TIAA Institute

Research Dialogue | Issue no. 231
September 2025

Fit to consume: How health shapes preferences for consumption

Abstract

Whether the marginal utility of consumption differs by health status is a key input to many economic models. However, evidence on the sign and magnitude of state dependence between consumption and health remains mixed. Using detailed panel data from Singapore that measures consumption and health each month, we find that health and consumption are complements. The marginal utility of consumption increases by 3.5% for a one-standard deviation increase in an index of health based on chronic conditions. Simulations illustrate our estimates of state dependence have quantitatively important implications for optimal retirement saving.

Adam Leive
UC-Berkeley and NBER
TIAA Institute Fellow

Jessica Ya Sun
Huazhong University of
Science and Technology

1. Introduction

How health affects the marginal utility of consumption is of primary importance to economic models of saving and insurance. If illness reduces the value of non-health consumption, then optimal health and disability insurance benefits should decline compared to a case where the value of a dollar is the same in good or bad health. Because health typically worsens with age, optimal retirement savings would also be lower. Conversely, if illness increases the value of non-health consumption, then optimal health, disability, and retirement benefits should be higher. Yet empirical evidence on whether the marginal utility of consumption depends on health ("state dependence") remains inconclusive.[1] Consequently, there's no consensus on how to model the sign or magnitude of state dependence between health and consumption.[2]

This paper provides novel evidence on how the marginal utility of consumption varies by health status using high-frequency survey data from Singapore. We first derive an empirically tractable formula for state dependence as a ratio of marginal utilities based on a simple life-cycle model of consumption and savings. State dependence is a function of the gap in nondurable consumption between health states, risk and time preferences, and parameters that govern uncertainty about future health and mortality. We then estimate state dependence using 55 waves of the Singapore Life Panel (2015–2019), which is a monthly survey containing detailed information on consumption, health, employment, and well-being. In addition to these monthly questions, the Singapore Life Panel (SLP) includes annual questions about assets and one-off modules about risk preferences and other topics relevant to aging. We construct an index of health by predicting self-assessed health status according to the diagnosis of chronic conditions and estimate how nondurable consumption responds to within-person changes in the health index over time.

We find strong evidence that health and consumption are complements. People reduce nondurable consumption if their health worsens, with a one-standard deviation decline in the health index corresponding to a 10% drop in nondurable consumption. Combining these results with our theory, the marginal utility of consumption is 3.5% higher for a one-standard deviation increase in the health index. We document heterogeneity in magnitude but not in sign of state dependence across several observable characteristics. Our estimates are robust to several alternative specifications. The categories of spending that decrease the most are consistent with our interpretation that health and consumption are complementary. While we document declines for nearly all spending components, we find strong evidence that food, dining out, and vacations fall sharply after illness.

Two supplemental approaches provide corroboration that consumption and health are complements. First, we estimate the regression suggested by Finkelstein et al. (2013) using data on life satisfaction that regresses a utility proxy on consumption, health, and the interaction between consumption and health. The sign of the interaction reveals the direction of state dependence. Second, we leverage random variation in lottery winnings, which are measured in three waves of the SLP. Motivated by the result from Kim and Oswald (2021) and Kim and Koh (2021) that lottery winnings increase life satisfaction and consumption, we test whether the effect of lottery winnings on these outcomes differ by baseline health status. In both cases, we find evidence consistent with our main result that the marginal utility of consumption is higher when healthy.

Using our main estimates, we illustrate the implications of state dependence for optimal retirement saving. We simulate a simple life-cycle model of consumption and savings that features changes in health status by age. This exercise yields three key findings. First, compared to state-independent utility, optimal retirement savings are 1.4% lower because the least healthy years of life occur in old age. Second, future cohorts should optimally save more than current cohorts due to technological advances and other factors that improve health at older ages. Third, ignoring state dependence will meaningfully underestimate risk aversion from observational data. The optimal consumption profile is less smooth under state-dependent utility because consumption is shifted earlier in life, and an econometrician estimating preference parameters from this profile will incorrectly infer the reason for these choices. Our model is intended to be illustrative, and a richer model with additional features may lead to slightly different quantitative magnitudes, but we expect the same qualitative conclusions regarding these three points.

1 Several studies find the marginal utility of consumption is higher when healthy (Finkelstein et al., 2013; Blundell et al., 2024), while others find the reverse (Lillard & Weiss, 1997; Ameriks et al., 2020) or document heterogeneity (Brown et al., 2016).

2 The modal assumption is that health and consumption are additively separable (see e.g., Hall & Jones, 2007; Hugonnier et al., 2013; De Nardi et al., 2016; Dobkin et al., 2018; Finkelstein et al., 2019; Seibold, 2021; Leive, 2022; Fonseca et al., 2023, among many others). Some studies, generally as part of robustness, analyze the case that the marginal utility of consumption is higher when healthy (Reichling & Smetters, 2015; Lieber & Lockwood, 2019; Hendren, 2021), others that it's higher when ill (Koijen et al., 2016), or both cases (Deshpande & Lockwood, 2022).

The setting of Singapore offers a useful context for studying this topic, and the richness of the SLP provides advantages relative to data used in previous work. One empirical challenge to measuring state dependence is how to determine whether consumption declines are due to preferences or due to incomplete insurance. Singapore is among the wealthiest countries in the world and covers its population through a robust system of Medical Savings Accounts (MSAs). If consumption decreases due to incomplete insurance, one would expect people with lower MSA balances to experience the largest consumption declines. Instead, we find that people with higher MSA balances cut their consumption by similar amounts in response to illness. This pattern is consistent with state dependence rather than incomplete insurance. In terms of data, the monthly frequency of the SLP is unparalleled among similar surveys focusing on aging-related questions.[3] We observe the median respondent 48 times, which improves our ability to track the dynamics between consumption and health.

Without such detailed data, prior research has generally needed to make additional assumptions to estimate state dependence between health and consumption. Studies using similar approaches with different datasets have reached different conclusions. Starting with Finkelstein et al. (2013), a series of papers have examined the correlation between utility proxies, chronic conditions, and permanent income. While Finkelstein et al. (2013) find the marginal utility of consumption is lower with more chronic conditions in the Health and Retirement Study (HRS), the sign and magnitude of state dependence vary between studies using the same methods and similar data from other countries (Wang & Wang, 2020; Simonsen & Kjaer, 2021; Bassoli, 2022). As an alternative approach, both Brown et al. (2016) and Ameriks et al. (2020) ask people hypothetical questions about how much money they prefer to allocate to states when they're disabled versus states when they're healthy. These studies also reach different conclusions. Ameriks et al. (2020) find people prefer more consumption when disabled than when healthy, while Brown et al. (2016) detect heterogeneity by age and type of disability, sometimes finding people favor less consumption when disabled and in other cases not rejecting state-independent utility.

Several studies have pursued more structural approaches, combining data on consumption or assets with detailed life-cycle models. This branch of research has also produced mixed results, with some finding a higher marginal utility of consumption when ill (Lillard & Weiss, 1997; Yogo, 2016), others finding the opposite (Koijen et al., 2016), and others finding little relationship (De Nardi et al., 2010). Most recently, Blundell et al. (2024) developed a semi-structural approach to test whether changes in consumption are driven more by changes in health or changes in resources. Using data from the HRS, they show that consumption declines after health shocks in the United States, and most of the response is explained by lower marginal utility of consumption when ill. We view our paper as complementary; Blundell et al. (2024) model a broader set of risks but don't estimate the magnitude of utility parameters, whereas we develop and estimate a theoretically grounded measure of state dependence while focusing on health risks.

The remainder of the paper proceeds as follows. We develop a formula for state dependence between consumption and health based on a stochastic life-cycle model in Section 2. We then describe the SLP survey and relevant aspects of Singapore's health system in Section 3. Section 4 presents our main results on state dependence using changes in consumption after health shocks. Section 5 shows the results of supplementary analyses using data on life satisfaction and lottery wins. We illustrate the quantitative importance of our estimates for optimal retirement saving by simulating a life-cycle model in Section 6. Section 7 briefly concludes by discussing external validity and the implications of our results for other social insurance programs.

3 The Health and Retirement Study (HRS), English Longitudinal Study of Ageing (ELSA), and the Survey of Health, Ageing, and Retirement in Europe (SHARE) each occur once every two years.

2. Theory

In this section, we develop a formula for health state dependence based on a stochastic life-cycle model of consumption and savings. An agent with preferences satisfying constant relative risk aversion (CRRA) chooses consumption to maximize their discounted expected utility of lifetime consumption. We assume medical costs are fully insured and that nondurable consumption doesn't affect health status. Utility is intertemporally separable, the future is discounted by the factor $0 < \beta < 1$, and the agent can borrow and save at interest rate r. Flow utility may vary by health states, with $u(c_{h,t})$ denoting the utility of consuming c_h in time t if healthy, and $v(c_{b,t})$ denotes the utility of consuming c_b in time t if unhealthy. While we assume the agent is either healthy or unhealthy, the framework naturally extends to any number of health states. The probability of being healthy in time $t + 1$ if healthy in time t is denoted by p.

Along the optimal consumption path, the agent equates the marginal utility of consumption when healthy in time t to the discounted expected utility of consumption in time $t + 1$ according to the Euler equation:

$$u'(c_{h,t}) = S\beta\,(1+r)[pu'\,(c_{h,t+1})+(1-p)\,v'(c_{h,t+1})] \qquad (1)$$

where S is the probability of survival to time $t + 1$ if healthy in time t. Omitting time subscripts (i.e., in steady state), Equation 1 can be rearranged as:

$$u'(c_h) = \frac{S\beta(1+r)(1-p)}{1-S\beta(1+r)p} v'(c_b) \equiv \delta v'(c_b) \qquad (2)$$

where δ is a scaling factor that depends on the parameters governing uncertainty and the dynamics of the decision problem. Following approaches used in previous studies to measure the value of insurance (Baily, 1978; Gruber, 1997; Chetty, 2006; Hendren, 2017; Fadlon & Nielsen, 2019; Coyne et al., 2024), we approximate marginal utility of consumption when healthy by taking a second-order expansion around c_b: $u'(c_h) \approx (c_b) + u''(c_b)(c_h - c_b) + \frac{1}{2}u'''(c_b)(c_h - c_b)^2$. The marginal utilities in both health states now correspond to the same consumption level, and combining with the right-hand side of Equation 2 and dividing by $u(c_b)$ yields:

$$1 + \frac{u''(c_b)}{u'(c_b)}(c_h - c_b) + \frac{u'''(c_b)}{u''(c_b)}(c_h - c_b)^2 = \delta \frac{v'(c_b)}{u'(c_b)} \quad (3)$$

Under CRRA utility, substitute $-\frac{u''(c_b)}{u'(c_b)} c_b = \gamma(c_b)$ as the coefficient of relative risk aversion, $-\frac{u'''(c_b)}{u''(c_b)} c_b = \gamma(c_b) + 1$ as the coefficient of relative prudence, and rearrange to write:

$$\frac{u'(c_b)}{v'(c_b)} = \frac{\delta}{1-\gamma(c_b)\left(\frac{c_h-c_b}{c_b}\right)\left[1-\frac{\gamma(c_b)+1}{2}\left(\frac{c_h-c_b}{c_b}\right)\right]} \qquad (4)$$

The ratio of marginal utilities is a function of the proportional change in consumption between health states $\frac{c_h-c_b}{c_b}$ risk aversion γ, and a term δ that incorporates time preferences, interest rates, and uncertainty in health and survival. A ratio over 1 indicates that the marginal utility of consumption is higher when healthy, while a ratio below 1 indicates that it's lower when healthy. $\frac{u'(c_b)}{v'(c_b)} = 1$ represents the special case of state-independent utility. In the next section, we describe the empirical setting and data used to estimate each component and then estimate in Section 4.

We close this section by noting that Equation 4 provides a general formula for state dependence—we have assumed CRRA preferences but not a specific functional form linking h and c. As we evaluate the quantitative magnitudes of our estimates in Sections 4 and 6, we impose additional structure on the utility function to illustrate specific formulations of how health and consumption are linked.

3. Setting and data

In this section, we first present a high-level overview of healthcare financing in Singapore to establish context. We then describe the Singapore Life Panel (SLP) and summarize the key variables used in our empirical analysis.

3.1 Background on Singapore's health system

Personal savings accounts are central to financing health care in Singapore.[4] Created in 1984, MediSave is the country's program of individual medical savings accounts (MSAs). While working, citizens and permanent residents make mandatory MSA contributions that range from 4% to 10.5% of salary, up to a ceiling.[5] Employers make equivalent matching contributions. Contributions are deductible from taxable income, and balances grow at a fixed interest rate of 5% annually. Individuals can also make additional tax-deductible contributions to their MediSave accounts up to the statutory maximum. The contribution limits are high—66,000 SGD (49,102 USD) in 2022—and exceed 10 times the limits for Health Savings Accounts in the United States.[6] All balances roll over each year.

MSA withdrawals can be made at any age to pay for deductibles and copayments for outpatient care, preventive care, prescription drugs, inpatient and other acute care, and long-term care. MediSave funds can also finance premiums for MediShield Life, which is supplemental health insurance that covers low-probability, high-cost care such as hospitalizations and expensive outpatient care like dialysis and chemotherapy. The deductible in MediShield ranges from 1,500 SGD to 3,000 SGD depending on the individual's age and the class of hospital ward. MediSave balances can also finance care for dependents, including a spouse, children, grandparents, and siblings. Unlike HSAs in the United States, MSA funds can only be used to finance healthcare expenses—withdrawals for other consumption are prohibited regardless of age.

MediSave covers approximately 97% of Singapore's population. Those with low incomes who can't afford MediShield are covered by the safety-net program MediFund, which is financed by a government investment fund and requires preapproval for care to be reimbursed. Collectively, MediSave, MediShield, and MediFund are referred to as the "3M system."

3.2 The Singapore Life Panel

The SLP is a longitudinal survey of a representative sample of Singaporean citizens age 50 to 70 in 2015. The baseline survey was conducted in July 2015 and surveyed more than 13,000 individuals, including both the representative respondent and the spouse. The survey design is similar to the Health and Retirement Study (HRS) in the United States, but with more frequent interviews and a slightly different set of questions.[7] Respondents are surveyed monthly, with core questions about income and spending, employment, chronic conditions, health status, and life satisfaction asked each month. Other questions are included quarterly, annually, or on a one-off basis. The monthly sample size is around 8,000, which corresponds to a response rate of 65%. We use 55 waves of the survey, spanning the period of July 2015 to December 2019.

Detailed expenditure data is collected monthly and records spending on 33 different items. Our main analysis focuses on nondurable consumption. In supplementary analyses, we also include durable consumption. As part of the latter measure, we construct a rental equivalent for homeowners to measure housing consumption and impute service flows for vehicle consumption for households who own vehicles. Appendix A details our consumption measures and the classification of items into nondurables and durables. Compared to other countries, in-kind transfers are less common in Singapore. There are rebates for utilities for some low- and middle-income households, and once a quarter, the SLP asks about the amount received. We include these amounts as part of utility consumption. Following the literature, we exclude education and insurance from consumption as these items are more likely to represent investments. We winsorize consumption at the 1st and 99th percentiles within each wave to remove the influence of outliers. We exclude healthcare spending from our main consumption measure and instead measure the impact of health shocks on healthcare spending in supplementary analyses. Information on (self-reported) healthcare spending is collected each month and divided into subcategories of prescription medication, outpatient spending, inpatient spending, and other medical spending. During the annual survey, which is longer than the monthly surveys, individuals are asked to report household assets, including balances in checking and savings accounts, MediSave, and retirement accounts.

4 For a detailed description of Singapore's health system and its history, see Haseltine (2013), Yin and He (2018), and information from the government available at https://www.moh.gov.sg/cost-financing/healthcare-schemes-subsidies/MediSave (accessed May 2, 2024).

5 The ceiling in 2022 was 5,760 SGD (4,284.84 USD) to 7,560 SGD (5,623.85 USD) depending on age.

6 The maximum individual contribution to Health Savings Accounts in the United States was $3,650 in 2022, by comparison.

7 See https://rosa.smu.edu.sg/singapore-life-panel/about-singapore-life-panelr (accessed May 2, 2024).

Each month, the SLP asks individuals if they've ever been diagnosed with the following seven chronic conditions: diabetes, hypertension, arthritis, psychiatric conditions, heart conditions, stroke, and cancer. By tracking changes to these responses over time, we construct indicators for the development of new chronic conditions. Self-assessed health is also measured monthly on a five-point scale, with responses ranging from Poor to Excellent. In Section 4, we present evidence that the chronic conditions in the SLP capture almost as much variation in health status as a larger set of conditions do.

In addition to these core questions, the SLP includes several other modules relevant to our analysis. In one wave, respondents are asked about hypothetical gambles that we use to set-identify risk aversion. We describe this question in more detail in the next section. Finally, life satisfaction is measured monthly, and lottery purchases and winnings are measured in three waves. We use these variables in supplemental analysis in Section 5 and defer their discussion until then.

3.3 Descriptive statistics of sample

We make minimal sample restrictions to construct our analytic sample. Starting from a sample of 13,353 individuals and 9,823 households, we drop respondents who only appear in wave 0 (the baseline wave) or wave 1 (a small test-wave), which drops 2.7% of the sample. We next restrict the sample to individuals age 45 to 75, reducing the remaining sample by 1.5%. We drop 34 people who are presented with different summary screens in the consumption module across waves (0.3% of the remaining sample). Finally, we exclude 14 individuals who live with a household member under the age of 22 at any time during the sample period because consumption may change when children leave the home. This restriction barely reduces the sample because most respondents are beyond the traditional age of parenthood. Our analytic sample includes 12,742 individuals in 9,547 households. The average number of waves per individual respondent is 33, and the median is 48.

Table 1 presents summary statistics for the analytic sample. In total, our data contains 433,139 observations at the individual-wave level. Panel A presents demographic information. The average age is 60 years, and 52.8% of the sample is female. The large majority (86.6%) identify as ethnically Chinese. Most respondents (69%) live in public housing provided by the Housing Development Board (HDB), rather than in private housing. Nearly 80% of the sample is married. In terms of educational attainment, 23% have completed primary schooling or less, 41% have completed secondary schooling, and 36% have completed post-secondary or higher education.

Panel B reports income, consumption, and assets. On average, monthly household nondurable consumption per capita is 1,053 SGD, and durable consumption is 1,743 SGD. By comparison, annual household income per capita is 4,991 SGD, on average, including both earnings and other sources. The average household has slightly more than 40,000 SGD in total assets per capita, with more than 26,000 SGD in MSA balances. Average monthly healthcare spending is 75 SGD per capita.

Panel C reports the number of diseases and self-assessed health at the individual level over the sample period. On average, respondents have one chronic condition in a given month. Diabetes, hypertension, arthritis, and psychiatric conditions comprise the majority of these, while heart conditions, stroke, and cancer are less common. Appendix Figure B.1 shows that the onset of diseases is gradual and fairly steady at both the individual and household level during the sample period. In terms of self-assessed health, 32.1% report being in fair health, 45.9% in good health, and 13.8% in very good health. The two extreme levels are rare, with 5.8% reporting being in poor health and 2.5% in excellent health.

TABLE 1. SUMMARY STATISTICS

	Mean	SD
Panel A. Demographics		
Age	61.21	5.96
Female	0.530	0.499
Chinese	0.866	0.340
Educational attainment: Primary or less	0.227	0.419
Educational attainment: Secondary	0.412	0.492
Educational attainment: Post-secondary or higher	0.361	0.492
Live in HDB housing	0.684	0.465
Married	0.792	0.406
Citizen	0.951	0.215
Household size	1.86	0.397
Panel B. Consumption, income, assets		
Nondurable consumption per capita (monthly)	1,053	1,057
Durable consumption per capita (monthly)	1,743	1,363
Healthcare consumption per capita (monthly)	74.57	345.55
Income	5,009	5,671
Household MediSave assets per capita	26,239	17,608
Household total assets per capita	40,259	89,404
Panel C. Chronic conditions and health status		
Number of chronic conditions	1.00	1.12
Heart disease	0.130	0.336
Stroke	0.027	0.163
Cancer	0.053	0.224
Diabetes	0.191	0.394
Hypertension	0.387	0.487
Psychiatric conditions	0.035	0.185
Arthritis	0.179	0.383
Self-assessed health		
Excellent	0.025	0.156
Very good	0.138	0.345
Good	0.458	0.498
Fair	0.321	0.467
Poor	0.058	0.234
Number of individuals	12,742	
Number of households	9,547	
Number of individual-waves	429,442	

Notes: Table presents descriptive statistics of the analytic sample. Statistics are weighted using the survey's sample weights. The HDB provides public housing in Singapore.

4. Estimates of state dependence from consumption changes

This section provides our main estimates of state dependence. We sequentially estimate each parameter in Equation 4, beginning with constructing a health index based on the incidence of chronic conditions. We then measure how consumption responds to changes in this index. Next, we describe how we use the SLP to internally calibrate other parameters—risk aversion and health status transitions—and, finally, we detail which parameters we externally calibrate.

4.1 Constructing a health index

The theory in Section 2 treated health as a binary measure. In reality, health has multiple dimensions. Our empirical measure of health status seeks to balance the conceptual simplicity of analyzing a small number of health states while considering the multiple factors that determine health. The prior literature has constructed empirical measures of health in various ways, including counting the number of chronic conditions (Finkelstein et al., 2013; Hosseini et al., 2022), using self-assessed health (De Nardi et al., 2024), predicting self-assessed health from self-reported conditions and limitations (Poterba et al., 2017; Blundell et al., 2024), and predicting mortality from conditions inferred by prescription drug consumption (Danesh et al., 2024). Leveraging our monthly frequency of both self-assessed health and (self-reported) diagnosis of chronic conditions, we construct an index that represents the probability of being in good, very good, or excellent health. Specifically, we estimate a logit model of reporting self-assessed health as being good or better as a function of the seven chronic conditions, their interactions with sex, and all two-way interactions of chronic conditions.

By predicting health status based on chronic conditions, we use objective information about disease diagnoses while also accounting for differences in the severity of the various conditions affecting health. For example, strokes and heart attacks are likely more severe than hypertension or diabetes diagnoses, and counting the number of chronic conditions would treat them as equivalent. The prediction model maps conditions onto health status in a way that weights their relative importance to health and allows for the possibility that interactions between conditions may also matter. As robustness, we consider alternative constructions of the health index, including a LASSO regression to avoid possible overfitting or simply using the count of chronic conditions. Regardless of how we specify the health index, we find qualitatively similar results for state dependence.

We choose to predict a binary measure of health status for several reasons. First, most people report being in one of two health states: 78% say they're in either good or fair health. Relatively few people report having excellent, very good, or poor health, as shown above in Table 1. Second, it's unclear that the difference between "poor" and "fair" health is necessarily the same as the difference between "excellent" and "very good" health, which complicates efforts to predict a continuous measure of health. Third, the binary measure of good or better health is conceptually simple, which facilitates its interpretation and application to various contexts.

Figure 1 visualizes the relationship between chronic conditions and the health index. Panel (A) presents the average index level according to the number of chronic conditions. The health index declines monotonically with the number of chronic conditions, and patterns for men and women are quite similar. Beyond this count, the specific conditions a person has and in what combinations produce different values of the health index. The whiskers in Panel (A) denote the standard deviation within men and women who have the same number of chronic conditions. As the number of conditions increases, so does the variability of the health index, driven by the influence of different combinations of conditions. Across the full sample, the standard deviation of the health index is 0.15.

To provide intuition for how each condition separately influences the health index, Panel (B) plots the coefficient estimates from a linear regression of the health index on that chronic condition, without any other controls. The coefficients can be interpreted as the contribution of developing that condition, accounting for the other conditions a person already has. Hypertension exerts the smallest change in the health index, which is not surprising given that the condition can often be asymptomatic. Diabetes, cancer, heart conditions, and arthritis have relatively similar contributions—reducing the index between 0.2 and 0.25—with occasionally small differences by gender. Psychiatric conditions correspond to the largest reductions among these seven.

We quantitatively assess the importance of these conditions by running a linear regression of observed health (good or better) against the health index. The seven conditions capture 9.8% variation in the binary measure of health status (R^2 = 0.098). We then repeat the same exercise in the U.S. Medical Expenditure Panel Survey (MEPS), where we observe health status and a larger set of chronic conditions.[8] The seven conditions in the SLP represent about half of the "priority conditions" designated by the U.S. Agency for Healthcare Research and Quality and measured in the MEPS, with the others being attention deficit hyperactivity disorder, angina, asthma, emphysema, high cholesterol, heart disease, and heart attacks. In the MEPS, a health index with

8 We restrict the MEPS to those with health insurance and use surveys from 2016 to 2019 to match the SLP's sample period.

the same conditions (excluding psychiatric conditions, which are not recorded) capture 12.3% of the variation in health status. Using all conditions in the MEPS to construct a health index, we find the larger set of conditions explains 15.0% of variation in health status. We view the improvement in R^2 from 0.123 to 0.150 as modest and interpret this exercise as evidence that the conditions measured in the SLP capture important dimensions of health status relative to a larger set that could be measured, even if many aspects of health remain unobserved.

FIGURE 1. RELATIONSHIP BETWEEN HEALTH INDEX AND CHRONIC CONDITIONS

(A) Mean health index by number of conditions

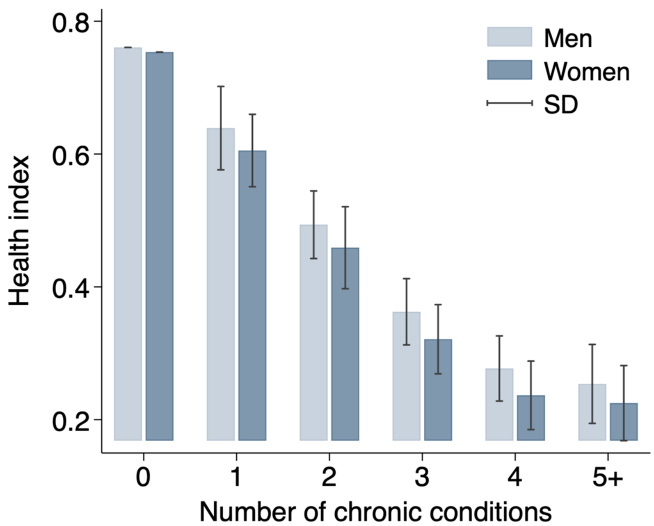

(B) Effect of conditions on health index

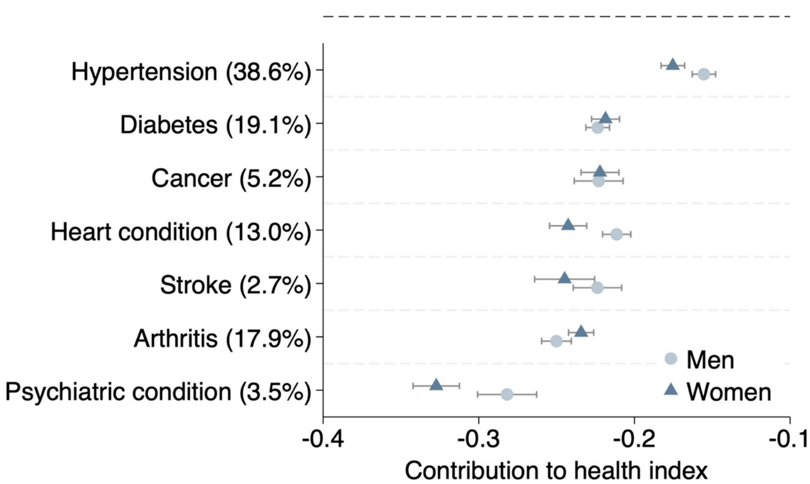

Notes: Figure shows statistics for how chronic conditions relate to the health index. Panel (A) plots the average of the health index for individuals with 1, 2, 3, 4, and 5 or more chronic conditions, separately by men and women, with whiskers denoting the standard deviation for that group. Panel (B) plots the regression estimates from a linear regression of the index against each of the chronic conditions, separately by men and women, with whiskers denoting 95% confidence interval on the difference relative to not having that chronic condition.

4.2 Changes in consumption between health states

We run the following regression to estimate the change in consumption between health states:

$$\ln(C_{ijt}) = \alpha_i + \tau_t + \lambda H_{ijt} + \eta X_{ijt} + \epsilon_{ijt} \qquad (5)$$

where $\ln(C_{ijt})$ measures log household consumption per capita for individual i of household j in month t, and H_{ijt} denotes the health index for individual i in month t. The coefficient λ measures the percentage change in consumption that appears in Equation 4 from Section 2. By including an individual fixed effect α_i and a wave fixed effect τ_t in Equation 5, λ is identified from changes within individuals over time. Time-varying covariates X_{ijt} are age and indicators for household size. Regressions are weighted using the survey's sample weights, and standard errors are clustered at the household level.[9]

TABLE 2. CONSUMPTION AND HEALTH INDEX

	(1)
Health index	0.103
	(0.041)
Individual fixed effects	Yes
Wave effects	Yes
R^2	0.775
N	10,989
NT	410,354

Notes: Table plots regression results of estimating Equation 5 via OLS. The health index is calculated as the predicted probability of being in good health based on a logit model that includes indicators for chronic conditions, interactions between conditions and gender, and all two-way interactions between conditions. The OLS regression also includes individual fixed effects, wave fixed effects, age, and indicators for household size and are weighted using survey sample weights. N denotes the number of unique individuals, and NT denotes the number of individual-waves.

Table 2 presents the results. On average, consumption is 10.3% higher when healthy compared to when unhealthy (column 1). However, this estimate corresponds to moving from one extreme of the health index to the other, which constitutes out-of-sample variation. Instead, a more meaningful interpretation is a one-standard deviation change in the index, which is equal to 0.15. The corresponding decline in consumption is then predicted to be about 1.5%. As a graphical analogue, Figure 2 plots a binned scatterplot of consumption against the health index using the methods of Cattaneo et al. (2024), controlling for individual fixed effects, month effects, and indicators for household size. The fitted line is a third-order polynomial and shows the relationship is linear over much of the range of the health index, which supports the specification in Equation 5.

As discussed earlier, an important assumption with this approach is that individuals are adequately insured against the cost of illness. We assess this assumption by testing whether consumption responses are the same for people with different initial levels of MSA balances per capita. One might expect those with smaller balances would reduce consumption more than those with higher balances if incomplete insurance explained the changes. Appendix Table B.1 shows that the estimates are similar if we split the sample by baseline level of MSA assets in 2015. Moreover, the average MSA balance among those in the bottom half exceeds 12,000 SGD, which is 13 times larger than annual healthcare spending in the full sample. We therefore interpret the change in consumption in Table 2 as reflecting preferences rather than insufficient resources.

The sources of the spending declines provide further evidence consistent with this interpretation. We document declines in consumption across most categories of nondurables, with the strongest evidence for reductions in food and beverages, dining out, and vacations (Appendix Figure B.3).

9 A large literature estimates similar regressions to study the completeness of insurance. Many studies have examined whether households in developing countries are protected against shocks (see e.g., Townsend, 1994; Gertler & Gruber, 2002; Mohanan, 2013; Liu, 2016; Garcia-Mandico et al., 2021). The macro literature estimates "pass through coefficients" between consumption and income shocks using similar approaches (Blundell et al., 2008; Kaplan & Violante, 2010; Blundell et al., 2024).

FIGURE 2. BINNED SCATTERPLOT: LOG CONSUMPTION VS. HEALTH INDEX

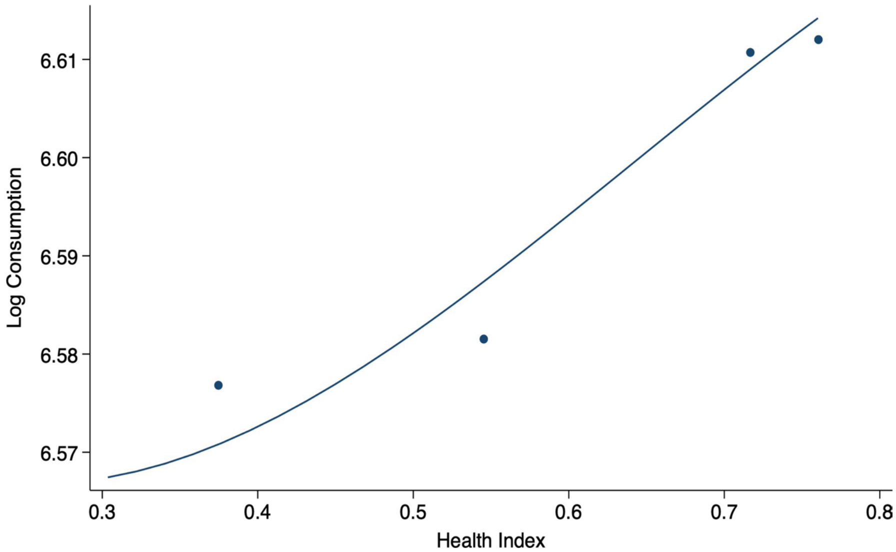

Notes: Figure shows binned scatterplots of log nondurable consumption per capita against the health index, controlling for individual fixed effects, month effects, and indicators for household size. The number of equally sized bins is chosen based on data-driven methods of Cattaneo et al. (2024). The line plots a third-order global polynomial, and the regression is weighted using the survey's sample weights.

4.3 Other internally calibrated parameters

Risk aversion: In one wave, the SLP asks a series of questions about hypothetical gambles related to permanent income. These questions follow Barsky et al. (1997) and set-identify risk aversion for each survey respondent. Respondents are first asked:

> *"Suppose that you were the only income earner in the family, and you have a good job guaranteed to give you your current income every year. Imagine you had the opportunity to take a new, equally good job, with a 50% chance it will double your income and a 50% chance that it will cut your income by a third. Would you take the new job?"*

Respondents who answer "Yes" are subsequently asked if they would accept the job if there were a 50-50 chance it would double their income or cut it in half, while those who answer "No" are instead asked if they would accept the job if there were a 50-50 chance it would double their income or cut it by 20%. Assuming CRRA utility, the answers to these choices bound each respondent's coefficient of relative risk aversion to be within one of the following intervals: $[0, 1)$, $[1, 2)$, $[2, 3.76)$, $[3.76, \infty)$. Our main estimates take the midpoint of the relevant interval, and we show state dependence estimates using the upper or lower bounds as robustness. We assume a lower bound of 0 (risk neutrality) for a person who answers "Yes" to both questions. For a person who rejects both gambles, we assume an upper bound of 5 to avoid implausibly high levels of risk aversion.[10] We assume that risk aversion is fixed during the sample period.

There's substantial heterogeneity in risk aversion. Slightly more than a third of respondents reject both gambles. A similar fraction accepts one or both gambles, with roughly equal shares in each of the three lowest risk aversion intervals. Finally, about one-quarter respond with "Don't know" to both questions. Fortunately, the survey includes other questions that proxy for risk preferences, which we use to impute responses for these individuals. In particular, the survey asks *"Are you generally a person who tries to avoid taking risks or one who is fully prepared to take risks?"* measured on a scale from 0 to 10 and asks similar questions for financial risks and health risks. These questions, which are commonly used in surveys but don't allow the same structural interpretation, are correlated with responses to hypothetical gambles about permanent income. For survey respondents who answer both sets of questions, we run linear regressions of the CRRA bounds against indicators for each possible response to the supplemental risk preference questions, along with indicators for female, ethnicity, education, and a quartic in age. We use these predictions to impute responses for respondents reporting "Don't know" to the gambles, or who didn't respond in that wave. In supplementary analysis, we show that our

10 A CRRA coefficient of 5 implies a person is roughly indifferent between accepting the gamble if there were an equal chance of doubling their salary or reducing it by 15%.

results are robust to excluding the 36% of observations with imputed risk aversion.

Health status transitions: We estimate the transition probabilities between health states by regressing the current health index against its one-month lag:

$$H_{it} = \kappa_1 + \kappa_2 H_{it-1} + e_{it} \quad (6)$$

The conditional probability of being healthy next month if healthy this month is $\kappa_1 + \kappa_2$. This exercise assumes that people have rational expectations about future changes in their health. As indirect support for this assumption, Appendix Figure B.2 shows there's a tight link between subjective probabilities of ever developing a chronic condition and the empirical probabilities during our sample window.[11]

4.4 Externally calibrated parameters

The SLP lacks questions that would allow us to measure the parameters governing time preferences and risk-free interest rates. We assume an annual discount factor of 0.975, which translates into a monthly discount factor of $\beta = 0.998 = 0.975^{1/12}$. We assume a gross interest rate of 1.02 annually based on the yield on 10-year Singapore government bonds during the sample period. This translates into a monthly gross interest rate of $1+r = 1.0016 = 1.02^{1/12}$. Finally, we calibrate age- and sex-specific survival probabilities using life tables published by the Singapore government.

4.5 Results

We now have all the inputs to calculate state dependence using Equation 4 from Section 2. Table 3 collects each of the inputs and presents our main estimates of state dependence.

We estimate $u'/v' = 1.27$, indicating that the marginal utility of consumption is 27% higher when the probability of being healthy is 1 compared to when it's zero. We calculate a 90% confidence interval on the ratio of [1.07, 1.43] and a 95% confidence interval of [1.04, 1.46] using a block-bootstrap procedure. To implement the bootstrap, we first sample individuals with replacement to construct a dataset with the same number of individuals as the original dataset, and then use all waves of those sampled individuals to estimate σ for that bootstrap sample. We repeat this process 1,000 times and then take the $\frac{\alpha}{2}$ and $100 - \frac{\alpha}{2}$ quantiles of that distribution as the $(100 - \alpha)\%$ confidence interval.

As in the case of Table 2, it's more instructive to consider a change in the health index that corresponds to within-sample variation. We can also translate these estimates into a precise magnitude by specifying the functional form of utility along with a baseline measure of health. As one example, consider $(1+\sigma h)\frac{c^{1-\gamma}}{1-\gamma}$ as the utility function, which is the formulation used in Palumbo (1999) and De Nardi et al. (2010). Then $\sigma = \frac{u'}{v'} - 1 = 0.27$ and a one-standard deviation increase in the health index starting from the sample mean corresponds to a 3.5% increase in the marginal utility of consumption.[12]

TABLE 3. STATE DEPENDENCE ESTIMATES

	Estimate	Notes
$\frac{c_h - c_b}{c_b}$, % consumption change between health states	0.103	Table 2
γ, Risk aversion	2.85	Section 4.3
p, Probability healthy next month if healthy	0.966	Section 4.3
S_h, Survival probability if healthy (monthly)	0.999	Externally calibrated
β, Discount factor (monthly)	0.997	Externally calibrated
1+r, Gross interest rate (monthly)	1.002	Externally calibrated
$\frac{u'(c)}{v'(c)}$, State dependence	1.27	
90% CI	[1.07, 1.43]	
95% CI	[1.04, 1.46]	

Notes: Table presents the components of the formula for u'/v' and the corresponding source, as described in the text. The last three rows show the estimated magnitude of u'/v' and its 90% and 95% confidence intervals, calculated by block bootstrapping 1,000 samples.

11 It's not surprising that the empirical probabilities are smaller than the subjective probabilities because we only observe four years of data, while the subjective probabilities ask about developing the disease at any time in the future.

12 $\frac{1+0.27\times(\bar{h}+0.15)}{1+0.27\times\bar{h}} = 1.035$ where $\bar{h} = 0.62$ is the sample mean of the health index and 0.15 is the standard deviation.

Appendix Figure B.4 plots estimates of $\frac{u'}{v'}$ for a variety of observables, including age, education, gender, marital status, and other characteristics. Estimates are consistently positive and statistically distinguishable from zero for most groups. Estimates are slightly larger for those with higher socioeconomic status, measured by higher education and living in non-HDB housing. To examine how estimates vary by asset levels, we split the sample by median of MSA wealth in the baseline wave. The magnitudes are similar, consistent with the assumption that our findings likely capture preferences rather than incomplete insurance. Collectively, we interpret the limited heterogeneity shown in these results to suggest that the complementarity between health and consumption is likely a general pattern.

Robustness: We probe robustness by considering a range of other empirical specifications and sub-samples in Figure 3. We continue to find strong evidence that health and consumption are complements. We estimate values of $\frac{u'}{v'}$ of 1.21 and 1.29 using the lower bounds and upper bounds of risk aversion, respectively. We obtain similar results if instead of measuring consumption per capita, we adjust consumption based on an equivalence scale that weights the second and subsequent adults half as much as one adult to account for economies of scale in household consumption. Results are also similar if we use alternative formulations of the health index, including using a LASSO, omitting interactions of chronic conditions, or taking the average of health indices within households. So far, we've assumed that individuals can freely adjust their consumption in response to illness based on their preferences. However, those who are hospitalized may exhibit a drop in consumption because they are not at home, not because of state dependence. To assess this possibility, we exclude observations in which the individual is hospitalized in the current or previous wave.[13] The results are smaller, but still statistically distinguishable from zero. Next, we restrict the sample to individuals who report no chronic conditions at retirement to examine a sub-sample who are not adjusting their labor supply in response to this set of illnesses. We again obtain positive estimates of state dependence. We find similar patterns if we omit those who ever report a psychiatric condition, which may be less exogenous compared to the other conditions. Finally, we obtain similar results when restricting the sample to those who respond in waves 50 or later, verifying that attrition—due to mortality or other reasons—does not affect the results.

The theory in Section 2 ignored bequest motives. To proxy for the strength of bequest motives, we split the sample by whether the household has living children. The responses of those without children would capture state dependence with weaker bequest motives. However, the small sample size of those without children makes those estimates noisy and inconclusive.

FIGURE 3. ROBUSTNESS: STATE DEPENDENCE

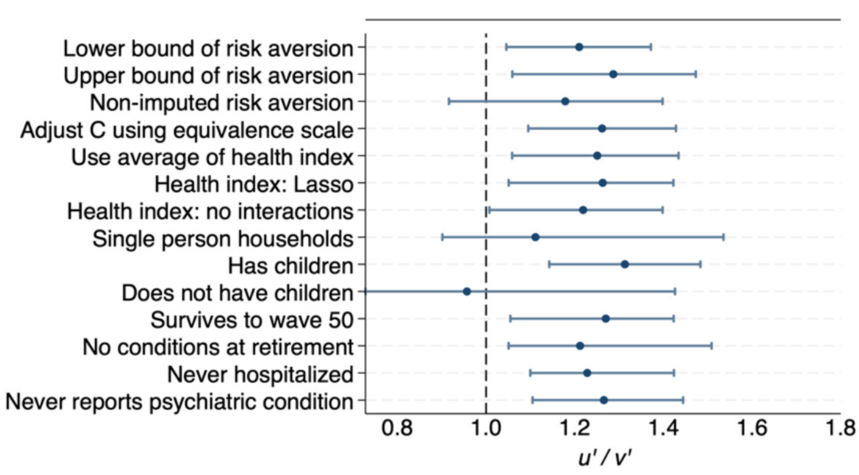

Notes: Figure plots estimates of state dependence for different sub-samples. Whiskers denote 90% confidence intervals calculated by block bootstrapping 1,000 samples for each group. For visual clarity, we top-code the upper bound of the confidence interval at 3.5.

13 We define hospitalization as having a positive amount of inpatient spending that month.

5. Supplementary results: Subjective well-being and lottery wins

This section presents additional estimates on state dependence between health and consumption using two alternative approaches. First, we estimate state dependence following the framework of Finkelstein et al. (2013) using data on life satisfaction as a proxy for utility. Second, we leverage random variation in lottery winnings. This second approach extends the main findings in Kim and Oswald (2021) and Kim and Koh (2021) to test whether the effects of an extra dollar of income on consumption and life satisfaction vary by health status. In both cases, we find evidence that health and consumption are complements. We focus less on the implied magnitudes of state dependence from these analyses compared to those in Section 4 because the latter are more tightly linked to theory. We instead concentrate on the qualitative results from these supplementary approaches and view them as providing corroborating evidence on the sign of state dependence.

5.1 Approach 1: Analysis using data on subjective well-being

Each month, the SLP includes the question *"How satisfied are you with your life as a whole these days?" with five responses ranging from "very dissatisfied" to "very satisfied."* Life satisfaction is a standard question previously used by economists to measure well-being (Frey & Stutzer, 2002; Oswald & Wu, 2010, 2011; Benjamin et al., 2014). In describing the ideal empirical approach using data on utility proxies, Finkelstein et al. (2013) write:

> *"If we could observe information on health, consumption, and a proxy for utility, we could simply regress the utility proxy on consumption, health, and the interaction of consumption and health, and the coefficient on the interaction term would give an estimate of state-dependent utility. In practice, however, we know of no panel dataset with sufficient sample size that contains information on consumption, health, and utility proxies."*

The recent creation of the SLP allows us to implement this specification. We estimate the following regression:

$$LifeSatisfaction_{it} = \varphi_1 H_{it} \times \ln(C_{it}) + \varphi_2 \ln(C_{it}) + \varphi_3 H_{it} + \alpha_i + \tau_t + \lambda X_{it} + e_{it} \quad (7)$$

where i indexes individuals, t indexes survey waves, and life satisfaction proxies for utility. Our primary specification is a fixed-effects ordered logit because the dependent variable is measured on a Likert scale. We also estimate linear probability models (LPMs), in which the dependent variable is an indicator for respondents who report being either "very satisfied" (5% of sample) or "satisfied" (50%) as opposed to "neither satisfied nor dissatisfied" (37%), "dissatisfied" (6%), or "very dissatisfied (2%)." Because utility is a continuous ordinal measure rather than a binary one, we believe there are also conceptual rationales to prefer the ordered logit over the LPM.[14] X_{ijt} includes a quartic in age and indicators for household size, which may vary over time. Finally, we include individual fixed effects α_i and wave fixed effects τ_t in both the ordered logit and LPMs. Standard errors are clustered at the household level, and regressions are weighted using the survey's sample weights. We also demean all variables to interpret coefficients relative to sample averages rather than for zero consumption or a zero probability of being healthy.

The key coefficient is φ_1. If $\varphi_1 > 0$, the marginal utility of consumption is higher in better health. If $\varphi_1 < 0$, the marginal utility of consumption is lower in better health. The null hypothesis is $\varphi_1 = 0$, which represents state-independent utility. By controlling for individual fixed effects and wave fixed effects, φ_1 is identified by changes in life satisfaction within individuals over time as their health status and consumption vary.

14 We estimate the fixed effects ordered logit developed by Baetschmann et al. (2015, 2020), which extends the conditional maximum likelihood estimator for an ordered dependent variable. The thresholds between categories of life satisfaction are allowed to vary across individuals.

5.1.1 Results

Table 4 presents the results of estimating Equation 7, with column 1 showing the estimates from the fixed effects ordered logit and column 2 showing those from the LPM. Before discussing state dependence, we note there's a positive and statistically significant relationship between consumption and life satisfaction, as shown by the estimates of φ_2. There's also a positive relationship between the health index and life satisfaction φ_3. Turning to the main coefficient of interest, we find strong evidence that the marginal utility of consumption is higher in good health. The estimates on φ_1 are positive and statistically significant across specifications.

We also consider alternative measures for satisfaction and health status. An important concern with subjective measures of well-being like life satisfaction is that it doesn't equal utility. As a supplemental measure, Appendix Table B.4 replicates Table 4 using a weighted average that incorporates other satisfaction questions in addition to life satisfaction as a whole. We first construct indicators for whether people report being satisfied or very satisfied with the following aspects of their life: social contacts and family life, daily activities and job, household income, overall economic situation, and health. These questions are also asked each month.

We then calculate a weighted satisfaction measure by assigning weights to each of these aspects, drawing on the results in Benjamin et al. (2014). For example, life satisfaction as a whole accounts for 19% of this weighted average, while satisfaction with daily activities and work accounts for 10%. The estimates of state dependence are slightly smaller than the results in Table 4, which use life satisfaction as a whole and again provide evidence that health and consumption are complements. Appendix Table B.5 presents results that use each person's number of chronic conditions instead of their predicted health status. These regressions also indicate that an extra dollar of consumption increases life satisfaction more when healthy than when ill.

TABLE 4. LIFE SATISFACTION REGRESSIONS

	Ordered Logit (1)	LPM (2)
φ_1, Health index ×Log consumption	0.381	0.027
	(0.101)	(0.012)
φ_2, Log consumption	0.099	0.015
	(0.016)	(0.002)
φ_3, Health index	1.239	0.143
	(0.235)	(0.027)
R^2		0.660
N	8,878	10,988
NT	353,163	410,220

Notes: Table presents results from estimating Equation 7. Log consumption is the natural logarithm of nondurable consumption. Column 1 presents results from the fixed effects ordered logit, and Column 2 presents results from the linear probability model. Both regressions include individual fixed effects and wave fixed effects. Standard errors are clustered at the individual level and regressions are weighted using the survey's sample weights.

TABLE 5. REGRESSIONS USING LOTTERY WINNINGS

	Dep var: Life satisfaction			Dep var: Log consumption		
	All (1)	Below median health (2)	Above median health (3)	All (1)	Below median health (2)	Above median health (3)
Lottery winnings/1,000	0.063	0.039	0.093	0.022	0.019	0.023
	(0.016)	(0.018)	(0.024)	(0.005)	(0.006)	(0.007)
Lottery purchases/1,000	-0.012	-0.007	-0.017	0.012	0.010	0.015
	(0.004)	(0.005)	(0.007)	(0.002)	(0.003)	(0.002)
N	5,524	2,990	3,397	5,281	2,867	3,196
NT	11,295	5,707	5,588	11,251	5,692	5,559

Notes: Table presents results from estimating Equation 7. Log consumption is the natural logarithm of nondurable consumption. Column 1 presents results from the fixed effects ordered logit, and Column 2 presents results from the linear probability model. Both regressions include individual fixed effects and wave fixed effects. Standard errors are clustered at the individual level and regressions are weighted using the survey's sample weights.

5.2 Approach 2: Analysis using lottery winnings

As another supplemental approach to measure state dependence, we exploit random variation in lottery winnings. Playing the lottery is very common in Singapore: 53% of the sample report they purchased lottery tickets in the last year. Previous research using the SLP shows that lottery winnings increase life satisfaction, self-assessed health, and consumption (Kim & Oswald, 2021; Kim & Koh, 2021). Motivated by these findings, we test for heterogeneity in the effect of lottery winnings on consumption and life satisfaction by (prior) health status. If health and consumption are complements, then lottery winnings should increase consumption and life satisfaction more when healthy than when unhealthy.

Information on lottery winnings and lottery spending is measured in three waves (numbers 16, 28, and 52). In our sample, the empirical probability of winning any prize money is 46% across these waves. Among those who play, annual ticket spending is 2,362 SGD and annual winnings are 460 SGD (999 SGD among winners), on average. Both ticket spending and winnings are highly skewed. The statistics in our sample are similar to those reported in Kim and Oswald (2021) and Kim and Koh (2021), which used data from waves 16 and 28. As expected, the amount of lottery winnings, conditional on lottery spending, is generally unrelated to observable characteristics (Appendix Figure B.5).

We estimate the following specification, which closely mirrors Kim and Oswald (2021):

$$y_{ijt} = \alpha + \eta_1 \ln(L_{ijt}) + \eta_2 \ln(M_{ijt}) + \mu_t + \xi X_{ijt} + u_{ijt} \quad (8)$$

where y_{ijt} denotes life satisfaction or log nondurable consumption, L_{ijt} denotes lottery winnings for individual i in household j in wave t, and M_{ijt} denotes lottery spending (both measured in thousands). As discussed in Kim and Oswald (2021), controlling for lottery spending is critical to estimate the effect of lottery winnings on outcomes because more ticket purchases increase the chance of winning. Ticket purchases may affect life satisfaction through psychological channels and may affect consumption by crowding out other spending. Covariates (X_{ijt}) include indicators for race, marriage, education, household size, and a quartic in age.[15] Unlike our earlier analyses, this specification omits individual fixed effects because we have only three waves with lottery information. Instead, we compare life satisfaction and consumption between people who (randomly) win more money than others but who have similar levels of lottery spending and health conditions. We split the regression by those above or below median values of the health index in the last wave before lottery purchases. We expect the coefficient η_1 to be larger for those who were in better health beforehand.

The results are presented in Table 5. First, column 1 replicates the results documented in Kim and Koh (2021) and Kim and Oswald (2021) that lottery winnings increase life satisfaction, conditional on lottery spending. Columns 2 and 3 show this effect is higher among those in better health. In column 4, we replicate the finding from Kim and Koh (2021) that lottery winnings increase consumption. Columns 5 and 6 show this effect is also higher for those in better health. We find similar results when using the number of chronic conditions instead of the health index to split the sample (Appendix Table B.6). Overall, we view these results using random lottery winnings as corroborating evidence of state dependence: health and consumption are complements.

15 We depart from Kim and Oswald (2021) by omitting other income, employment, home ownership, and private transfers because these variables may be outcomes that change as a result of winning the lottery.

6. Quantitative implications: Optimal retirement saving

These results have implications for a range of economic decisions and policy design, including health insurance and retirement saving. In this section, we show how our estimated magnitudes of state dependence meaningfully change conclusions about the level of optimal retirement saving. We solve a life-cycle model under different assumptions about state dependence and the profile of health shocks. The model builds on the setup developed in Section 2 and is intentionally parsimonious to focus on the role of state dependence. We do not model institutional details like taxes and government transfers to keep insights more general. We only allow for uncertainty in survival, abstracting from the role of uncertainty in health status, unemployment, or other consumption shocks.[16] The results should therefore be viewed as illustrative of how state dependence between health and consumption can potentially influence optimal saving decisions.

6.1 Model structure and implementation

Agents have a per-period utility function that depends on their health index h according to the parameter σ:

$$U(c,h) = (1+\sigma h)\frac{c^{1-\gamma}}{1-\gamma} \qquad (9)$$

where γ is the coefficient of relative risk aversion. Each period lasts five years. The agent's life cycle is divided between working years (ages 25–65) and retirement (ages 65–100), which is assumed to be anticipated and exogenous. Agents survive to the next period according to a probability that decreases with age and die with certainty by age 100. To illustrate the most basic implications of state-dependent utility, we assume health depreciates deterministically over time.

When working, agents receive labor earnings y that grow deterministically: $y_t + 1 = gy_t$. When retired, agents have zero labor earnings and withdraw from savings to finance consumption. Agents can borrow and save at the gross rate $1 + r$ and can't die in debt. They're endowed with initial assets A_0, which evolve according to:

$$A_{t+1} = (1+r)(A_t + y_t - c_t) \qquad (10)$$

We represent the dynamic programming problem recursively. Denoting β as the discount factor, the decision problem is given

$$V_t(A_t, h_t) = \max_{c_t}\left\{(1+\sigma h_t)\frac{c_t^{1-\gamma}}{1-\gamma} + \beta E_t[V_{t+1}(A_{t+1}, h_{t+1})]\right\}$$
(11)

subject to (10) and $A_T \geq 0$.

Our baseline model sets $\sigma = 0.27$ to be consistent with Table 3, since $\sigma = \frac{u'(c)}{v'(c)} - 1$. We illustrate the influence of state dependence by comparing the results to when σ is instead zero (state-independent utility). We continue to use the same parameters as Table 3 for γ, β, r (adjusted to five-year periods). We also use the same life tables as before to calibrate survival probabilities and match the average of the health index by age to that in the SLP. We extrapolate the health index to earlier and later ages using the profile of health observed in the U.S. National Health Interview Survey (NHIS), adjusted for differences in levels between the two surveys. We assume individuals earn a $50,000 salary at age 25 that grows 10% every five years and that $A_0 = 0$.

6.2 Results: Benchmark comparisons

Optimal retirement savings—defined as assets at age 65—are about 1.4% lower with state-dependent utility compared to the benchmark of state-independent utility. This calculation holds all parameters and inputs the same except for σ. Since health declines as people age, consumption is tilted earlier in the life cycle, and saving rates are lower while working. We view these magnitudes as moderately sized.

Another way to illustrate the importance of state dependence is to consider how ignoring it would affect estimates of other preference parameters. One advantage of our data is that risk aversion is estimated separately based on well-established survey questions on hypothetical gambles over permanent income. However, it's common for research to estimate risk aversion by matching observed choice data to predicted behavior from a structural model. We find that ignoring state dependence will meaningfully underestimate risk aversion because the optimal consumption profile is less smooth under state-dependent utility as consumption is shifted earlier in life. We numerically solve for the level of risk aversion assuming the (wrong) model of state-independent utility that would best rationalize the optimal consumption profile generated from the (true) model of state-dependent utility. We calculate a CRRA parameter 2.67, which is lower than the mean of 2.85 from the hypothetical gambles. An econometrician estimating preference parameters from this profile will incorrectly infer the reason for these choices as lower risk aversion, rather than a higher marginal utility of consumption when healthy. This underestimate would have implications for the evaluation of counterfactual policies. For example, a lower estimate of risk aversion would understate the value of annuities and insurance to individuals.

16 In future versions, we'll allow some of these variables to be stochastic and to influence survival probabilities.

FIGURE 4. TRENDS IN SELF-REPORTED HEALTH BY COHORT

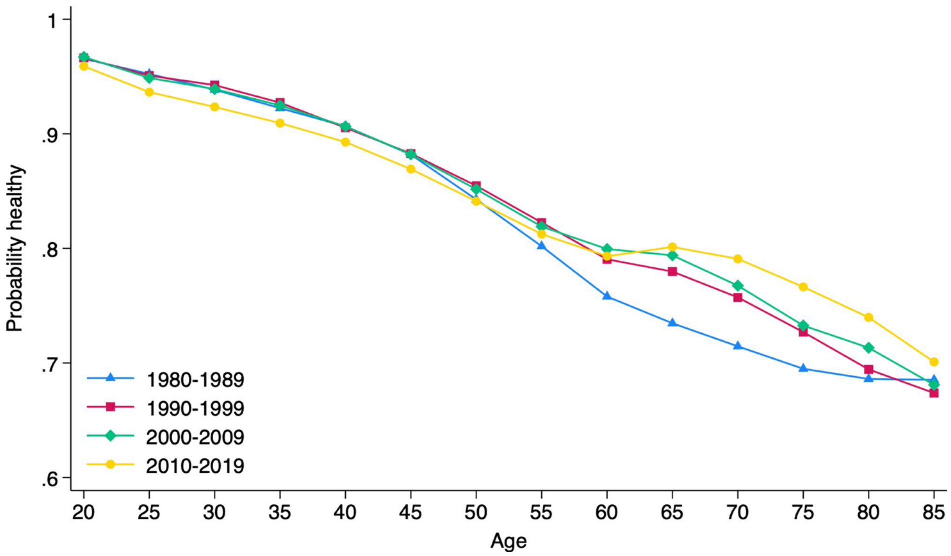

Notes: Figure shows the percentage of respondents who report being in good, very good, or excellent health by five-year age groups in the National Health Interview Survey (NHIS). Each line aggregates annual surveys from that particular decade. Means are calculated using the survey's sample weights.

6.3 Results: Improvements in health over time

We next illustrate how the complementarity between health and consumption increases optimal savings rates if more years in old age are spent in better health. For this exercise, we calculate optimal saving under different health profiles while holding σ and all other parameters fixed. Over the last six decades, there have been marked improvements in the length and quality of life, both in the United States and elsewhere. Reductions in mortality and morbidity have been especially pronounced among those with higher educational attainment and wealth (Leive & Ruhm, 2021; Bavafa et al., 2023). Continued progress in medical technology and other factors that improve health in old age will have important implications for saving decisions. To provide magnitudes of the health improvements, Figure 4 plots the predicted share of people reporting being healthy by five-year age bin in the National Health Interview Study, stratified by survey decade. These rates have been adjusted for race, gender, and education to purge compositional differences in these factors over time. The health profiles of different cohorts are fairly similar until age 55, but older cohorts subsequently experience sharper declines in health compared to younger cohorts.

For example, 84% of 50-year-olds report being healthy in both the 1980–1989 surveys and the 2010–2019 surveys. By contrast, 79% of 50-year-olds report being healthy in the 2010–2019 surveys compared to 71% in the 1980–1989 surveys. Health has depreciated at a slower rate in younger cohorts.[17]

We calculate optimal retirement savings under two assumptions about health profiles, which are similar to the differences between period and cohort life tables when calculating life expectancy. One profile uses that of the 1980s surveys (the "1980 period" profile). A 50-year-old in 1980 expects their health status to be similar in 30 years to that of 80-year-olds alive in 1980. The second profile follows a cohort forward in time, tracking their health with those measured in future surveys of the same cohort (the "1980 cohort" profile). Heuristically, the 50-year-old in 1980 expects they'll be on the yellow line at age 80 in 2010, not the blue line. Optimal retirement savings are 0.3% higher for the "1980 cohort" compared to the "1980 period" profile. This increase is fairly modest but should again be interpreted as an illustration of this particular model.

17 These differences are plausibly understated due to selection in mortality: A larger fraction of less healthy people may survive to a given age in later cohorts compared to earlier ones.

7. Discussion

In this paper, we've provided novel evidence on how the marginal utility of consumption varies by health status. We first develop a formula for state dependence between health and utility based on the Euler equations in a stochastic life-cycle consumption-savings model. We then estimate state dependence using 55 waves of a monthly panel survey that measures consumption, health, and other household characteristics in Singapore. We find that the marginal utility of consumption is higher in good health. A one-standard deviation decrease in a person's health index corresponds to a 3.5% reduction in their marginal utility of consumption. Our estimates are robust to a number of alternative specifications. We also present corroborating evidence of the complementarity between health and consumption using two supplementary approaches that use data on life satisfaction and lottery winnings, building upon Finkelstein et al. (2013), Kim and Koh (2021), and Kim and Oswald (2021).

Our paper has several limitations. Like prior research, we rely on data from a single country. Most survey respondents are between ages 50 and 70 given the survey's sampling frame. These features naturally raise questions about generalizability to other settings. Singapore's population is relatively small—slightly larger than that of Denmark—but it's reasonably diverse along several dimensions, including race, religion, language spoken, and income. In our setting, we found relatively little evidence of heterogeneity across observable characteristics of individuals. In terms of measuring health, we're limited to the seven chronic conditions recorded each month. While these conditions represent important ones that are standard in other surveys, there are naturally important dimensions of health that we miss. Generalizing our findings to other conditions requires assuming that the chain of influence from conditions to self-reported health and from self-reported health to consumption is similar for other conditions.

Our estimated magnitude of state dependence has important implications for insurance and saving decisions. We illustrate how incorporating state dependence shifts the consumption profile earlier in life and lowers optimal retirement saving. Ignoring state dependence can also affect inferences on other preference parameters if these are not measured separately. For example, studies that match observed choices to predicted choices from a structural model will underestimate risk aversion because the optimal consumption profile is less smooth under state dependence. A lower value of risk aversion would then understate the value of annuities and other forms of insurance. Future research could explore these issues in greater detail as well as other implications of the complementarity between health and consumption.

References

Ameriks, J., Briggs, J., Caplin, A., Shapiro, M. D., & Tonetti, C. (2020). Long-term-care utility and late-in-life saving. *Journal of Political Economy*, *128*(6), 2375–2451. https://doi.org/10.1086/706686

Armstrong, G., Cho, C., Garner, T. I., Matsumoto, B., Munoz, J., & Schild, J. (2022). Building a consumption poverty measure: Initial results following recommendations of a federal interagency working group. *AEA Papers and Proceedings*, *112*, 335–339. https://doi.org/10.1257/pandp.20221041

Baetschmann, G., Ballantyne, A., Staub, K., & Winkelmann, R. (2020). feologit: A new command for fitting fixed-effects ordered logit models. *The Stata Journal*, *20*(2), 253–275. https://doi.org/10.1177/1536867X20930984

Baetschmann, G., Staub, K. E., & Winkelmann, R. (2015). Consistent estimation of the fixed effects ordered logit models. *Journal of the Royal Statistical Society, Series A*, *178*(3), 685–703. http://www.jstor.org/stable/43965757

Baily, M. N. Some aspects of optimal unemployment insurance. (1978). *Journal of Public Economics*, *10*(3), 379–402. https://doi.org/10.1016/0047-2727(78)90053-1

Barsky, R. B., Juster, F. T., Kimball, M. S., & Shapiro, M. D. (1997). Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study. *Quarterly Journal of Economics*, *112*(2), 537–579. https://doi.org/10.1162/003355397555280

Bassoli, E. (2023). An empirical analysis of health-dependent utility on SHARE and ELSA data. *Italian Economic Journal*, *9*, 1217–1261. https://doi.org/10.1007/s40797-022-00208-0

Bavafa, H., Mukherjee, A., & Welch, T. (2023). Inequality in the golden years: Wealth gradients in disability-free and work-free longevity in the United States. *Journal of Health Economics*, *92*. https://doi.org/10.1016/j.jhealeco.2023.102820

Benjamin, D. J., Heffetz, O., Kimball, M. S., & Szembrot, N. (2014). Beyond happiness and satisfaction: Toward well-being indices based on stated preference. *American Economic Review*, *104*(9), 2698–2735. https://doi.org/10.1257/aer.104.9.2698

Blundell, R., Borella, M., Commault, J., & De Nardi, M. (2024). Old age risks, consumption, and insurance. *American Economic Review*, *114*(2), 575–613. https://doi.org/10.1257/aer.20220555

Blundell, R., Pistaferri, L., & Preston, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, *98*(5), 1887–1921. https://doi.org/10.1257/aer.98.5.1887

Brown, J., Goda, G. S. & McGarry, K. (2016). Heterogeneity in state-dependent utility: Evidence from strategic surveys. *Economic Inquiry*, *54*(2), 847–861. https://doi.org/10.1111/ecin.12313

Cattaneo, M. D., Crump, R. K., Farrell, M. H., &Feng, Y. (2024) On binscatter. *American Economic Review*, *114*(5), 1488–1514. https://doi.org/10.1257/aer.20221576

Chetty, R. (2006). A general formula for the optimal level of social insurance. *Journal of Public Economics*, *90*(10–11), 1879–1901. https://doi.org/10.1016/j.jpubeco.2006.01.004

Coyne, D., Fadlon, I., Ramnath, S. P., & Tong, P. K. (2024). Household labor supply and the value of Social Security survivors benefits. *American Economic Review*, *114*(5), 1248–1280. https://doi.org/10.1257/aer.20190813

Danesh, K., Kolstad, J. T., Parker, W. D., & Spinnewijn, J. (2024). *The chronic disease index: Analyzing health inequalities over the lifecycle.* (Working Paper 32577). National Bureau of Economic Research. https://doi.org/10.3386/w32577

De Nardi, M., French, E., & Jones, J. B. (2010). Why do the elderly save? The role of medical expenses. *Journal of Political Economy 118*(1), 39–75. https://doi.org/10.1086/651674

De Nardi, M., French, E., & Jones, J. B. (2016). Medicaid insurance in old age. *American Economic Review*, *106*(11), 3480–3520. https://doi.org/10.1257/aer.20140015

De Nardi, M., Pashchenko, S., & Porapakkarm, P. (2024). The lifetime costs of bad health. *Review of Economic Studies, 92*(3), 1987–2026. https://doi.org/10.1093/restud/rdae080

Deshpande, M. & Lockwood, L. M. (2022). “Beyond health: Nonhealth risk and the value of disability insurance. *Econometrica*, *90*(4), 1781–1810. https://doi.org/10.3982/ECTA19668

Dobkin, C., Finkelstein, A., Kluender, R., & Notowidigdo, M. J. (2018). The economic consequences of hospital admissions. *American Economic Review*, *108*(2), 308–352. https://doi.org/10.1257/aer.20161038

Fadlon, I. & Nielsen, T. H. (2019). Household labor supply and the gains from social insurance. *Journal of Public Economics*, *171*, 18–28. https://doi.org/10.1016/j.jpubeco.2018.01.010

Finkelstein, A., Hendren, N., & Luttmer, E. F. P. (2019). "The value of Medicaid: Interpreting results from the Oregon Health Insurance Experiment. *Journal of Political Economy*, *127*(6), 2836–2874. https://doi.org/10.1086/702238

Finkelstein, A., Luttmer, E. F. P., & Notowidigdo, M. J. (2013). What good is wealth without health? The effect of health on the marginal utility of consumption. *Journal of the European Economic Association*, *11*(s1), 221–258. https://doi.org/10.1111/j.1542-4774.2012.01101.x

Fonseca, R., Langot, F., Michaud, P.-C., & Sopraseuth, T. (2023). Understanding cross-country differences in health status and expenditures: Health prices matter. *Journal of Political Economy, 131*(8), 1949–1993. https://doi.org/10.1086/724113

Frey, B. S. & Stutzer, A. (2002). What can economists learn from happiness research? *Journal of Economic Literature*, *40*(2), 402–435. https://doi.org/10.1257/002205102320161320

Garcia-Mandico, S., Reichert, A., & Strupat, C. (2021). The social value of health insurance: Results from Ghana. *Journal of Public Economics*, *194.* https://doi.org/10.1016/j.jpubeco.2020.104314

Gertler, P. & Gruber, J. (2002). Insuring consumption against illness. *American Economic Review*, *92*(1), 51–70. https://doi.org/10.1257/000282802760015603

Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *American Economic Review*, *87*(1), 192–205. https://www.jstor.org/stable/2950862

Hall, R. E. & Jones, C. I. (2007). The value of life and the rise in health spending. *Quarterly Journal of Economics*, *122*(1), 39–72. https://doi.org/10.3386/w10737

Haseltine, W. A. (2013). *Affordable Excellence: The Singapore Healthcare Story.* Brookings Institution Press.

Hendren, N. (2017). Knowledge of future job loss and implications for unemployment insurance. *American Economic Review*, *107*(7), 1778–1823. https://doi.org/10.1257/aer.20151655

Hendren, N. (2021). Measuring ex ante welfare in insurance markets. *Review of Economic Studies, 88*(3), 1193–1223. https://doi.org/10.1093/restud/rdaa015

Hosseini, R., Kopecky, K. A., & Kai, Z. (2022). The evolution of health over the life cycle. *Review of Economic Dynamics*, *45*, 237–263. https://doi.org/10.1016/j.red.2021.07.001

Hugonnier, J., Pelgrin, F., & St-Amour, P. (2013). "Health and (other) asset holdings. *Review of Economic Studies*, *80*(2), 663–710. https://doi.org/10.1093/restud/rds033

Kaplan, G. & Violante, G. L. (2010). "How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics*, *2*(4), 53–87. https://doi.org/10.1257/mac.2.4.53

Kim, S. & Koh, K. (2021). "The effects of income on health: Evidence from lottery wins in Singapore. *Journal of Health Economics*, *76*, https://doi.org/10.1016/j.jhealeco.2020.102414

Kim, S. & Oswald, A. J. (2021). Happy lottery winners and lottery-ticket bias. *Review of Income and Wealth*, *67*(2), 317–333. https://ink.library.smu.edu.sg/soe_research/2438

Koijen, R. S. J., Philipson, T. J., & Uhlig, H. (2016). Financial health economics. *Econometrica*, *84*(1), 195–242. https://doi.org/10.3982/ECTA11182

Koijen, R. S. J., Van Nieuwerburgh, S., & Yogo, M. (2016). Health and mortality delta: Assessing the welfare cost of household insurance choice. *Journal of Finance*, *71*(2), 957–1009. https://www.jstor.org/stable/43869122

Leive, A. (2022). Health insurance design meets saving incentives: Consumer responses to complex contracts. *American Economic Journal: Applied Economics*, *14*(2), 200–227. https://doi.org/10.1257/app.20200135

Leive, A. & Ruhm, C. (2021). "Has mortality risen disproportionately for the least educated? *Journal of Health Economics*, *79*. https://doi.org/10.1016/j.jhealeco.2021.102494

Lieber, E. M. J. & Lockwood, L. M. (2019) Targeting with in-kind transfers: Evidence from Medicaid home care. *American Economic Review*, *109*(4), 1461–1485. https://doi.org/10.3386/w24267

Lillard, L. A. & Weiss, Y. (1997). "Uncertain health and survival: Effects on end-of-life consumption. *Journal of Business & Economic Statistics*, *15*(2), 254–268. https://www.jstor.org/stable/1392309

Liu, K. (2016). “Insuring against health shocks: Health insurance and household choices. *Journal of Health Economics*, *46*, 16–32. https://doi.org/10.1016/j.jhealeco.2016.01.002

Meyer, B. M. & Sullivan, J. X. (2013). Consumption and income inequality and the Great Recession. *American Economic Review, 103*(3), 178–183. https://doi.org/10.1257/aer.103.3.178

Meyer, B. M. & Sullivan, J. X. (2023). Consumption and income inequality in the United States since the 1960s. *Journal of Political Economy*, *131*(62), 247–284. https://doi.org/10.1086/721702

Mohanan, M. (2013). Causal effects of health shocks on consumption and debt: Quasi- experimental evidence from bus accident injuries. *Review of Economics and Statistics*, *95*(2), 673–681. https://doi.org/10.1162/REST_a_00262

Oswald, A. J. & Wu, S. (2010). Objective confirmation of subjective measures of human well-being: Evidence from the USA. *Science*, *327*(5965), 576–579. https://doi.org/10.1126/science.1180606

Oswald, A. J. & Wu, S. (2011). Well-being across America. *Review of Economics and Statistics*, *93*(4), 1118–1134. https://doi.org/10.1162/REST_a_00133

Palumbo, M. (1999). Uncertain medical expenses and precautionary saving near the end of the life cycle. *Review of Economic Studies*, *66*(2), 395–421. https://doi.org/10.1111/1467-937X.00092

Poterba, J. M., Venti, S. F., & Wise, D. A. (2017). The asset cost of poor health. *Journal of the Economics of Ageing*, *9*, 172–184. https://doi.org/10.1016/j.jeoa.2017.02.001

Reichling, F. & Smetters, K. (2015). Optimal annuitization with stochastic mortality and correlated medical costs. *American Economic Review*, *105*(11), 3273–3320. https://doi.org/10.1257/aer.20131584

Seibold, A. (2021). Reference points for retirement behavior: Evidence from German pension discontinuities. *American Economic Review*, *111*(4), 1126–1165. https://doi.org/10.1257/aer.20191136

Simonsen, N. F. & Kjaer, T. (2021). New evidence of health state dependent utility of consumption: A combined survey and register study. *Danish Centre for Health Economics Discussion Paper*. http://dx.doi.org/10.2139/ssrn.3867245

Townsend, R. M. (1994). Risk and insurance in village India. *Econometrica*, *62*(3), 539–591. https://doi.org/10.2307/2951659

Wang, X. & Wang, C. (2020). How does health status affect marginal utility of consumption? Evidence from China. *International Journal of Environmental Research and Public Health*, *17*(7), 1–20. https://doi.org/10.3390/ijerph17072234

Yin, J. D.-C. & He, A. J. (2018). Health insurance reforms in Singapore and Hong Kong: How the two ageing Asian tigers respond to health financing challenges? *Health Policy*, *122*(7), 693–697. https://doi.org/10.1016/j.healthpol.2018.04.012

Yogo, M. (2016). Portfolio choice in retirement: Health risk and the demand for annuities, housing, and risky assets. *Journal of Monetary Economics*, *80*, 17–34. https://doi.org/10.1016/j.jmoneco.2016.04.008

Appendix A. Construction of consumption measures

This appendix describes the expenditure items recorded in the SLP and our methods to calculate consumption. We adhere closely to standard methods used in the literature when possible (Meyer and Sullivan 2013, 2023; Armstrong et al. 2022). We include the following items as nondurables: utilities and other fuels, communications, domestic and housekeeping services, food and beverages, dining and/or drinking out, tobacco, clothing, footwear, jewelry, watches, accessories, personal care products, entertainment, sports, hobbies, vacations, road use fees, petrol, public transportation, and any other spending not elsewhere reported. We use spending on these items as our measure of nondurable consumption. We classify the following items as durables: housing, vehicles, furniture and furnishings, television, DVD/Blu-ray player and recorder, refrigerator, microwave, vacuum cleaner, washing machine, clothes dryer, and air conditioners. We construct monthly consumption flows for housing and vehicles as described below. We lack information on product characteristics of other durables to construct consumption flows for them and so include expenditure on them instead in our measure of durable consumption.

Each month, respondents saw one of five spending summary screens. Which screen was shown was randomized. The five screens are: (1) subtotals with drop-down and grand total; (2) subtotals with drop-down and no grand total; (3) long list with grand total; (4) long list with no grand total; and (5) summary sub-screens and no grand total. Our regressions include indicators for which summary screen the respondent saw in each wave.

Rental equivalence of housing for homeowners

For renters, we use expenditure on rent as housing consumption. For owners, we calculate a consumption flow by constructing the rental equivalent of their home. We do so by merging in external government-provided data on the average rent in HDB housing by city and the number of rooms. We assign households the rent reported in the government data corresponding to their type of housing, number of rooms, and city. The SLP records the following different types of housing: Housing Development Board (HDB) flat; Housing and Urban Development Company (HUDC) flat; Design, Build and Sell Scheme (DBSS) or Executive Condominium (EC) flat; private condo; landed house detached; and other housing. The survey records the number of rooms (ranging from 1 to 5) or if it's larger than five rooms ("Executive").

Service flow of vehicles

For households who rent a vehicle, we use vehicle spending as the service flow for vehicle consumption. For households who own a vehicle, we impute a service flow for vehicle consumption. The SLP doesn't include information on the characteristics of the vehicle or how long ago it was purchased, so we're unable to construct service flows based on vehicle type and depreciation. Instead, we impute a service flow based on vehicle expenditure among non-owners who have similar observables to vehicle owners. In particular, we regress vehicle spending on characteristics of the respondent (a quartic polynomial in age and income, indicators for education, race, town, household size) and survey wave. This imputation assumes that households who don't own a vehicle rent similar vehicles as households with the same characteristics that own in the same geographic areas. Vehicle ownership is reported annually, rather than monthly. If households report a change in ownership status, we assume that it occurred midway through the year when calculating service flows on a monthly basis.

Appendix B. Additional results

FIGURE B.1 NUMBER OF DISEASES BY SURVEY WAVE

(A) Individual-level

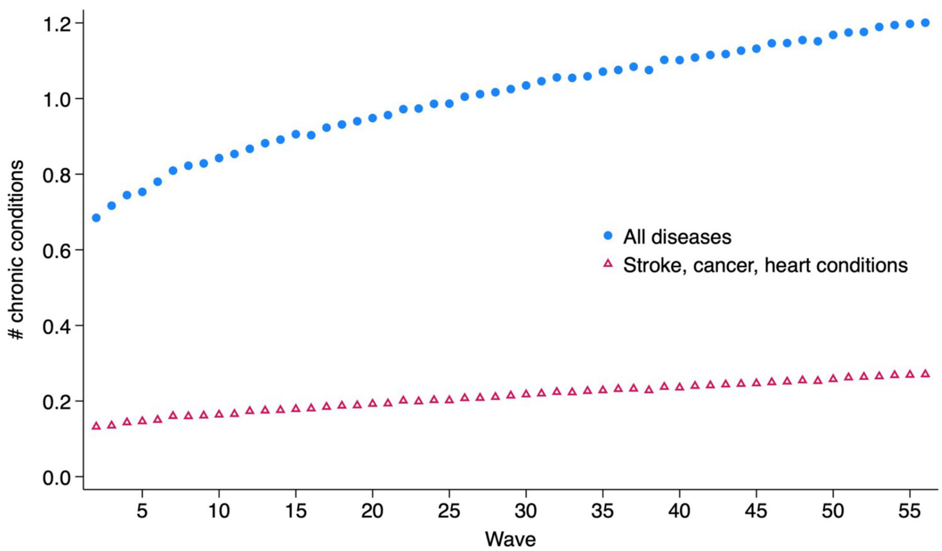

(B) Household-level

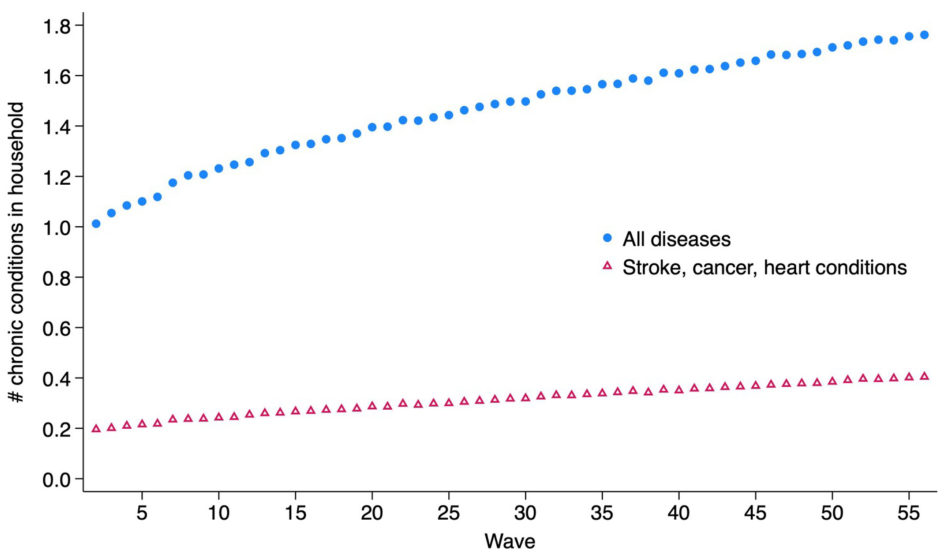

Notes: Figure shows the average number of diseases by wave across individuals (Panel A) and the average total number of diseases by households (Panel B). The circles denote the mean for all diseases and the hollow triangles denote the mean for heart conditions, stroke, or cancer. Means are calculated using individual-level sample weights, which are fixed over time. Sample is not restricted to be a balanced panel, so the mean rate sometimes declines between successive quarters due to changes in sample composition.

FIGURE B.2 CORRELATION BETWEEN SUBJECTIVE AND EMPIRICAL PROBABILITIES OF DISEASE

(A) Hypertension

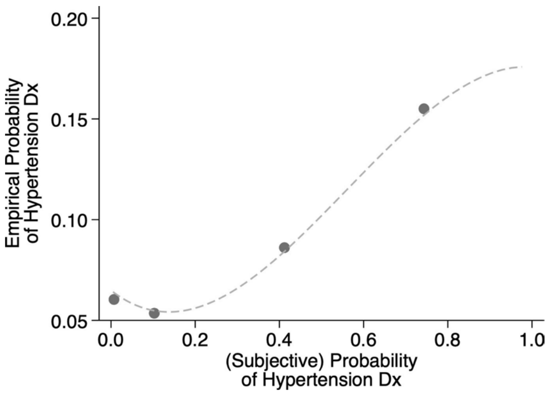

B) Diabetes

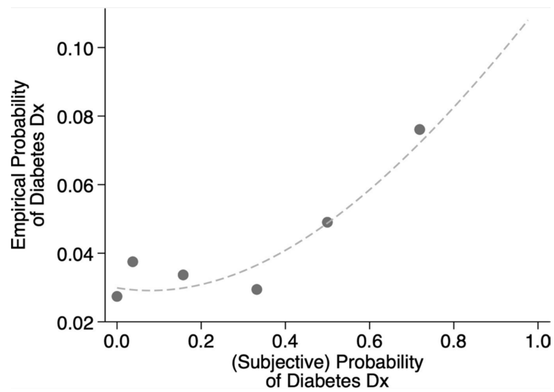

(C) Cancer

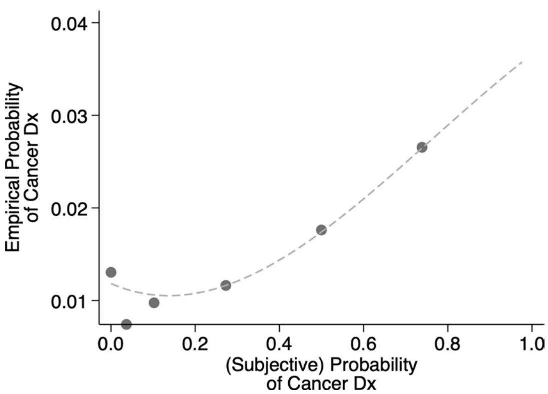

(D) Heart conditions

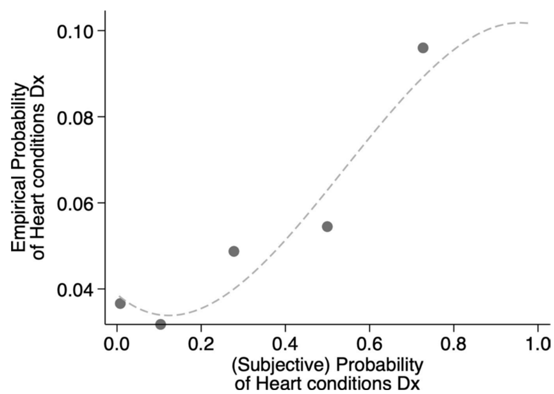

(E) Stroke

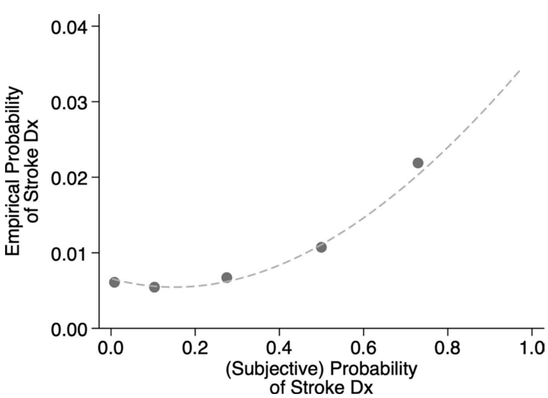

(E) Arthritis

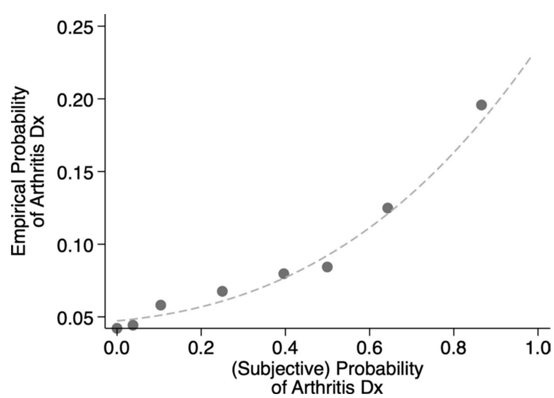

Notes: Figure plots binned scatterplots of a regression of an indicator for ever being diagnosed with a disease during the sample period against the subjective probability of ever being diagnosed with the disease. The number of equally sized bins are chosen based on data-driven methods of Cattaneo et al. (2024). The survey did not include a question about subjective probabilities of developing psychiatric conditions.

TABLE B.1 CONSUMPTION AND HEALTH INDEX BY MSA ASSETS

	(1)	(2)
Health index	0.101	0.095
	(0.070)	(0.048)
Individual fixed effects	Yes	Yes
Wave effects	Yes	Yes
MSA balances at baseline	Below median	Above median
R^2	0.751	0.745
N	4,957	4,691
NT	190,052	200,299

Notes: Table plots regression results of estimating Equation 5 via OLS, stratified by baseline level of MSA assets. The health index is calculated as the predicted probability of being in good health based on a regression model. The health index is calculated as the predictions from a logit model with indicators for chronic conditions, interactions between conditions and gender, and all two-way interactions between conditions. Column (1) is restricted to respondents with 2015 MSA assets below the median, and Column (2) is restricted to respondents with 2015 MSA assets above the median. All regressions include individual fixed effects, wave fixed effects, age, and indicators for household size and are weighted using survey sample weights. N denotes the number of unique individuals, and NT denotes the number of individual-waves.

FIGURE B.3 POISSON REGRESSIONS: CHANGE IN CONSUMPTION COMPONENTS

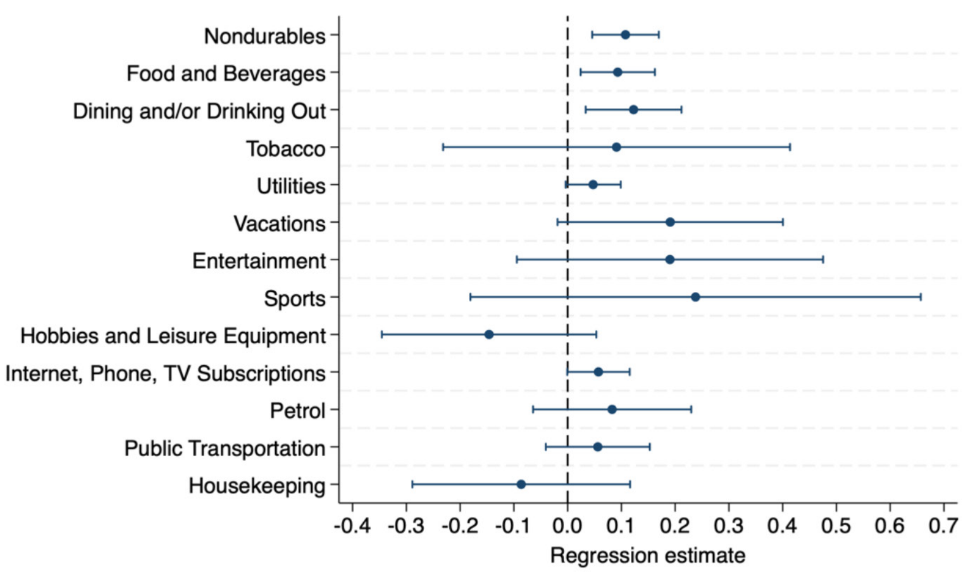

Notes: Figure plots Poisson regression estimates of components of nondurable consumption against the health index. Vacations include all transportation, accommodation, and recreational expenses on trips. Entertainment includes tickets to movies, sporting events, museums, etc. Hobbies and leisure equipment include photography, stamps, newspapers, magazines, books, camping, gardening, pets, electronic entertainment (e-magazines, e-books, streaming services), etc.

TABLE B.2 CONSUMPTION AND NUMBER OF CHRONIC CONDITIONS

	(1)
Number of chronic conditions	-0.011
	(0.005)
Individual fixed effects	Yes
Wave effects	Yes
R^2	0.789
N	11,001
NT	411,167

Notes: Table plots regression results of estimating Equation 5 via OLS where the independent variable is the number of chronic conditions. Regressions also include indicators for household size. N denotes the number of unique individuals and NT denotes the number of individual-waves.

FIGURE B.4 HETEROGENEITY IN STATE DEPENDENCE

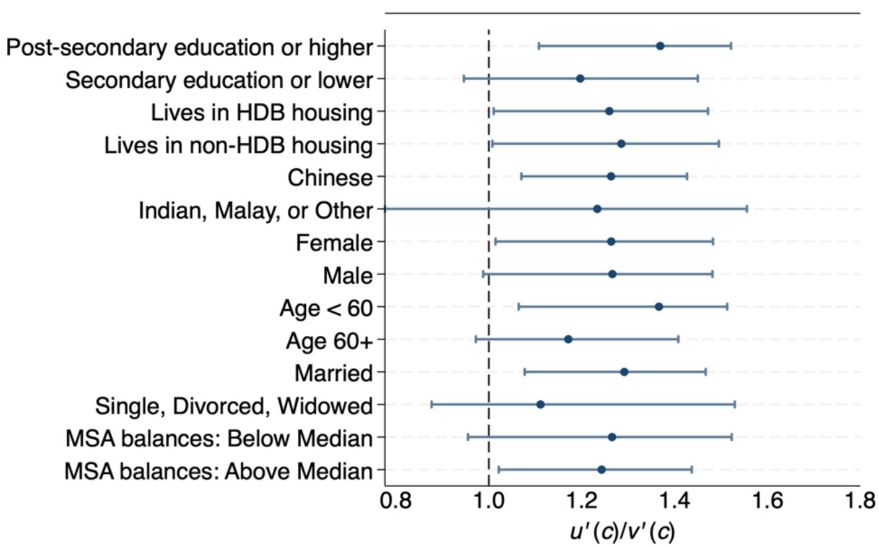

Notes: Figure plots estimates of state dependence for different sub-samples. Whiskers denote 90% confidence intervals calculated by block bootstrapping 250 samples for each group. For visual clarity, we top-code the upper bound of the confidence interval at 3.5.

TABLE B.3 RELATIONSHIP BETWEEN HEALTH INDEX AND OTHER OUTCOMES

	Positive health spending (1)	Log health spending (2)	Working (3)	Retired (4)
Health Index	-0.399	-0.734	0.081	-0.068
	(0.027)	(0.055)	(0.082)	(0.023)
Individual fixed effects	Yes	Yes	Yes	Yes
Wave effects	Yes	Yes	Yes	Yes
R^2	0.452	0.419	0.815	0.847
N	11,344	9,039	11,066	11,066
NT	*420,421*	*172,614*	*409,410*	*409,410*

Notes: Table plots regression results of estimating Equation 5 via OLS for an indicator of positive health spending (column 1), log of health spending (column 2), an indicator for working (column 3), and an indicator for retired (column 4). Regressions also include indicators for household size. Each column presents the results for a different dependent variable. N denotes the number of unique individuals, and NT denotes the number of individual-waves.

FIGURE B.5 BALANCE TESTS: LOTTERY WINNINGS BY OBSERVABLE CHARACTERISTICS

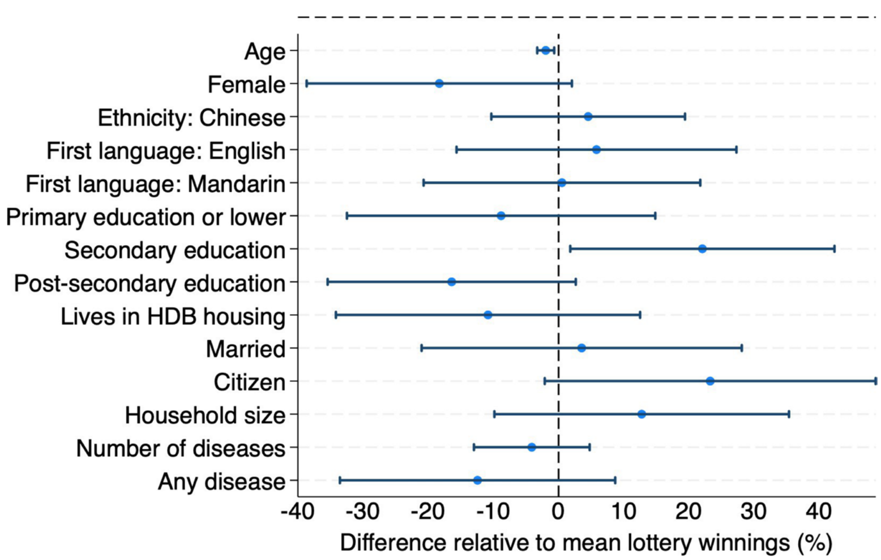

Notes: Figure shows regression estimates of lottery winnings in levels against observable characteristics, controlling for lottery purchases, and fixed effects for survey wave. Estimates are divided by mean lottery winnings (including zeros) and multiplied by 100 for interpretability. Whiskers denote 95% confidence intervals, calculated using standard errors clustered by individual respondents.

TABLE B.4 ROBUSTNESS: STATE-DEPENDENCE RESULTS, WEIGHTED SATISFACTION

	(1)
β_1 Health index ×Log consumption	0.019
	(0.010)
β_2 Log consumption	0.011
	(0.001)
β_3 Health index	0.154
	(0.022)
R^2	0.783
N	10,983
NT	409,062

Notes: Table shows results from estimating Equation 7 by OLS. The dependent variable in each regression is a weighted average of satisfaction measures. To construct weighted satisfaction, we assign the coefficients from Table 2 of Benjamin et al. (2014) as weights to the satisfaction questions measured in the SLP as follows: (1) coefficient = 0.31 (Aspect: "How satisfied you are with your life") assigned to SLP question: Life satisfaction as a whole; (2) average of coefficient = 0.37 (Aspect: "The quality of your family relationships") and coefficient = 0.16 (Aspect: "Your sense of community, belonging, and connection with other people") to the SLP question "How satisfied are you with your social contacts and family life?"; (3) average of coefficient = 0.10 (Aspect: "Overall quality of experience at work") and coefficient = 0.22 ("How rewarding the activities in your life are") to the SLP question "How satisfied are you with your daily activities and, if you are working, your job?"; (4) coefficient = 0.11 (Aspect: "Your material standard of living") to the SLP question "How satisfied are you with the total income of your household?"; (5) coefficient = 0.34 (Aspect: "Your financial security") to the SLP question "How satisfied are you with your overall economic situation?"; (6) coefficient = 0.41 (Aspect: "Your health") to the SLP question "How satisfied are you with your health?"; The mean of the dependent variable measuring weighted satisfaction is 0.55. Major diseases are stroke, heart conditions, or cancer. All diseases are major chronic conditions plus diabetes, arthritis, hypertension, and psychiatric conditions. Standard errors are clustered by individuals in parentheses. Regressions are weighted using survey sample weights. N denotes the number of unique individuals, and NT denotes the number of individual-waves.

TABLE B.5 LIFE SATISFACTION REGRESSIONS, NUMBER OF CHRONIC CONDITIONS

	Ordered Logit (1)	LPM (2)
β_1 Number of chronic conditions ×Log consumption	-0.051	-0.005
	(0.013)	(0.002)
β_2 Log consumption	0.086	0.013
	(0.016)	(0.002)
β_3 Number of chronic conditions	-0.171	-0.018
	(0.035)	(0.004)
R^2		0.651
N	8,855	10,957
NT	353,954	411,117

Notes: Table presents results from estimating Equation 7 using the number of chronic conditions instead of the health index. Log consumption is the natural logarithm of nondurable consumption. Column 1 presents results from the fixed effects ordered logit, and Column 2 presents results from the linear probability model. Both regressions include individual fixed effects and wave fixed effects. Standard errors are clustered at the individual level, and regressions are weighted using the survey's sample weights.

TABLE B.6 REGRESSIONS USING LOTTERY WINNINGS, NUMBER OF CHRONIC CONDITIONS

	Dep var: Life satisfaction			Dep var: Log consumption		
	All (1)	Below median health (2)	Above median health (3)	All (1)	Below median health (2)	Above median health (3)
Lottery winnings/1,000	0.021	0.032	0.015	0.016	0.016	0.013
	(0.006)	(0.007)	(0.007)	(0.005)	(0.007)	(0.007)
Lottery purchases/1,000	-0.005	-0.007	-0.003	0.012	0.013	0.011
	(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.002)
R^2	0.016	0.016	0.025	0.235	0.232	0.234
N	5,524	2,393	2,990	5,524	2,393	2,990
NT	11,295	4,348	5,707	11,257	4,341	5,693

Notes: Table presents results from estimating Equation 8 using the number of chronic conditions. The median conditions refers to the median number of chronic conditions. Columns 1–3 present results for life satisfaction, and columns 4–6 present results for log consumption estimated, both via OLS (some ordered logits for life satisfaction failed to converge). All regressions include wave fixed effects. Log consumption is the natural logarithm of nondurable consumption.

About the authors

Adam Leive is an assistant professor at UC-Berkeley's Goldman School of Public Policy. Leive is an economist who studies consumer behavior in health insurance, retirement saving and other social insurance programs. His research on Health Savings Accounts received the 2022 TIAA Paul A. Samuelson Award for outstanding scholarly writing on lifelong financial security.

Leive is a Faculty Research Fellow at NBER, a TIAA Institute Fellow, a J-PAL affiliate, and a member of the G53 Network on Financial Literacy and Personal Finance. He earned his PhD from the University of Pennsylvania's Wharton School and his BA from Princeton University's School of Public and International Affairs.

Jessica Ya Sun is a health economist whose work employs empirical methods to examine policy-relevant issues in population aging, healthcare markets, and social insurance. Her research focuses on the behavioral and welfare implications of public policies, with particular emphasis on retirement security and health insurance. Recently, she has also investigated the effects of centralized procurement policies on patients' medical burdens and the impact of anti-corruption campaigns on healthcare delivery. Her work has been published in the *Journal of Economic Behavior & Organization* and *Health Economics*. She holds a PhD in Economics from Singapore Management University.

About the TIAA Institute

The TIAA Institute helps advance the ways individuals and institutions plan for financial security and organizational effectiveness. The Institute conducts in-depth research, provides access to a network of thought leaders, and enables those it serves to anticipate trends, plan future strategies, and maximize opportunities for success.

To learn more, visit tiaainstitute.org.

Join the conversation online: @TIAAInstitute

We are grateful for helpful comments on an earlier draft from Misha Dworsky, Evan Saltzman, and seminar participants at UC-Berkeley Demography, UC-Berkeley Public Health, Universidade Federal de Minas Gerais, and National Tax Association Meetings. The Singapore Ministry of Education (MOE) Academic Research Fund Tier 3 [grant numbers MOE2019- T3-1-006] provided financial support for the collection, analysis, and interpretation of data. This research was supported by funding from the TIAA Institute.
TIAA Institute is a division of Teachers Insurance and Annuity Association of America (TIAA), New York, NY.

4842102-A

TIAA Institute

(09/25)