

Fit to Consume: How Health Shapes Preferences for Consumption*

Adam Leive
UC-Berkeley
& NBER

Jessica Ya Sun
Huazhong University
of Science and Technology

September 2026

Abstract

Whether the marginal utility of consumption differs by health is relevant to many economic models, but evidence on its sign and magnitude remains mixed. Using monthly panel data from Singapore, we show through multiple approaches that health and consumption are complements. First, healthier people have a higher MPC from lottery winnings. Second, we derive a sufficient-statistics formula and estimate that the marginal utility of consumption falls 5% for a one-standard-deviation decline in health. Third, life satisfaction data provide corroborating evidence. Calibrations show state dependence meaningfully affects optimal insurance and saving and, if ignored, biases estimates of risk aversion.

JEL Codes: I10, I31, G52, D15

*Leive: leive@berkeley.edu. Sun (Corresponding author): jessicasunya@hust.edu.cn. We are grateful for helpful comments from Misha Dworsky, Itzik Fadlon, Amy Finkelstein, Seonghoon Kim, Erzo F.P. Luttmer, Matt Notowidigdo, Andrew Oswald, Kathleen McGarry, Evan Saltzman and seminar participants at Stanford University, UC-Berkeley, UC-Irvine, SoCCAM, Fudan University 7th Annual Conference on the Economics of Social Security, Sichuan University Workshop on Frontier Research in Big Data and Health Economics, Universidade Federal de Minas Gerais, National Tax Association Meetings, and TIAA Institute/Pension Research Council Symposium. 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. The content, findings and conclusions are the responsibility of the author(s) and do not necessarily represent the views of TIAA or the TIAA Institute.

1 Introduction

The primary goal of insurance and saving is to transfer resources from states of the world where they are less valuable to states where they are most valuable. By smoothing the marginal utility of consumption across states and over time, these instruments help agents maximize expected utility in the face of income losses or other financial shocks. However, a person’s health represents another state that may affect their utility of consumption (Zeckhauser 1970; Arrow 1974; Cook and Graham 1977), with implications for optimal policy design. If illness reduces the value of non-medical consumption, for instance, then optimal health insurance benefits and retirement saving should be lower than if utility is state-independent. Conversely, if the marginal utility of consumption is higher when ill, optimal levels of insurance benefits and savings should rise.

Prior research has pursued various approaches to evaluate how the marginal utility of consumption depends on health, and empirical evidence remains inconclusive.¹ Finkelstein, Luttmer and Notowidigdo (2013) find that marginal utility declines in poor health by examining correlations between utility proxies, chronic conditions, and permanent income using the Health and Retirement Study (HRS), but other studies using similar methods and data from other countries find the opposite (Wang and Wang, 2020; Bassoli, 2022; Amo-Agyei, 2025). A second approach elicits hypothetical allocations of money across health states: Ameriks et al. (2020) find people prefer more consumption when disabled, while Brown, Goda and McGarry (2016) find the opposite or do not reject state-independent utility, with results varying by age and type of disability. Structural life-cycle models combined with data on consumption or assets have also produced divergent estimates, some finding a higher marginal utility when ill (Lillard and Weiss 1997; Yogo 2016), others finding the opposite (Kojen, Van Nieuwerburgh and Yogo 2016b; Blundell et al. 2024), and others finding little relationship (De Nardi, French and Jones 2010). Consequently, the literature has not reached consensus on how to model state dependence between health and consumption.²

We make progress on this question by pursuing multiple empirical approaches using monthly panel data from Singapore. We use 55 waves of the Singapore Life Panel (SLP), a monthly survey with detailed information on consumption, health, employment, and well-

¹See Finkelstein, Luttmer and Notowidigdo (2009) for a discussion of the trade-offs and assumptions of different approaches.

²The modal assumption is that health and consumption are additively separable (see e.g., Hall and Jones 2007; Hugonnier, Pelgrin and St-Amour 2013; Finkelstein, Hendren and Luttmer 2019; Fonseca et al. 2023). Some studies, often as part of robustness, analyze the case that the marginal utility of consumption is higher when healthy (Reichling and Smetters 2015; Lieber and Lockwood 2019; Hendren 2021), others that it is higher when ill (Kojen, Philipson and Uhlig 2016a), or both possibilities (Deshpande and Lockwood 2022).

being. We construct an index of health by predicting self-assessed health status as a function of the diagnosis of chronic conditions. The SLP also contains modules on risk preferences that allow us to set-identify risk aversion for each respondent, following methods established in Barsky et al. (1997). We first document the direction of state dependence with reduced-form analysis and then quantify its magnitude with two optimization-based formulas.

We begin with motivating evidence that establishes the sign of state dependence. Playing the lottery is common in Singapore, and winnings provide random variation in windfall income. The SLP measures lottery winnings and purchases in three waves. Building on prior work showing that lottery winnings increase consumption in this population (Kim and Koh 2021), we test whether the marginal propensity to consume (MPC) out of winnings varies with baseline health. We find the MPC is over three times larger for people in above-median health, indicating that health and consumption are complements. This test has the virtue of exploiting clean variation in windfalls and not requiring agent optimization, but it does not deliver the magnitude of state dependence, which requires a model of consumption decisions.

To recover the magnitude, we derive an empirically-tractable formula for health state dependence as a ratio of marginal utilities based on the Euler equations of a life-cycle model of consumption and savings. State dependence is a function of the gap in (non-medical) consumption between health states, risk and time preferences, and parameters that govern uncertainty about future health and mortality. Using all 55 survey waves, we estimate the formula’s key input—how consumption responds to within-person changes in health—and find that a one-standard-deviation decline in the health index corresponds to a 1.6% drop in non-durable consumption. Combining this estimate with our theory, the marginal utility of consumption is 5% lower for a one-standard-deviation decline in health. We document heterogeneity in magnitude but not in sign of state dependence across observable characteristics, and our estimates are robust to several alternative specifications and extensions.

The crucial assumption of the Euler equation approach is that agents optimize their consumption across states and over time. This assumption is powerful as it allows us to rely on well-measured consumption data rather than to evaluate all margins of behavior that might change in response to a health shock. For example, our formula is robust to the possibility that spouses adjust their labor supply after a health shock because this response is already reflected in household consumption. In contrast, our formula for state dependence would not be valid if households follow rules-of-thumb in their consumption decisions or have present-biased time preferences that violate the Euler equation. We also consider the role of bequest motives and show how modifying our formula to include incidental bequests as in

[Lockwood \(2018\)](#) leads to slightly larger magnitudes compared to our main results.

As a complement to the Euler equation approach, we derive a within-state first-order condition formula for state dependence that does not require inter-temporal optimization. This approach exploits the optimality condition that consumers equate marginal utilities across goods within a period, comparing how state-dependent consumption responds to changes in health relative to how a state-independent “outside good” changes. Because it operates within a period, the within-state approach requires no assumptions about discount rates, interest rates, or survival probabilities, and is robust to concerns about the precise timing of health shocks relative to saving decisions. Instead, the researcher must identify a good whose marginal utility does not depend on health. We use charitable giving and gifts, although the prevalence of zeros in this category requires a modified functional form of utility with an additional parameter. The within-state estimates are larger in magnitude, and the two approaches reinforce each other because they rely on different assumptions.

A final piece of empirical evidence comes from data on subjective well-being. We estimate the specification suggested by [Finkelstein, Luttmer and Notowidigdo \(2013\)](#), in which life satisfaction (a proxy for utility) is regressed on consumption, health, and their interaction, along with individual and time fixed effects. The sign of the interaction reveals the direction of state dependence, and the results again indicate that the marginal utility of consumption is higher when healthy.

Combining our estimated magnitudes of state dependence with simple calibrations, we quantify how the parameter affects optimal insurance and saving decisions. Starting with health insurance, we revisit the standard result that risk-averse individuals should fully insure if premiums are actuarially fair. By contrast, state dependence induces partial insurance; using a textbook model of how much insurance to buy against a prospective loss, a one standard deviation decrease in the health index results in people optimally purchasing 12.3% less insurance. For retirement saving, we simulate a simple life-cycle model of consumption and savings that features changes in health status by age. Compared to state-independent utility, optimal savings at retirement are 1.4% lower because the least healthy years of life occur in old age. Building off these illustrative examples, we show how ignoring state dependence can meaningfully underestimate risk aversion if the parameter is not measured separately (as it is in our context) and the true model is state-dependent utility. This insight is relevant for a range of empirical settings, where it is common to estimate risk aversion by matching observed choices to those predicted by a structural model. A model with state-independent utility would rationalize the insurance and saving choices in our calibrations with a lower value of risk aversion, rather than a higher marginal utility of consumption when healthy.

The setting of Singapore offers a useful context for studying state dependence, and the richness of the SLP provides advantages relative to data used in previous work. One empirical challenge is determining whether a drop in consumption is due to preferences or due to incomplete insurance. Singapore is among the wealthiest countries in the world and covers its population through an established 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. In terms of data, the SLP’s monthly frequency 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. Another advantage of the SLP is that we can separately estimate risk aversion over large financial stakes for the same population used to estimate consumption changes, which is important to align in sufficient-statistics frameworks (Chetty 2006) and because state dependence is sensitive to assumptions about risk aversion (Finkelstein, Luttmer and Notowidigdo 2009). Finally, the SLP enables us to implement multiple empirical approaches using the same data, which strengthens our conclusion about the complementarity between health and consumption.

The paper proceeds as follows. Section 2 describes Singapore’s health system and the SLP. Section 3 establishes the sign of state dependence using variation in lottery winnings. Section 4 derives a sufficient-statistics formula that delivers its magnitude, and Section 5 estimates each of the formula’s inputs. Section 6 evaluates robustness to both the empirical specification and the formula’s assumptions. Section 7 corroborates the sign of our findings using life satisfaction data. Section 8 illustrates the implications of our estimates for health insurance, retirement saving, and the estimation of risk aversion. Section 9 briefly concludes.

2 Setting and Data

This section describes the institutional setting and data that underlie our empirical analysis. We first provide an overview of healthcare financing in Singapore, which is based on medical savings accounts. Along with the country’s high income level, this institutional feature plays an important role for interpreting our results: a central empirical challenge is distinguishing whether consumption falls in bad health because of preferences or because of incomplete insurance, and account balances will later allow us to test between these explanations. We then describe the Singapore Life Panel (SLP), focusing on how we measure consumption and health at monthly frequency. Finally, we construct an index of health status from chronic

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

conditions, which serves as our measure of health throughout the paper.

2.1 Background on Singapore’s health system

Personal savings accounts are central to financing healthcare 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, which was 71,500 SGD ($\approx$ \$55,600 USD) in 2024.⁶ 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 cannot afford MediShield are covered by the safety-net program Medifund, which is financed by a government investment fund and requires pre-approval for care to be reimbursed. Collectively, MediSave, MediShield, and MediFund are referred to as the “3M system.”

2.2 The Singapore Life Panel (SLP)

The SLP is a longitudinal survey of a representative sample of Singaporean citizens aged 50 to 70 in 2015. The baseline survey was conducted in July 2015 and surveyed more than

⁴For 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).

⁵The ceiling in 2024 was 6,528 SGD ($\approx$ 5,100 USD) to 8,568 SGD ($\approx$ 6,700 USD) depending on age.

⁶This limit is called the “Basic Health Sum” and increases each year as new cohorts turn 65. For more information about MSA contribution rates and limits, see <https://www.cpf.gov.sg/service/article/is-there-a-maximum-amount-that-i-can-save-in-my-medisave-account> (accessed December 26, 2025).

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 generally around 8,000, which corresponds to a response rate of 65%. We use 55 waves of the survey, spanning the period July 2015 to December 2019.

Detailed expenditure data is collected monthly and records spending on 33 different items. Our analysis focuses on non-durable consumption. In supplementary analyses, we also examine durable consumption, for which we construct a rental equivalent for home owners and impute service flows for households who own vehicles. Appendix A details our consumption measures and the classification of items into non-durables 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 utilities 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 health care spending from our main consumption measure and instead estimate the impact of health shocks on health care 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.

Each month, the SLP asks individuals if they have 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.

A key strength of the SLP is that its modules enable us to test for state dependence using multiple, complementary approaches on the same individuals. In addition to the core questions described above, lottery purchases and winnings are measured in three waves. In

⁷See <https://rosa.smu.edu.sg/singapore-life-panel/about-singapore-life-panelr> (accessed May 2, 2024) for more details about the survey.

one wave, respondents are asked about hypothetical gambles that we use to set-identify risk aversion. We describe these variables in detail in [Section 3](#) and [Section 5](#).

2.3 Sample selection and descriptive statistics

We make minimal restrictions to construct our analytic sample: we drop respondents who appear only in the baseline or test waves, restrict to ages 45 to 75, and exclude a small number of individuals who are presented with different summary screens in the consumption module across waves or who live with a household member under age 22, since consumption may change when children leave the home. Together, these restrictions reduce the sample by 4.6%, yielding an analytic sample of 12,742 individuals in 9,547 households ([Appendix Table B.1](#) reports the number excluded at each step). 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 429,442 observations at the individual-wave level. Panel A presents demographic information. The average age is 61 years and 53% of the sample is female. The large majority (86.6%) identify as ethnically Chinese. Most respondents (68.4%) 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 education, 23% have completed primary schooling or less, 41% secondary schooling, and 36% post-secondary or higher.

Panel B reports income, consumption, and assets. On average, monthly household non-durable consumption per capita is 976 SGD and durable consumption is 1,694 SGD. By comparison, monthly household income per capita is 5,009 SGD, on average, including both earnings and other sources. The average household has approximately 40,261 SGD in total assets per capita (excluding housing), with 26,240 SGD in MSA balances. Average monthly out-of-pocket healthcare spending is 46 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, 45.8% 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
<i>Panel A. Demographics</i>		
Age	61.208	5.956
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	0.480
Live in HDB housing	0.684	0.465
Married	0.792	0.406
Citizen	0.951	0.216
Household size	1.861	0.397
<i>Panel B. Consumption, Income, Assets</i>		
Non-durable consumption per capita (monthly)	976	897
Durable consumption per capita (monthly)	1,694	1,292
Health care consumption per capita (monthly)	46	138
Income	5,009	5,672
Household MediSave assets per capita	26,240	17,607
Household total assets per capita	40,261	89,404
<i>Panel C. Chronic Conditions and Health Status</i>		
Number of chronic conditions	1.002	1.124
Heart disease	0.130	0.336
Stroke	0.027	0.163
Cancer	0.053	0.224
Diabetes	0.191	0.393
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. Household total assets per capita exclude housing wealth.

2.4 Constructing an index of health

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 measures of health in various ways, including counting the number of chronic conditions (Finkelstein, Luttmer and Notowidigdo 2013; Hosseini, Kopecky and Kai 2022), using self-assessed health (De Nardi, Pashchenko and Porapakkarm 2025), predicting self-assessed health from self-reported conditions and limitations (Poterba, Venti and Wise 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 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. For example, strokes and heart attacks are likely more debilitating 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 matter. As robustness, we consider alternative constructions of the health index, including a LASSO regression to avoid possible overfitting or simply counting the number of chronic conditions.⁸ 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 are 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 is unclear that the difference between “poor” and “fair” health is 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.


⁸Difficulties with Activities of Daily Living (ADLs) are only measured in three survey waves. To assess how much additional information ADLs provide conditional on chronic conditions, we re-estimate the health index using those waves, including and excluding ADLs and all two-way interactions. Including ADLs increases the R^2 from 0.096 to 0.114 relative to a model with only chronic conditions. Since this increase in explained variation is modest but the sample size falls by 95%, we do not pursue models that include ADLs.

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, with similar patterns for men and women. 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. The coefficients are 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 is often asymptomatic. Diabetes, cancer, heart conditions, and arthritis reduce the index by 0.2 to 0.25, occasionally with small differences by gender. Psychiatric conditions correspond to the largest reductions among these seven.

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 intervals on the difference relative to not having that chronic condition. The fraction of people reporting that condition is displayed in parentheses on the y-axis.

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% of the 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.⁹ 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 using only the same conditions recorded in the SLP (excluding psychiatric conditions, which are not recorded) captures 12.3% of the variation in health status. Using all observed 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.

3 Motivating Evidence from Lottery Wins

Before presenting a formal model, we establish the direction of state dependence using an approach that requires minimal assumptions: we test whether the marginal propensity to consume (MPC) out of random lottery winnings differs by health status. If health and consumption are complements, the MPC should be higher when healthy than when unhealthy. The reverse is true if health and consumption are substitutes, while equal MPCs across health states indicate that consumption is state-independent. We implement this test by comparing the MPC between people whose baseline health index is above or below the sample median.

Our test is similar in spirit to survey questions that ask respondents to allocate a fixed amount of money between a disabled state and a non-disabled state (Brown et al. 2016; Ameriks et al. 2020). A lottery win confronts people with an analogous allocation problem, as the winner chooses between consuming the windfall in their current state of health and transferring it to other states through saving or leaving a bequest. The advantage of the lottery windfall is that stakes are real and health states have been experienced, so no respondent must forecast preferences in an unfamiliar state. The disadvantage to this revealed-preference test is that observed behavior may also reflect constraints rather than preferences, which the survey questions remove by construction. We assess this issue by incorporating baseline assets in supplemental analysis. Differences in expected survival could also generate differences in the MPC. If people in worse health expect to spread a windfall over fewer remaining years, their MPC would be higher.

⁹We restrict the MEPS to those with health insurance and use surveys from 2016–2019 to match the SLP’s sample period.

Because this survival effect works against finding a higher MPC when healthy, our test is conservative for detecting complementarity.

The Singapore setting is well-suited to this test. Playing the lottery is common, with 53% of the sample reporting they purchased lottery tickets in the last year. Previous research using the SLP has studied the effect of lottery winnings on life satisfaction (Kim and Oswald 2021) and self-assessed health (Kim and Koh 2021). As a mechanism for their health effects, Kim and Koh (2021) show that lottery winnings increase consumption; our focus is to test whether this consumption response differs by baseline health.

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 1,981 SGD and annual winnings are 409 SGD (886 SGD among winners), on average (Appendix Table B.2). 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), who used data from waves 16 and 28. Observable characteristics are balanced conditional on lottery spending, supporting the assumption of conditional random assignment: the p -value from the omnibus test that all covariates jointly equal zero is 0.22 (Appendix Table B.3).

We start by estimating the MPC from lottery wins using the following specification:

$$C_{it} = \alpha + \eta_1 L_{it} + \eta_2 M_{it} + \mu_t + \xi X_{it} + u_{it} \quad (1)$$

where C_{it} denotes non-durable consumption for individual i in wave t , L_{it} denotes lottery winnings and M_{it} denotes lottery spending. η_1 represents the marginal propensity to consume from lottery winnings. As discussed in Kim and Oswald (2021), controlling for lottery spending is critical to estimate the effect of lottery winnings because more ticket purchases increase the chance of winning. And ticket purchases may affect consumption by crowding out other spending. Covariates (X_{it}) include indicators for the frequency of playing the lottery, race, marriage, education, household size, citizenship, gender, and a quartic in age. We split the regression by those above or below median values of the health index in the first wave the person is observed. Our second specification includes interactions between baseline health and lottery measures to test for a difference in the MPCs:

$$C_{it} = \zeta_0 + \zeta_1 L_{it} + \zeta_2 M_{it} + \zeta_3 (L_{it} \times D_i) + \zeta_4 (M_{it} \times D_i) + \zeta_5 D_i + \omega_t + \rho X_{it} + e_{it} \quad (2)$$

where D_i equals one if individual i 's health index in their first observed wave is above the sample median and zero otherwise. The coefficient ζ_3 measures the difference in the MPC

from lottery winnings between those in better health and those in worse health at baseline.

Table 2 – MPC from Lottery Winnings, by Baseline Health

Dependent variable: Non-durable consumption				
	All	Below Median	Above Median	All
	(1)	Health	Health	(4)
	(2)	(3)		
Lottery winnings	0.217*** (0.069)	0.087 (0.084)	0.324*** (0.101)	0.106 (0.084)
Lottery purchases	0.236*** (0.032)	0.243*** (0.048)	0.226*** (0.036)	0.256*** (0.048)
Winnings $\times$ Above median health				0.220* (0.128)
Purchases $\times$ Above median health				−0.034 (0.054)
Above median baseline health				343.543 (236.238)
N	5,523	2,711	2,812	5,523
NT	11,297	5,517	5,780	11,297

Notes: Table presents estimates of Equation 1 (columns 1–3) and Equation 2 (column 4). The dependent variable is non-durable consumption in SGD aggregated over the 12 months corresponding to the lottery module. Lottery winnings and purchases are measured in SGD, so the coefficient on winnings is a marginal propensity to consume. Health is measured at baseline. All regressions include wave fixed effects and covariates (indicators for race, marriage, education, household size, frequency of playing the lottery, and a quartic in age). Standard errors clustered by household in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The results are presented in Table 2. Column 1 shows the MPC out of winnings is 0.217 in the full sample. Splitting by baseline health, the MPC is 0.087 and statistically insignificant for those in below-median health (column 2), compared to 0.324 and highly significant for those in above-median health (column 3). Column 4 confirms this pattern by interacting lottery winnings and spending with an indicator for above-median baseline health: the interaction with lottery winnings is positive (0.220, $p = 0.086$), indicating that the propensity to consume out of a windfall is larger when healthy. Because analytic standard errors can be unreliable when prize distributions are highly skewed, we follow prior literature by also assessing significance using randomization inference (Cesarini et al. 2016; Picchio et al. 2018; Kim and Koh 2021). We randomly reassign winnings among players in the same wave and decile of lottery ticket spending and re-estimate Equation 2 on 5,000 placebo samples. The actual interaction coefficient is larger in magnitude than all but 4.3% of the

placebo estimates ([Appendix Figure B.2](#)). This p -value of 0.043 from the permutation test reinforces the results from [Table 2](#).¹⁰

The heterogeneity in consumption responses by health status does not appear to reflect financial constraints. If constraints drove the difference in MPCs, people with fewer resources should respond differently to windfalls regardless of their health. To evaluate this possibility, we let the MPC vary directly with baseline resources by interacting the lottery measures with indicators for above-median financial assets and above-median MSA balances, measured in the first wave. If baseline health merely proxied for constraints, its interaction with winnings should shrink toward zero once these terms are included. Instead, the coefficient rises slightly to 0.25 and is marginally significant at the 5 percent level with both sets of asset interactions included, while the asset interactions themselves are smaller and statistically insignificant ([Appendix Table B.6](#)).

Because lottery players may differ from the broader population, we also re-estimate the MPC regressions weighting each player by the inverse of an estimated propensity to play the lottery. The results from this reweighting exercise ([Appendix Table B.7](#)) are similar to [Table 2](#), indicating that selection into playing the lottery based on observable characteristics does not drive the difference in MPCs by health status.

The higher MPC when healthy provides evidence that health and consumption are complements. This estimate does not by itself recover a preference parameter, which motivates the model we develop next.

4 Theory

The lottery evidence establishes the sign of state dependence but not its magnitude, which requires a model of consumption decisions. In this section, we derive a sufficient-statistics formula for health state dependence from the Euler equation of a life-cycle model of consumption and savings, with full details in [Appendix C](#). The formula expresses state dependence as a function of three sets of measurable inputs: the change in consumption between health states, risk aversion, and parameters governing time preferences and uncertainty about future health and survival. [Section 5](#) then estimates each input using the SLP. At the end of this section, we describe a complementary formula based on within-state first-order conditions from a two-good model, which is also derived in [Appendix C](#) and does not require inter-temporal assumptions.

An agent chooses consumption today to maximize their discounted expected utility

¹⁰We find similar patterns for durable and total consumption, with a larger MPC for those in better health, though the interaction terms are less precisely estimated, with permutation p -values of 0.17 and 0.06, respectively ([Appendix Table B.4](#), [Appendix Table B.5](#), [Appendix Figure B.3](#), and [Appendix Figure B.4](#)).

of lifetime consumption. Within each period, the agent enters with wealth determined by the previous period's saving decision, observes their current health state, and then chooses consumption and saving. The agent's wealth evolves according to $W_t = R(W_{t-1} + y_{t-1} - c_{t-1})$, where y is income, c is consumption, and $R = 1 + r$ is the gross interest rate. Health shocks arrive between periods, after saving has been chosen. We assume medical costs are fully insured and that non-medical consumption does not affect health status.¹¹ Utility is intertemporally separable and the future is discounted by the factor $0 < \beta < 1$. 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 $v(c_{b,t})$ denoting the utility of consuming c_b in time t if unhealthy.¹² The probability of being healthy in time t if the agent is healthy in time $t - 1$ is denoted by p and represents the agent's belief at the time of the saving decision.

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

$$u'(c_{h,t-1}) = S\beta R[p u'(c_{h,t}) + (1 - p)v'(c_{b,t})] \quad (3)$$

where S is the agent's belief about their probability of survival to time t if healthy in time $t - 1$. We make two approximations that allow us to write state dependence as a simple formula based on this equation. First, we assume $u'(c_{h,t-1}) \approx u'(c_{h,t})$, which is a reasonable approximation when time periods are short.¹³ Equation 3 can then be rearranged as:

$$u'(c_{h,t-1}) = \frac{S\beta R(1 - p)}{1 - S\beta R p} v'(c_{b,t}) \equiv \delta v'(c_{b,t}) \quad (4)$$

where δ is a scaling factor that depends on the parameters governing uncertainty and the dynamics of the decision problem. Next, we follow approaches used in previous studies to measure the value of insurance (see e.g., Baily 1978; Gruber 1997; Chetty 2006; Hendren 2017; Fadlon and Nielsen 2019; Coyne et al. 2024) and approximate the marginal utility of consumption when healthy by taking a second-order expansion around $c_{b,t}$:

$$u'(c_{h,t-1}) \approx u'(c_{b,t}) + u''(c_{b,t})(c_{h,t-1} - c_{b,t}) + \frac{1}{2}u'''(c_{b,t})(c_{h,t-1} - c_{b,t})^2 \quad (5)$$

The marginal utilities in both health states now correspond to the same consumption level,

¹¹In Section 6, we consider the robustness of this framework to binding budget constraints and the possibility that some consumption may represent investments in health as in Grossman (1972).

¹²While we assume the agent is either healthy or unhealthy, the framework naturally extends to any number of health states.

¹³In Appendix C, we report that the median approximation error is 0 under CRRA utility.

and combining with the right-hand side of Equation 4 and dividing by $u'(c_{b,t})$ yields:

$$1 + \frac{u''(c_{b,t})}{u'(c_{b,t})}(c_{h,t-1} - c_{b,t}) + \frac{1}{2} \frac{u'''(c_{b,t})}{u'(c_{b,t})}(c_{h,t-1} - c_{b,t})^2 = \delta \frac{v'(c_{b,t})}{u'(c_{b,t})} \quad (6)$$

Finally, under the assumption that preferences satisfy constant relative risk aversion (CRRA) locally, we substitute $-\frac{u''(c_{b,t})}{u'(c_{b,t})}c_{b,t} = \gamma$ as the coefficient of relative risk aversion, $-\frac{u'''(c_{b,t})}{u''(c_{b,t})}c_{b,t} = \gamma + 1$ as the coefficient of relative prudence, and define the proportional change in consumption between health states as $\frac{\Delta c_t}{c_{b,t}} \equiv \frac{c_{h,t-1} - c_{b,t}}{c_{b,t}}$ to obtain:

$$\frac{u'(c_{b,t})}{v'(c_{b,t})} = \frac{\delta}{1 - \gamma \frac{\Delta c_t}{c_{b,t}} \left[1 - \frac{\gamma+1}{2} \frac{\Delta c_t}{c_{b,t}} \right]} \quad (7)$$

The ratio of marginal utilities $\frac{u'}{v'}$ is a function of the proportional change in consumption between health states in successive periods Δc_t , 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 is lower when healthy. $\frac{u'(c)}{v'(c)} = 1$ represents the special case of state-independent utility.

One potential concern with this approach is that it requires several inter-temporal assumptions, both about individual parameters and the timing of health shocks relative to saving decisions. An alternative approach exploits within-state first-order conditions rather than inter-temporal ones. In Appendix C, we develop a two-good model in which consumers allocate expenditure between a health-dependent consumption good and an outside good whose marginal utility does not depend on health status. The first-order conditions within a period yield a state-dependence formula that shares the core inputs of consumption changes and risk aversion ($\frac{\Delta c}{c}$ and γ), but replaces the inter-temporal parameter δ with quantities related to the response of the outside good to changes in health. Because this approach operates entirely within a period, it requires no assumptions about discount rates, interest rates, survival probabilities, or the timing of health shocks relative to saving decisions. It is therefore robust to concerns about whether wealth is predetermined at the time of a health shock.

We focus on the Euler equation approach in the main text for two reasons. First, the within-state approach requires the researcher to identify a consumption good whose marginal utility is plausibly invariant to health. We implement this using charitable giving and other gifts, but the prevalence of zeros in this category necessitates modifying the functional form to Stone-Geary utility, which requires calibrating an additional parameter and so introduces more modeling choices. Second, the Euler equation approach yields a parsimonious formula

whose inputs can be estimated or calibrated without taking a stand on which goods are state-independent. Nonetheless, we view the two approaches as providing a useful check on each other because they make different identifying assumptions. As we show in [Appendix C](#), the estimates based on the within-state first-order conditions are larger in magnitude than the Euler equation estimates, reinforcing the empirical findings we present next.

5 Estimates of State Dependence from Consumption Changes

This section provides our main estimates of state dependence using all 55 survey waves. We sequentially estimate each parameter in [Equation 7](#), starting with how consumption responds to changes in the health index. We then describe how we use the SLP to internally calibrate risk aversion and health status transitions and, finally, we detail which parameters are externally calibrated. Combining these inputs, our benchmark estimate implies that the marginal utility of consumption falls by about 5% for a one standard deviation decline in the health index.

5.1 Changes in Consumption between Health States

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

$$\ln(C_{ijt}) = \alpha_i + \tau_t + \lambda H_{ijt} + \eta X_{ijt} + \epsilon_{ijt} \quad (8)$$

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 $\frac{\Delta c_t}{c_{b,t}}$ that appears in [Equation 7](#) from [Section 4](#). By including an individual fixed effect α_i and a wave fixed effect τ_t in [Equation 8](#), 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.¹⁴

[Table 3](#) presents the results. On average, consumption is 10.8 percent higher when healthy compared to when unhealthy. 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

¹⁴A 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 and Gruber 2002](#); [Mohan 2013](#); [Liu 2016](#); [Garcia-Mandico, Reichert and Strupat 2021](#)). The macro literature estimates “pass-through coefficients” between consumption and income shocks using similar approaches ([Blundell, Pistaferri and Preston 2008](#); [Kaplan and Violante 2010](#); [Blundell et al. 2024](#)).

Table 3 – Consumption and Health Index

	(1)
Health index	0.108** (0.043)
Individual fixed effects	Yes
Wave effects	Yes
R^2	0.749
N	12,742
NT	411,471

Notes: Table reports regression results of estimating Equation 8 via OLS. The health index is calculated as the predicted probability of being in good or better 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. Regressions are weighted using survey sample weights. Standard errors clustered at the household level. N denotes the number of unique individuals and NT denotes the number of individual-waves. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

about 1.6 percent. 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 fourth-order polynomial and shows that the relationship is linear over much of the range of the health index, which supports the specification in Equation 8.¹⁵ The monthly frequency of the SLP also allows us to examine the timing of the consumption response. Appendix Figure B.7 presents an event study of consumption relative to the diagnosis of any chronic condition. There is no evidence of anticipatory adjustment or differential pre-trends, and consumption declines only after diagnosis.


5.2 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

¹⁵We present results that replace non-durable consumption with other outcomes in Appendix B since they are not used to estimate state dependence. Breaking non-durables into components, we find the strongest evidence for reductions in food, beverages, and dining out, as well as large but noisy declines in vacations, entertainment, and sports (Appendix Figure B.5). Both durable consumption and total consumption also increase with the health index (Appendix Figure B.6), but the magnitudes are smaller and not statistically different from zero. Finally, Appendix Table B.8 shows that declines in the health index increase health spending and the probability of retirement. The change in household income is smaller than the consumption change and is less precisely estimated.

Figure 2 – Binned Scatterplot: Log Consumption vs. Health Index



Notes: Figure shows a binned scatterplot of log non-durable consumption per capita against the health index, controlling for individual fixed effects, month effects, and indicators for household size, using 20 quantile-spaced bins. The line plots a 4th-order global polynomial and the regression is weighted using the survey’s sample weights.

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 is 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 is 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 a robustness check. 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 7 to avoid implausibly high levels of risk aversion. We assume that risk aversion is fixed during the sample period.

There is substantial heterogeneity in risk aversion. Slightly more than a third of respondents reject both gambles.¹⁶ 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

¹⁶By contrast, Barsky et al. (1997) find about two-thirds reject both gambles.

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 do not 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 did not respond in that wave. The average CRRA parameter in our sample is 3.64.

Health status transitions: We assume rational expectations about health risks. As indirect support for this assumption, subjective probabilities of developing a chronic condition align closely with the empirical probabilities in our sample window ([Appendix Figure B.8](#)). We estimate the transition probabilities between health states by regressing the current health index against its one-month lag, including individual fixed effects, wave effects, and age effects:

$$H_{it} = \kappa_1 + \kappa_2 H_{it-1} + \kappa_3 age_{it} + \alpha_i + \tau_t + e_{it} \quad (9)$$

The conditional probability of being healthy next month if healthy this month is $\kappa_1 + \kappa_2$, which we estimate to be 0.966.

5.3 Externally calibrated parameters

The SLP lacks questions that would allow us to measure time preferences or 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 2 percent annual interest rate based on the yield on 10-year Singapore government bonds during the sample period. This translates into a monthly gross interest rate of $R = 1.0016 = 1.02^{1/12}$. Finally, we calibrate age- and sex-specific survival probabilities using life tables published by the Singapore government. We assess the robustness of our estimates to lower and higher values of each of these parameters in [Section 6](#).

5.4 Results

We now have all the inputs to calculate state dependence using [Equation 7](#) from [Section 4](#). [Table 4](#) collects each of the inputs and presents our main estimates of state dependence. We estimate $\frac{u'(c)}{v'(c)} = 1.38$, indicating that the marginal utility of consumption is 38% higher when the probability of being healthy is 1 compared to when it is zero. We calculate a 95% confidence interval of [1.07, 1.58] and a 90% confidence interval of [1.11, 1.56] 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 each component of [Equation 7](#) and calculate $\frac{u'(c)}{v'(c)}$ 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.

Table 4 – State Dependence Estimates

	Estimate	Notes
$\frac{\Delta c_t}{c_{b,t}}$, consumption Δ between health states	0.108	Table 3
γ , coefficient of relative risk aversion	3.64	Section 5.2
p , probability healthy next month if healthy	0.966	Section 5.2
S_h , survival probability if healthy (monthly)	0.999	Externally calibrated
β , discount factor (monthly)	0.998	Externally calibrated
R , gross interest rate (monthly)	1.002	Externally calibrated
$\frac{u'(c)}{v'(c)}$, State dependence	1.38	
90% CI	[1.11, 1.56]	
95% CI	[1.07, 1.58]	

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

As before, it is more instructive to consider a smaller 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, French and Jones \(2010\)](#). Then $\sigma = \frac{u'}{v'} - 1 = 0.38$ and a one-standard deviation decrease in the health index starting from the sample mean corresponds to a roughly 5% decline in the marginal utility of consumption.¹⁷ This magnitude

¹⁷This decline is calculated as $\frac{1+0.38 \times (\bar{h}-0.15)}{1+0.38 \times \bar{h}} = 0.954$, where $\bar{h} = 0.62$ is the sample mean of the health index and 0.15 is the standard deviation.

is smaller than those estimated by [Finkelstein, Luttmer and Notowidigdo \(2013\)](#) and [Koijen, Philipson and Uhlig \(2016a\)](#), possibly due to the monthly frequency of our data, a different empirical approach, or the Singapore setting. We view the most likely explanation to be different methodologies because, as described in [Section 7.1](#), we obtain a larger magnitude than [Finkelstein, Luttmer and Notowidigdo \(2013\)](#) when replicating their approach using our data.

To explore heterogeneity, [Figure 3](#) plots estimates of $\frac{u'(c)}{v'(c)}$ for a variety of observables, including age, education, gender, marital status, and other characteristics. Estimates are consistently above 1 and statistically significant for most groups. Estimates are slightly larger for those with higher socioeconomic status, measured by higher education and living in non-HDB housing. Collectively, we interpret the limited heterogeneity shown in these results to suggest that the complementarity between health and consumption is likely a general pattern.


As discussed earlier, a crucial 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. If incomplete insurance explained our estimates, we would expect that those with smaller balances would reduce consumption more than those with higher balances. The last three rows split the sample by different asset levels in the baseline wave. The magnitudes are similar between those with MSA assets above and below the median level.¹⁸ The average MSA balance among those in the bottom half exceeds 12,000 SGD, which is an order of magnitude larger than annual health care spending in the full sample. The final row shows that the estimate of state dependence decreases slightly if we exclude the 7.5% of households with zero MSA assets in the baseline wave. This analysis reinforces our interpretation that the change in consumption in [Table 3](#) likely reflects preferences rather than insufficient resources.

6 Robustness

In this section, we evaluate the robustness of our estimates in two ways. First, we vary the empirical specification, examining the sensitivity of our results to alternative ways of measuring health, risk preferences, and consumption, as well as to the externally calibrated parameters. Our qualitative conclusions are unchanged: across these specifications, the marginal utility of consumption remains higher in better health. Second, we assess the assumptions underlying the sufficient-statistics formula. The formula in [Section 4](#) requires

¹⁸[Appendix Table B.9](#) presents the OLS regressions of consumption changes for these groups.

Figure 3 – Heterogeneity in State Dependence



Notes: Figure plots estimates of state dependence for different subsamples. The last three rows divide the sample based on MSA balances in the baseline wave. Whiskers denote 90% confidence intervals calculated by block bootstrapping 400 samples for each group.


that agents optimize consumption according to the Euler equation. We consider the leading reasons this condition could fail, including binding budget constraints, present bias, and rules of thumb, as well as margins omitted from the model, such as spousal labor supply, bequests, and consumption that improves health.

6.1 Alternative specifications and calibrations

We continue to find strong evidence that health and consumption are complements under alternative modeling assumptions, as shown in Figure 4. We estimate values of $\frac{u'(c)}{v'(c)}$ of 1.22 and 1.48 using the lower bounds and the upper bounds of risk aversion, respectively. Using non-imputed measures of risk aversion yields similar magnitudes but noisier estimates due to smaller sample sizes. We also 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. Finally, the results are robust to using 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 have 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, rather than because of state dependence. To assess this possibility, we exclude observations in which the individual is

Figure 4 – Robustness: Alternative Empirical Specifications and Sub-Samples



Notes: Figure plots estimates of state dependence for alternative measurements of risk, consumption, or health or for different subsamples. Whiskers denote 90% confidence intervals calculated by block bootstrapping 100 samples for each group.

hospitalized in the current or previous wave.¹⁹ The results are slightly smaller and still statistically distinguishable from zero. Next, we restrict the sample to individuals who report no chronic conditions at retirement to examine a subsample that does not adjust their own labor supply in response to this set of illnesses. We again obtain positive estimates of state dependence. We also obtain similar results when restricting the sample to those who respond in wave 50 or later, verifying that attrition—due to mortality or other reasons—does not affect the results.

We next probe robustness to different values of the three externally calibrated parameters by considering lower and higher values of the discount factor, the interest rate, and survival probabilities. [Appendix Table B.10](#) reports estimates of $\frac{u'(c)}{v'(c)}$ for 27 combinations of low, baseline, and high values of these parameters, with annual discount factors between 0.95 and 0.99, interest rates between 0 and 4 percent, and annual survival probabilities between 0.976 and 0.994. The estimated state dependence ranges from 1.20 to 1.50. As a complementary exercise, we calculate how far each parameter would need to move, holding the other two at their baseline values, for state dependence to disappear ($\frac{u'(c)}{v'(c)} = 1$). The required values lie outside any plausible calibration: an annual discount factor of 0.83, a gross interest rate of 0.87 (a real return of -13 percent per year), or an annual survival probability of 0.85—an order of magnitude higher mortality risk than in the Singapore life tables. We conclude that our finding of positive state dependence is not driven by the externally calibrated inputs.

¹⁹We define a hospitalization as reporting positive inpatient spending.

6.2 Robustness of sufficient statistics formula

Spousal labor supply: Spouses may alter labor supply in response to a partner’s health shock. Our formula is robust to such responses under the assumption of household efficiency and if consumption is well-measured. Intuitively, any change in spousal labor supply should be captured by the change in consumption. Labor supply is just another decision that is being optimized over, and so there is no new information in that response that does not already appear in consumption. Previous research has shown how researchers can use changes in spousal labor supply to recover the gap in marginal utilities for various shocks, rather than using changes in consumption (Fadlon and Nielsen 2019). Empirically, Fadlon and Nielsen (2021) find spouses adjust their labor supply after a fatal illness but not a non-fatal one using Danish administrative data. Our results are therefore robust to adjustments in spousal labor supply.

Bequests: Our theory in Section 4 omits bequests, treating the utility of leaving assets to heirs upon death as zero. Prior research has considered a number of different ways of modeling bequests (see e.g., Lockwood 2012). In Appendix D, we show how modifying the Euler equation to incorporate incidental bequests based on the specification and results in Lockwood (2018) and Liu, Song and Zhu (2024) leads state dependence to be slightly higher than our main estimates. While omitting incidental bequests understates the magnitude of state dependence, the same is not true if people have strategic bequest motives as in Bernheim, Shleifer and Summers (1985). If parents receive (or expect to receive) informal care from children in exchange for bequests, and a deterioration in health increases the likelihood that care is needed, then part of the consumption drop may reflect the desire to preserve assets to exchange for care.²⁰ Using incentive-compatible survey questions among older adults in Singapore, Liu, Song and Zhu (2024) find evidence against strategic bequests and in support of altruism. Consequently, our approach that abstracts from bequests likely yields a lower bound on the magnitude of state dependence.

Budget Constraints: As discussed earlier, one concern is that people may reduce consumption due to insufficient assets, rather than because of state-dependent preferences. If people would prefer to consume more when ill but are unable to do so because of budget constraints, then the Euler equation will not hold and our formula is not valid. Such constraints may indeed matter in many contexts—including in the United States (Blundell et al. 2024)—but they are likely less important in Singapore, which is among the world’s

²⁰We tried restricting the analysis to households without children, but doing so cuts 86% of the sample and yields extremely noisy estimates.

richest countries and has a well-functioning safety net. As direct evidence, we showed earlier in [Figure 3](#) that state dependence estimates are similar for those with different levels of MSA assets and robust to excluding those with zero MSA balances. A related concern is that illness reduces household income, so that consumption falls because resources decline rather than because of preferences. However, household income does not change significantly with the health index ([Appendix Table B.8](#)), indicating that the consumption response is not driven by contemporaneous income losses. The subsample of individuals with no chronic conditions at retirement provides further evidence. Because these individuals develop conditions only after leaving the labor force, illness cannot reduce their earnings, and the positive estimates of state dependence in this group cannot reflect income losses. These patterns suggest that our results are unlikely to be explained by budget constraints that prevent the agent from smoothing consumption across states.

Present bias: The Euler equation would differ from [Equation 3](#) if, instead of exponential discounting, agents have the quasi-hyperbolic discounting function: $1, \tilde{\beta}\tilde{\delta}, \tilde{\beta}\tilde{\delta}^2, \dots, \tilde{\beta}\tilde{\delta}^t$ ([Harris and Laibson 2001](#)). If the short-run discount factor $\tilde{\beta}$ is below 1, preferences are inconsistent over time as the agent over-consumes today relative to the preferences of their future self. In addition to the level of $\tilde{\beta}$, the degree to which such preferences pull consumption earlier in time also depends on whether the agent recognizes their self-control problem and their amount of liquid wealth.

We empirically assess this issue by using survey responses to a series of one-off questions that ask about procrastination and savings regret. In wave 34, the SLP asks the following questions and we interpret the responses in bold as those consistent with present bias:

Q: How often do you delay things you aren't really interested in?

Response options: Never; Sometimes; **Most of the time**; Always

Q: How often do you put things off you are not good at?

Response options: Never; Sometimes; **Most of the time**; Always

Q: Please think back to when you were around 45 years old. Suppose you could redo your spending and saving from then to now, which of the following would you do?

Response options: **Spend less and save more over the years**; Spend and save about the same over the years; Spend more and save less over the years

For each question, we exclude all observations for any respondent who provides one of the answers in bold. This restriction leaves us with sub-samples of respondents less likely

to be influenced by present bias. The estimated magnitudes of state dependence range between 1.29 and 1.38 in these sub-samples ([Appendix Figure B.9](#)). Since we find similar estimates for state dependence in [Table 4](#), we view this as suggestive evidence that present bias does not explain our results. In addition, the monthly frequency of our data is helpful in distinguishing state dependence from present bias: we leverage sharp changes between consumption and health to estimate state dependence, whereas present bias alone would predict a consumption profile that declines more smoothly over time.

Rules of Thumb: [Equation 3](#) would not be valid if people follow rules of thumb rather than solving the dynamic optimization problem. For example, some common rules might be: consume the same fraction of income each period; consume permanent income; or consume cash-on-hand if it is less than expected income as in [Deaton \(1992\)](#). In studying the utility losses from such rules compared to a benchmark of life-cycle optimization, [Winter, Schlafmann and Rodepeter \(2012\)](#) show that the consumption profiles can sometimes differ substantially, especially around the ages in our sample that overlap with retirement.

Consumption as Investment in Health Production: Our model assumes non-medical consumption does not affect health, and our regressions exclude medical and pharmacy spending for this reason. It is instructive to assess how feedback from consumption to health would affect our analysis, as in canonical models of health production ([Grossman 1972](#)). The direction of the resulting bias depends on how the investment operates. Investment may take the form of purchasing healthier but more expensive goods. For example, someone managing a chronic condition has a stronger incentive to switch to organic foods, which would raise spending in bad health. Under this channel, the decline in consumption we observe understates the drop in marginal utility, and our estimates are a lower bound on state dependence. However, investment may instead take the form of abstention; a newly diagnosed patient may be advised to cut back on dining out, alcohol, and unhealthy food. If these items are relatively expensive, following such advice reduces measured spending in bad health for reasons unrelated to preferences. Under this channel, our estimates would overstate state dependence. The composition of the consumption response helps assess the importance of abstention. If it were the primary driver, declines should be concentrated in the categories subject to medical advice. While food, beverages, and dining out do decline, we also estimate large—albeit noisier—declines in vacations, entertainment, and sports ([Appendix Figure B.5](#)), categories with no plausible connection to dietary guidance, suggesting abstention is not the primary driver of the decline.

7 Supplementary Results using Subjective Well-Being Data

The preceding robustness checks maintained the assumption that agents optimize consumption. This section presents an additional estimate of state dependence between health and consumption using an approach that does not require this assumption. We estimate state dependence following the framework of [Finkelstein, Luttmer and Notowidigdo \(2013\)](#) using data on life satisfaction as a proxy for utility, and again find evidence that health and consumption are complements. We focus less on the implied magnitude of state dependence from this analysis compared to that in [Section 5](#) because the latter is more tightly linked to theory. We instead concentrate on the qualitative result from this supplementary approach and view it as providing corroborating evidence on the sign of state dependence.

7.1 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 that has previously been used by economists to measure well-being ([Frey and Stutzer 2002](#); [Oswald and Wu 2010, 2011](#); [Benjamin et al. 2014](#)). In describing the ideal empirical approach using data on utility proxies, [Finkelstein, Luttmer and Notowidigdo \(2013\)](#) 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 creation of the SLP allows us to implement this specification. We estimate the following regression:

$$LifeSatisfaction_{it} = \varphi_1 H_{it} \times \ln(C_{it}) + \varphi_2 \ln(C_{it}) + \varphi_3 H_{it} + \alpha_i + \tau_t + \lambda X_{it} + e_{it} \quad (10)$$

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%). Since 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.²¹ X_{it} includes a quartic polynomial in age and indicators for household size, which may vary over time. 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. Finally, we de-mean all variables to interpret coefficients relative to sample averages rather than relative to zero consumption or a zero probability of being healthy.

The key coefficient is φ_1 . If $\varphi_1 > 0$, the marginal utility of consumption (as proxied by life satisfaction) 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.

Table 5 presents the results of estimating Equation 10, 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 is a positive and statistically significant relationship between consumption and life satisfaction, as shown by the estimates of φ_2 . There is 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.

An important concern with subjective measures of well-being is that they do not equal utility, and state dependence estimates will be biased if the function mapping utility to these measures is misspecified. Finkelstein, Luttmer and Notowidigdo (2013) address this issue by performing a nonparametric test of state dependence, which we replicate in Appendix Table B.11. To implement the test, we first construct bins of mean levels of life satisfaction and mean levels of the health index, with the means calculated within each person over all waves. We then run OLS regressions of Equation 10 within each bin and collect the t -statistics on the interaction term. Under the null of state-independent utility, the CDF of each t -statistic is drawn from a uniform distribution between 0 and 1. A mean CDF significantly different from 0.5 rejects the null in favor of state-dependent utility, with an average above 0.5 indicating that marginal utility is higher when healthy. Under a range

²¹We 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.

Table 5 – Life Satisfaction Regressions

	Ordered Logit (1)	LPM (2)
φ_1 , Health index $\times$ Log consumption	0.334*** (0.102)	0.025** (0.012)
φ_2 , Log consumption	0.079*** (0.018)	0.013*** (0.002)
φ_3 , Health index	1.236*** (0.242)	0.142*** (0.028)
R^2		0.650
N	8,855	10,957
NT	552,460	411,094

Notes: Table presents results from estimating [Equation 10](#). Log consumption is the natural logarithm of non-durable 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 household level and regressions are weighted using the survey’s sample weights. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

of different choices for bin sizes, we consistently find statistically significant evidence of this sign of state dependence.

As additional robustness, we also consider an alternative utility proxy and measure of health status. [Appendix Table B.12](#) replicates [Table 5](#) 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. We then calculate a weighted satisfaction measure by assigning weights to each of these aspects, drawing on the results in [Benjamin et al. \(2014\)](#). The estimate on the interaction term is again positive.

Finally, we measure health by the count of chronic conditions, instead of the health index, and again find similar patterns ([Appendix Table B.13](#)). Since this specification is closest to [Finkelstein, Luttmer and Notowidigdo \(2013\)](#), we compare our estimated magnitudes to theirs as a way to investigate differences stemming from our data and setting. Our OLS regression estimates the ratio of the interaction term to log consumption as 0.38, which is twice the size of their estimate. This comparison suggests that our data and setting do not account for the smaller magnitudes we obtained in [Section 5](#), and instead points to different methods as the most likely explanation.

8 Quantitative Implications

State dependence has implications for a range of economic decisions and policy design. In this section, we show how our estimated magnitudes in [Section 5](#) meaningfully change conclusions about the level of health insurance and retirement saving. We then show how ignoring state dependence can lead to incorrect estimation of risk preferences when matching observed choices to those predicted from a standard model that assumes state-independent utility.

8.1 Optimal health insurance coverage

We first illustrate the implications of state dependence for health insurance using a simple model of how much insurance to buy. Compared to the textbook result that risk-averse individuals will choose to fully insure the loss if the premium is actuarially fair, state dependence induces them to partially insure.

Model: A risk-averse individual faces a prospective financial loss L with probability π . In this state, their health index is h_l . With probability $1 - \pi$, no loss occurs and their health index is $h_g \geq h_l$. Agents have a utility function that depends on their health index h according to the parameter σ : $U(c, h) = (1 + \sigma h) \frac{c^{1-\gamma}}{1-\gamma}$. They have wealth W and choose the amount k of insurance to buy at the actuarially fair price π per dollar of coverage to maximize their expected utility:

$$\max_k E[U] = (1 - \pi) \left[(1 + \sigma h_g) \frac{(W - \pi k)^{1-\gamma}}{1 - \gamma} \right] + \pi \left[(1 + \sigma h_l) \frac{(W - L + (1 - \pi)k)^{1-\gamma}}{1 - \gamma} \right] \quad (11)$$

The first-order condition for this problem can be rewritten after some algebra as:

$$L - k^* = \frac{(\psi - 1)(W - \pi L)}{\pi + \psi(1 - \pi)} \quad (12)$$

where $\psi \equiv \left(\frac{1 + \sigma h_g}{1 + \sigma h_l} \right)^{1/\gamma}$ captures the ratio of marginal utilities between states.²²

The left-hand side of [Equation 12](#) is the amount by which the loss is partially insured, measured in dollars. The right-hand side comprises three components. First, $\psi - 1$ determines the sign of the equation and incorporates the strength and direction of state dependence. The special case of full insurance ($k^* = L$) occurs if there is state-independent

²²Appendix E presents the derivation of the more general case that includes an insurance load, of which this is a special case.


utility ($\sigma = 0$) or if there is no change in health status between states ($h_g = h_l$), in which case $\psi = 1$. $\psi > 1$ if health and consumption are complements ($\sigma > 0$) and $h_g > h_l$. The term is larger for higher values of σ (stronger complementarity between health and consumption), and the greater the difference $h_g - h_l$ (the change in health between states). The second term in the numerator $W - \pi L$ represents wealth net of the expected loss. Someone with more wealth would choose to buy less insurance coverage than a person with lower wealth, which is a standard theoretical prediction. Finally, the term $\pi + \psi(1 - \pi)$ is a scaling factor that averages the relative value of a dollar in each state, weighted by their probabilities.

Illustrative Calculations: To assess magnitudes, we calculate how the optimal levels of insurance coverage vary with state dependence based on [Equation 12](#). These calculations assume $W = \$100,000$, $L = \$10,000$, $\pi = 0.1$, and $\gamma = 3.64$. We perform calculations assuming that in addition to the financial loss, the health index declines by 1 standard deviation (from 0.75 to 0.6) or two standard deviations (from 0.75 to 0.45). These magnitudes were chosen to represent moderate reductions in health and a meaningful financial loss relative to wealth. Our baseline model sets $\sigma = 0.38$ to be consistent with [Table 4](#), since $\sigma = \frac{u'(c)}{v'(c)} - 1$. [Figure 5](#) presents the results over a range of state dependence from 0 to 0.6 for comparison. The figure illustrates how a stronger complementarity between consumption and health moves the optimal choice further from full insurance. For $\sigma = 0.38$ and a 1 standard deviation decline in health, the optimal insurance k is \$8,772, equivalent to 12.3% less than full insurance. For twice as large a decline in health, the optimal level of insurance is \$7,499, equivalent to 25.0% less than full coverage. We interpret this level of partial insurance to be substantial, but not implausibly large.

8.2 Optimal retirement saving

We next illustrate the implications for retirement saving by solving 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 4](#) 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. The results should therefore be viewed as illustrative of how state dependence can potentially influence optimal saving decisions.

Figure 5 – Optimal Health Insurance: Partial Coverage



Notes: Figure plots the calculations for the optimal level of insurance coverage k as a function of the state dependence parameter σ , based on Equation 12. The calculations assume a wealth equal to \$100,000, CRRA coefficient = 3.64, financial loss = \$10,000, health in the no-loss state = 0.75, health in the loss state = 0.6 (1 SD decrease) or = 0.45 (2 SD decrease). The dashed vertical line denotes the level of state dependence estimated in Table 4.

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

$$U(c, h) = (1 + \sigma h) \frac{c^{1-\gamma}}{1-\gamma} \quad (13)$$

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 for saving, we assume health depreciates deterministically over time. This assumption rules out precautionary motives to save due to the variance of health risks and instead focuses solely on changes in the level of health.

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 $R = 1 + r$, and cannot die in debt. They are endowed with initial assets A_0 , which evolve according to:

$$A_{t+1} = R(A_t + y_t - c_t) \quad (14)$$

We represent the dynamic programming problem recursively. Denoting β as the

discount factor, the decision problem is given by:

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

subject to (14) and $A_T \geq 0$.

Our baseline model sets $\sigma = 0.38$ and 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 4 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 $A_0 = 0$.


Results—benchmark comparisons and improvements in health over time:

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.

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 and Ruhm 2021; Bavafa, Mukherjee and Welch 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 6 plots the predicted share of people reporting being healthy by five-year age bin in the National Health Interview Survey, 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. In contrast, 79% of 70-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.²³

Figure 6 – Trends in Self-Reported Health by NHIS Cohort



Notes: Figure shows the percentage of respondents who report being in good, very good, or excellent health by 5-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.

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 will 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 specific to this particular model.

8.3 Implications for estimating risk preferences

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

²³These 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.

hypothetical gambles over permanent income. Without such information, it is common for researchers to estimate risk aversion by matching observed choice data to predicted behavior from a structural model. Below we show how ignoring state dependence will meaningfully underestimate risk aversion, which could have implications for the evaluation of counterfactual policies. For example, a lower estimate of risk aversion would understate the value of annuities and insurance.

In the case of health insurance, we showed above that state dependence induces partial insurance even if premiums are actuarially fair. In a standard model with state-independent utility and no liquidity constraints, risk-loving preferences would be needed to rationalize this result. This is an extreme example in which state-independent utility gets the sign of a key parameter wrong. Now consider the more realistic case in which there is a $z\%$ load above expected cost, so the premium is no longer actuarially fair. The corresponding formula for the optimal amount of coverage with state dependence (derived in Appendix E) is:

$$L - k^* = \frac{(W - L\alpha)(\psi\Lambda - 1)}{\alpha + \psi\Lambda(1 - \alpha)} \quad (15)$$

where $\alpha = \pi(1 + z)$ is the premium including the load and $\Lambda \equiv [\frac{\alpha(1-\pi)}{\pi(1-\alpha)}]^{1/\gamma}$. Assuming a wealth of \$100,000, CRRA coefficient $\gamma = 3.64$, a 10% chance of a \$10,000 loss that reduces the health index by one standard deviation (as in [Section 8.1](#)), and a 5% insurance load, the optimal coverage is \$7,318 if $\sigma = 0.38$. To explain this choice of coverage under a model with state-independent utility, we set $\sigma = 0$ and solve for γ . Doing so yields a risk aversion coefficient of 1.98. This level of risk aversion is substantially below 3.64 in the state-dependent model that generates the same choice of coverage k .

In the case of retirement saving, ignoring state dependence will also underestimate risk aversion because the optimal consumption profile is less smooth under state-dependent utility as consumption is shifted earlier in life. We numerically calculate the level of risk aversion assuming a model of state-independent utility that would best rationalize the optimal consumption profile generated from the model of state-dependent utility. We perform this calculation for the benchmark case from [Section 8.2](#) in which state dependence reduces optimal wealth at retirement by 1.4% compared to a model with state-independent utility. To match the observed consumption profile, we calculate a CRRA parameter of 3.35 when assuming state-independent utility. As in the case of health insurance, an econometrician who assumes state-independent utility and estimates preference parameters from this consumption profile will incorrectly infer the reason for these choices as lower risk aversion, rather than a higher marginal utility of consumption when healthy.

9 Discussion

In this paper, we have provided novel evidence on how the marginal utility of consumption varies by health status. We develop a formula for state dependence between health and consumption 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 decline in a person's health index by one standard deviation reduces their marginal utility of consumption by 5%. Our estimates are robust to a number of alternative specifications. We also derive a within-state first-order condition formula that does not require inter-temporal assumptions and yields results slightly larger in magnitude. We also present corroborating evidence from two approaches that do not require agent optimization: reduced-form estimates showing that the marginal propensity to consume out of lottery winnings is larger in better health, and a replication of the life-satisfaction approach of [Finkelstein, Luttmer and Notowidigdo \(2013\)](#). Both reinforce our qualitative conclusion that health and consumption are complements.

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 small—slightly larger than that of Denmark—but it is reasonably diverse along several dimensions, including race, religion, language spoken, and income. In our setting, we find relatively little evidence of heterogeneity across observable characteristics of individuals. In terms of measuring health, we are 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 health conditions requires assuming that the way changes in self-reported health map to changes in consumption would be similar for such conditions.

Our estimated magnitude of state dependence has important implications for insurance and saving decisions. Even if health insurance premiums are actuarially fair, state dependence induces partial insurance. In the case of life-cycle saving, state dependence shifts the optimal consumption profile earlier in life when people are healthier, reducing the amount of wealth at retirement. Using these two examples, we show how ignoring state dependence can also affect the estimation of other preference parameters if these are not measured separately. Matching observed choices to predicted choices from a model that assumes state-independent utility will meaningfully underestimate risk aversion. A lower value of risk aversion would then understate the value of insurance and

influence the design of counterfactual policies. Future research could explore these issues in more detail along with other implications of the complementarity between health and consumption.

References

- Ameriks, John, Joseph Briggs, Andrew Caplin, Matthew Shapiro, and Christopher Tonetti, “Long-Term-Care Utility and Late-in-Life Saving,” *Journal of Political Economy*, 2020, 128 (6), 2375–2451.
- Amo-Agyei, Silas, “The Effect of Health on the Marginal Utility of Consumption: Evidence Based on Household Consumption Expenditure,” *Journal of Economic Behavior & Organization*, 2025, 229, 106809.
- Armstrong, Grayson, Caleb Cho, Thesia Garner, Brett Matsumoto, Juan Munoz, and Jake Schild, “Building a Consumption Poverty Measure: Initial Results Following Recommendations of a Federal Interagency Working Group,” *AEA Papers and Proceedings*, 2022, 112, 335–339.
- Arrow, Kenneth, “Optimal Insurance and Generalized Deductibles,” *Scandinavian Actuarial Journal*, 1974, 1, 1–42.
- Baetschmann, Gregori, Alexander Ballantyne, Kevin Staub, and Rainer Winkelmann, “feologit: A new command for fitting fixed-effects ordered logit models,” *The Stata Journal*, 2020, 20 (2), 253–275.
- , Kevin Staub, and Rainer Winkelmann, “Consistent estimation of the fixed effects ordered logit models,” *Journal of the Royal Statistical Society, Series A*, 2015, 178, 685–703.
- Baily, Martin, “Some Aspects of Optimal Unemployment Insurance,” *Journal of Public Economics*, 1978, 10 (3), 379–402.
- Barsky, Robert, F. Thomas Juster, Miles Kimball, and Matthew Shapiro, “Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study,” *Quarterly Journal of Economics*, 1997, 111 (2), 537–579.
- Bassoli, Elena, “An Empirical Analysis of Health-Dependent Utility on SHARE and ELSA Data,” *Italian Economic Journal*, 2022.
- Bavafa, Hessam, Anita Mukherjee, and Tyler Welch, “Inequality in the golden years: Wealth gradients in disability-free and work-free longevity in the United States,” *Journal of Health Economics*, 2023, 92, 102820.
- Benjamin, Daniel, Miles Kimball, Ori Heffetz, and Nichole Szembrot, “Beyond Happiness and Satisfaction: Toward Well-Being Indices Based on Stated Preference,” *American Economic Review*, 2014, 104 (9), 2698–2735.
- Bernheim, B. Douglas, Andrei Shleifer, and Lawrence H. Summers, “The Strategic Bequest Motive,” *Journal of Political Economy*, 1985, 93 (6), 1045–1076.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston, “Consumption Inequality and Partial Insurance,” *American Economic Review*, 2008, 98 (5), 1887–1921.
- , Margherita Borella, Jeanne Commault, and Mariacristina De Nardi, “Old Age Risks, Consumption, and Insurance,” *American Economic Review*, 2024, 114 (2), 575–613.
- Brown, Jeffrey, Gopi Shah Goda, and Kathleen McGarry, “Heterogeneity in State-Dependent Utility: Evidence from Strategic Surveys,” *Economic Inquiry*, 2016, 54, 847–861.
- Cattaneo, Mattias, Richard Crump, Max Farrell, and Yingjie Feng, “On Binscatter,” *American Economic Review*, 2024, 114 (5), 1488–1514.
- Cesarini, David, Erik Lindqvist, Robert Östling, and Björn Wallace, “Wealth, health, and child development: Evidence from administrative data on Swedish lottery players,” *Quarterly Journal of Economics*, 2016, 131 (2), 687–738.
- Chetty, Raj, “A general formula for the optimal level of social insurance,” *Journal of Public Economics*, 2006, 90, 1879–1901.
- Cook, Philip and Daniel Graham, “The Demand for Insurance and Protection: The Case of

- Irreplaceable Commodities,” *Quarterly Journal of Economics*, 1977, 91, 143–156.
- Coyne, David, Itzik Fadlon, Shanthi Ramnath, and Patricia Tong, “Household Labor Supply and the Value of Social Security Survivors Benefits,” *American Economic Review*, 2024, 114 (5), 1248–1280.
- Danesh, Kaveh, Jonathan Kolstad, William Parker, and Johannes Spinnewijn, “The Chronic Disease Index: Analyzing Health Inequalities Over the Lifecycle,” *NBER Working Paper* 32577, 2024.
- De Nardi, Mariacristina, Eric French, and John B. Jones, “Why do the Elderly Save? The Role of Medical Expenses,” *Journal of Political Economy*, 2010, 118, 39–75.
- , Svetlana Pashchenko, and Ponpoje Porapakkarm, “The Lifetime Costs of Bad Health,” *Review of Economic Studies*, 2025, 92, 1987–2026.
- Deaton, Angus, “Household Saving in LDCs: Credit Markets, Insurance, and Welfare,” *Scandinavian Journal of Economics*, 1992, 94 (2), 253–273.
- Deshpande, Manasi and Lee M. Lockwood, “Beyond Health: Nonhealth Risk and the Value of Disability Insurance,” *Econometrica*, 2022, 90 (4), 1781–1810.
- Fadlon, Itzik and Torben Nielsen, “Household Labor Supply and the Gains from Social Insurance,” *Journal of Public Economics*, 2019, 171, 18–28.
- and —, “Family Labor Supply Responses to Severe Health Shocks: Evidence from Danish Administrative Records,” *American Economic Journal: Applied Economics*, 2021, 13, 1–30.
- Finkelstein, Amy, Erzo Luttmer, and Matthew Notowidigdo, “Approaches to Estimating the Health State Dependence of the Utility Function,” *American Economic Review: Papers and Proceedings*, 2009, 99 (2), 116–121.
- , —, and —, “What Good is Wealth without Health? The Effect of Health on the Marginal Utility of Consumption,” *Journal of the European Economic Association*, 2013, 11 (s1), 221–258.
- , Nathaniel Hendren, and P. Luttmer Erzo F.}, “The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment,” *Journal of Political Economy*, 2019, 127 (6), 2836–2874.
- Fonseca, Raquel, François Langot, Pierre-Carl Michaud, and Thepthida Sopraseuth, “Understanding Cross-Country Differences in Health Status and Expenditures: Health Prices Matter,” *Journal of Political Economy*, 2023, 131 (8), 1949—1993.
- Frey, Bruno and Alois Stutzer, “What Can Economists Learn from Happiness Research?,” *Journal of Economic Literature*, 2002, 40 (2), 402–435.
- Garcia-Mandico, Silvia, Arndt Reichert, and Christoph Strupat, “The Social Value of Health Insurance: Results from Ghana,” *Journal of Public Economics*, 2021, 194, 104313.
- Gertler, Paul and Jonathan Gruber, “Insuring Consumption Against Illness,” *American Economic Review*, 2002, 92 (1), 51–70.
- Grossman, Michael, “On the Concept of Health Capital and the Demand for Health,” *Journal of Political Economy*, 1972, 80 (2), 223–255.
- Gruber, Jonathan, “The Consumption Smoothing Benefits of Unemployment Insurance,” *American Economic Review*, 1997, 87 (1), 192–205.
- Hall, Robert E. and Charles I. Jones, “The Value of Life and the Rise in Health Spending,” *Quarterly Journal of Economics*, 2007, 122 (1), 39–72.
- Harris, Christopher and David Laibson, “Dynamic Choices of Hyperbolic Consumers,” *Econometrica*, 2001, 69, 935–957.
- Haseltine, William, “Affordable Excellent: The Singapore Healthcare Story,” *Brookings Institution Press*, 2013.

- Hendren, Nathaniel, “Knowledge of Future Job Loss and Implications for Unemployment Insurance,” *American Economic Review*, 2017, 107 (7), 1778–1823.
- , “Measuring Ex Ante Welfare in Insurance Markets,” *Review of Economic Studies*, 2021, pp. 1193–1223.
- Hosseini, Roozbeh, Karen Kopecky, and Zhao Kai, “The evolution of health over the life cycle,” *Review of Economic Dynamics*, 2022, 45, 237–263.
- Hugonnier, Julien, Florian Pelgrin, and Pascal St-Amour, “Health and (Other) Asset Holdings,” *Review of Economic Studies*, 2013, 80 (2), 663–710.
- Kaplan, Greg and Giovanni L. Violante, “How Much Consumption Insurance Beyond Self-Insurance?,” *American Economic Journal: Macroeconomics*, 2010, 2, 53–87.
- Kim, Seonghoon and Andrew Oswald, “Happy Lottery Winners and LT Bias,” *Review of Income and Wealth*, 2021, 67 (2), 317–333.
- and Kanghyock Koh, “The effects of income on health: Evidence from lottery wins in Singapore,” *Journal of Health Economics*, 2021, 76, 102414.
- Koijen, Ralph S. J., Tomas J. Philipson, and Harald Uhlig, “Financial Health Economics,” *Econometrica*, 2016, 84 (1), 195–242.
- Koijen, Ralph, Stijn Van Nieuwerburgh, and Motohiro Yogo, “Health and Mortality Delta: Assessing the Welfare Cost of Household Insurance Choice,” *Journal of Finance*, 2016, 71 (2), 957–1009.
- Leive, Adam and Christopher Ruhm, “Has Mortality Risen Disproportionately for the Least Educated?,” *Journal of Health Economics*, 2021, 79, 102494.
- Lieber, Ethan and Lee Lockwood, “Targeting with In-Kind Transfers: Evidence from Medicaid Home Care,” *American Economic Review*, 2019, 109 (4), 1461–1485.
- Lillard, Lee and Yoram Weiss, “Uncertain Health and Survival: Effects on End-of-Life Consumption,” *Journal of Business Economic Statistics*, 1997, 15 (2), 254–268.
- Liu, Haoming, Changcheng Song, and Shenghao Zhu, “Intentional Bequest Motives and the Choice of Annuity,” *Working Paper*, 2024.
- Liu, Kai, “Insuring Against Health Shocks: Health Insurance and Household Choices,” *Journal of Health Economics*, 2016, 46, 16–32.
- Lockwood, Lee, “Bequest motives and the annuity puzzle,” *Review of Economic Dynamics*, 2012, 15, 226–243.
- , “Incidental Bequests and the Choice to Self-Insure Late-Life Risks,” *American Economic Review*, 2018, 108 (9), 2513–2550.
- Meyer, Bruce and James Sullivan, “Consumption and Income Inequality and the Great Recession,” *AEA Papers and Proceedings*, 2013, 110, 178–183.
- and —, “Consumption and Income Inequality in the United States since the 1960s,” *Journal of Political Economy*, 2023, 131 (62), 247–284.
- Mohanan, Manoj, “Causal Effects of Health Shocks on Consumption and Debt: Quasi-Experimental Evidence from Bus Accident Injuries,” *Review of Economics and Statistics*, 2013, 95 (2), 673–681.
- Oswald, Andrew and Stephen Wu, “Objective confirmation of subjective measures of human well-being: Evidence from the USA,” *Science*, 2010, 327 (5965), 576–579.
- and —, “Well-being across America,” *Review of Economics and Statistics*, 2011, 93 (4), 1118–1134.
- Palumbo, Michael, “Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle,” *Review of Economic Studies*, 1999, 66, 395–421.
- Picchio, Matteo, Sigrid Suetens, and Jan C. van Ours, “Labour supply effects of

- winning a lottery,” *Economic Journal*, 2018, 128 (611), 1700–1729.
- Poterba, James, Steven Venti, and David Wise, “The asset cost of poor health,” *Journal of the Economics of Ageing*, 2017, 9, 172–184.
- Reichling, Felix and Kent Smetters, “Optimal Annuitization with Stochastic Mortality and Correlated Medical Costs,” *American Economic Review*, 2015, 105 (11), 3273–3320.
- Townsend, Robert, “Risk and Insurance in Village India,” *Econometrica*, 1994, 62 (3), 539–591.
- Wang, Xiaoyu and Chunan Wang, “How Does Health Status Affect Marginal Utility of Consumption? Evidence from China,” *International Journal of Environmental Research and Public Health*, 2020, 17 (2234), 1–20.
- Winter, Joachim, Kathrin Schlafmann, and Ralf Rodepeter, “Rules of Thumb in Life-cycle Saving Decisions,” *Economic Journal*, 2012, 122 (560), 479–501.
- Yin, Jason Dean-Chen and Alex Jingwei He, “Health Insurance Reforms in Singapore and Hong Kong: How the Two Ageing Asian Tigers Respond to Health Financing Challenges?,” *Health Policy*, 2018, 122 (7), 693–697.
- Yogo, Motohiro, “Portfolio choice in retirement: Health risk and the demand for annuities, housing, and risky assets,” *Journal of Monetary Economics*, 2016, 80, 17–34.
- Zeckhauser, Richard, “Medical Insurance: A Case Study of the Tradeoff Between Risk Spreading and Appropriate Incentives,” *Journal of Economic Theory*, 1970, 2, 10–26.

Online 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 non-durables: 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 non-durable 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; (5) summary sub-screens and no grand total.

Rental equivalence of housing for home owners

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; Other housing. The survey records the number of rooms (ranging from 1 to 5) or if it is larger than 5 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 does not include information on the characteristics of the vehicle or how long ago it was purchased, so we are 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 do not 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.


Online Appendix B: Additional Results

Table B.1 – Sample Selection

	Individuals	Households
Full SLP sample	13,353	9,822
Excluding respondents observed only in waves 0 or 1	13,003	9,647
Excluding ages below 45 or above 75	12,800	9,582
Excluding inconsistent consumption summary screens	12,756	9,556
Excluding households with a member under age 22	12,742	9,547

Notes: Table reports the number of unique individuals and households remaining after sequentially applying each sample restriction. Wave 0 is the baseline wave and wave 1 was a small test wave. The third restriction excludes individuals who are presented with different summary screens in the consumption module across waves. The final row constitutes the analytic sample used throughout the paper.

Figure B.1 – Number of Diseases by Survey Wave




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, and so the mean rate sometimes declines between successive quarters due to changes in sample composition.

Table B.2 – Descriptive Statistics of Lottery Variables

	Mean	SD	Median	P90	P95	P99	NT
Lottery winnings	409	1,379	0	1,000	2,000	8,500	11,261
Lottery purchases	1,981	3,899	520	5,200	8,320	26,000	11,470
Winnings, conditional on winning	886	1,922	200	2,000	5,000	12,000	5,215

Notes: Table presents descriptive statistics of annual lottery winnings and ticket purchases, in SGD, among respondents who report playing the lottery and positive ticket spending. Statistics pool the three lottery modules (waves 16, 28, and 52) and use survey weights. Winnings and purchases are winsorized at the 99th percentile, matching the values used in the regressions in Table 2. NT is the number of person-wave observations.

Figure B.2 – Permutation Test for the Lottery Winnings $\times$ Health Interaction



Notes: Figure shows the null distribution of the interaction between lottery winnings and above-median baseline health from randomization inference. In each of 5,000 placebo samples, lottery winnings are randomly reassigned among players in the same survey wave and within-wave decile of ticket spending, which preserves the dependence of winnings on ticket purchases, and the interaction specification in column 4 of Table 2 is re-estimated. The histogram plots the resulting placebo interaction coefficients. The vertical line marks the actual estimate (0.220). The actual estimate is larger in magnitude than all but 4.3% of the placebo estimates.

Table B.3 – Balance Test: Lottery Winnings and Baseline Characteristics

	Dep var: Lottery winnings (1)
Lottery purchases	0.0921*** (0.0105)
Age	6844.6 (4462.1)
Age ² /100	-16781.3 (11156.0)
Age ³ /1,000	1816.5 (1231.7)
Age ⁴ /10,000	-73.33 (50.68)
Female	-12.66 (29.82)
Ethnicity: Chinese	-57.74 (64.24)
Primary education or lower	-64.19 (61.76)
Secondary education	16.02 (60.61)
Post-secondary education	-20.24 (62.04)
Married	-39.76 (79.91)
Citizen	45.40 (58.48)
Household size	44.70 (87.35)
Omnibus p -value	0.221
N	5,304
NT	11,298

Notes: Table presents a regression of annual lottery winnings in SGD on observable characteristics among lottery players, conditional on lottery ticket purchases, indicators for the frequency of playing, and wave fixed effects. Age enters as a quartic polynomial. The omnibus p -value tests that all covariate coefficients (excluding purchases and play frequency) jointly equal zero. Standard errors clustered by respondent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.4 – MPC from Lottery Winnings for Durable Consumption, by Baseline Health

Dependent variable: Durable consumption				
	All	Below Median	Above Median	All
	(1)	Health	Health	(4)
	(1)	(2)	(3)	(4)
Lottery winnings	0.204*	0.052	0.316**	0.076
	(0.116)	(0.162)	(0.159)	(0.160)
Lottery purchases	0.228***	0.289***	0.185***	0.297***
	(0.043)	(0.073)	(0.043)	(0.073)
Winnings $\times$ Above median health				0.252
				(0.219)
Purchases $\times$ Above median health				−0.124
				(0.082)
Above median baseline health				967.958**
				(396.597)
N	5,523	2,711	2,812	5,523
NT	11,297	5,517	5,780	11,297


Notes: Table presents estimates of [Equation 1](#). The dependent variable is durable consumption in SGD aggregated over the 12 months corresponding to the lottery module. Lottery winnings and purchases are measured in SGD, so the coefficient on winnings is a marginal propensity to consume. Columns 2 and 3 split the sample by baseline (first-wave) health; column 4 interacts winnings and purchases with an indicator for above-median baseline health. All regressions include wave fixed effects and covariates (indicators for race, marriage, education, household size, frequency of playing the lottery, and a quartic in age). Standard errors clustered by household in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.5 – MPC from Lottery Winnings for Total Consumption, by Baseline Health

Dependent variable: Total consumption				
	All	Below Median	Above Median	All
	(1)	Health	Health	(4)
	(1)	(2)	(3)	(4)
Lottery winnings	0.421*** (0.161)	0.139 (0.217)	0.640*** (0.225)	0.182 (0.213)
Lottery purchases	0.465*** (0.067)	0.533*** (0.110)	0.411*** (0.067)	0.553*** (0.111)
Winnings $\times$ Above median health				0.472 (0.301)
Purchases $\times$ Above median health				−0.158 (0.122)
Above median baseline health				1,311.501** (562.635)
<i>N</i>	5,523	2,711	2,812	5,523
<i>NT</i>	11,297	5,517	5,780	11,297


Notes: Table presents estimates of Equation 1. The dependent variable is total consumption (non-durable plus durable) in SGD aggregated over the 12 months corresponding to the lottery module. Lottery winnings and purchases are measured in SGD, so the coefficient on winnings is a marginal propensity to consume. Columns 2 and 3 split the sample by baseline (first-wave) health; column 4 interacts winnings and purchases with an indicator for above-median baseline health. All regressions include wave fixed effects and covariates (indicators for race, marriage, education, household size, frequency of playing the lottery, and a quartic in age). Standard errors clustered by household in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure B.3 – Permutation Test for the Lottery Winnings $\times$ Health Interaction: Durable Consumption



Notes: Figure shows the null distribution of the interaction between lottery winnings and above-median baseline health from randomization inference, with annual durable consumption per capita as the dependent variable. In each of 5,000 placebo samples, lottery winnings are randomly reassigned among players in the same survey wave and within-wave decile of ticket spending, which preserves the dependence of winnings on ticket purchases, and the interaction specification in column 4 of [Table B.4](#) is re-estimated. The histogram plots the resulting placebo interaction coefficients. The vertical line marks the actual estimate (0.252). The actual estimate is larger in magnitude than all but 16.7% of the placebo estimates.

Figure B.4 – Permutation Test for the Lottery Winnings $\times$ Health Interaction: Total Consumption



Notes: Figure shows the null distribution of the interaction between lottery winnings and above-median baseline health from randomization inference, with annual total consumption (non-durable plus durable) per capita as the dependent variable. In each of 5,000 placebo samples, lottery winnings are randomly reassigned among players in the same survey wave and within-wave decile of ticket spending, which preserves the dependence of winnings on ticket purchases, and the interaction specification in column 4 of [Table B.5](#) is re-estimated. The histogram plots the resulting placebo interaction coefficients. The vertical line marks the actual estimate (0.472). The actual estimate is larger in magnitude than all but 6.3% of the placebo estimates.

Table B.6 – Robustness: MPC from Lottery Winnings, including Asset Interactions

Dependent variable: Non-durable consumption			
	Main (1)	Non-missing Baseline Assets (2)	+ Baseline Asset Interactions (3)
Lottery winnings	0.106 (0.084)	0.130 (0.089)	−0.001 (0.104)
Winnings × Above median health	0.220* (0.128)	0.206 (0.133)	0.251* (0.130)
Winnings × Above median financial assets			0.122 (0.138)
Winnings × Above median MSA balance			−0.119 (0.146)
Lottery purchases	0.256*** (0.048)	0.257*** (0.052)	0.172*** (0.041)
Purchases × Above median health	−0.034 (0.054)	−0.049 (0.057)	−0.043 (0.052)
Purchases × Above median financial assets			0.032 (0.060)
Purchases × Above median MSA balance			0.065 (0.063)
<i>N</i>	5,304	4,967	4,967
<i>NT</i>	11,297	10,778	10,778


Notes: Column 1 repeats the specification in column 4 of [Table 2](#). Column 2 estimates the same specification on the sample with non-missing asset information. Column 3 adds interactions of lottery winnings and purchases with indicators for above-median financial assets and above-median MSA balances. The asset indicators equal one if the household’s per-capita financial assets or MSA balance, measured in the baseline wave, is above the survey-weighted median, and indicator main effects are included but not reported. *N* counts respondents in each estimation sample. All regressions include wave fixed effects and covariates (indicators for race, marriage, education, household size, frequency of playing the lottery, and a quartic in age). Standard errors clustered by household in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.7 – Robustness: MPC from Lottery Winnings, Reweighted to the Full Population

	Dep var: Annual Non-durable Consumption			
	All	Below Median Health	Above Median Health	Interaction
	(1)	(2)	(3)	(4)
Lottery winnings	0.224*** (0.075)	0.093 (0.091)	0.340*** (0.111)	0.112 (0.090)
Lottery purchases	0.243*** (0.035)	0.270*** (0.053)	0.209*** (0.036)	0.281*** (0.055)
Winnings $\times$ Above median health				0.231 (0.142)
Purchases $\times$ Above median health				-0.072 (0.061)
Above median baseline health				387.015 (259.113)
N	5,304	2,633	2,671	5,304
NT	11,297	5,517	5,780	11,297


Notes: Table re-estimates the [Table 2](#) regressions weighting each player by the survey weight times the inverse of a propensity score for playing the lottery, so that players match the covariate distribution of all respondents in the lottery waves. The propensity score is a logit of lottery participation on a quartic in age and indicators for gender, ethnicity, education (relative to primary or lower), marriage, citizenship, household size, and quintiles of baseline financial assets per capita, with wave fixed effects. Weights are trimmed at the 99th percentile. All regressions include wave fixed effects and demographic controls. Standard errors clustered by household in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure B.5 – Poisson Regressions: Change in Consumption Components




Notes: Figure plots Poisson regression estimates of components of non-durable consumption against the health index, with whiskers denoting 90% confidence intervals. 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. We estimate Poisson regressions for this analysis to accommodate zero values for some spending components in some months.

Figure B.6 – Binned Scatterplots: Log Durable Consumption and Total Consumption



Notes: Figure shows binned scatterplots of log durable consumption per capita (Panel A) and log total consumption per capita (Panel B) against the health index, controlling for individual fixed effects, month effects, and indicators for household size for 20 quantiles based on the methods of Cattaneo et al. (2024). The line plots a 4th-order global polynomial and the regression is weighted using the survey's sample weights.

Figure B.7 – Event Study of Consumption Response to Diagnosis of Chronic Conditions




Notes: Figure plots event-study coefficients of log non-durable consumption per capita around the diagnosis of any chronic condition, with whiskers denoting 95% confidence intervals. Event time is measured in months relative to diagnosis, with month -1 as the omitted reference period. The regression includes individual fixed effects, wave fixed effects, indicators for household size, age, an indicator for each respondent's first interview, and is weighted using survey sample weights; standard errors are clustered at the household level. Never-diagnosed respondents serve as controls. Observations more than 12 months before or 36 months after diagnosis are pooled in the binned endpoint categories. The p -value of the joint test that the pre-diagnosis coefficients (months -12 to -2) equal zero is 0.46.

Table B.8 – Relationship Between Health Index and Other Outcomes

	Positive Health Spending (1)	Log Health Spending (2)	Working (3)	Retired (4)	Income (5)
Health Index	-0.400*** (0.028)	-0.677*** (0.078)	0.080*** (0.026)	-0.068*** (0.024)	0.066 (0.043)
Individual fixed effects	Yes	Yes	Yes	Yes	Yes
Wave effects	Yes	Yes	Yes	Yes	Yes
R^2	0.421	0.398	0.810	0.844	0.874
N	12,742	10,677	12,386	12,386	12,268
NT	420,421	166,956	409,410	409,410	378,838

Notes: Table reports regression results of estimating Equation 8 for an indicator of positive health spending (column 1), log of health spending (column 2), an indicator for working (column 3), an indicator for retired (column 4), and income per capita (column 5). Columns 1–4 are estimated via OLS; column 5 is estimated via Poisson pseudo-maximum likelihood to account for zeros in income, and its reported R^2 is a pseudo R^2 . Income is winsorized at the 1st and 99th percentiles within each wave. Regressions also include indicators for household size. N denotes the number of unique individuals and NT denotes the number of individual-waves. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure B.8 – Correlation Between Subjective and Empirical Probabilities of Disease




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 is 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. It is 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.

Table B.9 – Consumption and Health Index by MSA Assets

	(1)	(2)
Health index	0.122* (0.074)	0.097* (0.050)
Individual fixed effects	Yes	Yes
Wave effects	Yes	Yes
MSA balances at baseline	Below median	Above median
R^2	0.727	0.719
N	4,956	4,695
NT	190,854	200,670

Notes: Table reports regression results of estimating Equation 8 via OLS, stratified by baseline level of MSA assets. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure B.9 – Sub-sample Analysis: Questions Related to Preferences to Delay and Regret



Notes: Figure plots estimates of state dependence for sub-samples that exclude respondents who report preferences that are consistent with present bias. Whiskers denote 90% confidence intervals calculated by block bootstrapping 1,000 samples for each group.

Table B.10 – Sensitivity to Externally Calibrated Parameters

β_{annual}	R_{annual}	S_{annual}	$u'(c)/v'(c)$
0.950	1.00	0.9763	1.20
0.950	1.00	0.9930	1.24
0.950	1.00	0.9940	1.24
0.950	1.02	0.9763	1.25
0.975	1.00	0.9763	1.26
0.950	1.02	0.9930	1.30
0.950	1.02	0.9940	1.30
0.950	1.04	0.9763	1.30
0.990	1.00	0.9763	1.31
0.975	1.00	0.9930	1.31
0.975	1.00	0.9940	1.32
0.975	1.02	0.9763	1.32
0.950	1.04	0.9930	1.35
0.950	1.04	0.9940	1.36
0.990	1.00	0.9930	1.36
0.990	1.00	0.9940	1.36
0.990	1.02	0.9763	1.37
0.975	1.02	0.9930	1.38
0.975	1.02	0.9940	1.38
0.975	1.04	0.9763	1.38
0.990	1.02	0.9930	1.43
0.990	1.02	0.9940	1.43
0.990	1.04	0.9763	1.44
0.975	1.04	0.9930	1.44
0.975	1.04	0.9940	1.45
0.990	1.04	0.9930	1.50
0.990	1.04	0.9940	1.50

Notes: The table reports sensitivity of $\delta = S\beta R(1-p)/(1-S\beta Rp)$, and of the implied state-dependence ratio $u'(c)/v'(c) = \delta/D$. The denominator $D = 1 - \gamma\lambda + \frac{1}{2}\gamma(\gamma+1)\lambda^2$ is held fixed at the [Table 4](#) value estimated from the consumption response and risk-aversion data. Rows show all 27 combinations of $\beta_{annual} \in \{0.95, 0.975, 0.99\}$, $R_{annual} \in \{1.00, 1.02, 1.04\}$, and $S_{annual} \in \{0.9763, 0.9930, 0.9940\}$, where the middle value of each set is the baseline calibration. All parameters are displayed at annual frequency; δ is computed at monthly frequency. The transition probability is held fixed at its baseline value $p = 0.966$. Rows are sorted by the implied ratio $u'(c)/v'(c)$. The bolded row is the baseline calibration of β, R, S .

Table B.11 – Nonparametric tests from Finkelstein, Luttmer and Notowidigdo (2013)

	(1)	(2)	(3)	(4)
Bin width for mean life satisfaction	0.100	0.125	0.125	0.050
Bin width for mean health index	0.100	0.125	0.050	0.125
Number of cells	70	48	104	119
Number of usable cells	64	48	97	110
Number of individuals in usable cells	7040	7046	7044	7043
Average number of observations per cell used	4501	6007	2826	2484
Average number of individuals per cell used	110	147	72	64
Mean of CDFs of t -statistics (F)	0.690	0.710	0.648	0.609
p -value of test that $F = 0.5$	< 0.001	< 0.001	< 0.001	< 0.001
Fraction of estimated interaction terms > 0	0.734	0.771	0.680	0.636

Notes: Table reports results of the nonparametric test of state dependence. Each column shows the results of a test implemented for the choice of bin widths presented in the first two rows of the table. Usable cells include at least one individual. For each bin, we estimate Equation 10 via OLS where the dependent variable is the binary measure of life satisfaction. The bottom three rows present the results of the test, which compare the mean of the cumulative distribution of the t -statistics of the interaction term to 0.5, which is the mean under the null of state independence.

Table B.12 – Robustness: State-Dependence Results, Weighted Satisfaction

	(1)
φ_1 , Health index $\times$ Log consumption	0.020** (0.010)
φ_2 , Log consumption	0.009*** (0.002)
φ_3 , Health index	0.153*** (0.022)
R^2	0.777
N	10,953
NT	409,918

Notes: Table shows results from estimating Equation 10 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. Standard errors are clustered by household in parentheses. Regressions are weighted using survey sample weights. N denotes the number of unique individuals and NT denotes the number of individual-waves. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.13 – Life Satisfaction Regressions, Number of Chronic Conditions

	Ordered Logit (1)	LPM (2)
φ_1 , Number of Chronic Conditions $\times$ Log consumption	-0.056*** (0.014)	-0.005*** (0.002)
φ_2 , Log consumption	0.095*** (0.016)	0.015*** (0.002)
φ_3 , Number of Chronic Conditions	-0.171*** (0.035)	-0.019*** (0.004)
R^2		0.651
N	8,878	10,988
NT	550,393	410,201

Notes: Table presents results from estimating Equation 10 using the number of chronic conditions instead of the health index. Log consumption is the natural logarithm of non-durable 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Online Appendix C: Derivation of State-Dependence Based on Sufficient-Statistics Formulas

This appendix presents two derivations of the state-dependence formula. The first, in Section C.1, derives the formula from the Euler equation as presented in Equation 7. The second, in Section C.2, derives an alternative formula that exploits within-state first-order conditions from a two-good model and does not require the Euler equation or the stationarity approximation.

C.1 Euler equation derivation

The Euler equation for an agent who is healthy in period $t - 1$ is:

$$u'(c_{h,t-1}) = S\beta R [p \cdot u'(c_{h,t}) + (1 - p) \cdot v'(c_{b,t})] \quad (\text{C.1})$$

where S is the survival probability, β is the discount factor, $R = 1 + r$ is the gross interest rate, p is the probability of remaining healthy, and u' and v' denote marginal utility in the healthy and sick states respectively. We make the approximation $u'(c_{h,t-1}) \approx u'(c_{h,t})$, which we empirically evaluate below, and then substitute for $u'(c_{h,t})$ on the right-hand side:

$$u'(c_{h,t-1}) = S\beta R [p \cdot u'(c_{h,t-1}) + (1 - p) \cdot v'(c_{b,t})] \quad (\text{C.2})$$

Collecting terms in $u'(c_{h,t-1})$:

$$u'(c_{h,t-1}) [1 - S\beta R p] = S\beta R (1 - p) \cdot v'(c_{b,t}) \quad (\text{C.3})$$

Solving for $u'(c_{h,t-1})$:

$$u'(c_{h,t-1}) = \frac{S\beta R(1-p)}{1-S\beta Rp} \cdot v'(c_{b,t}) \equiv \delta \cdot v'(c_{b,t}) \quad (\text{C.4})$$

To express the marginal utilities in both health states at a common consumption level, we take a second-order Taylor expansion of $u'(c_{h,t-1})$ around $c_{b,t}$ and substitute it into [Equation C.4](#):

$$u'(c_{b,t}) + u''(c_{b,t})(c_{h,t-1} - c_{b,t}) + \frac{1}{2}u'''(c_{b,t})(c_{h,t-1} - c_{b,t})^2 = \delta \cdot v'(c_{b,t}) \quad (\text{C.5})$$

Dividing both sides by $u'(c_{b,t})$:

$$1 + \frac{u''(c_{b,t})}{u'(c_{b,t})}(c_{h,t-1} - c_{b,t}) + \frac{1}{2} \frac{u'''(c_{b,t})}{u'(c_{b,t})}(c_{h,t-1} - c_{b,t})^2 = \delta \cdot \frac{v'(c_{b,t})}{u'(c_{b,t})} \quad (\text{C.6})$$

Under CRRA utility, the coefficient of relative risk aversion is $\gamma = -\frac{u''(c)}{u'(c)}c$ which implies $\frac{u''(c_{b,t})}{u'(c_{b,t})} = -\frac{\gamma}{c_{b,t}}$. For the third derivative, the coefficient of relative prudence under CRRA utility equals $-\frac{u'''(c)}{u''(c)}c = \gamma + 1$. This gives $u'''(c) = -u''(c) \cdot \frac{\gamma+1}{c}$. Substituting the expression for $u''(c)$:

$$\frac{u'''(c_{b,t})}{u'(c_{b,t})} = \frac{u'''(c_{b,t})}{u''(c_{b,t})} \cdot \frac{u''(c_{b,t})}{u'(c_{b,t})} = \frac{-(\gamma+1)}{c_{b,t}} \cdot \frac{-\gamma}{c_{b,t}} = \frac{\gamma(\gamma+1)}{c_{b,t}^2} \quad (\text{C.7})$$

Substituting into [Equation C.6](#) and defining $\frac{\Delta c_t}{c_{b,t}} \equiv \frac{c_{h,t-1} - c_{b,t}}{c_{b,t}}$:

$$1 - \gamma \frac{\Delta c_t}{c_{b,t}} + \frac{\gamma(\gamma+1)}{2} \left(\frac{\Delta c_t}{c_{b,t}} \right)^2 = \delta \cdot \frac{v'(c_{b,t})}{u'(c_{b,t})} \quad (\text{C.8})$$

Finally, factor $\gamma \frac{\Delta c_t}{c_{b,t}}$ on the left-hand side, take the reciprocal, and rearrange to write the ratio of marginal utilities when healthy relative to when unhealthy:

$$\frac{u'(c_{b,t})}{v'(c_{b,t})} = \frac{\delta}{1 - \gamma \frac{\Delta c_t}{c_{b,t}} \left[1 - \frac{\gamma+1}{2} \frac{\Delta c_t}{c_{b,t}} \right]} \quad (\text{C.9})$$

which is the formula presented in the text as [Equation 7](#).

The approximation $u'(c_{h,t-1}) \approx u'(c_{h,t})$ is an assumption about stationarity that allows us to simplify the formula for state dependence. We empirically evaluate the approximation error $\varepsilon = u'(c_{h,t-1}) - u'(c_{h,t})$ using our data by first calculating consumption in sequential months for respondents without chronic conditions in either month. The median difference in consumption is zero and the mean difference is 11 SGD, which is small relative to an average non-durable consumption around 1,000 SGD.

C.2 Within-period derivation using a state-independent outside good

As a complementary approach, we derive an alternative formula for state dependence that does not rely on the Euler equation or the stationarity approximation. Instead, we exploit within-state first-order conditions from a model in which the agent consumes two goods: a state-dependent consumption good c_t and a state-independent outside good x_t . In our empirical application, we use charitable giving, spending for others, and gifts to family and friends as a set of items more likely

to be state-independent to stand in for x_t , though the derivation applies to any good whose utility is invariant to health status.

Utility is additively separable across goods and health states. In each period, the agent derives utility from c_t and x_t :

$$\begin{aligned} \text{Healthy: } U_{h,t} &= u(c_{h,t}) + f(x_{h,t}) \\ \text{Sick: } V_{b,t} &= v(c_{b,t}) + f(x_{b,t}) \end{aligned} \quad (\text{C.10})$$

where $f(\cdot)$ is identical across health states, reflecting the assumption that the utility from the outside good is invariant to health status. As in Section C.1, $u(\cdot)$ and $v(\cdot)$ satisfy CRRA with curvature γ . Because the outside good may be zero for some households (e.g., those who do not give to charity in a given month), we adopt the following specification for f to ensure that marginal utility is defined at zero:

$$f(x) = \frac{(x + \psi)^{1-\gamma_x}}{1 - \gamma_x} \quad (\text{C.11})$$

where $\psi > 0$ is a threshold parameter and $\gamma_x > 0$ is the curvature. The parameter $\psi > 0$ serves a similar role as the consumption floor parameter in the literature on bequests to model how bequests are luxury goods (Lockwood 2018). This yields $f'(x) = (x + \psi)^{-\gamma_x}$, which is finite at $x = 0$.

Conditional on health status, the agent equates marginal utilities across goods: $u'(c_{h,t}) = f'(x_{h,t})$ when healthy and $v'(c_{b,t}) = f'(x_{b,t})$ when unhealthy. Because the outside good's utility is the same function f in both states, these conditions imply that the shadow value of wealth in each state is equated across both goods. Taking the ratio of the two conditions:

$$\frac{u'(c_{h,t})}{v'(c_{b,t})} = \frac{f'(x_{h,t})}{f'(x_{b,t})} \quad (\text{C.12})$$

This relationship holds within a single period and does not require any inter-temporal assumptions. We approximate each side by a second-order Taylor expansion and suppress the time subscripts below because the expression holds within a particular time period. Expanding $u'(c_h)$ around c_b and dividing by $u'(c_b)$, with $\frac{\Delta c}{c_b} \equiv (c_h - c_b)/c_b$:

$$\frac{u'(c_h)}{u'(c_b)} \approx 1 - \gamma \frac{\Delta c}{c_b} + \frac{1}{2} \gamma(\gamma + 1) \left(\frac{\Delta c}{c_b} \right)^2 \quad (\text{C.13})$$

following the same steps as in Section C.1. Multiplying and dividing the left-hand side of Equation C.12 by $u'(c_b)$:

$$\frac{u'(c_h)}{v'(c_b)} = \frac{u'(c_h)}{u'(c_b)} \cdot \frac{u'(c_b)}{v'(c_b)} \quad (\text{C.14})$$

Substituting Equation C.13 and rearranging Equation C.12:

$$\frac{u'(c_b)}{v'(c_b)} = \frac{1}{1 - \gamma \frac{\Delta c}{c_b} \left[1 - \frac{\gamma+1}{2} \frac{\Delta c}{c_b} \right]} \cdot \frac{f'(x_h)}{f'(x_b)} \quad (\text{C.15})$$

Applying a similar Taylor approximation for f' and denoting $\frac{\Delta x}{x_b + \psi} \equiv \frac{x_h - x_b}{x_b + \psi}$ yields:

$$\frac{f'(x_h)}{f'(x_b)} \approx 1 - \gamma_x \frac{\Delta x}{x_b + \psi} \left[1 - \frac{\gamma_x + 1}{2} \frac{\Delta x}{x_b + \psi} \right] \quad (\text{C.16})$$

The key difference from the consumption component is that the proportional change is measured relative to $x_b + \psi$ rather than x_b , reflecting the shifted origin to accommodate zeros.

Substituting Equation C.16 into Equation C.15:

$$\frac{u'(c_b)}{v'(c_b)} \approx \frac{1 - \gamma_x \frac{\Delta x}{x_b + \psi} \left[1 - \frac{\gamma_x + 1}{2} \frac{\Delta x}{x_b + \psi} \right]}{1 - \gamma \frac{\Delta c}{c_b} \left[1 - \frac{\gamma + 1}{2} \frac{\Delta c}{c_b} \right]} \quad (\text{C.17})$$

The denominator is the same expression that appears in the Euler equation derivation and captures how state-dependent consumption changes between health states. The numerator captures the corresponding change in the outside good, where the proportional change $\frac{\Delta x}{x_b + \psi}$ is attenuated relative to $\frac{\Delta x}{x_b}$ by the threshold parameter. This formula requires estimating the proportional changes in consumption ($\frac{\Delta c}{c_b}$) and the outside good ($\frac{\Delta x}{x_b + \psi}$) between health states, along with the curvature parameters γ and γ_x and the threshold parameter ψ , but does not require estimating or calibrating health transitions, survival rates, discount rates, or interest rates.

Table C.1 presents the results. We use charitable giving, spending for others, and gifts to family and friends as the state-independent outside good. We estimate $\frac{\Delta c}{c_b}$ from a regression of log consumption on the health index with individual fixed effects, and $\frac{\Delta x}{x_b + \psi}$ via Poisson pseudo-maximum likelihood regression of $(x + \psi)$ on the health index with individual fixed effects. We calibrate $\gamma_x = \gamma$ and present results for three values of ψ , set to $0.5\times$, $1\times$, and $2\times$ mean giving among givers, which is equal to 287 SGD.

Table C.1 – State Dependence Estimates: Within-State First-Order Condition Approach

	$\psi = 0.5 \times \bar{x}$	$\psi = 1 \times \bar{x}$	$\psi = 2 \times \bar{x}$
<i>Inputs</i>			
$\frac{\Delta c}{c_b}$, consumption change between health states		0.108	
γ , coefficient of relative risk aversion		3.64	
γ_x , curvature of outside good		3.64	
$\frac{\Delta x}{x_b + \psi}$, outside good change between health states	-0.079	-0.051	-0.030
Extensive margin (health $\rightarrow$ Pr(give))		-0.055	
<i>Results</i>			
$\frac{u'(c)}{v'(c)}$, State dependence	1.75	1.65	1.56
90% CI	[1.18, 2.09]	[1.21, 1.99]	[1.20, 1.84]
95% CI	[1.10, 2.14]	[1.14, 2.04]	[1.14, 1.88]

Notes: Table presents the components of the formula for $\frac{u'(c)}{v'(c)}$ based on the within-state first-order condition approach described in Appendix C.2. Each column corresponds to a different value of the Stone-Geary subsistence parameter ψ , calibrated as a multiple of mean giving among givers ($\bar{x}$). The consumption change $\frac{\Delta c}{c_b}$ is estimated from a regression of log consumption on the health index with individual fixed effects. The outside good change $\frac{\Delta x}{x_b + \psi}$ is estimated via Poisson pseudo-maximum likelihood regression of $(x + \psi)$ on the health index with individual fixed effects. Confidence intervals are calculated by block bootstrapping 1,000 samples at the individual level.

Online Appendix D: Euler Equation with Bequest Motives

This appendix extends the stochastic lifecycle model of consumption and savings to incorporate a bequest motive. In this framework, an agent with constant relative risk aversion (CRRA) preferences chooses consumption to maximize the discounted expected utility of lifetime consumption and bequests. We maintain the baseline assumptions: medical costs are fully insured, non-medical consumption does not affect health status, utility is inter-temporally separable, the future is discounted by factor $0 < \beta < 1$, and the agent can borrow and save at the gross interest rate $R \equiv 1 + r$. The key modification is the inclusion of utility derived from wealth remaining at the time of death. The agent's decision problem balances the utility of current consumption against the expected utility of future consumption if alive and the utility of leaving a bequest if dead. The optimal consumption path is characterized by the following modified Euler equation:

$$u'(c_{h,t-1}) = S\beta R[p u'(c_{h,t}) + (1-p)v'(c_{b,t})] + (1-S)\beta R w'(A_t) \quad (\text{D.1})$$

where $w'(A_t)$ is the marginal utility of leaving a bequest of value A_t and the first term on the right-hand side represents the expected marginal utility from future consumption conditional on survival as in the main text. As before, we rearrange Equation (D.1) under the assumption that the marginal utility of consuming c_h is constant over time: $u'(c_{h,t-1}) = u'(c_{h,t})$.

$$u'(c_{h,t-1}) = \frac{S\beta R(1-p)}{1-S\beta R p} v'(c_{b,t}) + \frac{(1-S)\beta R}{1-S\beta R p} w'(A_t)$$

Defining the scaling factors δ (as in the baseline model without bequests) and D (the bequest factor) as:

$$\delta \equiv \frac{S\beta R(1-p)}{1-S\beta R p}, \quad D \equiv \frac{(1-S)\beta R}{1-S\beta R p}$$

we can express the relationship more compactly:

$$u'(c_{h,t-1}) = \delta v'(c_{b,t}) + D w'(A_t) \quad (\text{D.2})$$

We assume the marginal utility of bequests, $w'(A_t)$, is proportional to the marginal utility of consumption in the unhealthy state: $w'(A_t) = \theta v'(c_{b,t})$, where $\theta \geq 0$ represents the strength of the bequest motive. Substituting this into Equation (D.2) yields:

$$u'(c_{h,t-1}) = (\delta + \theta D) v'(c_{b,t}) \equiv \delta^B v'(c_{b,t}) \quad (\text{D.3})$$

Following the methodology in the main text, we approximate the marginal utility of healthy consumption via a second-order Taylor expansion of $u'(c_{h,t-1})$ around $c_{b,t}$. Combining this with Equation (D.3) and substituting the definitions for relative risk aversion (γ) and prudence ($\gamma + 1$), with $\frac{\Delta c_t}{c_{b,t}} \equiv \frac{c_{h,t-1} - c_{b,t}}{c_{b,t}}$ as defined in the main text, we obtain the following modified expression for state dependence:

$$\frac{u'(c_{b,t})}{v'(c_{b,t})} = \frac{\delta^B}{1 - \gamma \frac{\Delta c_t}{c_{b,t}} \left[1 - \frac{\gamma+1}{2} \frac{\Delta c_t}{c_{b,t}} \right]} \quad (\text{D.4})$$

The estimated ratio of marginal utilities between health states is directly proportional to the augmented scaling factor δ^B . Note that $\delta^B > \delta$ if the bequest motive $\theta > 0$ since $D > 0$. And since δ^B increases with θ , a stronger desire to leave bequests implies a higher degree of complementarity between health and consumption.

Our model in the main text corresponds to the case where $\theta = 0$. To assess how the presence of a bequest motive would increase our estimates, we calculate the implied state-dependence ratio across a range of values for θ . We utilize the monthly parameters detailed in Table 4 of the main text: survival probability $S = 0.999$, discount factor $\beta = 0.998$, gross interest rate $R = 1.0016$, and health transition probability $p = 0.966$. These yield $\delta \approx 0.960$ and $D \approx 0.028$. Table D.1 presents estimates of state dependence as θ varies. If the marginal utility of bequests is equal to the marginal utility of consumption in the bad state ($\theta = 1$), then state dependence would be only slightly higher (1.42) than the 1.38 estimated without a bequest motive. If the marginal utility of a bequest is five times the marginal utility of consumption when ill ($\theta = 5$), then state dependence increases to 1.58.

Table D.1 – Sensitivity of State Dependence to the Strength of Bequest Motive (θ)

Bequest Motive (θ)	Scaling Factor (δ^B)	State Dependence $\frac{u'(c)}{v'(c)}$
0 (Baseline)	0.960	1.38
0.5	0.974	1.40
1.0	0.989	1.42
2.0	1.017	1.46
5.0	1.102	1.58

Calibration based on Lockwood (2018) and Liu et al. (2024): We calibrate θ by mapping this parameter to the structural bequest motive developed in Lockwood (2018) and estimated in the context of Singapore by Liu et al. (2024). In these models, the utility of consumption c is $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ and utility from a bequest of assets A is specified as:

$$w(A) = \left(\frac{\phi}{1-\phi} \right)^\gamma \frac{\left(\frac{\phi}{1-\phi} \bar{c} + A \right)^{1-\gamma}}{1-\gamma} \quad (\text{D.5})$$

where $\bar{c}$ is the consumption threshold above which people leave bequests and ϕ represents the bequest intensity, equal to the marginal propensity to bequeath for those with consumption above $\bar{c}$.

In our framework, θ represents the ratio of the marginal utility of bequests to the marginal utility of consumption in the bad health state: $\theta = \frac{w'(A)}{v'(c_b)}$. Differentiating Equation D.5 yields:

$$\theta = \frac{\left(\frac{\phi}{1-\phi} \right)^\gamma \left(\frac{\phi}{1-\phi} \bar{c} + A \right)^{-\gamma}}{c_b^{-\gamma}} = \left(\frac{\frac{\phi}{1-\phi} c_b}{\frac{\phi}{1-\phi} \bar{c} + A} \right)^\gamma \quad (\text{D.6})$$

We calibrate this value using parameter estimates from Table 7 in Liu et al. (2024) and data from the SLP. The two key parameters are $\phi = 0.924$ and $\bar{c} = 3,984$ SGD annually, which we convert to 332 SGD monthly to match the frequency of our data. The estimate of $\phi = 0.924$ in Singapore is only slightly below the estimate of 0.96 from Lockwood (2018) in the US. We take the other inputs from the SLP. In their structural estimation, Liu et al. (2024) externally calibrate $\gamma = 3$, which is slightly below our estimate of 3.64. We calculate median household assets per capita $A = 10,000$ SGD and average non-durable consumption for those in the bottom quartile of the health index $c_b = 933$ SGD. These inputs yield $\theta = 0.46$:

$$\theta = \left(\frac{\frac{0.924}{1-0.924} \cdot 933}{\frac{0.924}{1-0.924} \cdot 332 + 10000} \right)^{3.64} \approx (0.808)^{3.64} \approx 0.46. \quad (\text{D.7})$$

This magnitude indicates that the marginal utility of leaving a bequest is 46% as large as the marginal utility of consumption when ill. As shown in [Table D.1](#), this exercise suggests that state dependence would be slightly higher (1.40 vs. 1.38) if our model incorporated incidental bequests.

Online Appendix E: Optimal Insurance with State Dependence

This appendix presents the derivations used in [Section 8.1](#) and [Section 8.3](#) of the main text. There are two states of the world. In the good state (g), which occurs with probability $(1 - \pi)$, there is no financial loss and the health index is h_g . In the loss state (l), which occurs with probability π , a loss of size L occurs and the health index is h_l . We assume $h_l \leq h_g$. Starting with wealth W , the agent can buy k dollars of insurance coverage at a price of α per dollar. The price is based on the loss probability and a load factor $z \geq 0$: $\alpha = \pi(1 + z)$. In the case of actuarially fair insurance, $z = 0$ so $\alpha = \pi$. With these prices, consumption in the two states is as follows:

$$C_g = W - \alpha k \quad (\text{E.1})$$

$$C_l = W - L + k - \alpha k = W - L + (1 - \alpha)k \quad (\text{E.2})$$

Preferences satisfy Constant Relative Risk Aversion (CRRA) and the parameter σ captures the strength of state dependence based on health status h_s . Utility in state s is

$$u(C_s) = (1 + \sigma h_s) \frac{C_s^{1-\gamma}}{1-\gamma} \quad (\text{E.3})$$

The individual chooses the amount of insurance coverage, k , to maximize their expected utility, $E[U]$.

$$\max_k E[U] = (1 - \pi) \left[(1 + \sigma h_g) \frac{(W - \alpha k)^{1-\gamma}}{1-\gamma} \right] + \pi \left[(1 + \sigma h_l) \frac{(W - L + (1 - \alpha)k)^{1-\gamma}}{1-\gamma} \right] \quad (\text{E.4})$$

The first-order condition is:

$$\pi(1 - \alpha)v'(C_l) = \alpha(1 - \pi)u'(C_g) \quad (\text{E.5})$$

which yields the ratio of marginal utilities:

$$\frac{u'(C_g)}{v'(C_l)} = \frac{\pi(1 - \alpha)}{\alpha(1 - \pi)} \quad (\text{E.6})$$

Substituting the marginal utility functions and rearranging yields:

$$\frac{(1 + \sigma h_g)C_g^{-\gamma}}{(1 + \sigma h_l)C_l^{-\gamma}} = \frac{\pi(1 - \alpha)}{\alpha(1 - \pi)} \quad (\text{E.7})$$

$$\left(\frac{C_g}{C_l} \right)^{-\gamma} = \left(\frac{1 + \sigma h_l}{1 + \sigma h_g} \right) \left[\frac{\pi(1 - \alpha)}{\alpha(1 - \pi)} \right]$$

$$\frac{C_g}{C_l} = \left(\frac{1 + \sigma h_g}{1 + \sigma h_l} \right)^{1/\gamma} \left[\frac{\alpha(1 - \pi)}{\pi(1 - \alpha)} \right]^{1/\gamma}$$

Define $\psi \equiv \left(\frac{1 + \sigma h_g}{1 + \sigma h_l} \right)^{1/\gamma}$, which captures the ratio of marginal utilities between states, and $\Lambda \equiv \left[\frac{\alpha(1 - \pi)}{\pi(1 - \alpha)} \right]^{1/\gamma}$, which captures the price wedge from loading relative to an actuarially fair premium. Plugging in the expressions for C_g, C_l and solving for k^* yields

$$\begin{aligned} W - \alpha k^* &= \psi \Lambda (W - L + (1 - \alpha) k^*) \\ W - \alpha k^* &= \psi \Lambda W - \psi \Lambda L + \psi \Lambda (1 - \alpha) k^* \\ \psi \Lambda L - \psi \Lambda W + W &= \alpha k^* + \psi \Lambda (1 - \alpha) k^* \\ W(1 - \psi \Lambda) + \psi \Lambda L &= k^* [\alpha + \psi \Lambda (1 - \alpha)] \\ k^* &= \frac{W(1 - \psi \Lambda) + \psi \Lambda L}{\alpha + \psi \Lambda (1 - \alpha)} \end{aligned} \tag{E.8}$$

To measure the degree of under-insurance, we subtract L from both sides and rearrange:

$$\begin{aligned} k^* - L &= \frac{W(1 - \psi \Lambda) + \psi \Lambda L}{\alpha + \psi \Lambda (1 - \alpha)} - L \\ &= \frac{W(1 - \psi \Lambda) + \psi \Lambda L - L[\alpha + \psi \Lambda (1 - \alpha)]}{\alpha + \psi \Lambda (1 - \alpha)} \\ &= \frac{W - W\psi \Lambda + \psi \Lambda L - L\alpha - L\psi \Lambda (1 - \alpha)}{\alpha + \psi \Lambda (1 - \alpha)} \\ &= \frac{W - W\psi \Lambda + \psi \Lambda L - L\alpha - L\psi \Lambda + L\psi \Lambda \alpha}{\alpha + \psi \Lambda (1 - \alpha)} \\ &= \frac{W - W\psi \Lambda - L\alpha + L\psi \Lambda \alpha}{\alpha + \psi \Lambda (1 - \alpha)} \\ &= \frac{W(1 - \psi \Lambda) - L\alpha(1 - \psi \Lambda)}{\alpha + \psi \Lambda (1 - \alpha)} \\ &= \frac{(W - L\alpha)(1 - \psi \Lambda)}{\alpha + \psi \Lambda (1 - \alpha)} \\ &= \frac{(L\alpha - W)(\psi \Lambda - 1)}{\alpha + \psi \Lambda (1 - \alpha)} \end{aligned} \tag{E.9}$$

Alternatively, we can express this as the difference from L as in [Section 8.1](#):

$$L - k^* = \frac{(W - L\alpha)(\psi \Lambda - 1)}{\alpha + \psi \Lambda (1 - \alpha)} \tag{E.10}$$

When insurance is actuarially fair, $\Lambda = 1$ and the equation becomes [Equation 12](#) of the main text.