixed effects from equation (4) on permanent income and demographic controls:

$$\hat{\theta}_i = \beta_4 \bar{Y}_i^{\beta_2} + X_{it} \Psi_2 + \eta_{it} . \quad (5)$$

Our empirical estimate of β_1/β_4 gives the proportional change in the marginal utility of *permanent income* resulting from moving from good to poor health. The underlying structural parameter of interest is the proportional change in the marginal utility of *consumption* due to falling into poor health evaluated at a constant consumption level, i.e., γ_1 as defined by in equation (1). These two are the same, given the key assumption in our baseline model that second-period consumption is proportional to permanent income or, equivalently and more generally, that consumption in the sick state is pre-determined. Violation of this assumption is one way to invalidate our second identifying assumption; we explore the validity of the pre-determined consumption assumption in Section 6.2.3 below.

Identification of β_4 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. This is a considerably stronger assumption than what is required for identification of β_1 . We emphasize, however, that it is needed only to provide one way of scaling (interpreting) our key parameter β_1 ; it is not fundamental to our identification of the existence of state-dependent utility.

4.3 Data and baseline specification

We use all cohorts in the first seven waves of the HRS and limit the sample to individuals (and their spouses) aged 50 and older who are not in the labor force and who have health insurance. The resulting sample of 45,447 person-years consists of an average of about 4 observations on 11,514 unique individuals. Appendix B provides more detail on our sample and variable definitions. Table 3 presents some descriptive statistics. The sample is 63 percent female and 87 percent white with an average age of 72. About three-fifths of the person-years are married.

To operationalize equations (4) and (5) 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) as the average across all waves of their total annual household income, adjusted for household composition, and a 5 percent annual draw down of current financial wealth; 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 (see, e.g., Smith 1999), we consider the following 7 diseases that are asked consistently over time: hypertension, diabetes, cancer, heart disease, chronic long disease, stroke, and arthritis.⁸ 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 baseline 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. On average, 87 percent of person-years respond in the affirmative. The within-person standard

⁸ 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.

deviation in the response (which is our key left hand side variation) is 0.28.

Finally, we operationalize the general regression equation (4) for our baseline specification by assuming that the mapping $g(\cdot)$ is linear and that $\beta_2 \rightarrow 0$, or, equivalently, the coefficient of relative risk aversion $\alpha \rightarrow 1$.⁹ We will relax both assumptions in Section 6.1.

Our baseline specification is therefore the following linear probability model:

$$HAPPY_{it} = \beta_1 NUM_DISEASE_{it} \times \log(\bar{Y}_i) + \beta_3 NUM_DISEASE_{it} + X_{it} \Psi_1 + \theta_i + \varepsilon_{it}. \quad (6)$$

We estimate (6) using de-meaned values of \bar{Y}_i so that the coefficient on $NUM_DISEASE_{it}$ (β_3) describes the relationship between an individual's change in sickness and change in utility at the sample average level of permanent income. The estimated fixed effects ($\hat{\theta}_i$) denote the individual's average utility when $NUM_DISEASE_{it} = 0$. We include in X_{it} controls for time-varying individual characteristics that might be correlated with changes in utility and health: household size, an indicator for whether the individual is married, a quadratic in age, and a fixed effect for each of the seven waves of the HRS.

We operationalize equation (5) in the same fashion, estimating:

$$\hat{\theta}_i = \beta_4 \log(\bar{Y}_i) + X_{it} \Psi_2 + \eta_{it}. \quad (7)$$

We use the same time-varying X 's from (6) and also include additional time-invariant individual controls (a constant, race and gender) which are absorbed in (6) by the individual fixed effects; we do this so as not to confound the relationship between permanent income and SWB in (7) with demographics that are correlated with permanent income.

5. Main Results

5.1 Baseline results

Table 4 shows our baseline estimates of equation (6) and the auxiliary regression equation (7). Our utility proxy looks sensible: it is decreasing in number of diseases ($\beta_3 < 0$) and increasing in permanent income ($\beta_4 > 0$). The coefficient β_3 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

⁹ While, as a practical matter, we use a range of values of permanent income and therefore need to make an assumption about curvature (i.e., α), in principle this is not essential for our empirical strategy. In principle, we could estimate state-dependent utility in our setting without making an assumption about curvature. This can be seen in Figure 1 and imagining we only observed permanent income in two places (e.g., low and high).

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 their health state over time. The coefficient β_4 of 0.048 indicates that a 10 percent 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 percent). 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 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-percent level, which implies that we reject the null of state-independent utility.

The bottom panel of Table 4 reports several ways to interpret our estimate of β_1 . The ratio β_1/β_4 of -17.9% indicates that for a healthy person (i.e., someone with no diseases), acquiring one chronic disease is associated with a 17.9 percent decline in marginal utility.¹⁰ Likewise, a one-standard-deviation increase in the number of chronic diseases is associated with an 11.2 percent decline in marginal utility for a previously healthy individual, with a 95 percent confidence interval ranging from 2.7 percent to 16.8 percent. 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 3 diseases, and 0.012 with 4 diseases (only 3.2% of our sample has more than 4 diseases); this is the empirical counterpart to the stylized picture in Figure 1B of utility curves “fanning out.”

Our finding of negative state dependence implies that the optimal level of health insurance and the optimal fraction of earnings saved for retirement are lower than indicated by the standard formula that assumes no state dependence. To calibrate the models in Section 2, which were based on a binary health variable, we reran our model with the variable *NUM_DISEASE* replaced by an indicator variable for being in the sick state, which we defined as having the median

¹⁰ 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_4 and $\sigma\beta_1/\beta_4$ in this table and in all subsequent tables.

number of diseases (2) or more. This yields (in results not shown) an estimate of γ_1 (i.e., β_1/β_4) of -0.29 (bootstrapped p-value = 0.132), which means that the marginal utility of those in the sick state (2 or more diseases) is 29 percent lower than the marginal utility of those in the healthy state (less than 2 diseases). The stylized models presented in Section 2 suggest that with this amount of state dependence, the optimal level of health insurance is roughly 20 to 45 percentage points lower than it would be in the absence of state dependence. Likewise, the models suggest that the optimal fraction of earnings saved for retirement is about 1 percentage point lower (or about 4 percent lower) than it would be with a state-independent utility function. While these calculations should be viewed as merely illustrative, they suggest that the magnitude of state dependence we have detected may have a non-trivial effect on important economic phenomena.

5.2 Additional analysis

Our approach yields an estimate of the average effect of deteriorating health on the marginal utility of consumption in a representative sample of the elderly and near elderly. This is the economically relevant parameter for savings and health insurance decisions; indeed, we consider it a strength of our approach that it yields estimates of the average effect of common health conditions in the population on the marginal utility of consumption. However, because the marginal utility of consumption may not change with the onset of each disease in the same way, we examine the effect on marginal utility of each disease separately.

Table 5 interacts each of the seven disease dummies with the log of permanent income and includes all seven interaction terms and the seven disease dummies in one regression. The first seven columns give the estimates on the interaction term, the disease dummy for each of the seven diseases, and log permanent income. Not surprisingly, the precision of the estimates is often considerably worse than the baseline number of diseases variable. 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 = 0.131). In the final column, we show that the prevalence-weighted sum of the seven interaction terms from this specification is statistically significant and that the magnitude (-10.5%) is very similar to our baseline result of -11.2%.¹¹ In unreported results, we also examine whether the magnitude of the drop in marginal utility from

¹¹ Prevalence-weighting is based on the person-years in the baseline sample that have the disease dummy turned on.

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 = 0.360).

Table 6 reports the results from several additional sensitivity analyses. In this and all subsequent robustness tables, Column 1 replicates our baseline results from Table 4. Subsequent columns always report results for one specified change relative to this baseline. To facilitate comparability of the magnitude of state-dependent utility across these and later analyses, the bottom row reports the implied percent change in marginal utility for a healthy person associated with a one-standard-deviation decline in health (i.e., $\sigma\beta_1/\beta_4$). This provides a scale-free way of comparing different estimates.

Column 2 shows that the results are not sensitive to excluding the demographic controls (X_{it}). Column 3 restricts the analysis to individuals who are always single. Since three-fifths of our sample is married, our estimates are potentially confounded by correlations in health changes within a couple and by any effects that spousal health has on one's own marginal utility. As shown in Column 3, the point estimate of state dependence is still negative among single individuals. However, based on just 30 percent of the original sample, the estimate is no longer statistically significant. Column 4 shows that the estimate of β_1 is unaffected by adding additional covariates for spousal health and the interaction of spousal health with log permanent income. Interestingly, the results suggest that while a deterioration in spousal health has a similar impact on an individual's utility as a deterioration in own health, a deterioration in spousal health has no detectible effect on an individual's marginal utility.

Oswald and Powdthavee (2007) show that individuals partly adapt to disability; the onset of disability reduces happiness more in the short run than in the longer run. While their finding concerns adaptation of the level of utility, it raises the question of whether marginal utility also adapts to health shocks. In column 5, we test for such habituation effects by adding a two-year lag of the number of diseases as well as this lag interacted with log permanent income as regressors. We find that the coefficient on permanent income interacted with lagged number of diseases is small and statistically insignificant. In other words, the decline in marginal utility after a negative health shock does not appear to diminish over time.

Finally, Table 7 investigates the sensitivity of our results to alternative measures of our key

variables. Columns 2 and 3 show that we continue to estimate negative and statistically significant state dependence (i.e., $\beta_1 < 0$) if we replace our permanent income measure \bar{Y} with, respectively, education and wealth, which are other reasonable proxies for consumption opportunities; in both columns, the magnitude of our estimate of state dependence (i.e., $\sigma\beta_1/\beta_4$ shown in the bottom row) is slightly larger than in the baseline estimate. Columns 4 through 7 show that we continue to obtain negative and (usually) statistically significant estimates of state dependence if, instead of our baseline measure of the number of chronic diseases, we use other standard measures of health, including (respectively) limitations to activities of daily living (ADLs), limitations to instrumental activities of daily living (IADLs), other functional limitations (OFLs), and a health index measure in the spirit of Dor et al. (2006) in which we sum the three limitation measures and the individual's reported pain score.

The last three columns of Table 7 report results for alternative utility proxies. In addition to the baseline utility proxy (the subjective well-being question “Much of the past week I felt happy [yes or no]?”), the HRS contains seven other items from Radloff's (1977) CES-D depression scale. These items have a similar format but instead of “I felt happy” substitute “I enjoyed life”, “I felt sad”, “I felt lonely”, “I felt depressed”, “I felt that everything I did was an effort”, “my sleep was restless”, and “I could not get going”. We code these 0/1 measures such that 1 corresponds to higher utility and define a *CESD-8* variable as the sum of the answers over these eight questions. We also follow Smith et al. (2005) by defining a subjective well-being measure *CESD-4* that consists of the sum of answers to the first four items from the Radloff scale; these focus more on happiness and less on the feelings more typically associated with depression or stress.

Columns 8 and 9 of Table 7 report results of estimating equations (6) and (7) using *CESD-8* and *CESD-4* respectively as our utility proxy. Both have desirable properties for a utility proxy in that they both decline with worsening health (i.e., $\beta_3 < 0$) and increase with permanent income (i.e., $\beta_4 > 0$). Most importantly, both indicate a decline in the marginal utility of permanent income associated with deteriorating health, i.e., $\beta_1 < 0$, though this decline is only statistically significant for *CESD-8*. The bottom row of Table 7 shows that the magnitude of the estimated state dependence (i.e., $\sigma\beta_1/\beta_4$) is somewhat smaller than in our baseline, although it lies within

the baseline's 95 percent confidence interval.¹² In column 10, we draw on a similar sample of individuals in a different data set – the British Household Panel Survey (BHPS) – to perform the analysis with subjective life satisfaction (measured on a 7-point scale), a commonly used alternative subjective well-being measure; appendix B describes these data in more detail. We continue to estimate negative state-dependent utility (i.e., $\beta_1 < 0$) with a magnitude (i.e., $\sigma\beta_1/\beta_4$) of -14.0%, which is quite similar to our baseline estimate of -11.2% using different data, from a different country, with a different SWB measure. However, in the considerably smaller sample, our estimate of state-dependent utility is no longer statistically significantly different from zero.

6. Threats to identifying assumptions

Threats to validity come from violations of the identifying assumptions that allow us to interpret the estimate of β_1 as a test of state dependence. In Section 6.1, we evaluate the first identifying assumption, that there are no errors in our specification of the mapping $g(\cdot)$ from latent cardinal utility to the utility proxy that vary with health differentially by permanent income. In Section 6.2 we explore in detail several possible threats to the second identifying assumption that, conditional on sickness S , control variables X , and fixed effects θ , there are no omitted determinants of utility that vary with health differentially by permanent income. While it is never possible to directly test the identifying assumptions, our reading of the results in this section is that they greatly alleviate concerns that these types of potential biases are responsible for our finding of negative state-dependent utility.

6.1 Differential errors in the mapping from cardinal utility to subjective well-being

There is growing evidence that measures of self-reported well-being are meaningful measures of utility: 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 (e.g., increasing in income and health, decreasing in divorce and unemployment); happier people are less likely to commit suicide; subjective well-being correlates with activity in the brain's left-prefrontal cortex, which is known from clinical and experimental evidence to

¹² Smith et al. (2005) compare the impact of moving from no ADL limitations to at least two ADL limitations on CESD-4 by household net worth. They find that those with below-median net worth experience a significantly larger drop in subjective well-being as a result of acquiring two ADL limitations than those with above-median net worth. Because their sample also includes individuals in the labor force and those without health insurance, these estimates could be driven by negative consumption shocks as a result of the onset of disability rather than by positive state-dependent utility. Indeed, the estimate on our interaction term also becomes positive if we add to our sample individuals in the labor force and individuals without health insurance, but we argue that the interaction term in this case no longer estimates state dependence because it is biased by the direct effect of disability on consumption.

process pleasure.¹³ As noted earlier, in our data we find the (sensible) properties that SWB increases with permanent income and declines with adverse health events.

Such findings suggest that subjective well-being data can be useful in economic research, especially for questions that cannot be credibly answered using observed choices such as in Di Tella, MacCulloch, Oswald (2001), Gruber and Mullainathan (2005), or Luttmer (2005). 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 et al. (2003) and Conlin et al. (2007) show that individuals suffer from “projection bias”: people put 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.

At the same time, however, there is doubtless considerable measurement error in SWB measures as utility proxies. It is well documented that answers to these questions can be sensitive to wording, framing, question order, social desirability, or how the respondent processes the question (see, e.g., Bertrand and Mullainathan, 2001). Indeed, in one study, researchers were able to elicit a significant movement in self-reported happiness with subjects’ lives as a whole by experimentally manipulating whether subjects found a dime (Schwarz and Strack, 1991).

Random (even non-mean-zero) measurement error in SWB does not bias our estimate of state dependence. As discussed in Section 4.2, the key assumption needed about our utility proxy for the validity of our analysis is that errors in our specification of the mapping $g(\cdot)$ from latent cardinal utility to the utility proxy cannot vary systematically with the interaction of health and permanent income. In other words, a given *change* in true underlying 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 $g(\cdot)$ may cause this assumption to fail, suppose that the true mapping from underlying 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

¹³ See Frey and Stutzer (2002), Di Tella and MacCulloch (2006), or Krueger et al. (2008) for more extensive discussions of the validity and usefulness of happiness data in economic research.

drop from 2 to 1; in this example the true 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. The fundamental cause of this mistake is that the derivative of the true mapping increases with utility (which, in turn, depends on income), so that a given change in true utility from a health shock maps into a larger change in the utility proxy for a high-income individual than for a low-income individual.

In the example above, the misspecification of $g(\cdot)$ led to an incorrect inference 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. Alternatively, consider a high-income individual and a low-income individual who both have mean utility proxies (over time) of 3 and who both experience the onset of a disease. Suppose that the disease causes utility to drop from 5 to 1 for the high-income individual, but only from 4 to 2 for the low-income individual; in this example the true utility function is characterized by negative state dependence. It is easy to see that *any* monotonically increasing mapping from true utility to the utility proxy will lead to a larger drop in the utility proxy for the high-income individual than for the low-income individual. More generally, when we compare the effect of the onset of a disease within bins of high-income and low-income individuals with same mean levels of utility, misspecification of $g(\cdot)$ will not lead to incorrect inference about the sign of state dependence. 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 on average high-income individuals have higher utility levels 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 and, more importantly, suggests that the effects of potential misspecification of $g(\cdot)$ may not be significant in our setting.

The results from the nonparametric test (described in 6.1.1) show clear and consistent evidence of negative state dependence. The (considerable) advantage of the nonparametric

approach is that it is robust to arbitrary monotonically increasing mappings from true utility to the utility proxy. However this flexibility comes at a cost: the nonparametric approach cannot be used to estimate the magnitude of state dependence, only its sign. Therefore, in Section 6.1.2 we explore a range of alternative parametric assumptions and estimates of $g(\cdot)$, and we are reassured that they all yield similar quantitative estimates of the magnitude of state dependence.

6.1.1 Nonparametric test of the sign of state dependence

The logic behind the nonparametric test is the insight, explained in the preceding example, that when we 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 misspecification of $g(\cdot)$ will not lead to incorrect inference about the sign of state dependence. The reason for this insensitivity is that within such bins, high- and low-income individuals are on average at the same range of the $g(\cdot)$ function. 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.

For each bin based on mean values of the utility proxy and mean values of the number of diseases, we run an OLS regression of the utility proxy on number of diseases and an interaction of the number of diseases with permanent health.¹⁴ Since the number of observations in these narrowly defined bins is low, the estimates of the coefficients on the interaction term will be imprecise. Still, under the null hypothesis of no state dependence, the t-statistics of these coefficients will 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 will be 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 $\frac{1}{2}$. 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.

Table 8 shows the results of the nonparametric test. Column 1 shows our baseline case in which we have sorted individuals into 10 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 percent) of these t-statistics is negative, which provides a first indication of negative state dependence. Moreover, the mean of the CDFs

¹⁴ 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.

of the t-statistics is 0.402, which is significantly different from the expected value of $\frac{1}{2}$ under the null hypothesis of no state dependence ($p = 0.003$). Columns 2 through 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 true cardinal utility to the utility proxy.

6.1.2 Robustness to alternative parametric assumptions and estimates of the mapping function

As discussed above, the nonparametric test allows us to flexibly estimate the sign of health state dependence but does not provide information on the magnitude of state dependence. In this section therefore we examine the robustness of our quantitative estimates to alternative parametric assumptions or estimates of $g(\cdot)$. Table 9 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 examine the robustness of our estimates to an alternative common specification of the mapping. Specifically, we estimate a probit model rather than the linear probability model in the baseline specification.¹⁵ In the linear probability model, we implicitly assume that the probability that an individual responds “Yes” to the happiness question is a cardinal measure of true utility. In the probit specification, the implicit assumption is instead that the latent variable in the probit model is a cardinal measure of true utility. Column 2 presents results from a fixed effects probit model. In the table, we report marginal effects evaluated at the mean, but we use the original probit coefficients to calculate the magnitude of state dependence (i.e., $\sigma \beta_1 / \beta_4$).¹⁶ Our sample size is reduced because we cannot include individuals who never change their response to the happiness question, but we still estimate a statistically significant decline in marginal utility that is roughly twice as large as the baseline result (-28.7% versus -11.2%). The larger magnitude from the fixed effects probit specification appears to be a result of the different sample used rather than the alternative functional form; column 3 demonstrates that running the baseline (linear) specification on the same sample used in the fixed effects probit yields very similar results. Thus, we are reassured to find similar estimates for state dependence for two common but quite different specifications of the mapping

¹⁵ Our exploration in Section 5.2 of the robustness of our findings to alternative utility proxies (see Table 7 columns 8 through 10) can be seen as another way of exploring the robustness of our findings to alternative mappings from latent utility to the utility proxy.

¹⁶ Because of the large number of fixed effects in the nonlinear model, we use a “zig-zag” approach to maximizing the log-likelihood function (Heckman and MaCurdy 1980).

function of cardinal utility to subjective well-being.

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 (5) by nonlinear least squares; the coefficient of relative risk aversion, α , is given by one minus the exponent on permanent income (i.e., $\alpha = 1 - \beta_2$). 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 over time.¹⁷ Naturally, our estimate of α will be affected by our assumption regarding the mapping function $g(\cdot)$; this provides one way of gauging the plausibility of our assumed mapping.

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 3.5.¹⁸ Our estimates of negative state dependence remain 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. 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. Hence, both the linear and the probit mappings imply latent utility measures that contain plausible levels of curvature. 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 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. Then, using this estimate of the mapping $g(\cdot)$, we estimate the sign and

¹⁷ We then use our estimate of β_2 from the cross-sectional equation (5) in our panel equation (4). We iterate until β_2 converges.

¹⁸ Our estimates 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.

magnitude of health state dependence. We find that imposing CRRA values between 1 and 5 and flexibly estimating $g(\cdot)$ yields very similar estimates of state dependence.

We flexibly estimate the mapping $g(\cdot)$ by allowing the latent variable in the probit model to be a monotonically increasing ninth-degree polynomial function of our measure of cardinal utility, v :

$$\Pr(HAPPY_i) = g(v_i) = \Phi(h(v_i)),$$

where $\Phi(\cdot)$ denotes the standard normal cumulative density function and $h(\cdot)$ is the flexible polynomial. We impose on v the curvature that corresponds to a coefficient of relative risk aversion of α by specifying $v_i = \pi_1 \bar{Y}_i^{1-\alpha} / (1-\alpha) + \pi_0$. The function $h(\cdot)$ is identified cross-sectionally because we only have cross-sectional variation in permanent income \bar{Y}_i . Of course, the estimate $\hat{h}(\cdot)$ will depend on our assumed degree of risk aversion. Next, we define our mapping from our cardinal utility measure v to the utility proxy $HAPPY$ as $\hat{g}(\cdot) = \Phi(\hat{h}(\cdot))$. We use the estimated mapping $\hat{g}(\cdot)$ and set $\beta_2 = 1-\alpha$ when we estimate equation (4), which identifies state dependence by the interaction between permanent income and health in a panel regression with individual fixed effects. Appendix C provides further technical details of this procedure.

Columns 6 through 8 show the estimates for state-dependence when we select the mapping $g(\cdot)$ such that the coefficient of relative risk aversion in latent cardinal utility equals 1, 3, and 5, respectively. We choose these three values of relative risk aversion since they span the most commonly used values in the literature. 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 through 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 -20% and -29% for these three mappings, which is similar to the estimate of -29% in the standard fixed effects probit regression in column 2. The estimate is statistically significant at the 5 percent level for the mapping corresponding to a coefficient of relative risk aversion of 3, and for the other two mappings, the p-value of our estimate of state dependence lies around 10 percent.

Overall, we believe the evidence presented in section 6.1 greatly alleviates concerns that our results are primarily due to errors in the specification of the mapping from true cardinal utility to the utility proxy that are correlated with the interaction between permanent income and health,

though 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. This overlap ensures that in the semiparametric specifications the comparisons between high- and low-income individuals to a large extent take place in the same range of the mapping function rather than at different ranges. When high- and low-income individuals are in the same range of the mapping function, any errors in the mapping function will not be strongly correlated with permanent income or with the interaction of permanent income and health.

6.2 Omitted determinants of utility that vary with health by permanent income

We explore three possible threats to our second identifying assumption that, conditional on sickness S , control variables X , and fixed effects θ , there are no omitted determinants of utility that vary with health differentially by permanent income. These possible threats are: (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 consumption.

6.2.1 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 we find (in unreported regressions) that the consumption path of the poor declines (in percentage terms) relative to that of the rich over time; this would in fact bias us against our finding negative state dependence.¹⁹

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. We therefore allowed the effect of permanent income to vary not only with number of

¹⁹ See Appendix B for a description of the consumption data. An alternative way to investigate this issue would be to add an interaction of permanent income with time. Unfortunately, the high collinearity between time and the onset of a disease makes it hard to disentangle the two effects; not surprisingly, our estimate of the interaction of permanent income with health becomes insignificant (not shown).

diseases but also with marital status and with household size. The estimate of the interaction term of permanent income and number of diseases remains similar in magnitude to our baseline estimate, but is no longer statistically significant at conventional levels (not shown).²⁰

6.2.2 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 there is no bias from differential rates of disease occurrence for individuals of different permanent income. This rate simply affects the frequency with which we observe the sick state.

The existing evidence suggests that if any reporting differences by socio-economic status (SES) exist, they would likely bias against our finding of negative state dependence. Banks et al. (2006) compare the education-disease gradient for individuals aged 40 to 70 in the 1999-2002 National Health and Nutrition Examination Survey (NHANES) based on self-reported health measures and on biological measures.²¹ For hypertension, the gradients using the two different health measures are virtually indistinguishable; for diabetes, there is some evidence of under-reporting by individuals of lower education (Banks et al., 2006 Table 4). Our own analysis of the HRS shows that conditional on reporting that a doctor has told them they have a particular disease, individuals of higher permanent income are less likely to report conditions that indicate a more severe form of the disease (see Table 10). This suggests that the threshold for reporting a disease is higher for the poor, since 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 under-reporting 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.

6.2.3 What if consumption is not predetermined?

²⁰ Note that if the effective interest rate or discount rate varies by permanent income, it would not cause bias in our estimate of state dependence as long as second-period consumption is some monotonically increasing function of lifetime income. See Appendix D for details.

²¹ Both the NHANES and HRS self-reported measures are based on the question “has a doctor ever told you that you have X” rather than respondent’s subjective assessment, which may mitigate potential differential reporting.

We discuss the effect of relaxing our key modeling assumption from Section 4.1 that consumption in the sick state is pre-determined. As we explain, any resulting bias would likely be against our finding of negative state-dependent utility. We also present complementary empirical evidence that is consistent with the assumption of pre-determined consumption in our baseline sample as well as the empirical effects of relaxing it.

Relaxing the assumption of pre-determined consumption:

Consumption would not be pre-determined if health has a mechanical effect on the resources available to consume. For example, health may affect labor income, out-of-pocket medical expenditures, household production, and longevity. In addition, consumption would not be pre-determined if optimizing individuals are able to re-allocate consumption in response to their state-dependent preferences; this could occur if health shocks are anticipated or transitory, relatives provide informal insurance, or there is an outside and future consumption good whose marginal utility is state independent, such as bequests as in Lillard and Weiss (1997).

When consumption is not pre-determined, health can affect marginal utility of consumption for two distinct reasons: its direct effect on the slope of the utility function (i.e., the state-dependent parameter of interest), and its indirect effect through its effect on consumption. Appendix D provides a formal treatment of the effect of extending the baseline model to allow for health to directly affect consumption. Here, we limit ourselves to an informal discussion designed to convey the intuition behind the bias that may result.




Figure 2: Estimates of state dependence when consumption varies with health

Figure 2 illustrates the natural case in which consumption falls with health, due for example

to reduced income or increased medical expenses. Because our regression specification does not correct for the consumption drop (we proxy consumption 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 consumption associated with the health decline. Curvature is captured by the parameter α , the coefficient of relative risk aversion.

Consider first the case in which the drop in consumption is proportional to consumption in the healthy state. If $\alpha=1$, the slope of the utility curve is inversely proportional to consumption. 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 2. Since the estimate of the slope of the utility curve in poor health is unbiased, the estimate of γ_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 2006b) many papers estimate a substantially higher level of relative risk aversion (e.g., Gertner 1993, Cohen and Einav 2007). If $\alpha>1$, then the marginal utility of consumption falls more than proportionally with consumption, resulting in a smaller drop in utility for high-consumption individuals than for those at lower levels of consumption. This will cause the slope of the estimated utility curve for poor health to be biased upward, leading to a positive bias in γ_1 and biasing against our finding of negative state-dependent utility.

Other cases follow the same basic intuition. For example, the drop in consumption may be less than proportional to consumption in the healthy state if there are absolute expenses associated with poor health, such as medical expenses. As a result of these absolute expenses, the drop in utility associated with poor health is smaller at higher consumption levels, creating even more positive bias in our estimate of γ_1 , assuming $\alpha \geq 1$.²² Of course, if deterioration in health causes an increase in consumption – for example because of reduced life expectancy for individuals who are not completely annuitized – 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. We show below that any such bias is likely to be small in practice.

²² 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 higher consumption individuals (which seems empirically less relevant).

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.²³ In such cases, individuals would adjust their savings (re-allocate consumption) to equate marginal utility across current and future periods (and/or the outside good). For example, for negative state-dependent utility (i.e., $\gamma < 0$), people who know they will become sick will save less than they otherwise would have for future consumption. As a result of such re-allocation, our estimate of state-dependent utility would be biased toward zero (for $\alpha > 1$) and would be unbiased for $\alpha = 1$, whether or not the true state dependence were positive or negative.

Is consumption pre-determined in our sample? Some suggestive evidence

The evidence in Table 11 is consistent with our assumption that, in our sample of elderly individuals not in the labor force and with health insurance, consumption is pre-determined and cannot be adjusted in response to health events. We examine how current income and consumption change as health deteriorates. Since both income and consumption are household measures, our health measure becomes the number of chronic diseases that the respondent and his or her spouse report as a fraction of the maximum number of diseases possible (7 for singles and 14 for couples). We now include household fixed effects instead of individual fixed effects, and continue to include wave fixed effects and controls for time-varying household characteristics (specifically, household size, a quadratic in average household age, and a dummy for marital status). Column 1 indicates that a one-standard-deviation increase in the number of household diseases is associated with marginally statistically significant but economically insignificant 0.9 percent increase in current household income.

More informatively, we are also able to examine whether *consumption* changes with adverse health events. The Consumption and Activities Mail Survey (CAMS) – a small topical module administered to about 30% of households in the HRS for only three waves (i.e., about 10 percent of our person-years) – allows us to construct a broad-based measure of total consumption as well as non-durable consumption; Appendix B provides more detail. 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 insignificant 1.3 percent increase in total consumption and 1.5

²³ 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. It is less clear the extent to which they are anticipated. 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.”

percent increase in non-durable consumption. There is therefore no evidence of a statistically significant change in consumption or income following an adverse health event. Moreover, the resultant downward bias in our estimate of γ if we use the point estimate of the 1.5 percent increase in consumption following a one-standard-deviation decline in health is too small to explain our finding of negative state dependence.²⁴

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, it 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. In columns 2, 4 and 6, we find that neither the income nor the consumption response differs significantly by level of permanent income. The point estimate on the interaction term between the number of diseases and permanent income is 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, thus biasing against our finding of negative state dependence.

The most plausible reason why consumption might increase following an adverse health event – and thus bias 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 12 compares the results in the full sample (column 1) to results in a sample limited to those with more than 50 percent of their permanent income annuitized through Social Security or defined-benefit pensions (column 2) or more than 75 percent annuitized (column 3).²⁵ The implied

²⁴ As we derive in Appendix C, the bias can be computed given an assumption about the curvature of utility and the proportionality of the effect of health on consumption. We assume any increase in consumption following a negative health event is proportional to consumption when healthy. Equation (8) in Appendix C shows that with no state-dependent utility, our estimate of β_1/β_4 will be $-(\alpha-1)n$, where α is the coefficient of relative risk aversion and n is the proportional change in net resources from a health shock. For log utility ($\alpha = 1$) there is no bias. If we assume $\alpha = 3$, 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 et al. 2006) then the point estimate of a 1.5 percent increase in non-durable 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.

²⁵ The percent of permanent income annuitized is calculated by summing the total household income from IRAs, pensions, SSDI, and social security and averaging this total across all waves. This average total is then divided by the total household permanent income.

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 become only marginally significant.

As previously discussed, most empirically relevant scenarios suggest that the direct effect of health would be to decrease consumption and therefore bias against our finding of negative state-dependent utility. These scenarios are less of a concern because we can remain confident that our ability to reject the null of $\beta_1 = 0$ is interpretable as negative state-dependent utility. However, these potential sources of bias are relevant for the magnitude of our estimate. Two likely causes of decreased consumption when health deteriorates are the effect of health on labor market income and medical expenditures. In part for these reasons, we excluded individuals who were in the labor force and individuals without health insurance from our baseline sample. Columns 4 and 5 show the results from adding in, respectively, individuals without health insurance and individuals in the labor force; we continue to restrict attention to individuals aged 50 and over. The inclusion of individuals without health insurance has no detectable effect on our estimate of β_1 , perhaps because those individuals comprise less than 15 percent of our sample. By contrast, adding in individuals in the labor force has a substantial effect in the expected direction. Specifically, the negative coefficient on β_1 falls by half and is no longer statistically significant at conventional levels (p-value = 0.118). This is exactly the sign of the bias we would expect if deteriorating health reduces labor market earnings for those in the labor force, thus reducing resources and hence potential consumption.²⁶

Finally, we limit the sample to individuals who are at least 65 years old, which is the Medicare eligibility age. In this sample, 19.5 percent are in the labor force. By defining the sample based on age, we limit potential concerns about the sample selection on retirement and health insurance decisions. Column 6 shows that our estimate for state dependence is very similar for this sample.

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

²⁶ In a similar vein, we also examined a further restriction of the baseline sample to individuals whose spouses were also not in the labor force. The baseline estimates were not affected (not shown).

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 of 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 health deterioration is associated with a statistically significant decline in the marginal utility of consumption. Our central estimate is that a one-standard-deviation increase in the number of chronic diseases of an individual is associated with an 11 percent 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 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 1 percentage point (or about 4 percent).

Our findings also raise several important questions for future work. We estimate the average effect on marginal utility from the onset of different 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.

Appendix A: Derivation of formulas used in calibration

Optimal insurance calibration:

An agent faces probability p of receiving a health shock. If the agent receives the health shock, she receives insurance for a fraction b of her health expenses, H . When she does not receive a health shock, she pays an insurance premium τ . The agent maximizes utility given permanent income \bar{Y} :

$$\begin{aligned} \max_{C_h, C_s, H} & (1-p)u_h(C_h) + p(u_s(C_s) + \Psi(H)) \\ \text{s.t.} & \bar{Y} - \tau - C_h \geq 0 \\ & \bar{Y} - C_s - (1-b)H \geq 0 \end{aligned}$$

where $\Psi(H)$ is a function that captures the effect of health spending on utility, and C_h and C_s denote, respectively, consumption in the healthy and sick state. Utility is state dependent, with $u_h(\cdot)$ denoting the utility function in the healthy state and $u_s(\cdot)$ denoting the utility function in the sick state.

The planner chooses b and τ to maximize agent utility subject to the budget-balance constraint (insurance premiums collected equal benefits paid in expectation). Following the derivation in Chetty (2006), the problem above has the following exact solution:

$$\frac{u'_s(C_s) - u'_h(C_h)}{u'_h(C_h)} = \frac{d \log H}{d \log b} = \varepsilon_{H,b}.$$

Our modification to the standard Baily/Chetty set-up involves introducing state-dependent utility of the following form: $u'_s(C) = (1 + \gamma_1)u'_h(C) = (1 + \gamma_1)u'(C)$. This results in the following exact solution for optimal insurance:

$$\frac{(1 + \gamma_1)u'(C_s) - u'(C_h)}{u'(C_h)} = \varepsilon_{H,b}.$$

The above formula makes it clear that with no moral hazard and no state-dependent utility, there will be full insurance and consumption will be equalized across states. Using the agent's budget constraint and assuming CRRA utility, we can solve for the optimal level of insurance, b^* , in the above expression:

$$\frac{(1 + \gamma_1)(C_h - (1-b)H)^{-\alpha} - C_h^{-\alpha}}{C_h^{-\alpha}} = \varepsilon_{H,b} \quad \Leftrightarrow$$

$$b^* = 1 + \frac{C_h}{H} \left(\left(\frac{1 + \gamma_1}{1 + \varepsilon_{H,b}} \right)^{1/\alpha} - 1 \right).$$

This is the formula that we use to create the table of optimal insurance values in the paper. We choose $\varepsilon_{H,b} = 0.2$ based on Manning et al. (1987). We approximate $C_h / H = 3$ based on data on the distribution of health spending and the distribution of annual household consumption. Since H is the incremental health spending associated with becoming sick, we approximate it using data from the 2000 Medical Expenditure Panel Survey (MEPS) on the difference in mean total medical spending for those whose medical spending is above the median ($\sim \$9,800$) and those whose medical spending is below the median ($\sim \$790$). Using the consumption data in the CAMS survey (described in more detail in Appendix B), we find that median consumption is about $\$25,000$.²⁷ We use median consumption divided by the difference in average health spending (between average spending for those above and below the median) to get our estimate for $C_h / H : \$25,000 / (\$9,800 - \$790) \approx 3$. Using these two values and allowing γ_1 and α to vary, we calculate each cell in the calibration table.

Optimal savings calibration:

We use a two-period model of optimal savings for retirement. The agent must decide how much of her first period (fixed) wage w to save for consumption in the second period. The agent faces a probability p of becoming sick in the second period. Savings yield e^{rT} , where r is the annual rate of return on savings and T is the number of years between periods. As before, state-dependent utility is introduced by allowing a proportional change in the marginal utility of consumption if the agent receives a health shock: $u'_s(C) = (1 + \gamma_1)u'_h(C)$. Finally, the agent discounts the second period by $e^{\delta T}$ and has CRRA per-period utility (with coefficient of relative risk aversion α). The full maximization problem of the agent is the following:

$$\begin{aligned} \max_s & u_h(C_1) + e^{-\delta T} (p u_s(C_2) + (1-p) u_h(C_2)) \\ \text{s.t. } & s = w - C_1 \\ & C_2 = e^{rT} s \end{aligned}$$

where s is the fraction of the first-period (fixed) wage to save for retirement and $u_h(C) = \frac{C^{1-\alpha}}{1-\alpha}$.

Unconstrained maximization of the above problem yields the following formula for the optimal level of savings:

$$s^* = \frac{w}{1 + (1 - p + p(1 + \gamma_1))^{-1/\alpha} e^{rT} e^{-T(r-\delta)/\alpha}}$$

²⁷ Since consumption in the CAMS is calculated on a household basis, we then convert it to an individual-level variable using the OECD adjustment for household composition described in the text.

Appendix B: Data Appendix

I. Health and Retirement Study

Our analysis uses data from all cohorts (and their spouses) in the first seven waves of the HRS. The original HRS cohort is surveyed in every even year starting in 1992. The AHEAD cohort is surveyed in 1993, 1995, 1996, 1998, 2000, 2002 and 2004. The War Baby and CODA cohorts are surveyed in 1998, 2000, 2002 and 2004. For more detail on the data and the sample see <http://hrsonline.isr.umich.edu/intro/index.html>. We use the RAND HRS data set, which is a “cleaned, easy-to-use, streamlined version” (<http://hrsoline.isr.umich.edu/meta/r/and/>), and merge on some additional variables that are needed.

Sample selection:

- Aged 50 and older. This restriction is only binding for spouses, since the HRS only sampled main respondents age 50 and older.
- Not in labor force: We define an individual as not in the labor force if they (1) self report that they are either retired or that the retirement question is “inapplicable” (presumably reflecting no serious prior labor market attachment) *and* (2) have annual earnings of less than \$5,000. Since the retirement question is not asked in the 1994/1995 waves, we include individuals in this wave if they meet the criteria in the prior wave.
- Have health insurance: We define an individual as having health insurance if she is covered by any private or public insurance.
- We require that the individual maintains her retirement status and insurance coverage while she is in the sample. Individuals who do not initially meet these criteria can enter our sample in subsequent waves if they subsequently meet the criteria, but we drop all spells in the sample that do not terminate with the last observation of the individual meeting the sample selection criteria.²⁸
- We exclude the bottom percentile of the permanent income (defined below) distribution from our analysis, given the potential sensitivity of the coefficient on the log of permanent income (see equation 6) to such outliers. In practice, including these individuals does not have a substantive effect on the results.
- Finally, we require that the individual appear in the baseline sample for more than one wave, and only use person-years where the key variables have non-missing values.

Variable definitions

- *Annual household income (adjusted for household composition)*: Total annual household income is the sum of household income from wages and salaries, capital income (business income, dividend and interest income, and other asset income), pensions, government transfers and other sources. We also add 5% of the household’s current financial wealth (that is, total household wealth not including housing or automobile) to this aggregate household income measure to account for the fact that elderly households may be spending down their accumulated financial savings; results are unaffected if we instead assume a 10% or 0% “drawdown” rate of financial wealth. We use the OECD

²⁸ As a specification check, we also define a sample where once an individual enters the sample, the individual remains in the sample indefinitely regardless of changes to health insurance and retirement status, and the results are extremely similar. As an additional specification check, we applied the sample criteria on a year-by-year basis, and again find very similar results.

adjustment for household size (Atkinson et al. 1995), dividing total household income by 1.7 if the respondent is married and living with a spouse in the same household in that wave.

- *Permanent income*: Average across all waves of annual household income (adjusted for household composition)
- *Measures of chronic disease*: The exact question is “Has a doctor ever told you that you had X.” These have been coded in the RAND data set to be absorbing.
- *Wealth measure* (used in Table 6 column 4 as an alternative measure of permanent income): The wealth measure used is constructed by averaging household wealth across all waves in which a household appears. The measure of wealth we use excludes net housing wealth and automobile wealth. It includes the sum of the net value of financial wealth (e.g., stocks, mutual funds, investment trusts, checking, savings, money markets, CD’s, T-bills) and other savings and assets minus non-housing and non-automobile debts. We limit the sample to households with more than \$1,000 in wealth, which results in a roughly 20% reduction in sample from baseline sample.

Consumption and Activities Mail Survey (CAMS)

The Consumption and Activities Mail Survey (CAMS) – a small topical module administered to about 30% of households in the HRS for three waves – allows us to construct a broad-based measure of total consumption, as well as non-durable consumption. The CAMS survey was mailed to 5,000 households selected at random from the 13,214 households in HRS 2000; they received 3,866 responds in 2001 and followed up with the respondent sample in 2003 and 2005 to form a household-level panel data set on consumption.

We use all three waves of CAMS, matching each to the preceding HRS survey years since the CAMS asks about consumption in the previous year. The survey asks about 6 “big-ticket” durable consumption items and 26 non-durable consumption categories that are modeled after the Consumer Expenditure Survey (CEX) and designed to encompass the exhaustive set of non-durable consumption categories in the CEX. We follow Hurd and Rohwedder (2005) to construct measures of total consumption and total non-durable consumption; they also provide more detail on the survey and the underlying data. We exclude medical expenditure from both consumption measures.

II. British Household Panel Survey (BHPS)

The British Household Panel Survey (BHPS) is an annual longitudinal survey covering “labor markets outcomes, income, savings and wealth, household organization, housing, consumption, social and political values, education, and training.”²⁹ The original sample contained about 5,500 households and about 10,000 individuals. Initially, the response rate was very high (>95%), but by year 2000 there was moderate attrition (65% of wave 1 respondents who were still alive and in United Kingdom in 2000 gave an interview in 2000, and 55% of respondents gave responses in all waves).

We use the 14 waves between 1990 and 2003, and we impose baseline sample restrictions similar to those that we used with the HRS sample. In particular, just as in the HRS, we limit the BHPS sample to age greater than 50 who are not in the labor force and are covered by health

²⁹ The BHPS is described in more detail here: <http://www.iser.essex.ac.uk/ulsc/bhps/quality-profiles/BHPS-QP-01-03-06-v2.pdf>

insurance.

The proxy for utility that we use in the BHPS is the answer to the following question: “How dissatisfied or satisfied are you with your life overall?” Individuals respond on a 7-point scale where 1 corresponds to “not satisfied at all” and 7 corresponds to “completely satisfied” (intermediate values are not explicitly defined). This 7-point scale may capture movements in subjective well-being better than our binary happiness measure in the HRS. A possible drawback of the BHPS measure, however, is that it asks about a global evaluation of one’s life, which may be a harder question to answer than one’s current feeling of happiness (and, therefore, perhaps more prone to judgment biases). Further, a global evaluation of happiness may be affected by past or future well-being rather than just current well-being. A more general drawback to using the BHPS data set is the considerably smaller sample size of individuals who meet the baseline sample restrictions (only about 4,500 individuals and 20,000 observations in the BHPS versus roughly 11,500 individuals and 45,000 observations in the HRS).

As with the HRS data set, we compute household permanent income by taking an average of total household income (including pension payments, annuity payments, disability payments, government transfers and other government allowances) across all waves. The primary health measure that we use in the BHPS sample is the individual’s total number of reported diseases, using the diseases that appear consistently in all waves (excluding depression, alcohol use, drug use, and epilepsy): problems with arms/legs/hands, difficulty hearing, difficulty seeing, skin problems, chest/breathing problems, heart problems and/or hypertension, stomach/digestion problems, diabetes, migraines, and other diseases.³⁰ Note that this set of diseases is somewhat different from the HRS set (see Table 3 for a description of diseases used to construct the HRS health measure).

³⁰ Cancer and Stroke (which are diseases that also appear in HRS) appear only in the last 3 waves.

Appendix C: Semiparametric estimator of the mapping $g(\cdot)$

We generalize the standard probit model by flexibly estimating a nonlinear, monotonic transfer function $h(v)$. In our application, this transfer function maps cardinal utility v to the latent variable in a probit model with a binary subjective well-being outcome variable, *HAPPY*:

$$HAPPY_i = \begin{cases} 1 & \text{if } h(v_i) > \varepsilon_i \\ 0 & \text{if } h(v_i) \leq \varepsilon_i \end{cases},$$

where ε_i is a standard normal error term. The transfer function $h(v)$ is specified as a ninth-order polynomial that is constrained to be monotonically increasing using the rearrangement technique of Chernozhukov, Fernandez-Val, and Galichon (2009). Without loss of generality, we normalize $h(0)=0$ and $h'(0)=1$. We impose latent utility v to have the amount of curvature that corresponds to a coefficient of relative risk aversion of α :

$$v_i = \pi_1 \frac{\bar{Y}_i^{1-\alpha}}{1-\alpha} + \pi_0,$$

where π_1 and π_0 are parameters to be estimated. The polynomial coefficients and π_1 and π_0 are estimated by maximizing the following log likelihood function:

$$\max_{h(\cdot), \pi_0, \pi_1} \sum_{i,t} \left(HAPPY_{it} \times \log(\Phi(h(\pi_1 \frac{\bar{Y}_i^{1-\alpha}}{1-\alpha} + \pi_0))) + (1 - HAPPY_{it}) \times \log(1 - \Phi(h(\pi_1 \frac{\bar{Y}_i^{1-\alpha}}{1-\alpha} + \pi_0))) \right),$$

where $\Phi(\cdot)$ denotes the standard normal cumulative density function. The outcome of this maximization problem is an estimated transfer function $\hat{h}(\cdot)$, which will depend on our choice of α .

Next, we define the mapping from our cardinal utility measure v to the utility proxy *HAPPY* as $\hat{g}(\cdot) = \Phi(\hat{h}(\cdot))$. We use the estimated mapping $\hat{g}(\cdot)$ and set $\beta_2 = 1-\alpha$ when we estimate equation (4), which identifies state dependence by the interaction between permanent income and health in a panel model with individual fixed effects. We estimate equation (4) by maximum likelihood.³¹ Finally, using our estimated fixed effects, we estimate equation (5) which identifies the marginal utility of permanent income (β_4).

We report bootstrapped standard errors clustered by individual for two reasons. First, this is a three-step estimator – the first step estimates $h(\cdot)$, the second step estimates β_1 , fixed effects (θ_s), and other parameters given $\hat{h}(\cdot)$, and the third step estimates β_4 given the fixed effect estimates. Second, we are most interested in the magnitude of state dependence ($\sigma\beta_1/\beta_4$) and bootstrapping allows us to take into account the covariance between β_1 and β_4 , which are estimated in two separate equations. A single iteration of the three-step estimator takes about 4 hours to run, so we only run 100 iterations to compute our bootstrapped standard errors. We report p-values based on asymptotic t-tests constructed from our point estimate and the bootstrapped standard errors.

³¹ As with our baseline fixed effects probit model, we maximize the log likelihood function using a “zig-zag” approach which iterates between maximizing over the fixed effects holding the other parameters fixed and maximizing over the other parameters using current estimates of the fixed effects.

Appendix D: Estimates of State Dependence when Consumption Varies with Health

We model consumption responses to health states by allowing the individual's income to be state dependent. In particular, let the individual receive *net* income $\tilde{N}(S)$ in period 2 (in addition to the permanent income Y received in period 1). We think of $\tilde{N}(S)$ consisting of medical insurance payments, effects of health on labor income and household production, informal transfers from friends and family that depend on health status, or resources that would have otherwise been used on an outside state-independent consumption good (such as bequests). $\tilde{N}(S)$ is lowered by medical expenses and other possible negative shocks to income available for consumption.

The lifetime budget constraint now becomes:

$$Y = C_1 + \frac{1}{1+r}(C_2 - \tilde{N}(S)).$$

Further, assume that $\tilde{N}(0) = 0$, and $\tilde{N}(1) = N$. Because C_2 is stochastic, we express expected utility as a function of C_1 (rather than of C_2 as we did in the basic model):

$$E[U] = \frac{1}{1-\alpha} C_1^{1-\alpha} + \frac{1}{1+\delta} \left((1-p) \frac{1}{1-\alpha} ((1+r)(Y - C_1))^{1-\alpha} + \gamma_0 p + p(1+\gamma_1) \frac{1}{1-\alpha} ((1+r)(Y - C_1) + N)^{1-\alpha} \right).$$

There is no closed form solution for C_1 , but denote the solution of this maximization problem by C_1^* . Next, consider the effect of relaxing the budget constraint in the second period by providing additional resources R in the second period. Since we are only interested in the derivative of utility with respect to R , we can set the level of R at zero. The realized marginal utility of additional resources in period 2 depends on health and is given by:

$$\left. \frac{\partial U}{\partial R} \right|_{S=0} = \frac{1}{1+\delta} ((1+r)(Y - C_1^*))^{-\alpha}$$

$$\left. \frac{\partial U}{\partial R} \right|_{S=1} = \frac{1+\gamma_1}{1+\delta} ((1+r)(Y - C_1^*) + N)^{-\alpha}$$

Since C_1^* is chosen before the random variable health status is realized, C_1^* is independent of health status for any individual. Next, we define net income shocks n as a fraction of second period consumption in the healthy state:

$$n = \frac{N}{(1+r)(Y - C_1^*)}$$

This allows us to express the marginal utility of resources in the sick state as a fraction of the marginal utility of resources in the healthy state:

$$\frac{\partial U}{\partial R} \Big|_{S=1} \Big/ \frac{\partial U}{\partial R} \Big|_{S=0} = (1 + \gamma_1)(1 + n)^{-\alpha}$$

This expression tells us that the ratio of the marginal utility of resources depends on two factors. First, it is proportional to the marginal utility of consumption in ill health to the marginal utility of consumption in good health when both marginal utilities are evaluated *at the same level of consumption*. Second, it depends on the effect of illness on the level of consumption. If the individual's net income falls in poor health ($n < 0$), her level of consumption will fall with illness. Since the elasticity of marginal utility with respect to consumption is $-\alpha$, this second factor will increase her marginal utility of resources as health declines.

This expression also gives insight into the optimal level of net state-dependent income. This income should depend on health such that the marginal utility of resources is equalized across states of the world. So, the optimal level of net state-dependent income is:

$$n^* = (1 + \gamma_1)^{1/\alpha} - 1 \approx \gamma_1 / \alpha$$

Thus, absent state dependence ($\gamma_1 = 0$), the optimal level of net state-dependent income is zero ($n^* = 0$), i.e., it is optimal for health expenditure shocks to be fully covered by insurance. However, if the marginal utility of consumption is lower in poor health ($\gamma_1 < 0$), then less than full insurance is first-best optimal. In other words, even in the absence of moral hazard, partial insurance is optimal in this case. Similarly, if $\gamma_1 > 0$, optimal insurance would more than cover medical expenditure.

Even though we cannot obtain a closed form solution for C_1 , it is clear that second-period resources from savings are increasing in permanent income. We parameterize this relationship as $(1+r)(Y-C_1) = \theta_0 Y^{\theta_1}$, with $\theta_0 > 0$ and $\theta_1 > 0$. Modeling second-period resources as a monotonically increasing function of permanent income also captures cases in which the effective interest rate, discount rate, or probability of diseases varies by permanent income. It follows that second-period indirect utility, $v(Y,S)$, equals:

$$v(Y,0) = \frac{1}{1-\alpha} (\theta_0 Y^{\theta_1})^{1-\alpha} = \frac{\theta_0^{1-\alpha}}{1-\alpha} Y^{(1-\alpha)\theta_1}, \text{ and}$$

$$v(Y,1) = \gamma_0 + \frac{1+\gamma_1}{1-\alpha} ((1+n)\theta_0 Y^{\theta_1})^{1-\alpha} = \gamma_0 + \frac{(1+n)^{1-\alpha} (1+\gamma_1) \theta_0^{1-\alpha}}{1-\alpha} Y^{(1-\alpha)\theta_1}.$$

Thus, running the regression given by equation (3) yields the following parameter estimates:

$$\beta_1 = \left(\frac{\theta_0^{1-\alpha}}{1-\alpha} \right) \left((1+n)^{1-\alpha} (1+\gamma_1) - 1 \right), \quad \beta_2 = (1-\alpha)\theta_1, \quad \beta_3 = \gamma_0, \text{ and} \quad \beta_4 = \frac{\theta_0^{1-\alpha}}{1-\alpha}.$$

As before, the parameter ratio β_1/β_4 allows us to estimate state dependence:

$$\beta_1 / \beta_4 = (1+n)^{1-\alpha}(1+\gamma_1) - 1 \approx \gamma_1 - (\alpha-1)n(1+\gamma_1). \quad (8)$$

This expression formalizes the intuition developed from Figure 2 concerning the bias from having net state-dependent income. If $\alpha = 1$, the ratio β_1/β_4 yields an unbiased estimate of the state dependence in the marginal utility of consumption, even in the presence of state-dependent net income shocks. For $\alpha > 1$, the bias has the opposite sign as the sign of the net state-dependent income shocks.

In certain cases, individuals may be able to choose the level of net state-dependent income. This may occur if there are well functioning (informal) insurance networks or if the individual has an outside good of which the utility does not depend on health. In that case, individuals would set n such that the marginal utility of resources is equalized across health states, so $n^* = (1+\gamma_1)^{1/\alpha} - 1$. The estimated ratio of the marginal utility of permanent income now becomes:

$$\beta_1 / \beta_4 = (1+n)^{1-\alpha}(1+\gamma_1) - 1 = ((1+\gamma_1)^{1/\alpha})^{1-\alpha}(1+\gamma_1) - 1 = (1+\gamma_1)^{1/\alpha} - 1$$

As before, there is no bias if $\alpha=1$. If $\alpha>1$, the estimate of state dependence is biased towards zero.

We can also model predictable or temporary health changes in this framework. Individuals who can predict health changes will adjust their savings such that the marginal utility of second-period consumption is equal to the marginal utility of first-period consumption. Such individuals can effectively self-insure, so we can think of them as selecting n such that the marginal utility of consumption is equalized across periods. However, the cost of the insurance is fully born by the person knowing that he will fall ill, so the analogy with insurance would only hold exactly if we could ignore this cost. Since the analogy does not hold exactly, we now derive the consumption decisions of individuals who can predict their future health status.

The first-order condition for those knowing they will fall ill is given by:

$$C_1^{-\alpha} - \frac{1+r}{1+\delta}(1+\gamma_1)((1+r)(Y-C_1))^{-\alpha} = 0,$$

while the first-order condition for those knowing they will remain healthy is given by:

$$C_1^{-\alpha} - \frac{1+r}{1+\delta}((1+r)(Y-C_1))^{-\alpha} = 0.$$

Both first order conditions have analytical solutions. Substituting these solutions into the second-period utility function yields indirect utility as a function of income and health status. As before, the indirect utility function is linear in $Y^{1-\alpha}$. The estimator β_4 is found as the coefficient on $Y^{1-\alpha}$ in the indirect utility function when healthy and β_1 is the additional effect of $Y^{1-\alpha}$ when sick.

When the discount rate equals the interest rate, the ratio of these parameters is given by:³²

$$\beta_1 / \beta_4 = \left(\frac{2+r}{1+r+(1+\gamma_1)^{1/\alpha}} \right)^{1-\alpha} (1+\gamma_1)^{1/\alpha} - 1.$$

As before, if $\alpha = 1$, then the ratio of the marginal utility of permanent income, β_1/β_4 , yields an unbiased estimate of γ_1 .

Taking a first-order Taylor series approximation around $\gamma_1=0$, yields the following expression for the bias term:

$$(\beta_1 / \beta_4 - \gamma_1) \approx - \left(\frac{(\alpha - 1)(1 + r)}{\alpha(2 + r)} \right) \gamma_1$$

This means that for $\alpha > 1$, the estimate of γ_1 is biased toward zero.

³² The expression for $\delta \neq r$ is lengthier and therefore harder to interpret, but qualitatively similar.

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(1): 123-129.
- Arrow, Kenneth. 1974. "Optimal Insurance and Generalized Deductibles." *Scandinavian Actuarial Journal* 1-42. Reprinted in *Collected Papers of Kenneth J. Arrow, Volume 3*, Cambridge MA, Harvard University Press, 1984, 212-260.
- Atkinson, A.B., L. Rainwater, and T. M. Smeeding. 1995. "Income Distribution in OECD Countries." *OECD Social Policy Studies* No. 18:
<http://www.oecd.org/dataoecd/61/52/35411111.pdf>
- Baily, Martin Neil. 1978. "Some Aspects of Optimal Unemployment Insurance." *Journal of Public Economics* 10(3): 379-402.
- 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(17): 2037-2045.
- Bernheim, B. Douglas, Andrei Shleifer, and Lawrence Summers. 1985. "The Strategic Bequest Motive." *Journal of Political Economy* 93(6): 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.
- Chernozhukov, Victor, Ivan Fernandez-Val, and Alfred Galichon. 2009. "Improving Point and Interval Estimates of Monotone Functions by Rearrangement," *Biometrika*, forthcoming.
- Chetty, Raj. 2006a. "A General Formula for the Optimal Level of Social Insurance." *Journal of Public Economics* 90(10-11): 1879-1901.
- Chetty, Raj. 2006b. "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. 2006. “Differential Mortality, Uncertain Medical Expenses, and the Saving of Elderly Singles.” NBER working paper no. 12554.

Di Tella, Rafael, and Robert MacCulloch. 2006. “Some Uses of Happiness Data in Economics.” *Journal of Economic Perspectives* 20(1): 25-46.

Di Tella, Rafael, Robert J. MacCulloch, and Andrew J. Oswald. 2001. “Preferences over Inflation and Unemployment: Evidence from Surveys of Happiness.” *American Economic Review* 91(1): 335-41.

Dor, Avi, Joseph Sudano, and David Baker. 2006. “The Effect of Private Insurance on the Health of Older, Working Age Adults: Evidence from the Health and Retirement Study.” *Health Services Research* 41(3): 759-787.

Engen, Eric M., William G. Gale and Cori E. Uccello. 1999. “The Adequacy of Household Savings.” *Brookings Papers on Economic Activity* 2: 65-187.

Evans, William N., and W. Kip Viscusi. 1991. “Estimation of State-Dependent Utility Functions Using Survey Data.” *Review of Economics and Statistics* 73(1): 94-104.

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(2): 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), forthcoming.

Frey, Bruno S. and Alois Stutzer. 2002. “What Can Economists Learn from Happiness Research?” *Journal of Economic Literature* 40(2): 402-435.

Gertner, Robert. 1993. “Game Shows and Economic Behavior: Risk-Taking on ‘Card Sharks’.” *Quarterly Journal of Economics* 108(2), 507-521.

Golosov, Mikhail and Aleh Tsyvinski. 2006. “Designing Optimal Disability Insurance: A Case for Asset Testing.” *Journal of Political Economy* 114(2): 257-269.

Gruber, Jonathan H. and Sendhil Mullainathan. 2005. “Do Cigarette Taxes Make Smokers

Happier?" *B.E. Journals in Economic Analysis and Policy: Advances in Economic Analysis and Policy* 5(1): 1-43.

Hall, Robert E., and Charles I. Jones. 2007. "The Value of Life and the Rise in Health Spending." *Quarterly Journal of Economics* 122(1): 39-72.

Heckman, James J. and Thomas E. Macurdy. 1980. "A Life Cycle Model of Female Labour Supply." *The Review of Economic Studies* 47(1): 47-74.

Hurd, Michael D. and Susan Rohwedder. 2005. "Consumption and Economic Well-Being at Older Ages: Income- and Consumption-based Poverty Measures in the HRS." Prepared for the 7th Annual Conference of the Retirement Research Consortium "Towards a Secure Retirement System." <http://www.mrrc.isr.umich.edu/publications/conference/pdf/cp05-Hurd.pdf>

Krueger, Alan B., Daniel Kahneman, David Schkade, Norbert Schwarz, and Arthur A. Stone. 2008. "National Time Accounting: The Currency of Life." Unpublished manuscript, Princeton University.

Layard, Richard, Guy Mayraz, and Stephen J. Nickell. 2008. "The Marginal Utility of Income." *Journal of Public Economics* 92(8-9): 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(2): 254-268.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2003. "Projection Bias in Predicting Future Utility." *Quarterly Journal of Economics* 118(4): 1209-1248.

Luttmer, Erzo F.P. 2005. "Neighbors as Negatives: Relative Earnings and Well-Being." *Quarterly Journal of Economics*, 120(3): 963-1002.

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.

Palumbo, Michael G. 1999. "Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle." *Review of Economic Studies* 66(2): 395-421

Oswald, Andrew J. and Nattavudh Powdthavee. 2007. "Does Happiness Adapt? A Longitudinal Study of Disability with Implications for Economists and Judges," *Journal of Public Economics*,

forthcoming.

Radloff, Lenore S. 1977. "The CES-D Scale: A Self-Report Depression Scale for Research in the General Population." *Applied Psychological Measurement* 1(3): 385-401.

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun. 2006. "Are Americans Saving 'Optimally' for Retirement?" *Journal of Political Economy* 114(4): 607-643.

Schwarz, Norbert, and Fritz Strack. 1991. "Evaluating One's Life: a Judgment Model of Subjective Well-Being." In *Subjective Well-Being: An Interdisciplinary Perspective*, ed. Fritz Strack, Michael Argyle, Norbert Schwarz, 27-47. Pergamon Press.

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(4): 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*, ed. David Wise, 213-237. University of Chicago Press.

Smith, Dylan M., Kenneth M. Langa, Mohammed U. Kabeto, and Peter A. Ubel. 2005. "Health, Wealth, and Happiness: Financial Resources Buffer Subjective Well-Being After the Onset of a Disability." *Psychological Science* 16(9): 663-666.

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.

Zeckhauser, Richard. 1970. "Medical Insurance: A Case Study of the Tradeoff Between Risk Spreading and Appropriate Incentives." *Journal of Economic Theory* 2(1): 10-26.

TABLE 3
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.)	
<u>Demographics</u>								
<i>NUM_WAVE</i>	45447	4.52	1.50	2	4	7	0	— Number of waves that respondent was interviewed.
\bar{Y} (Permanent income, \$)	45447	29224	33297	6236	20667	77285	0	— Permanent income constructed by taking the average across all waves of total household income plus a 5 percent 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).
<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	— Household size includes all residents of household (including spouse).
<u>Measures of Health</u>								
<i>NUM_DISEASE</i>	45447	1.95	1.30	0	2	4	0.63	— Sum of Yes/No "Has a doctor ever told you have <i>D</i> ?" (0-7)
<i>SPOUSE_NUM_DISEASE</i>	45447	1.03	1.30	0	0	4	0.71	— Sum of spouse's Yes/No "Has a doctor ever told you have <i>D</i> ?" (0-7)
<i>ADL_TOTAL</i>	45447	0.44	1.05	0	0	3	0.74	— Sum of Yes/No "Does anyone help you <i>A</i> ?" (0-6)
<i>IADL_TOTAL</i>	45384	0.41	0.89	0	0	2	0.67	— Sum of Yes/No "Are you able to <i>I</i> ?" (0-6)
<i>OFL_TOTAL</i>	45446	2.75	2.70	0	2	8	1.69	— Sum of "How difficult is <i>O</i> ?" (1 = Very or somewhat difficult) (0-10)
<i>HEALTH_INDEX</i>	45334	4.21	4.50	0	3	14	2.67	— Sum of severity of pain (0-2), ADL, IADL, and OFL (1-24)
<u>Measures of Utility</u>								
<i>HAPPY</i>	45447	0.87	0.34	0	1	1	0.28	— Yes/No "Much of the time the past week I felt happy?"
<i>CESD-8</i>	45447	6.32	2.01	2	7	8	1.38	— Sum of Yes/No "Much of the time the past week I felt/was <i>C</i> ?"
<i>CESD-4</i>	45447	3.38	1.05	1	4	4	0.79	— Subset of 4 out of 8 <i>CESD-8</i> questions (enjoy life, happy, sad, lonely).

Notes: Set of diseases: *D* = {hypertension, diabetes, cancer, heart disease, chronic lung disease, stroke, arthritis}. Set of Activities of Daily Living (ADLs): *A* = {dress, bathe or shower, walk across a room, eat (such as cutting up your food), get in and out of bed, use the toilet (including getting up and down)}. Set of Instrumental Activities of Daily Living (IADLs): *I* = {prepare hot meals, shop for groceries, make telephone calls, take medications, use a map, use a calculator}. Set of Other Functional Limitations (OFLs): *O* = {walk several blocks, walk one block, sit up for about 2 hours, get up from a chair, climb several flights of stairs, climb one flight of stairs, stoop/kneel/crouch, pick up a dime, extend your arms above shoulder level, push large objects like a living room chair}. Set of medications taken for disease: *M* = {hypertension, diabetes, chronic lung disease, heart disease, arthritis, stroke}. Set of CESD items: *C* = {depressed, everything I did was an effort, my sleep was restless, happy, lonely, enjoyed life, sad, could not 'get going'}. Spouse diseases set to 0 if the respondent is single.

TABLE 4
ESTIMATED MAGNITUDE OF STATE-DEPENDENT UTILITY

Dependent Variable: <i>HAPPY</i>	
<u>Panel A: Estimates</u>	
$NUM_DISEASE_{it} \times \log(\bar{Y}_i)$ (β_1)	-0.009 (0.004) [0.018]
$NUM_DISEASE_{it}$ (β_3)	-0.011 (0.003) [0.001]
$\log(\bar{Y}_i)$ (β_4)	0.048 (0.003) [0.000]
R^2	0.474
N	45447
Number of individuals	11514
<u>Panel B: Interpretation</u>	
% change in marginal utility for a 1 unit increase in $NUM_DISEASE_{it}$ (β_1/β_4) (bootstrapped standard error) (bootstrapped 95% CI) [bootstrapped p-value]	-17.9% (5.8%) (-4.2%, -26.9%) [0.018]
% change in marginal utility for a 1 standard deviation increase in $NUM_DISEASE_{it}$ ($\sigma\beta_1/\beta_4$) (bootstrapped standard error) (bootstrapped 95% CI) [bootstrapped p-value]	-11.2% (3.6%) (-2.7%, -16.8%) [0.018]
Marginal utility with 0 diseases (14.2% of sample)	0.048
Marginal utility with 1 disease (27.3% of sample)	0.039
Marginal utility with 2 diseases (28.4% of sample)	0.030
Marginal utility with 3 diseases (18.4% of sample)	0.021
Marginal utility with 4 diseases (8.9% of sample) (3.2% of sample has 5+ diseases)	0.012

Notes:

Panel A reports coefficients β_1 and β_3 from estimating equation (6) and coefficient β_4 from estimating equation (7). The dependent variable in equation (6) 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 (7) is the individual fixed effect estimated in equation (6). $NUM_DISEASE_{it}$ 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 have demeaned $\log(\bar{Y}_i)$. In addition to the covariates shown in the table, the estimates of equation (6) 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 (7) include wave fixed effects and a dummy for Non-white, a dummy for Female, Age, Age², Household size, and a dummy for whether the individual is single. Standard errors for β_1 , β_3 , and β_4 are in parentheses and are adjusted to allow for an arbitrary variance-covariance matrix for each individual over time; p-values are in brackets. Appendix B contains more detail on data and variable definitions; summary statistics are in Table 3.

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_4 and $\sigma\beta_1/\beta_4$ based on 10,000 bootstrap iterations, resampling individuals with replacement. The within-person standard deviation change in $NUM_DISEASE_{it}$ (σ) is 0.625.

TABLE 5
INDIVIDUAL DISEASE RESULTS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Blood pressure	Diabetes	Cancer	Lung Disease	Heart Disease	Stroke	Arthritis	Prevalence-weighted linear combination of (1)-(7)
$DISEASE_{it} \times \log(\bar{Y}_i)$ (β_1)	-0.020 (0.010) [0.045]	-0.013 (0.017) [0.425]	-0.003 (0.013) [0.802]	-0.045 (0.019) [0.017]	0.001 (0.011) [0.906]	-0.019 (0.018) [0.280]	0.007 (0.008) [0.350]	-0.008 (0.004) [0.032]
$DISEASE_{it}$ (β_3)	0.008 (0.008) [0.293]	0.004 (0.011) [0.694]	-0.025 (0.011) [0.020]	-0.042 (0.014) [0.003]	-0.016 (0.009) [0.072]	-0.001 (0.012) [0.904]	-0.017 (0.006) [0.009]	
$\log(\bar{Y}_i)$ (β_4)								0.047 (0.003) [0.000]
R ²								0.474
N								45447
Number of individuals								11514
% change in marginal utility for a 1 std. dev. change in $NUM_DISEASE_{it}$ ($\sigma\beta_1/\beta_4$)								-10.5% [0.029]
F-statistic that all interaction terms are equal								1.64
p-value of F-test, F(6,11513)								0.131

Notes: This table shows results from estimating a modified version of equations (6) and (7) in which seven disease dummies ($DISEASE_{it}$) indicating whether the respondent has the particular disease listed in the column heading are separately interacted with log permanent income in a single regression. See notes to Table 4 (Panel A) for more details on the estimating equations. The prevalence-weighted linear combination of the interaction terms (shown in column 8) gives an estimate of state-dependent utility that is comparable to the baseline specification (see Table 4). 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 $\sigma\beta_1/\beta_4$ is bootstrapped based on 10,000 iterations, resampling individuals with replacement.

TABLE 6
ALTERNATIVE SPECIFICATIONS

	(1)	(2)	(3)	(4)	(5)
	Baseline	No covariates	Always single	Own and spousal health	Habituation
$NUM_DISEASE_{it} \times \log(\bar{Y}_i)$ (β_1)	-0.009 (0.004) [0.018]	-0.008 (0.004) [0.023]	-0.003 (0.007) [0.632]	-0.008 (0.004) [0.022]	-0.011 (0.006) [0.062]
$NUM_DISEASE_{it}$ (β_3)	-0.011 (0.003) [0.001]	-0.012 (0.003) [0.000]	-0.010 (0.007) [0.142]	-0.010 (0.003) [0.002]	-0.010 (0.004) [0.029]
$SPOUSE_NUM_DISEASE_{it} \times \log(\bar{Y}_i)$				0.005 (0.004) [0.156]	
$SPOUSE_NUM_DISEASE_{it}$				-0.011 (0.003) [0.001]	
$NUM_DISEASE_{i,(t-1)} \times \log(\bar{Y}_i)$					-0.001 (0.006) [0.815]
$NUM_DISEASE_{i,(t-1)}$					-0.002 (0.005) [0.629]
$\log(\bar{Y}_i)$ (β_4)	0.048 (0.003) [0.000]	0.052 (0.003) [0.000]	0.049 (0.006) [0.000]	0.041 (0.003) [0.000]	0.048 (0.003) [0.000]
R ²	0.47	0.47	0.48	0.47	0.55
N	45447	45447	13437	45447	33910
Number of individuals	11514	11514	3411	11514	9032
Within-person std dev of $NUM_DISEASE_{it}$ (σ)	0.625	0.626	0.634	0.625	0.555
% change in marginal utility for a 1 std. dev. change in $NUM_DISEASE_{it}$ ($\sigma\beta_1/\beta_4$)	-11.2% [0.018]	-9.9% [0.022]	-4.2% [0.638]	-12.7% [0.022]	-12.3% [0.063]

Notes: Column (1) reports the results from the baseline specification in Table 4; see notes to Table 4 (Panel A) for more details. Subsequent columns report results from a single change relative to the baseline specification. Column (2) omits the covariates Age, Age², Household size, and Single from equation (6). Column (3) restricts the sample to individuals who are always single. Column (4) includes the total number of reported diseases of the spouse. Column (5) includes one-wave lags to test for adaptation/habituation. 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 $\sigma\beta_1/\beta_4$ is bootstrapped based on 10,000 iterations, resampling individuals with replacement.

TABLE 7
ALTERNATIVE MEASURES OF KEY VARIABLES

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Baseline	Permanent Income replaced by:		NUM_DISEASE _{it} replaced by:				HAPPY _{it} replaced by:		
		Education	Wealth	ADL TOTAL	IADL TOTAL	OFL TOTAL	HEALTH INDEX	CESD-8	CESD-4	LIFE_SAT (BHPS)
NUM_DISEASE _{it} × log(\bar{Y}_i) (β_1)	-0.009 (0.004) [0.018]	-0.002 (0.001) [0.016]	-0.004 (0.002) [0.052]	-0.005 (0.004) [0.201]	-0.010 (0.004) [0.020]	-0.002 (0.001) [0.086]	-0.002 (0.001) [0.060]	-0.038 (0.018) [0.040]	-0.016 (0.010) [0.107]	-0.032 (0.033) [0.323]
NUM_DISEASE _{it} (β_3)	-0.011 (0.003) [0.001]	-0.011 (0.003) [0.001]	0.029 (0.020) [0.139]	-0.022 (0.003) [0.000]	-0.022 (0.003) [0.000]	-0.010 (0.001) [0.000]	-0.009 (0.001) [0.000]	-0.150 (0.017) [0.000]	-0.050 (0.009) [0.000]	-0.053 (0.016) [0.001]
log(\bar{Y}_i) (β_4)	0.048 (0.003) [0.000]	0.010 (0.001) [0.000]	0.013 (0.001) [0.000]	0.031 (0.003) [0.000]	0.034 (0.003) [0.000]	0.033 (0.003) [0.000]	0.028 (0.003) [0.000]	0.613 (0.020) [0.000]	0.200 (0.010) [0.000]	0.143 (0.035) [0.000]
R ²	0.47	0.47	0.45	0.48	0.48	0.48	0.48	0.66	0.60	0.64
N	45447	45404	35698	45447	45384	45446	45334	45447	45447	19586
Number of individuals	11514	11501	8911	11514	11504	11513	11498	11514	11514	4676
Within-person std dev of NUM_DISEASE _{it} (σ)	0.625	0.626	0.631	0.738	0.669	1.686	2.673	0.625	0.625	0.620
% change in marginal utility for a 1 std. dev. change in NUM_DISEASE _{it} ($\sigma\beta_1/\beta_4$)	-11.2% [0.018]	-13.0% [0.011]	-16.8% [0.057]	-12.2% [0.204]	-19.2% [0.021]	-12.6% [0.081]	-17.4% [0.065]	-3.8% [0.045]	-5.1% [0.119]	-14.0% [0.329]

Notes: Column (1) reports the results from the baseline specification in Table 4; see notes to Table 4 (Panel A) for more details. Subsequent columns report results from alternative measures of key variables. Column (2) uses years of education (top-coded at 17 years) instead of permanent income. Column (3) uses log net worth excluding net housing wealth and net automobile wealth (see Appendix B and accompanying text) instead of permanent income. Columns (4) through (7) use other composite health measures described in Table 3 instead of total number of diseases. Columns (8) and (9) use other subjective well-being measures from the HRS. Column (10) uses a different data set (the British Household Panel Survey) and a 7-point measure of overall life satisfaction as the dependent variable; see Appendix B for more details. 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 $\sigma\beta_1/\beta_4$ is bootstrapped based on 10,000 iterations, resampling individuals with replacement.

TABLE 8
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, $NUM_DISEASE_{it} \times \log(\bar{Y}_i)$</i>					
Mean of CDFs of t-statistics, F	0.402	0.420	0.383	0.420	0.393
p-value of test that $F = 0.5$	[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 t-statistics, F	0.335	0.363	0.340	0.383	0.313
p-value of test that $F = 0.5$	[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 6.1. The sample includes all individuals with variation in the variable *HAPPY*; we drop all individuals with mean values of *HAPPY* of either 0 or 1. Individual mean values of *HAPPY* and mean values of *NUM_DISEASE* 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*, and their interaction. The cumulative distribution of the t-statistic on the interaction term is recorded, and we report the p-value from a two-sided test of whether the average of these recorded values is equal to 0.5 to test for state dependence. We do the same for the t-statistic on the number of diseases.

TABLE 9
ROBUSTNESS TO ALTERNATIVE MAPPINGS FROM CARDINAL UTILITY TO THE UTILITY PROXY

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Coefficient of relative risk aversion, α :	$\alpha=1$	$\alpha=1$	$\alpha=1$	α estimated		$\alpha=1$	$\alpha=3$	$\alpha=5$
Specification:	OLS	FE Probit	OLS (drop all 0s and all 1s)	OLS w/ NLLS	FE Probit w/ NLLS	Semiparametric Fixed Effects Probit		
$NUM_DISEASE_{it} \times \bar{Y}_i^{1-\alpha} / (1-\alpha)$ (β_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]	-22.097 (11.043) [0.045]	-0.016 (0.006) [0.005]	-14.787 (6.096) [0.015]	-535.557 (171.173) [0.002]
$NUM_DISEASE_{it}$ (β_3)	-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]	-0.115 (0.024) [0.000]	-0.132 (0.050) [0.008]	-0.143 (0.046) [0.002]
$\bar{Y}_i^{1-\alpha} / (1-\alpha)$ (β_4)	0.048 (0.003) [0.000]	0.059 (0.005) [0.000]	0.056 (0.005) [0.000]	1.253 (0.032) [0.000]	57.940 (2.559) [0.000]	0.051 (0.016) [0.001]	34.856 (10.701) [0.001]	1301.534 (386.552) [0.001]
Coefficient of relative risk aversion (α)				2.100 (0.148) [0.000]	3.464 (0.467) [0.000]			
R ²	0.474		0.17	0.47				
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 $NUM_DISEASE_{it}$ (σ)	0.625	0.638	0.638	0.625	0.638	0.638	0.638	0.638
% change in marginal utility for a 1 std. dev. change in $NUM_DISEASE_{it}$ ($\sigma\beta_1/\beta_4$)	-11.2% [0.018]	-28.7% [0.012]	-30.7% [0.016]	-11.2% [0.007]	-24.3% [0.002]	-20.3% [0.118]	-28.7% [0.011]	-26.3% [0.078]

Notes: Column (1) reports the results from the baseline specification in Table 4; see notes to Table 4 (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 our baseline (linear probability model) specification. Columns (6) through (8) report results from a semiparametric approach; see Appendix C 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 $\sigma\beta_1/\beta_4$ is bootstrapped base on 10,000 iterations, resampling individuals with replacement. Because of computational complexity, in columns (6) through (8) the p-value for $\sigma\beta_1/\beta_4$ is based on an asymptotic t-test using the t-statistic derived from the point estimate of $\sigma\beta_1/\beta_4$ divided by the bootstrapped standard error of this ratio, using 100 bootstrap iterations. In columns (6) through (8), all standard errors are bootstrapped standard errors based on 100 bootstrap iterations, resampling individuals with replacement.

TABLE 10
DISEASE NUMBER AND SEVERITY BY PERMANENT INCOME

	(1)	(2)	(3)	(4)	(5)	(6)
	Sample: has {lung disease, diabetes, stroke}					
Dependent Variable:	Number of diseases	Taking oxygen for lung disease	Severe diabetes, (1-2)	Severe diabetes dummy	Severe stroke, (1-7)	Severe stroke dummy
$\log(\bar{Y}_i)$	-0.231 (0.016) [0.000]	-0.037 (0.010) [0.000]	-0.028 (0.013) [0.031]	-0.025 (0.013) [0.047]	-0.194 (0.050) [0.000]	-0.037 (0.014) [0.010]
R ²	0.04	0.02	0.03	0.03	0.02	0.01
N	45447	4864	7927	7927	4387	4387
Mean of dependent variable	1.510	0.009	0.036	0.034	0.076	0.034
Within-individual std. dev. of dep. var.	0.625	0.275	0.281	0.257	1.370	0.433

Notes: Table reports results from a regression of the dependent variable shown in the column heading on , wave fixed effects Age, Age², Household size, and a dummy for whether the individual is single. In column (3), severe diabetes is defined as the sum of two dummy variables for whether the respondent takes insulin and whether the respondent has ever been hospitalized for kidney problems; in column (4), severe diabetes is defined if either dummy equals 1. In column (5), severe stroke is defined as the sum of 7 dummy variables for whether the respondent has vision problems, memory problems, speech problems, has seen a doctor recently, has general weakness from stroke, has therapy from stroke, or whether the respondent has other long-lasting problems from stroke; in column (6), severe stroke is defined if any of the 7 dummy variables equals 1. Standard errors, adjusted to allow for an arbitrary variance-covariance matrix for each individual over time, are in parentheses and p-values are in brackets.

TABLE 11
INCOME AND CONSUMPTION RESPONSE TO DISEASE

	(1)	(2)	(3)	(4)	(5)	(6)
	Current income		Total consumption		Non-durable consumption	
Household Fraction of Diseases _{it}	0.109 (0.056) [0.052]	0.106 (0.059) [0.069]	0.217 (0.228) [0.343]	0.227 (0.227) [0.319]	0.242 (0.245) [0.324]	0.255 (0.244) [0.297]
Household Fraction of Diseases _{it} × log(\bar{Y}_i)		-0.021 (0.083) [0.803]		0.286 (0.243) [0.239]		0.373 (0.257) [0.147]
R ²	0.79	0.79	0.79	0.79	0.79	0.79
N	36577	36577	4997	4997	4997	4997
Within-HH std. dev. change in Household Fraction of Diseases _{it}	0.080	0.080	0.061	0.061	0.061	0.061

Notes: Table reports results from a 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 Household Fraction of Diseases is the total number of diseases in the household divided by the maximum number possible (7 if single, 14 if married). 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 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.

TABLE 12
ALTERNATIVE SAMPLE RESTRICTIONS

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	<i>X</i> = Fraction of perm. income from SSI, SSDI, pensions, and annuities		Add in those without health insurance	Add in those in labor force	Age ≥ 65
		<i>X</i> > 50%	<i>X</i> > 75%			
$NUM_DISEASE_{it} \times \log(\bar{Y}_i)$ (β_1)	-0.009 (0.004) [0.018]	-0.013 (0.005) [0.012]	-0.013 (0.008) [0.090]	-0.010 (0.003) [0.004]	-0.004 (0.002) [0.118]	-0.007 (0.003) [0.034]
$NUM_DISEASE_{it}$ (β_3)	-0.011 (0.003) [0.001]	-0.013 (0.004) [0.002]	-0.014 (0.006) [0.014]	-0.011 (0.003) [0.000]	-0.008 (0.002) [0.002]	-0.012 (0.003) [0.000]
$\log(\bar{Y}_i)$ (β_4)	0.048 (0.003) [0.000]	0.074 (0.005) [0.000]	0.080 (0.007) [0.000]	0.049 (0.003) [0.000]	0.036 (0.002) [0.000]	0.033 (0.003) [0.000]
R^2	0.47	0.46	0.46	0.47	0.43	0.46
N	45447	28064	16450	52645	90405	48081
Number of individuals	11514	6814	4041	13189	19352	12380
Within-person std dev of $NUM_DISEASE_{it}$ (σ)	0.625	0.643	0.649	0.633	0.628	0.633
% change in marginal utility for a 1 std. dev. change in $NUM_DISEASE_{it}$ ($\sigma\beta_1/\beta_4$)	-11.2% [0.018]	-11.7% [0.013]	-10.8% [0.092]	-12.4% [0.003]	-6.8% [0.119]	-13.7% [0.035]

Notes: Column 1 reports the results from the baseline specification in Table 4; see notes to Table 4 (Panel A) for more details. Subsequent columns report results from re-estimating the baseline specification on different samples, as indicated in the column headings. 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 $\sigma\beta_1/\beta_4$ is bootstrapped based on 10,000 iterations, resampling individuals with replacement.