Consumption and Portfolio Choices under Health Risks: A Structural Life-Cycle Approach in Retirement

Xiaohe Li and Xinrui Xu

Supervisor:

Zoltán Rácz

Master in Finance

of

Stockholm School of Economics.

June 7, 2025

We would like to sincerely thank Zoltán Rácz for his invaluable guidance throughout the development of this thesis. His patient encouragement, insightful feedback, and thoughtful support—particularly regarding the structure and implementation of the code—have been crucial to the progress of our work. His rigorous approach and attention to detail have significantly shaped our understanding of the modeling process. We are deeply grateful for his time and consistent support at every stage of this thesis.

Abstract

We analyze the effect of health risk on consumption and saving policies in the post-retirement period by generating a life-cycle model with a stochastic health transition process and endogenous mortality risk. We consider both a benchmark model with only risk-free savings and an extended version that incorporates portfolio choice between risk-free and risky assets. The model is estimated using HRS data and the method of simulated moments. We find that healthier individuals tend to consume less and accumulate more wealth, driven by lower marginal utility and longer life expectancy, while those in poorer health tend to deplete wealth more rapidly. The share of risky assets declines with age, while health-related heterogeneity in portfolio choice at a given age is limited. The model provides a tractable framework for analyzing the economic impact of health deterioration and longevity risk in retirement, especially in light of population aging and the expanding retired population.

Keywords: Life-cycle model, retirement, health risk, endogenous mortality, portfolio choice

Contents

Chapter 1

Introduction

As populations age and life expectancy continues to rise for many countries, the understanding of the economic behavior of retirees becomes more and more important for both researchers and policymakers. Most of the existing life-cycle literature primarily focuses on the preretirement period and often treats the post-retirement period as a simplified extension of the former. However, the shift in the primary sources of uncertainty creates a substantial gap between the two phases. In particular, health deterioration and the associated mortality risks are key determinants of consumption and saving behavior among retirees, rather than the labor income risk in the pre-retirement phase. A growing number of studies on the post-retirement period incorporate health uncertainty in the model by highlighting the monetary consequences of medical expenditure risk. However, the direct impact of health status and longevity expectations on utility remains relatively underexplored in the existing literature.

Building on the existing post-retirement models of health-related saving behavior, we develop a more focused model to examine the influence of the health transition and corresponding mortality risk on the post-retirement individual behavior. Our model integrates the utility specification of health status and consumption from De Nardi et al. (2010) and modifies the health transition process introduced by Yogo (2016) to maintain tractability. We also extend the classic portfolio choice model generated by Cocco et al. (2005) to better reflect the health- and longevity-driven features specific to post-retirement behavior. This integrated approach allows us to better capture the influence of health dynamics on retirement behavior and asset allocation decisions, thereby presenting a more explicit understanding of the economic consequences of health deterioration and longevity expectations in the post-retirement period.

The rest of the thesis is organized as follows. Chapter 2 reviews the related literature. Chapter 3 presents the life-cycle model in retirement. Chapter 4 introduces the data source and outlines the methodology used to construct the simulated model. Chapter 5 demonstrates the results obtained from the simulated model and discusses several extensions. We present conclusions in Chapter 6.

Chapter 2

Literature Review

2.1 Life Cycle Model of Consumption and Saving

The life-cycle model was first proposed in the 1950s. Several papers suggest a similar idea with slightly different emphasis that individuals face a consumption and saving choice during their lifetime. Among them, Modigliani and Brumberg (1954) proposes the theory that individuals aim to solve the consumption-saving problem during their lifetime and maximize the utility subject to their budget constraints. In addition, Friedman (1957) suggests a related idea which is presented as the permanent income hypothesis (hereafter PIH). The PIH suggests the consumption levels for individuals depend not on their current income, but rather on their expected lifetime (or 'permanent') income.

The empirical data provide support for the PIH, among which Campbell and Mankiw (1989) argues the PIH holds for a portion of the population while the rest are hand-to-mouth and borrowing constrained. Zeldes (1989) empirically confirms the presence of the borrowing constraint by comparing the parameter linked to current income on consumption growth across groups with different wealth-to-income ratios.

2.2 Precautionary Savings and Uncertainty

However, the empirical data reveal a puzzling pattern: aggregate consumption closely follows income, while the correlation between the two at the individual level is weak. Therefore, Carroll and Samwick (1997) examines the empirical data and provides some of the first direct evidence that consumers facing greater income uncertainty tend to hold systematically higher levels of wealth. This study introduces the concept of the precautionary saving motive when individuals face uncertainty. Empirical studies (such as Gourinchas and Parker (2002)) also provide support to the precautionary saving concept.

Most of the literature investigates labor income risk as the primary source of uncertainty, as it primarily focuses on the pre-retirement period. For instance, Storesletten et al. (2004) investigates the uninsurable labor income risk by constructing an overlapping generations general equilibrium model and finds that the risk realized throughout the working years has a greater impact on household inequality. For a comprehensive survey of household behavior over lifetime, see Gomes et al. (2021) for an overview. In contrast, relatively few studies focus on household behavior in the post-retirement phase, where the labor income risk is no longer the primary source of uncertainty and households instead face risks related to health and mortality.

2.3 Modeling the Retirement Period

The rising academic research on the retirement period suggests the growing importance of retirement as the population ages and life expectancy continue to increase. In the classical literature, the retirement period is often treated as a simplified extension of the pre-retirement period. To address this limitation, several factors have been introduced to enrich the retirement model.

The main sources of the uncertainty during the retirement period come from health, medical expenditure, and lifespan. Palumbo (1999) incorporates health status into the utility function, distinguishing between two discrete states: healthy and unhealthy. French and Jones (2004) decomposes the out-of-pocket medical expense process similarly to the decomposition of the labor income before retirement. De Nardi et al. (2010) combines those two and finds that the mortality risk and expensive medical care are the key drivers of elderly savings. These studies incorporate health uncertainty and highlight the monetary consequences of medical expenditure risk.

Yet, the direct utility effect from health deterioration and the associated mortality risk is somewhat less explored in the literature. While several studies focus on the direct effect of health on wealth (e.g., Mahler and Yum (2024) and De Nardi et al. (2024)) or earnings (e.g., Hosseini et al. (2021)) over the entire life cycle, and others explore how health influences labor supply and retirement timing (e.g., Blundell et al. (2023) and French (2005)), few of them emphasize the post-retirement period.

Retirees could allocate their wealth across different asset classes to mitigate the effects caused by health risks and maximize lifetime utility. Cocco et al. (2005) provides a general solution to the portfolio allocation rule at different ages, imposing an exogenous mortality risk during retirement. In the presence of health and mortality risks, Yogo (2016) develops a rich model by allowing the asset allocation across bond, stock, and housing. Moreover, this paper treats health as ”health capital” and allows the health capital to change endogenously in response to health investment.

2.4 Model Foundations and Contributions

Building on the existing literature, our thesis generates a more focused and tractable model investigating the effect of health and related mortality risk. Specifically, we follow De Nardi et al. (2010) in modeling the direct utility effects of health status and consumption. We extend their framework by allowing asset allocation between risk-free and risky assets and adopting a richer specification of the health transition process, following Yogo (2016). This allows us to incorporate an endogenous mortality rate that varies with age and health status. We simplify the health and non-health investment opportunity set compared to Yogo's richer specification to maintain tractability while focusing on the role of health dynamics. In addition, we build on the classic portfolio choice model of Cocco et al. (2005) by enriching it with a more detailed structure during retirement, capturing the health- and longevity-driven features specific to post-retirement behavior.

Chapter 3

Model

3.1 Fundamental Concept of the Life-Cycle Model

The objective of the life-cycle model is to maximize the expected lifelong utility, based on current circumstances and future expectations. The life-cycle model describes how individuals should allocate their wealth between consumption, savings, and investment throughout their lives. The main idea of the model is as follows. Individuals tend to have relatively low savings and consume on labor income when they are young. During middle age, individuals begin to receive higher salaries and accumulate wealth. Subsequently, individuals consume from accumulated wealth as well as pension income after retirement. We focus solely on the post-retirement period and study the allocation policies of retired individuals when facing health risk. We introduce two models in this thesis.

1. **Benchmark model**: Individuals maximize their utility by choosing consumption corresponding to health risk and saving at the fixed risk-free rate.

2. **Model with Portfolio Choice**: Extending the benchmark model by incorporating the risky asset. Hence, individuals could choose the consumption and optimal risky share when allocating their savings.

3.2 Utility Function

We follow De Nardi et al. (2010) to model the retiree's flow utility as a combination of consumption level and health status:

$$u(c_t, h_t) = \delta(h_t)\frac{c_t^{1-\gamma}}{1-\gamma}$$

γ represents the coefficient of relative risk aversion, whereas c_t refers to the consumption of the retiree in period t. The utility function is based on the Constant Relative Risk Aversion (CRRA) specification, which captures diminishing marginal utility of consumption and risk preferences through the parameter γ. To incorporate the role of health in shaping utility, our model introduces a multiplicative weighting function $\delta(h_t)$. The specification keeps the standard CRRA structure while extending it to account for health-related variation in utility.

The Dependence of utility on health status is as follows.

$$\delta(h_t) = 1 + \delta h_t$$

The function models the direct effect of health status and consumption on the utility. δ is the health risk factor. A positive δ indicates that better health increases the marginal utility of consumption, leading healthier individuals to consume more. Conversely, a negative δ indicates that better health status reduces the marginal utility of consumption; thus, a healthier individual would consume less and save more. h_t refers to the health status of the retiree.

3.3 Budget constraints

The risk-free asset has a constant rate of return $(1 + r_f)$. Following Yogo (2016), the risky asset has a stochastic gross rate of return

$$R_{s,t+1} = \overline{R_s}\varepsilon_{s,t+1},$$

where $\log(\varepsilon_{s,t+1}) \sim \mathcal{N}\left(-\frac{\sigma_s^2}{2}, \sigma_s^2\right)$ is independently and identically distributed.

The retiree receives pension income Y every year from Social Security. Let x_t denote cash-on-hand, which is the total cash amount of the retiree in period t. The retiree consumes each period from the cash-on-hand, and the rest goes into the savings.

The intertemporal budget constraint for **Benchmark Model** is:

$$x_{t+1} = (1 + r_f)a_t + Y$$

In **Benchmark Model**, only investment in risk-free assets, such as treasury bonds, is allowed. r_f is the risk-free rate of return. a_t is the saving in period t. The rationale behind the function is that, the cash-on-hand level in period $t+1$ should equal the sum of the savings in period t, the risk-free return on savings in period t, and the pension income Y.

The intertemporal budget constraint for **Model with Portfolio Choice** is:

$$x_{t+1} = (1 - \alpha_t)(1 + r_f)a_t + \alpha_t R_{s,t+1} a_t + Y$$

In **Model with Portfolio Choice**, the retiree could invest in the risk-free asset as well as the risky asset. Similar to **Benchmark Model**, the core idea is that the cash-on-hand level in period $t+1$ should equal to the sum of the savings in period t, investment return on savings in period t, as well as the pension income Y. Since the retiree now has more investment options, she could invest a portion of the savings into the risky asset and the rest into the risk-free asset. Let $0 < \alpha_t < 1$ denote the percentage of the savings invested in the risky asset. Individuals cannot short the risky asset. Similarly, $0 < 1 - \alpha_t < 1$ represents the percentage of savings invested in the risk-free asset.

3.4 Objective Function

The objective function is defined recursively as

$$V(t,x_t,Y,h_t) = max\{\delta(h_t)\frac{c_t^{1-\gamma}}{1-\gamma} + \beta E_t(\mathbb{1}_{\{h_{t+1}\neq 0\}}V(t+1,x_{t+1},Y,h_{t+1}) + \mathbb{1}_{\{h_{t+1}=0\}}0)\}$$

where $\mathbb{1}_{\{h_{t+1}=0\}}$ denotes an indicator function which equals zero if a retiree dies in period $(t+1)$. And β is the subjective discount factor.

It is also commonly referred to as the value function. The value function of period t consists of the utility from period t and the expected value function of next period $t+1$. The first-order conditions and the corresponding Euler equations are derived from the Bellman equation and presented in Appendix A.

Chapter 4

Calibration

4.1 Data Source

4.1.1 Overview of the HRS

The Health and Retirement Study (HRS) is a longitudinal survey conducted by the University of Michigan. The HRS provides comprehensive data on the health, economic, and social conditions of individuals over 50 years old and their spouses in the United States. Participants are interviewed biennially starting from the year of 1992 until their death.

Since the HRS provides comprehensive data on health conditions and wealth status, it is particularly suitable for research on life-cycle behavior under health risk. To the best of our knowledge, most of the research papers in this area in the context of the United States use the HRS as their data source.

In our thesis, the HRS is utilized for the following key purposes:

1. **Health Status Transition**: The HRS provides self-reported health of interviewees. We use the data to estimate transition matrices of health status over time, which is critical for modeling health shocks in the life-cycle framework. Additionally, we embed the mortality risk into the health transition matrices, making the probability of death an endogenous factor in our model.

2. **Income and Wealth Distribution**: The HRS provides observations on retirement income (e.g., Social Security, pensions) and wealth status across different age groups. We use the data to generate the initial pension income and wealth distribution.

Given its unique combination of health, financial, and demographic data, the HRS serves as an essential empirical foundation for constructing and calibrating the retirement life-cycle model developed in our thesis.

4.2 Sample Description

Our sample includes primary respondents aged 67 and above, who were single (including widowed or divorced), and retired at the time of interview. All respondents in the sample have both positive income and wealth.

We focus on individuals aged 67 and above, as this corresponds to the legal retirement age in the United States. We select single respondents because the HRS dataset reports wealth at the household level, making it difficult to attribute individual wealth to married respondents. Equal division of household wealth by the number of household members may introduce bias into the model. Besides, respondents who were previously excluded since married may enter the sample as they are widowed or divorced.

4.3 Model Inputs

4.3.1 Health Transition Probabilities

The HRS provides the self-reported health status of each interviewee on a biennial basis, spanning from 1992, when the first interview was conducted, to the year of the interviewee's death. In the HRS, the self-reported health status is divided into 6 levels. A larger h_t^* represents a healthier status. If an interviewee becomes unreachable, the survey team will contact the interviewee's family members to check their situation.

$$h_t^* = \begin{cases} 0 & \text{Dead} \\ 1 & \text{Poor} \\ 2 & \text{Fair} \\ 3 & \text{Good} \\ 4 & \text{Very good} \\ 5 & \text{Excellent} \end{cases}$$

We collect the self-reported data of interviewees who meet our criteria in section 4.2 and cluster the data into different health status and age groups. Then we calculate the corresponding probabilities of the interviewees in each group, changing from the current health status to all 6 health statuses in the next period. For example, if the retiree's health status is 'Good' and her age is 85, according to **figure 4.1 (a)**, the probability of becoming 'Fair' in 87 is around 30%. Notably, since the interviews were conducted every two years, the interval between the periods in our model is also two years.

Figure 4.1 shows that if the retiree starts with a 'Poor' or 'Fair' health, in her early age she has a much higher probability to recover from it and become healthy again, while in her older days she faces great risk to stay in poor health condition, even die. A similar situation happens to a retiree with good health conditions. A retiree who starts with good health conditions is likely to stay healthy in her early years. However, in her later years, she faces a significantly higher probability of health deterioration.

4.3.2 Initial Wealth Distribution

We use the empirical data from the HRS survey to assign the initial wealth in our model. To simulate the initial wealth situation for individuals who have just retired

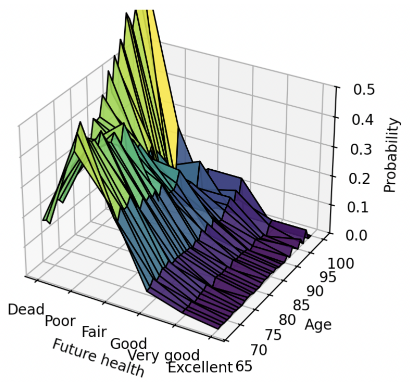

(a) Present health = 1 (Poor)

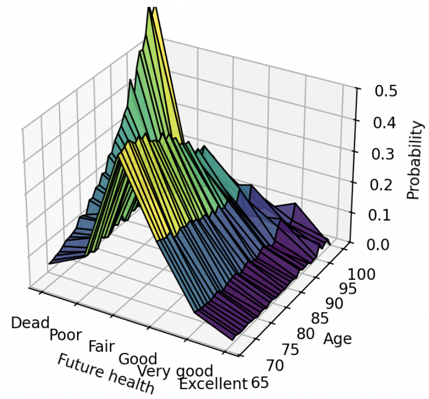

(b) Present health = 2 (Fair)

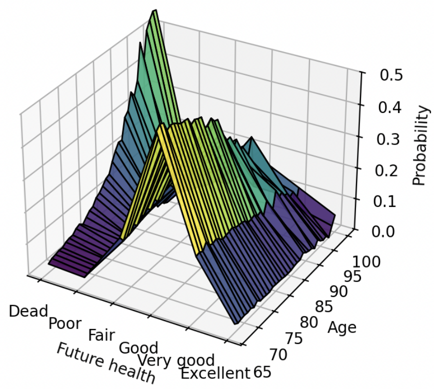

(c) Present health = 3 (Good)

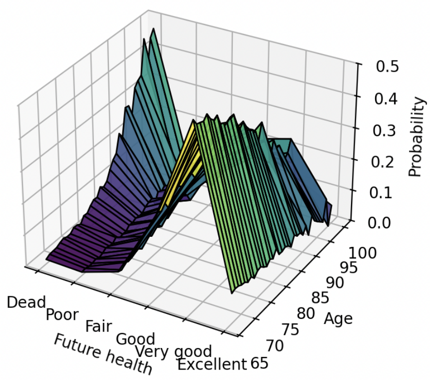

(d) Present health = 4 (Very Good)

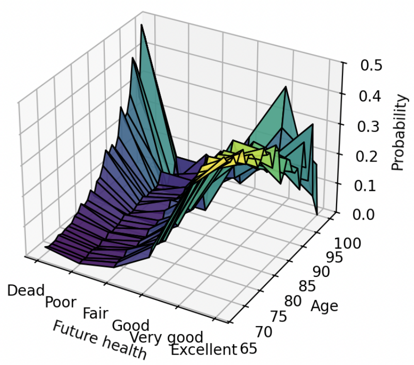

(e) Present health = 5 (Excellent)

Figure 4.1: Overall caption for five health transition matrices of different present health status.

from their jobs, we collect the wealth data of respondents at the age of 67 or 68 years old. Since we aim to study the effect of health on the life-cycle model, we divide the respondents into 5 health groups according to their self-reported health status and assign the initial wealth based on the corresponding wealth distribution separately.

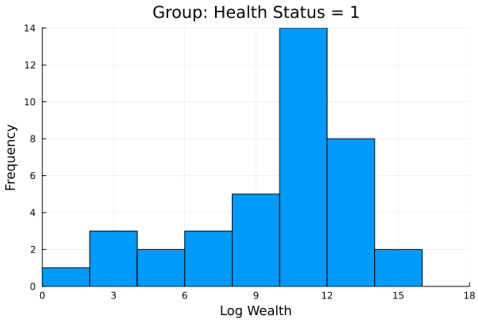

(a) Present health = 1 (Poor)

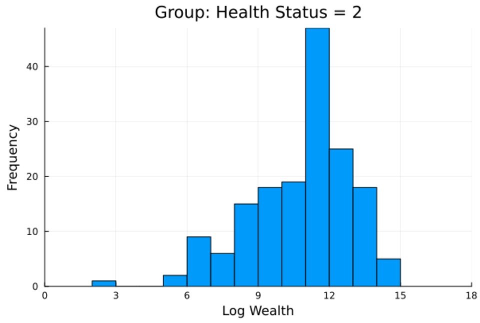

(b) Present health = 2 (Fair)

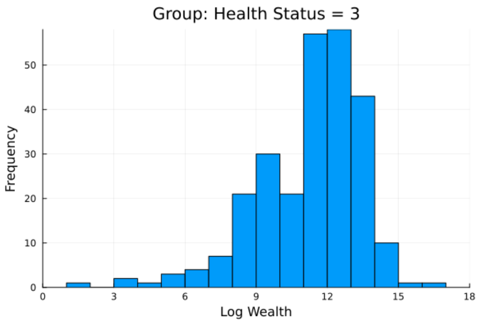

(c) Present health = 3 (Good)

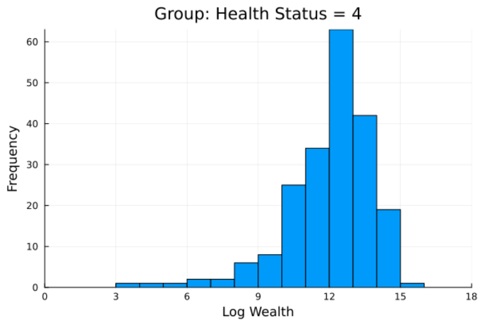

(d) Present health = 4 (Very Good)

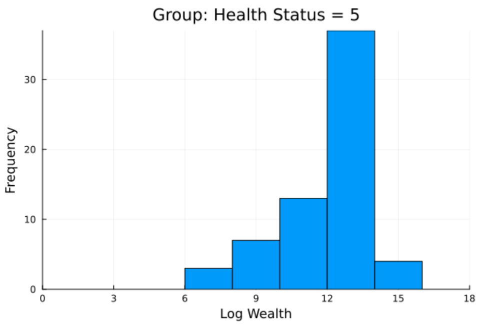

(e) Present health = 5 (Excellent)

Figure 4.2: Overall caption for wealth distributions of five health groups.

Figure 4.2 shows that each group exhibits distinct distributional characteristics. Group 1 skewed to the left as individuals with poorer health conditions tend to be trapped in worse financial conditions. Group 5 contains no individuals with log wealth below 6, which aligns with our expectations, as wealthier individuals are more likely to enjoy better health conditions. Group 2 and Group 4 exhibit the same patterns with a peak around the log wealth value of 12, frequency gradually declines on both sides of the mode, representing a relatively balanced dispersion of wealth within those health groups. Individuals in Group 3 clustered around the log wealth value of 12, suggesting that individuals with stable health conditions are often in stable financial situations as well. Notably, the sample sizes in Group 1 and Group 5 are relatively small, which is consistent with expectations, as extreme health states tend to be less prevalent in the general population.

4.3.3 Retirement Income

We assign the initial income based on the same overall income distribution from all respondents in the HRS data aged 67 or 68, as income is generally distributed more equally than wealth and is also computationally easier to work with. **Figure 4.3** shows the income distribution.

4.4 Model Simulation and Estimation Strategy

4.4.1 Dynamic Programming of the Model

The life-cycle model is solved numerically using dynamic programming — every individual faces stochastic health transition risks and wishes to maximize their value function from today's utility as well as the expected future value. In our model, every individual is assigned an initial wealth and a constant pension income throughout their retirement life. The health transition process for an individual is a Markov chain in which the health status in the next period depends only on the health status in the current period, without the effects from the earlier periods. Each individual considers the conditional expected future health status based on their most recent health status by taking the probability of transitioning to every health status (including the probability of death) into account. Moreover, one period

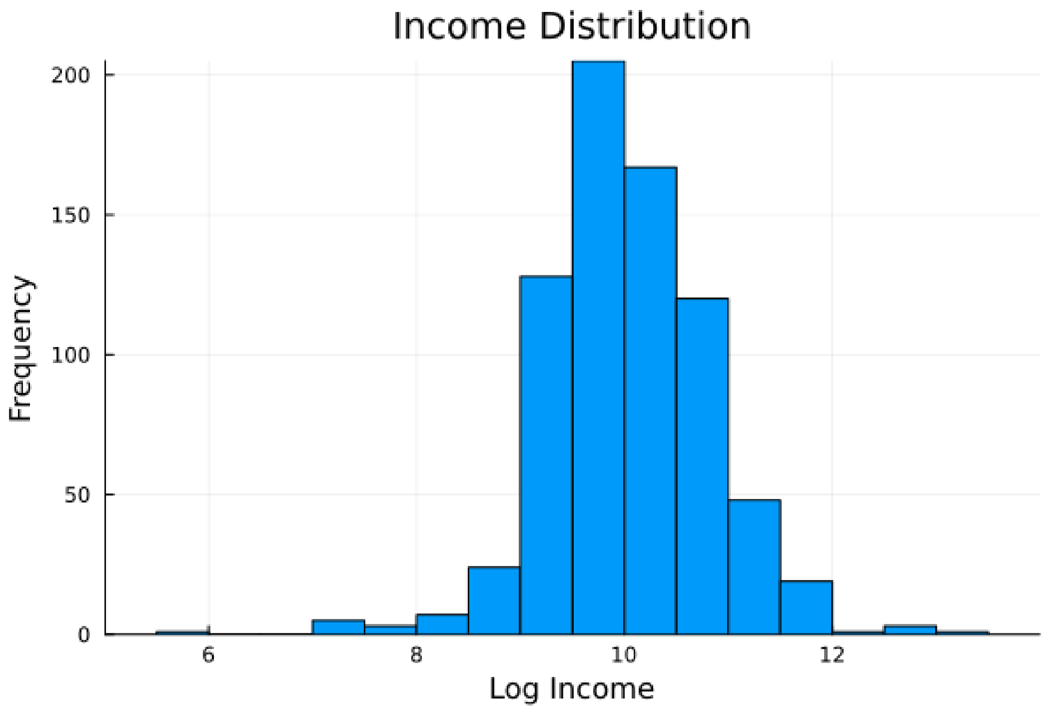

Figure 4.3: Income distribution

corresponds to two calendar years since the HRS data is biennial. We set the initial age of 68 years old, and the maximum age the individual can live up to is 100 years. Individuals make consumption and saving decisions based on their current age, health status, wealth, and income, which is constant over time but varies across individuals. Individuals would exit the model upon death.

In each period, an individual's cash-on-hand consists of the savings from the previous period and the pension income they receive in the current period. Borrowing from the future is not allowed in our model, so individuals must have non-negative savings. The cash-on-hand is either used for consumption or savings in the current period.

Following Carroll (2006), we choose to grid over saving (which is also 'wealth' in empirical data) rather than cash-on-hand to avoid root-finding steps in the solution of the Euler equation, greatly improving computational efficiency. The maximum saving grid is set to one million to cover most situations of wealth in the empirical data, and the number of points on the saving grid is 100. We use linear interpolation to solve for the wealth level that falls between two grid points. For the saving level that exceeds the upper grid (i.e., the super-rich individuals), the policy function is extrapolated linearly. Based on classic academic literature (e.g., Campbell and Mankiw (1989)), it is widely accepted that there is a considerable proportion of individuals who consume all the money they have without saving anything for the next period. Therefore, to capture the 'eat-everything' behavior of the poor individuals, we increase the accuracy of the policy function by adding fifty grids between the lowest cash-on-hand level corresponding to minimal saving and the lowest income level in case the lowest saving grid is too high for a relatively substantial fraction of individuals.

The policy function is solved backwards, as it is difficult to solve the stochastic process forward. Solving forward requires estimating the future value function and making an optimal decision simultaneously when facing uncertainty to calculate the expectation on the right-hand side of the Euler equation. By using backwards induction, the consumption in the next period is known, and therefore, it is more convenient to solve for the previous state when the expectation is known.

In the last period, the optimal policy for an individual is to consume everything and save nothing since it is assumed that there is no bequest motive in our model. After solving the last period for each saving grid, the consumption between the two periods can be linked using the Euler equation as mentioned in Appendix A. Therefore, by solving backwards, the policy function can be derived for the individual at any cash-on-hand level, conditional on different retirement periods of life, health status, and income level.

For the model with portfolio choice, the optimal share is retrieved by maximizing the expected future value function given the saving level. Shorting the risky assets is not allowed in our model hence the optimal portfolio share is between zero and one and is extrapolated by holding its value fixed at the boundary.

4.4.2 Simulation Procedure

We simulated the life-cycle path for 100,000 individuals. At the initial state, each individual is assigned an initial health status that reflects the empirical distribution of initial health statuses in HRS data. In addition, individuals' pension income follows the log normal distribution with mean and standard deviation calibrated to the HRS data.

Given the initial health distribution, individuals are then assigned the initial wealth based on the log distributions with different means and standard deviations under different health groups according to the empirical data. The wealth distributions vary across different health groups at age 68.

When simulating the subsequent health procedure, the individual would face different health transition probabilities that vary by different periods and current health statuses, as described in Section 4.3.1. In general, individuals face declining health conditions, which implies that older individuals are less healthy and might face a higher likelihood of health deterioration.

Following the setting in Yogo (2016), there is one simulated portfolio gross rate of return for all individuals where the return is calculated as the product of a constant and a log-normally distributed error term.

For each period, the cash-on-hand of an individual is calculated by adding savings from the previous period and pension income received in this period. Then, the optimal consumption and saving policy corresponding to a certain cash-on-hand level can be determined based on the policy function derived in Section 4.4.1.

4.4.3 The Method of Simulated Moments

The model is estimated following the approach from De Nardi et al. (2010). First, we calibrated parameters $(\beta, \gamma, r_f, \bar{R}_s, \sigma_s)$, which can be identified directly from empirical data or follow the commonly used values from existing academic literature. Second, the Method of Simulated Moments is adopted to estimate the health risk factor δ.

In the first step, we obtain the constant risk-free rate r_f and the stochastic gross return process of the risky asset following the assumption in Yogo (2016). The annual risk-free rate $r_f = 2.5\%$ is calibrated as the average real return on the one-year Treasury bond, adjusted for inflation using the CPI (consumer price index) from 1958 to 2008. In addition, the stochastic gross annual rate of return on the risky asset is calibrated as the product of a constant $(\bar{R}_s)$ and an error term $(\varepsilon_{s,t})$ which follows an independent and identical distribution $\log(\varepsilon_{s,t}) \sim \mathcal{N}\left(-\frac{\sigma_s^2}{2}, \sigma_s^2\right)$. These two parameters are calibrated as $\bar{R}_s = 1.065$ and $\sigma_s = 0.18$ using the gross real return on the Center for Research in Securities Price value-weighted stock index, adjusted for inflation using the CPI from 1958 to 2008. With this calibration, the annual equity premium for our simulated risky asset is 4%, which is slightly lower than the historical estimation of 4.5% but follows the widely accepted assumption in the existing life-cycle literature (Cocco et al., 2005). Moreover, the subjective annual discount factor $\beta = 0.96$ and the risk aversion parameter $\gamma = 5$ are also calibrated following Yogo (2016), and both values fall within the range of classical economics literature.

In the second step, we follow a standard Method of Simulated Moments (MSM) approach to estimate the health risk factor δ, conditional on the calibrated parameters in the first step. The purpose of the MSM is to minimize the distance metric between the moments from the empirical and simulated data by changing the model parameter (i.e., the health risk factor δ). We start by assuming each moment in the model has an identical weight when estimating the parameter, from which a first-stage estimated δ can be obtained. In addition, due to the instability of stochastic simulation, the optimal weighting matrix should be computed over multiple simulation runs to mitigate the impact of unstable moments. For example, the number of individuals in the excellent health group would be low when individuals are old, hence the simulated outcomes could have substantial variation across replications due to a relatively small subsample size. The optimal weighting matrix is computed as the inverse of the covariance across multiple simulation draws. Subsequently, we minimize the weighted distance metric to obtain the ”best match” between the simulated and empirical moments and derive the second-stage estimated δ as the final result. By adjusting using the optimal weighting matrix, greater weights would be assigned to the moments from the stable group, hence the impact of instability from stochastic simulation could be diminished.

4.4.4 Targeted Moments

To calibrate the health risk factor δ, we match the simulated moments with the empirical medians of wealth-to-income ratios at ages 70, 76, and 82 across five health groups. This choice reflects one of our key motivations: to understand how health status affects saving behavior and wealth accumulation, as well as how transitions across health states influence consumption decisions in the post-retirement period. Furthermore, we choose to use medians instead of means to reduce sensitivity to outliers and potential small subsample sizes. Moments from very old individuals are not selected as targets since the samples might be too small, and the fact that rich people live longer (De Nardi et al., 2009), also the bequest motive from the super-rich (De Nardi, 2004) could skew the data.

Chapter 5

Results

5.1 Estimation Results

The estimated result of the health risk factor δ is -0.18. The negative sign of the health risk factor indicates a consistent direction of the effect of health on consumption, as other existing literature (e.g., De Nardi et al. (2010)). An anticipated change from excellent health status (where $h_t = 5$) to a very good one ($h_t = 4$) would increase the optimal consumption by 22.87%, whereas a change from fair health status (where $h_t = 2$) to a poor one ($h_t = 1$) would increase the optimal consumption by 5.08%.

When the consumption level is held fixed, being in a better health condition lowers the marginal utility of consumption, hence, healthier individuals would choose to consume less and save more. In contrast, individuals with poorer health would consume at a higher level as a means of compensating for the disutility associated with their bad health. An alternative interpretation is that the higher consumption observed among less healthy individuals may serve as a proxy for medical expenditure. In this case, the negative δ is reasonable because individuals in poor health may be forced to allocate more resources to meet their medical needs. These two mechanisms may jointly contribute to the observed consumption pattern.

5.2 Policy Functions

5.2.1 Benchmark Model: Health Risk Only

5.2.1.1 Consumption Policy

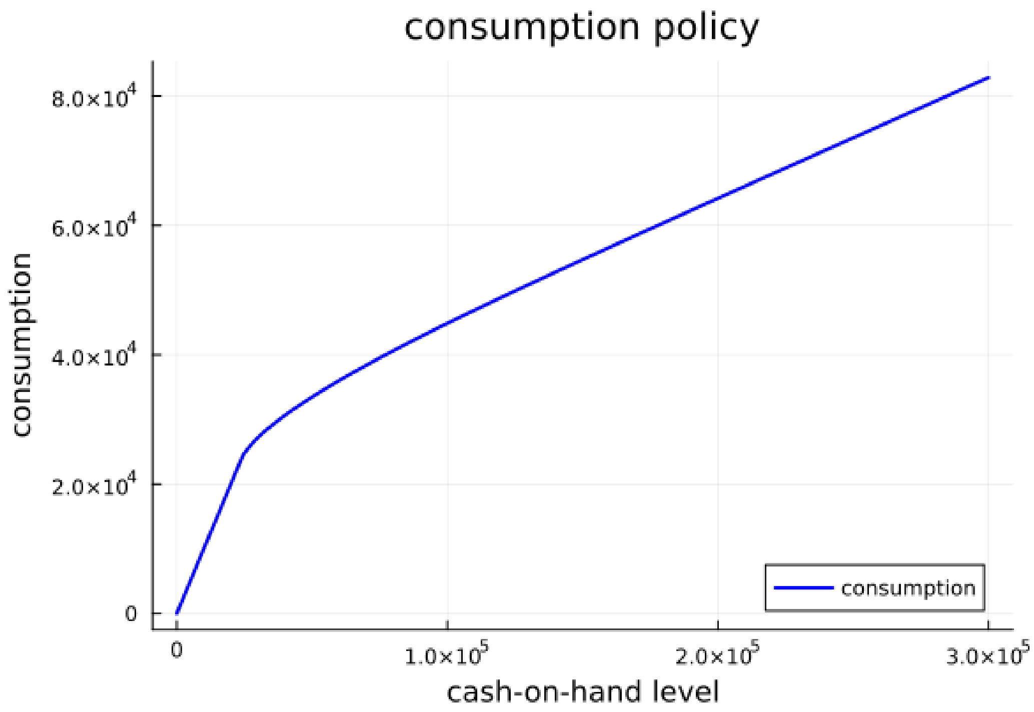

Figure 5.1: Consumption Policy

Figure 5.1 shows the optimal consumption policy for selected age, income level, and health status. Apart from the imposed hand-to-mouth behavior at low cash-on-hand level, the optimal consumption is an increasing and concave function. The shape of the consumption policy reflects that consumption does not grow proportionally as wealth increases. Taking into account other variables, individuals with a high level of cash on hand tend to save at an even higher rate and smooth consumption during their useful life in response to future uncertainty.

Figure 5.2 compares different consumption policies for individuals with the same health status and income level at the ages of 70 and 86. The optimal consumption level is lower for 70-year-old individuals. Theoretically, the expectation of health status of 86-year-old individuals is significantly lower than that of 70-year-old

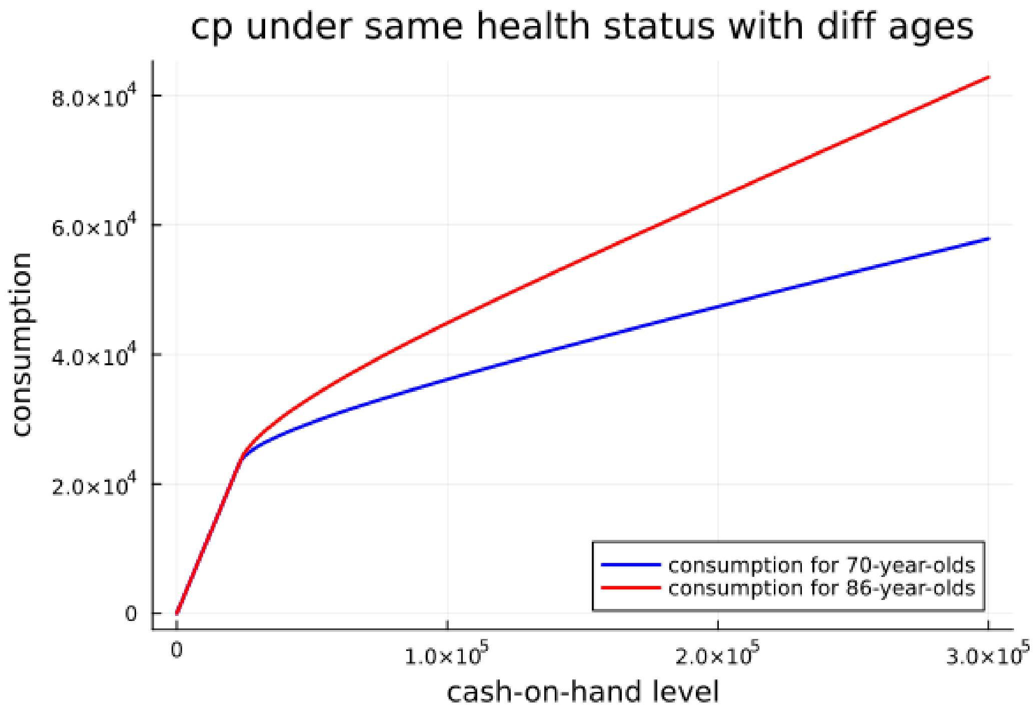

Figure 5.2: Consumption Policy for Different Ages

individuals. Based on the Euler equations, the anticipated change from good to bad health status should increase the consumption. Additionally, since our model generates an endogenous mortality rate linked to the health expectation, the 70-year-olds have a longer life expectancy than the 86-year-olds, holding other factors constant. Without gaining any utility from leaving a bequest to children, individuals tend to use up all (at least spend as much as possible) of their wealth within their limited lives. Therefore, it is reasonable for older individuals to consume more while holding the same cash-on-hand level.

Figure 5.3 illustrates the comparison between different health groups. With the same age and income level, healthier individuals would prefer to consume less and save more for the future. The mechanism behind the gap is similar to the rule we discussed when comparing the consumption policy for different age groups. In reality, a healthier retiree has a higher probability of maintaining good health. Then, according to the Euler equation, minor changes in potential health conditions would not move the optimal consumption significantly. Besides, the healthier retiree has a

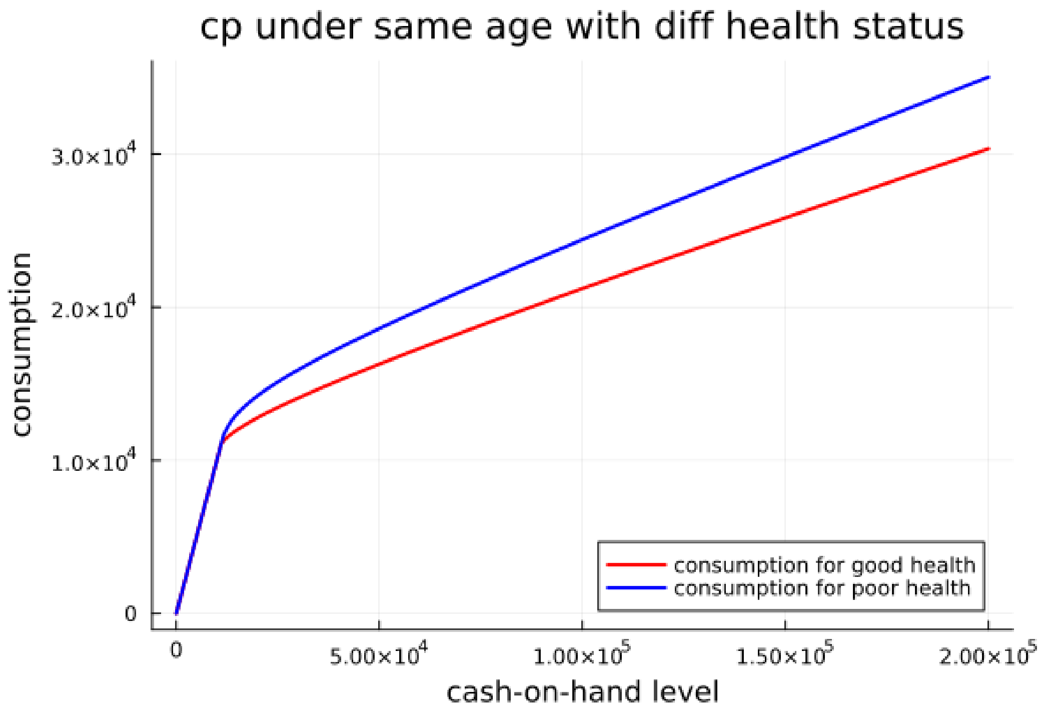

Figure 5.3: Consumption Policy for Different Health Groups

longer life expectancy than the unhealthier one, and a smoother consumption along the lifespan would be expected. Thus, the individual with a poorer health condition exhibits a higher consumption level within a shorter remaining lifespan, holding other factors the same.

5.2.1.2 Saving Policy

Figure 5.4 presents the optimal saving policy for selected ages, income levels, and health status. Similar to the consumption policy above, apart from the imposed hand-to-mouth behavior at low cash-on-hand levels, the optimal saving is an increasing function and convex since saving is defined as the difference between cash-on-hand and consumption. Individuals with a high cash-on-hand level tend to increase the rate of saving in response to future uncertainty.

Figure 5.5 reports the comparisons of saving policies between different ages and different health groups separately. We can draw the same conclusions as those obtained from analyzing the consumption policy that younger and healthier

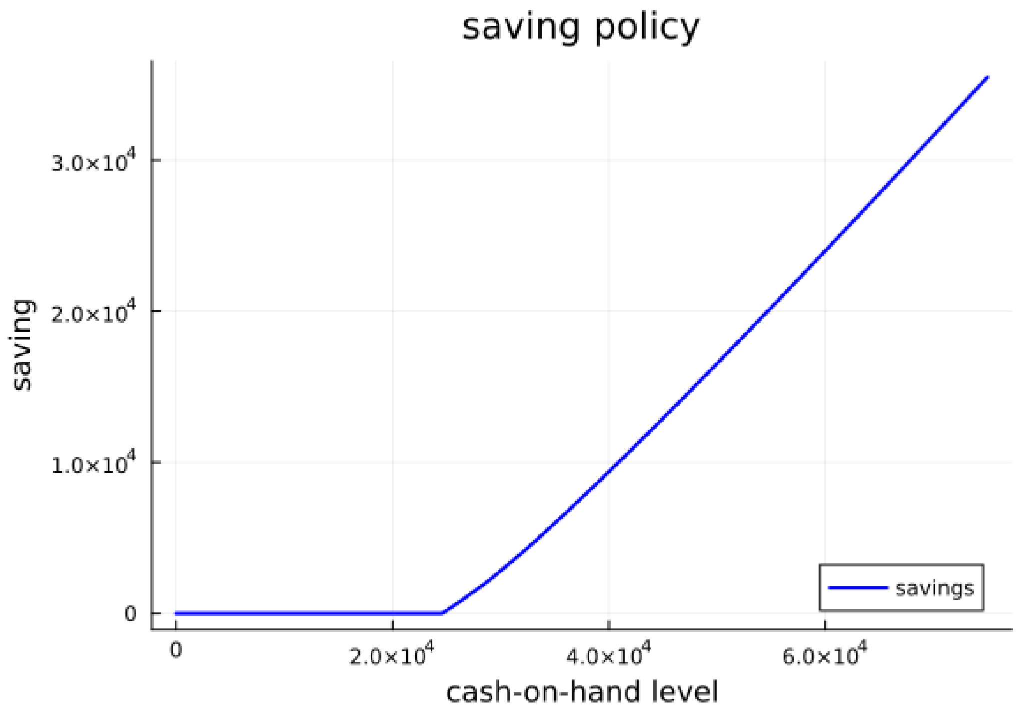

Figure 5.4: Saving Policy

individuals have stronger incentives to save at a higher rate, considering their longer life expectancy and lower risks of health deterioration.

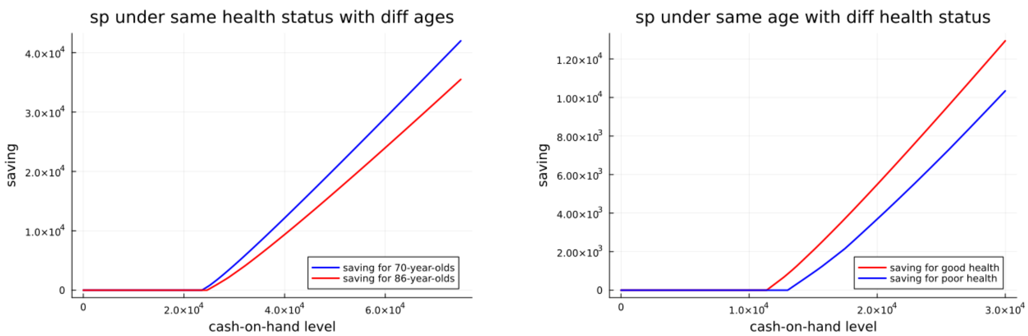

(a) Different Ages

(b) Different Health status

Figure 5.5: Comparison of Saving Policy

5.2.1.3 Value Function

As shown in Figure 5.6, the value function is an increasing and concave function of the cash-on-hand level. The concavity demonstrates that individuals are risk-averse when facing uncertainty and would be happier with a certain average cash-on-hand

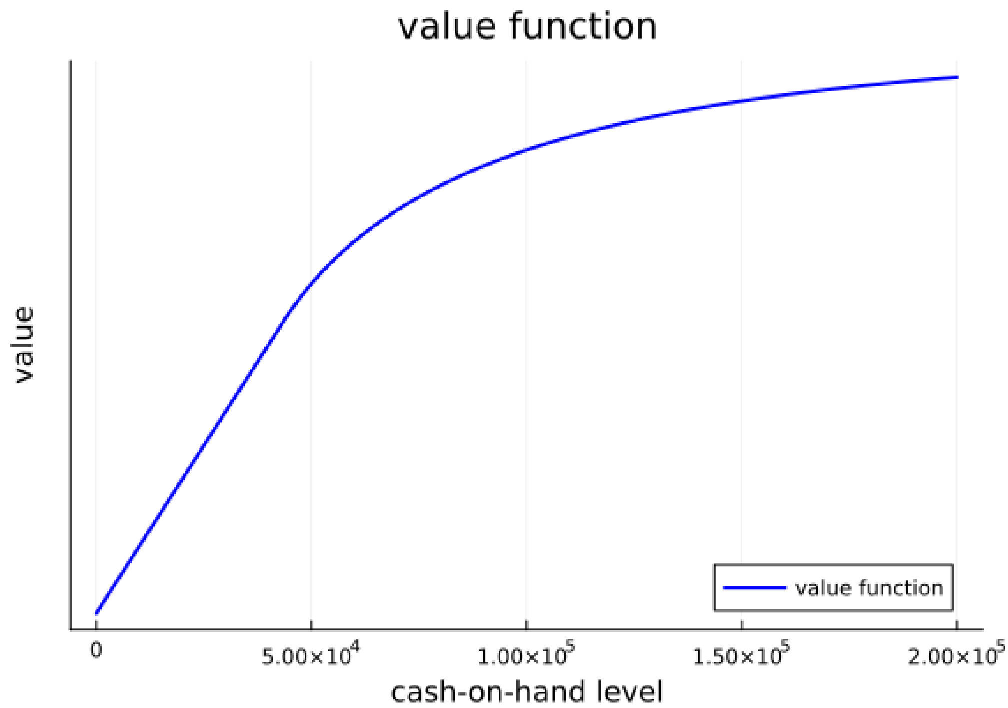

Figure 5.6: Value Function

rather than random utility from it. The properties of the value function are in line with results observed from empirical analyses (e.g., Carroll and Samwick (1997)).

According to the properties of the CRRA (Constant Relative Risk Aversion) utility function, individuals exhibit positive prudence, that is, for a given level of risk, there is a decreasing response as the consumption level rises. In other words, richer individuals show lower sensitivity when facing the same level of risk.

5.2.2 Model with Portfolio Choice

This model extends the benchmark model by introducing a risky asset. Individuals can choose the optimal share to invest in the risky asset and maximize the expected lifetime utility. The consumption, saving, and value functions in the portfolio choice model are very similar to the ones in the benchmark model, and the mechanisms behind them are the same. For a detailed illustration, refer to Appendix B.

5.2.2.1 Optimal Portfolio Share Policy

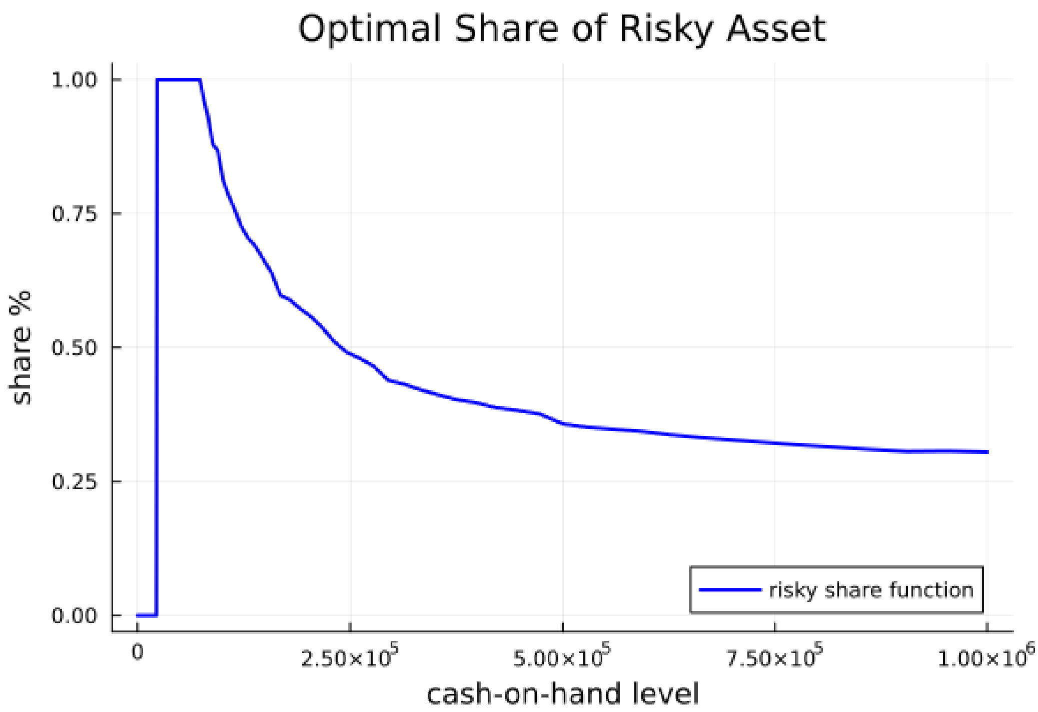

Figure 5.7: Optimal Share of Risky Asset

The optimal share of the risky asset invested in the next year as a percentage of wealth is presented in Figure 5.7, where the optimal risky share decreases with wealth. Apart from the hand-to-mouth individuals forced to save nothing in the risky asset, poor individuals would invest more in the risky asset for a higher expected return rate. Our result is in line with Merton's solution (Merton et al., 1971):

$$\frac{\alpha_t W_t}{W_t + PDV_t(FY_t)} = \frac{\mu}{\gamma \sigma_\eta^2}$$

α_t is the optimal risky share, W_t represents the financial wealth, $PDV_t(FY_t)$ is the human capital — calculated as the present value of the future income, μ and σ_η indicates the expected excess return and variance of the risky asset, and γ is the risk-aversion level of individuals. In Figure 5.7, μ, σ_η, and γ are the same across individuals. Since we plot the optimal share of the risky asset in a group with the same age and health status, the expected remaining lifetime is identical across the group. Consequently, the human capital part $PDV_t(FY_t)$ is a constant within the group. Therefore, wealthier individuals are holding less risky shares.

An intuitive interpretation of this mechanism would be to consider the future pension income as the substitution of the risk-free asset. Thus poor individuals would adopt a more aggressive strategy on risky assets because the ”risk-free asset” already constitutes a large share of their total wealth, and they are looking for a higher future return. Whereas pension income plays a less significant role for richer individuals among those who have the same income level. When wealth is sufficiently high, the optimal proportion of the risky asset in our model is slightly higher than the complete market solution derived by Merton (1969) and Samuelson (1975):

$$\alpha = \frac{\mu}{\gamma\sigma_\eta^2} = \frac{4\%}{5.0 \cdot 0.0374} = 21.4\%$$

A detailed derivation of σ_η^2 is provided in Appendix C.

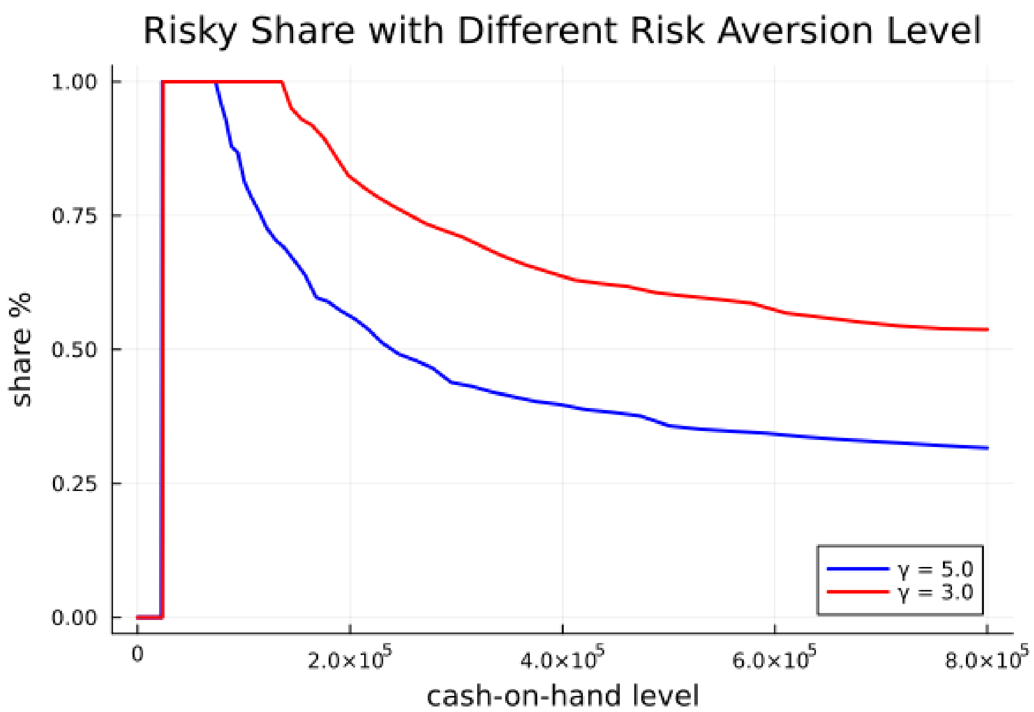

Figure 5.8: Optimal Share Comparison for Different Levels of Risk Aversion

As shown in Figure 5.8, individuals with a lower risk-averse level invest more in the risky asset since they pay less attention to the risk and prefer a higher expected future return. This result is also in line with Merton's solution. The group with low risk aversion serves here for comparison purposes only; we will not examine it in depth in the subsequent discussion.

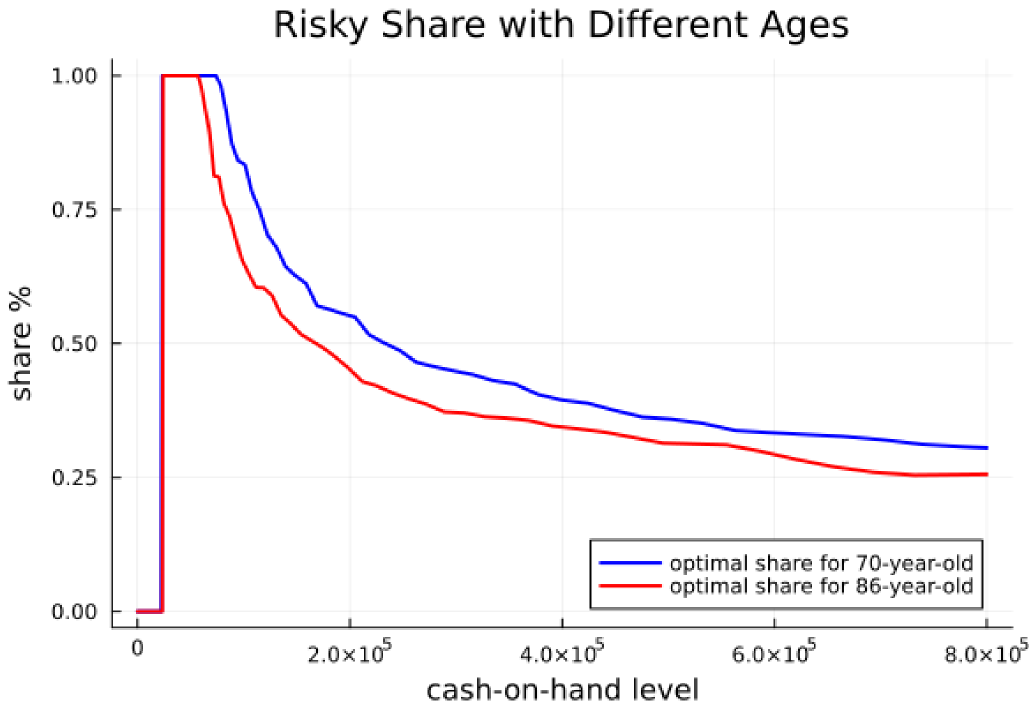

Figure 5.9: Optimal Share for Different Ages

Figure 5.9 compares the difference in optimal share for two age groups. Younger individuals always hold more share as percentages of wealth compared to older ones at the same cash-on-hand level. Again, this result aligns with the optimal portfolio rule since the human capital is higher for younger individuals. It is also consistent with the findings of Cocco et al. (2005), who show that the share of risky assets declines with age in a rather simple post-retirement setting. While Yogo (2016) demonstrates that the optimal shares of the risky asset across ages are relatively stable, the discrepancy may be due to the richer set of investment options available in his model, where the wealth can be allocated between bond, stock, and housing. When we focus solely on the risk-free asset, our results are consistent with Yogo's finding that the share of the risk-free asset increases as individuals age.

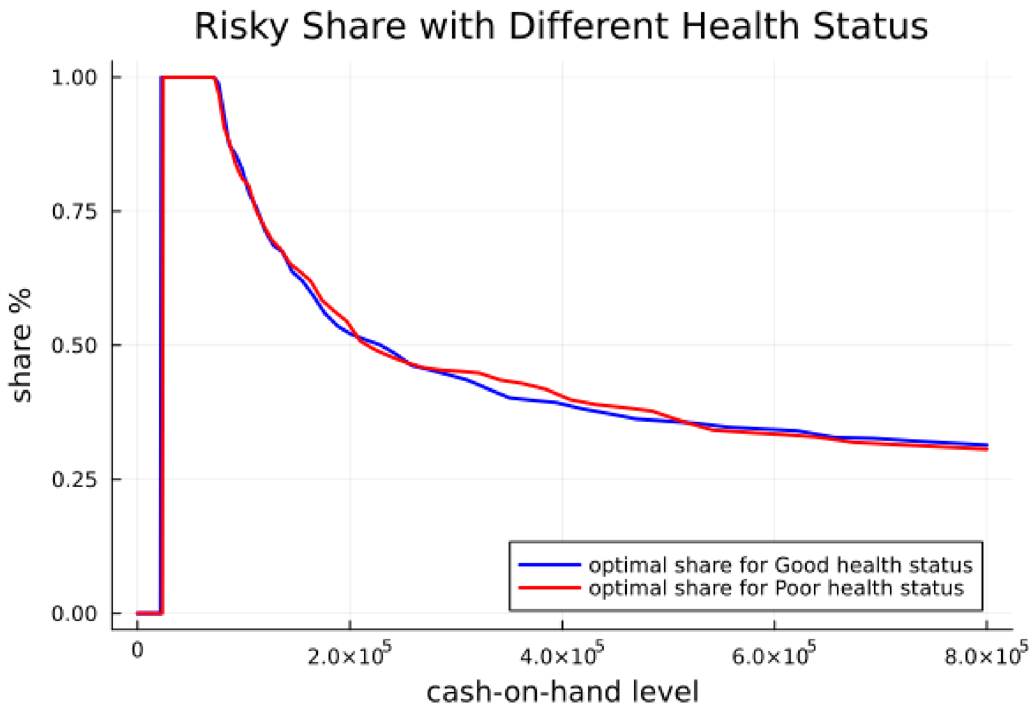

Figure 5.10: Optimal Share for Different Health Status

According to Figure 5.10, the difference in optimal risky share between the two health groups is marginal, and the corresponding difference in the risk-free asset share is similarly small, which is consistent with Yogo (2016). The optimal shares for the two health groups converge to the same value as wealth increases.

5.2.3 Comparison of Policy Functions from Different Models

In this section, we compare the difference between policy functions of the benchmark model with health risk only and the model with optimal portfolio choice, within the 68-year-old individuals who have a Good health status ($h_t = 3$) and a median level of income.

5.2.3.1 Consumption Policy Comparison

As shown in Figure 5.11 (a), the optimal consumption in the benchmark model is slightly lower than that in the optimal portfolio choice model. Since individuals would expect a higher rate of return on wealth when the risky asset is involved in the model, it allows a higher consumption and less saving to achieve the same level of expected future resources. Figure 5.11 (b) clearly reports the difference in consumption levels increases with wealth.

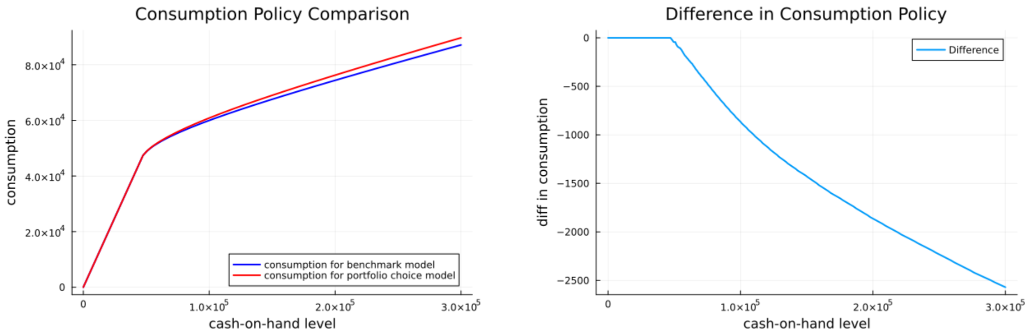

(a) Consumption Policy Comparison **(b)** Difference

Figure 5.11: Consumption Policy Comparison for Different Models

5.2.3.2 Saving Policy Comparison

Figure 5.12 (a) demonstrates the result of the comparison between saving policies for different models. Saving is the opposing choice of consumption, so individuals would save more in the benchmark model compared to the model with portfolio choice which requires less saving to maintain equivalent expected future resources. The group of individuals we choose to plot are rather healthy and in the very early stage of retirement; thus, they save at a relatively high proportion of wealth while consuming less than they save. As a result, the two saving policy lines nearly overlap. The difference between the two saving policy functions is clearly depicted in Figure 5.12 (b).

5.2.3.3 Value Function Comparison

Figure 5.13 compares the value functions under two different models. The value function obtained from the benchmark model is slightly lower than that from the model with portfolio choice, especially when the cash-on-hand level is high, indicating that allowing individuals' participation in the stock market would increase the utility in retirement life. The relative shape of the value function

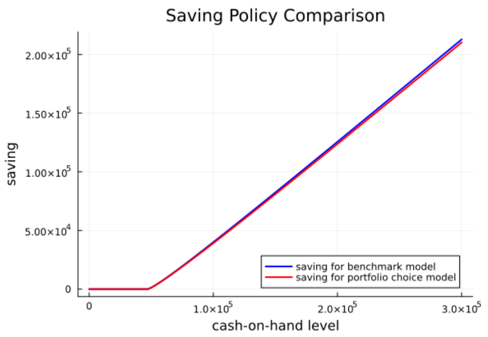

(a) Saving Policy Comparison

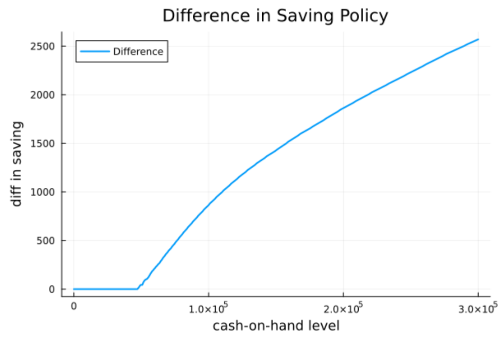

(b) Difference

Figure 5.12: Saving Policy Comparison for Different Models

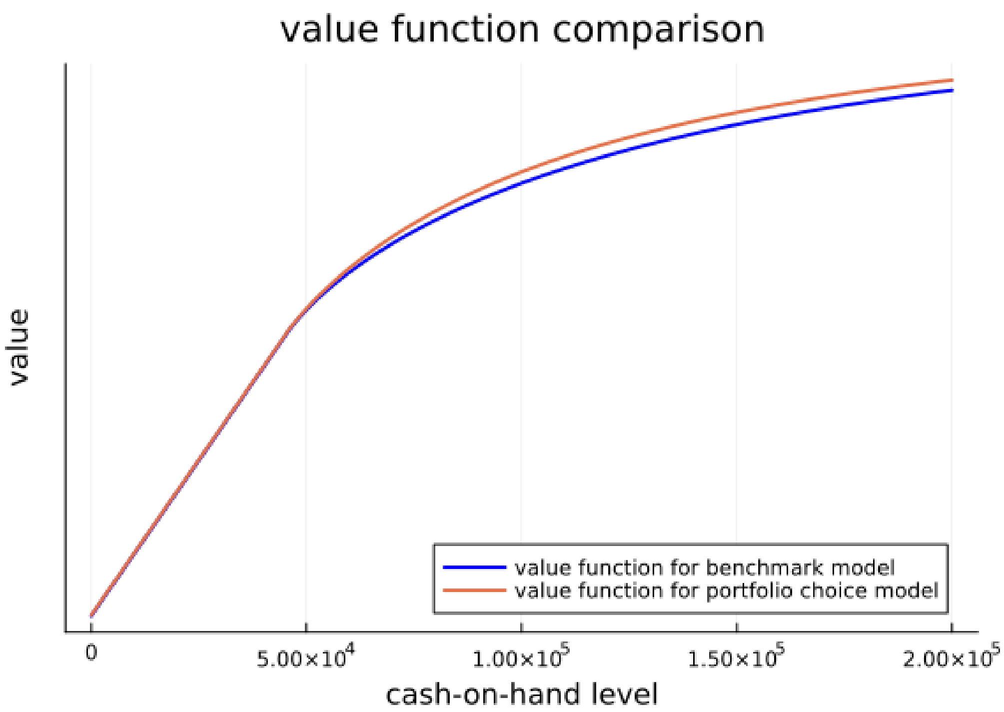

Figure 5.13: Difference in Value Function

might reflect more conservative behavior in the benchmark model, such as lower consumption and higher savings due to the absence of allocation flexibility, as demonstrated in Section 5.2.3.1 and Section 5.2.3.2. Individuals in the benchmark model could be regarded as those who choose not to participate in the stock market even when participation is allowed.

5.3 Comparison of Empirical and Simulated Moments

The comparisons of the medians of the wealth-to-income ratio between the simulated model and empirical data under different health statuses are demonstrated in Figure 5.14.

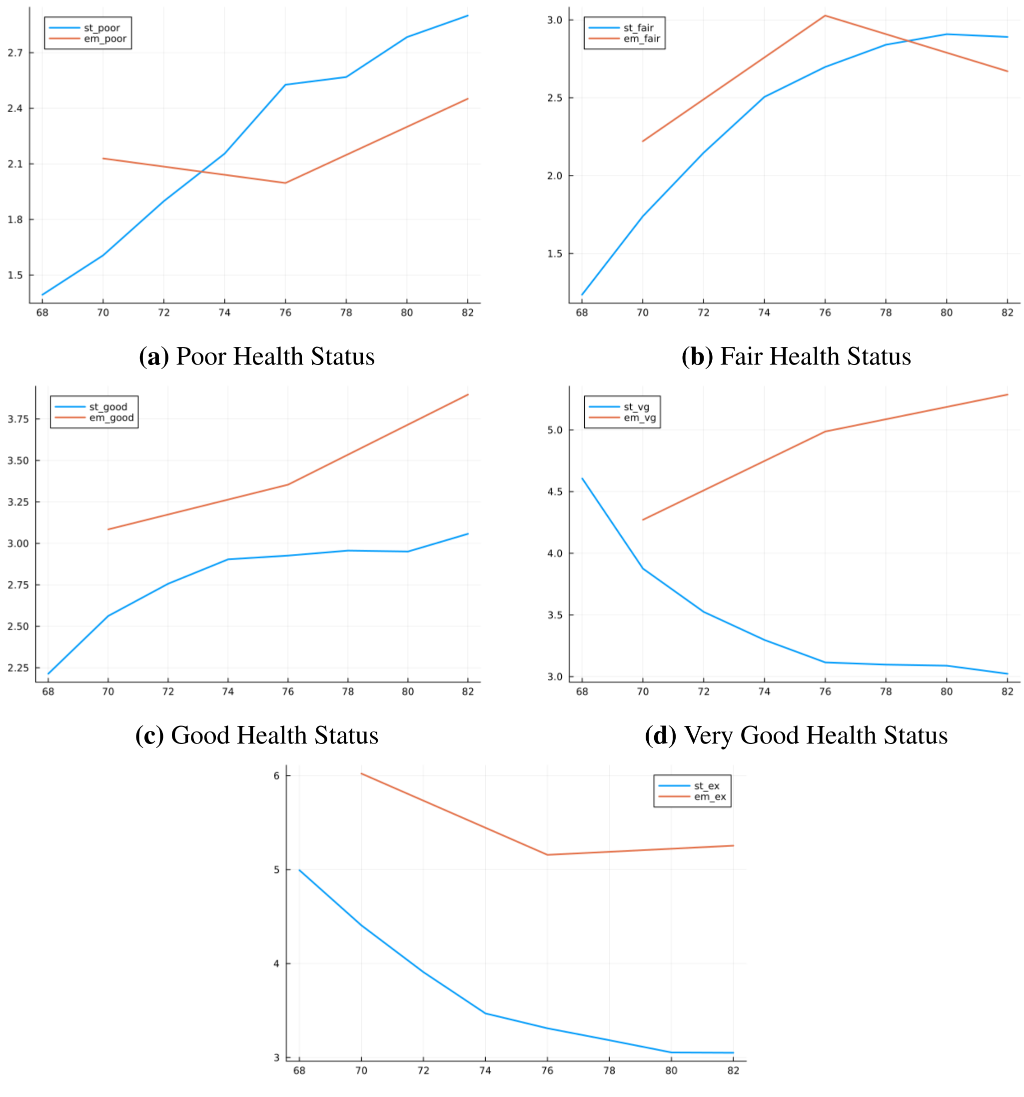

(a) Poor Health Status

(b) Fair Health Status

(c) Good Health Status

(d) Very Good Health Status

(e) Excellent Health Status

Figure 5.14: Comparison of Empirical and Simulated Wealth-to-Income Ratio

In general, our model provides a better fit among individuals who have a relatively worse health status. As shown in Figure 5.14 (a) and (b), the simulated wealth-to-income ratios for the Poor and Fair health groups reasonably match both the trend and magnitude of the empirical data. This relatively accurate fit may reflect the effectiveness of the endogenous mortality rate and health transition process, which capture two key dynamics: the high mortality risk among individuals in poor health, and the gradual health deterioration among initially healthier individuals. These forces jointly influence the composition of the observed less healthy groups over time. In particular, the relatively stable (or even increasing) wealth-to-income ratios in the poorer health groups may be driven not by individual wealth accumulation but by the exit of the less healthy and wealthy individuals through mortality and the inflow of moderately healthy and wealthy individuals whose health has deteriorated. Thus, once the model clarifies the effect of mortality-driven selection and health deterioration processes, the wealth decumulation driven directly by the health risk parameter δ can be more clearly identified.

For the Good health status, the simulated trend aligns reasonably well with the empirical data in the earlier stages, although there is some discrepancy in magnitude, as shown in Figure 5.14 (c). The similarity of trends may result from the dynamics mentioned above, whereas the discrepancy in magnitude could stem from differences in the initial assigned wealth-to-income ratios. Additionally, the empirical wealth-to-income ratio increases more steeply than the simulated one as individuals age. This may indicate that the effect of health transition risk on wealth decumulation is overstated in the model. Thus, the model may lack key incentives for individuals to save apart from the health risk, such as a bequest motive, as in Ameriks et al. (2011) or precautionary saving for future medical expenditures, as in De Nardi et al. (2010).

Our model fails to replicate the empirical data pattern accurately within the healthier groups. As in Figure 5.14 (d), the simulated medians exhibit an opposite direction from the empirical data in the Very Good health status, even under the condition that the two initial wealth-to-income ratios are similar. Figure 5.14 (e) indicates that the wealth decumulation tendencies in the simulated model are similar to the pattern in the empirical data in the Excellent health status. However, there is a substantial difference in the magnitudes of the wealth-to-income ratios. The downward trend in the Excellent health status suggests that the health risk parameter δ allows the model to successfully capture health-driven wealth decumulation. However, in the Excellent health group, the empirical wealth decumulation is slower than in the model, and in the Very Good group, wealth-to-income ratios even increase with age, suggesting that health status may be more persistent than the Markov process assumed in the model. The model may underestimate the strength of precautionary saving or bequest motives among healthier individuals, who tend to live longer and possibly plan for late-life medical costs or intergenerational transfers, as discussed previously in the context of the Good health group.

Potential mechanisms beyond health risk will be further discussed in Section 5.5.

5.4 Solution to the Life-Cycle Model

Table 5.1 reports the predicted values of portfolio share of risky and risk-free assets, wealth level (represented by wealth-to-income), consumption and health distribution in our model.

Panel A shows the proportion of risky assets as a percentage of total wealth. Individuals in relatively poor or good health tend to allocate a smaller share of their portfolio to risky assets. Within the poorer health group, the share of risky assets declines with age and converge to approximately 80% by the age of 92. At earlier ages, as discussed in Section 4.2, health status is closely associated with wealth levels, with less wealthy individuals more likely to be in poorer health, and vice versa. Consequently, individuals in better health may allocate less to risky assets, as pension income represents only a small portion of their total wealth, and they may prefer to diversify risks with other risk-free assets as well. On the other hand, individuals in poor health also tend to reduce their risky asset holdings due to the need to reserve funds for potential future medical expenses. The rationales behind Panel B are similar to those in Panel A.

Panel C presents the wealth level across different health conditions and ages, measured by the wealth-to-income ratio. We use savings in each period to represent the wealth. We follow the approach of Yogo (2016) to use the wealth-to-income ratio instead of absolute wealth values, as the absolute values tend to vary substantially across individuals and are therefore less suitable for cross-group comparisons. Moreover, given that income serves as the primary cash inflow contributing to wealth accumulation in our model, using the wealth-to-income ratio provides a more meaningful and standardized measure. Individuals in better health exhibit significantly higher wealth-to-income ratios at age 68, with the ratio declining as age increases. This pattern aligns with a core assumption of our model: with no bequest motive, individuals aim to deplete their wealth before death to maximize lifetime utility. As individuals age, they consume more of their accumulated wealth, which results in a gradual decline in the wealth-to-income ratio. In contrast, individuals in poorer health share an opposite pattern, which is inconsistent with the model's assumptions. One possible explanation is related to sample composition: at younger ages, the sample size is larger and includes many individuals in very poor health, often associated with low wealth. As age increases, those individuals are more likely to have died, and the remaining population in poor health may consist of previously healthier individuals who have experienced health deterioration, but retain relatively higher wealth. This dynamic may explain why, by age 92, the wealth-to-income ratio converges to approximately 3 across all health groups.

Panel D reports the consumption level as a percentage of cash-on-hand. As discussed in Section 3, cash-on-hand refers to the total spendable wealth available at the beginning of each period, and is therefore used in place of savings as a more relevant measure. Within each health group, the consumption level increases with age, which is consistent with the model's core assumption mentioned above. At any given age, individuals in worse health conditions exhibit higher consumption ratios, which aligns with economic theory that lower-wealth individuals are more likely to allocate a larger share of their resources to basic necessities. In contrast, wealthier individuals are more likely to reserve a greater portion of their wealth in the expectation of potential future health deterioration.

Table 5.1: Solutions to the portfolio choice model

Health status	**Age**				
	68	74	80	86	92
Panel A: Risky Asset (% of total wealth)					
Poor	79.6	79.4	77.0	77.1	77.0
Fair	86.5	80.0	78.4	80.0	78.5
Good	81.9	79.5	79.5	81.5	80.1
Very good	75.9	78.0	80.2	84.4	87.5
Excellent	75.8	77.8	79.3	82.4	89.1
Panel B: Risk-free Asset (% of total wealth)					
Poor	20.4	20.6	23.0	22.9	23.0
Fair	13.5	20.0	21.6	20.0	21.5
Good	18.1	20.5	20.5	18.5	19.9
Very good	24.1	22.0	19.8	15.6	12.5
Excellent	24.2	22.2	20.7	17.6	10.9
Panel C: wealth-to-income					
Poor	1.43	2.19	2.97	1.73	3.07
Fair	1.24	2.50	2.84	2.97	3.10
Good	2.31	2.91	3.01	2.96	3.08
Very good	4.71	3.34	3.01	3.02	2.93
Excellent	5.01	3.72	3.09	3.14	2.94
Panel D: Consumption (% of cash-on-hand)					
Poor	54.2	59.8	66.2	75.9	88.1
Fair	53.5	54.8	61.8	73.8	86.2
Good	43.8	48.4	57.0	69.7	84.3
Very good	31.2	41.9	51.6	64.8	80.6
Excellent	27.2	36.2	44.6	57.5	74.4
Panel E: Simulated Health distribution					
Poor	6.9	8.2	10.8	11.8	14.8
Fair	19.7	20.9	23.7	26.8	28.3
Good	32.5	34.1	33.8	33.3	30.7
Very good	30.8	28.9	25.9	22.1	20.6
Excellent	10.0	8.0	5.8	5.9	5.7

5.5 Discussions

From an overall perspective, the simulated data generates a lower wealth-to-income ratio than observed data. Moreover, the empirical medians of wealth-to-income ratios within each health group remain stable across different periods. In contrast, the simulated medians seem to converge to the same value for all health groups as age grows.

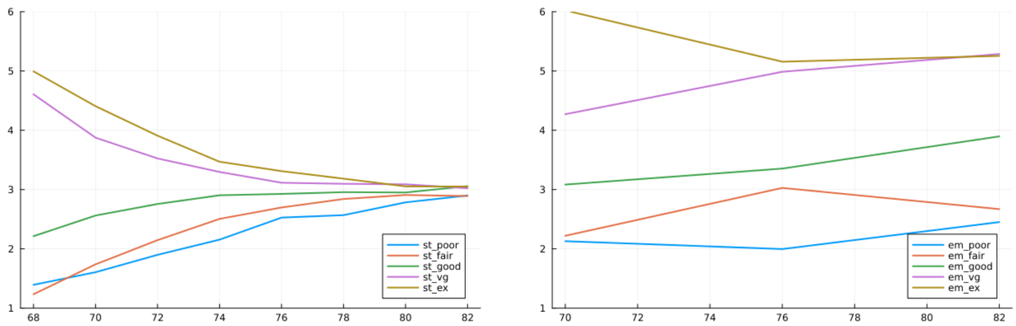

(a) Wealth-to-Income Ratio in Simulated Model **(b)** Wealth-to-Income Ratio in Empirical Data

Figure 5.15: Comparative Analysis of Simulated and Empirical Data

Since the health status for most individuals deteriorates over time (see Figure 4.1), there are three typical processes of health status changes. First, a large percentage of individuals move from an initial healthier status to a deteriorated one. Second, some of the individuals with a bad initial health status would exit the model earlier than their healthier counterparts. Third, the number of individuals with excellent health status declines over time, as the probability of health improvement is lower. Moreover, the composition of the healthier group (excellent and very good) is relatively stable.

5.5.1 Bequest Motive

Based on the mechanism listed above, the substantial differences between simulated and empirical data in the excellent and very good health groups (even the minor difference in the good health group) might be due to the lack of mechanisms for individuals to save beyond their life, that is, the lack of the bequest motive. Without the bequest motive, the optimal policy for each individual is to consume everything and save nothing before they die. Even for healthy individuals who have a relatively long life expectancy, the wealth should deteriorate slowly with age. The stability of the empirical wealth levels may suggest the presence of a bequest motive in the model, which is in line with prior research (Ameriks et al., 2011) highlighting bequest motives as a key determinant of late-life wealth accumulation.

5.5.2 Medical Expenditure

Medical expenditure is a commonly used component in most classical literature. Some studies treat it as an exogenous process (De Nardi et al., 2010), while others allow individuals to choose out-of-pocket medical spending based on their health status, age, and wealth (Yogo, 2016), with spending rising substantially as health deteriorates. These approaches offer a more precise account of how health influences consumption and saving behavior. In contrast, our model does not allow the medical expenditure mechanism, thus resulting in less flexibility for individuals to allocate their wealth when facing the health risk. Though the direct model of the medical expenditure process would clearly identify the change in wealth, focusing solely on medical expenses might underestimate the influence of actual health risks.

5.5.3 Heterogeneity

We assume that everybody follows the same health transition probabilities, whereas this is not always the case in reality. For example, we assume the health transition process follows a Markov chain. However, in reality, the health status could be persistent where the future health status depends not only on the current health status but also on the one from a more distant past. On the other hand, a sudden health shock in old age might deteriorate the health status dramatically.

Medically, the health transition process could be influenced by health investment, lifestyle behavior, genetic predispositions, and so on, let alone the unclear underlying mechanisms. The heterogeneity of health transition probabilities for individuals may create a bias between simulated and empirical data.

In addition to the heterogeneity of the health transition process, preference heterogeneity is another concern. The different risk-aversion levels and time preferences (the discount factor) would affect individuals' behavior significantly over the life cycle. We assume homogeneous preferences across individuals for computational efficiency.

Chapter 6

Conclusion

Our thesis develops a life-cycle model of retired individuals that incorporates health shocks and endogenous mortality risk into consumption, saving, and portfolio choice decisions. Our model extends the model developed by De Nardi et al. (2010), which assumes a binary health state and exogenous mortality. By doing so, our model enables a more nuanced analysis of the relationship between health dynamics and the financial behavior of retirees.

The result of the health risk factor suggests that individuals in better health conditions accumulate more wealth and exhibit lower consumption rates, due to longer life expectancy and lower marginal utility of consumption. Moreover, our findings replicate two key patterns documented in the existing literature: a declining risky asset share with age, and limited health-related heterogeneity in portfolio choice conditional on age.

While the model performs better in capturing the wealth decumulation of less healthy individuals, potentially due to its ability to capture the health dynamics, it shows some misalignment with the empirical wealth decumulation pattern from HRS. In particular, the model underestimates the wealth levels of healthier individuals at older ages, possibly due to the absence of some saving motives and health persistence.

Future research could enhance the model by adding the out-of-pocket medical expenditure mechanism to better capture the financial implications of health risk in retirement. In addition, introducing additional saving motives—such as a bequest motive or the option to purchase health-contingent financial products—would allow for a more realistic representation of precautionary saving behavior under health uncertainty.

Overall, this study highlights the importance of health dynamics in shaping post-retirement financial behavior and provides a tractable framework for evaluating retirement policy under health uncertainty.

References

Ameriks, J., Caplin, A., Laufer, S., and Van Nieuwerburgh, S. (2011). The joy of giving or assisted living? using strategic surveys to separate public care aversion from bequest motives. *The journal of finance*, 66(2):519–561.

Blundell, R., Britton, J., Dias, M. C., and French, E. (2023). The impact of health on labor supply near retirement. *Journal of Human Resources*, 58(1):282–334.

Campbell, J. Y. and Mankiw, N. G. (1989). Consumption, income, and interest rates: Reinterpreting the time series evidence. *NBER macroeconomics annual*, 4:185–216.

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics letters*, 91(3):312–320.

Carroll, C. D. and Samwick, A. A. (1997). The nature of precautionary wealth. *Journal of monetary Economics*, 40(1):41–71.

Cocco, J. F., Gomes, F. J., and Maenhout, P. J. (2005). Consumption and portfolio choice over the life cycle. *The Review of Financial Studies*, 18(2):491–533.

De Nardi, M. (2004). Wealth inequality and intergenerational links. *The Review of Economic Studies*, 71(3):743–768.

De Nardi, M., French, E., and Jones, J. B. (2009). Life expectancy and old age savings. *American Economic Review*, 99(2):110–115.

De Nardi, M., French, E., and Jones, J. B. (2010). Why do the elderly save? the role of medical expenses. *Journal of political economy*, 118(1):39–75.

De Nardi, M., Pashchenko, S., and Porapakkarm, P. (2024). The lifetime costs of bad health. *Review of Economic Studies*, page rdae080.

French, E. (2005). The effects of health, wealth, and wages on labour supply and retirement behaviour. *The Review of Economic Studies*, 72(2):395–427.

French, E. and Jones, J. B. (2004). On the distribution and dynamics of health care costs. *Journal of Applied Econometrics*, 19(6):705–721.

Friedman, M. (1957). The permanent income hypothesis. In *A theory of the consumption function*, pages 20–37. Princeton University Press.

Gomes, F., Haliassos, M., and Ramadorai, T. (2021). Household finance. *Journal of Economic Literature*, 59(3):919–1000.

Gourinchas, P.-O. and Parker, J. A. (2002). Consumption over the life cycle. *Econometrica*, 70(1):47–89.

Hosseini, R., Kopecky, K. A., and Zhao, K. (2021). How important is health inequality for lifetime earnings inequality? Technical report, Working paper.

Mahler, L. and Yum, M. (2024). Lifestyle behaviors and wealth-health gaps in germany. *Econometrica*, 92(5):1697–1733.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *The review of Economics and Statistics*, pages 247–257.

Merton, R. C. et al. (1971). Theory of rational option pricing.

Modigliani, F. and Brumberg, R. (1954). Utility analysis and the consumption function: An interpretation of cross-section data. *Franco Modigliani*, 1(1):388–436.

Palumbo, M. G. (1999). Uncertain medical expenses and precautionary saving near the end of the life cycle. *The Review of Economic Studies*, 66(2):395–421.

Samuelson, P. A. (1975). Lifetime portfolio selection by dynamic stochastic programming. *Stochastic optimization models in finance*, pages 517–524.

Storesletten, K., Telmer, C. I., and Yaron, A. (2004). Consumption and risk sharing over the life cycle. *Journal of monetary Economics*, 51(3):609–633.

Yogo, M. (2016). Portfolio choice in retirement: Health risk and the demand for annuities, housing, and risky assets. *Journal of monetary economics*, 80:17–34.

Zeldes, S. P. (1989). Consumption and liquidity constraints: an empirical investigation. *Journal of political economy*, 97(2):305–346.

Appendices

Appendix A

Euler Equations

The Euler Equation of the Benchmark model:

$$c_t^{-\gamma} = \frac{1}{1+\delta h_t}\beta(1+r_f)E_t(\mathbb{1}_{\{h_{t+1}\neq 0\}}(1+\delta h_{t+1})c_{t+1}^{-\gamma} + \mathbb{1}_{\{h_{t+1}=0\}}0)$$

The Euler Equation of the portfolio choice model:

$$\begin{aligned} c_t^{-\gamma} = \frac{\beta}{1+\delta h_t}\Big[&(1+r_f)(1-\alpha_t)\cdot E_t\left[\mathbb{1}_{\{h_{t+1}\neq 0\}}(1+\delta h_{t+1})c_{t+1}^{-\gamma} + \mathbb{1}_{\{h_{t+1}=0\}}0\right] \\ &+R_{s,t+1}\alpha_t\cdot E_t\left[\mathbb{1}_{\{h_{t+1}\neq 0\}}(1+\delta h_{t+1})c_{t+1}^{-\gamma} + \mathbb{1}_{\{h_{t+1}=0\}}0\right]\Big] \end{aligned}$$

Appendix B

Optimal Consumption and Saving Policy for Model with Portfolio Choice

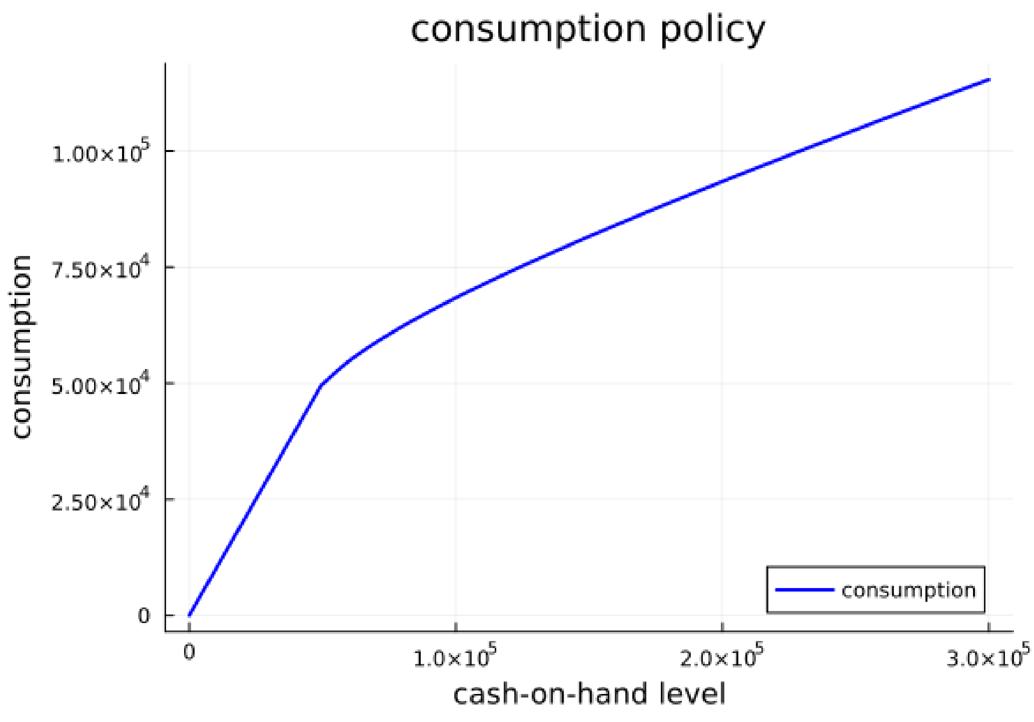

Figure B.1: Consumption Policy

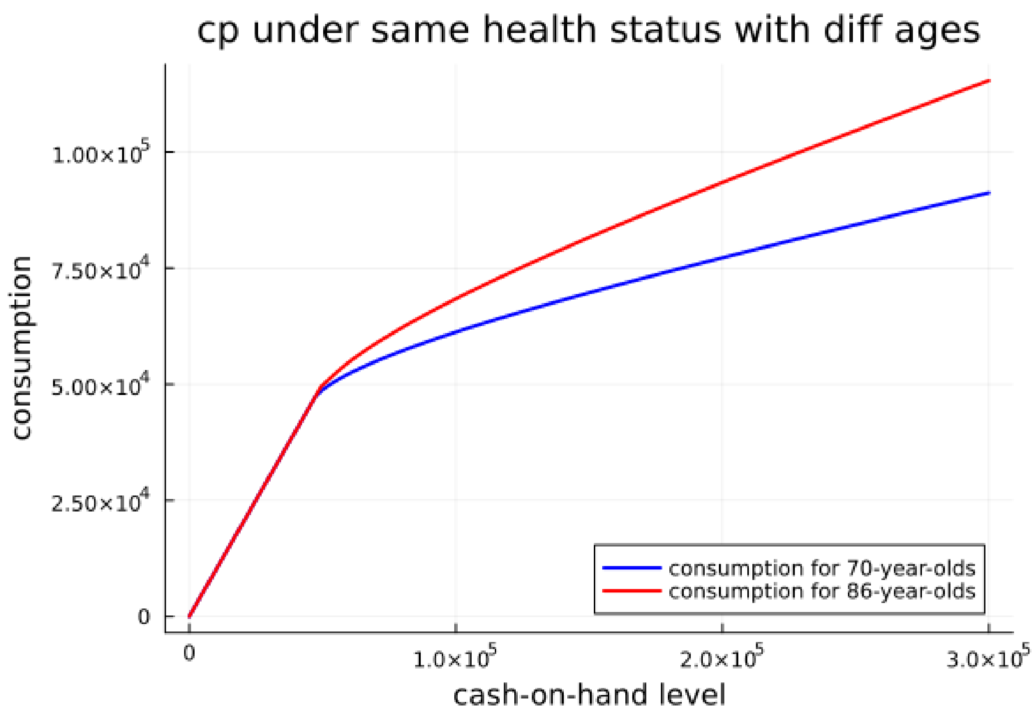

Figure B.2: Consumption Policy for Different Ages

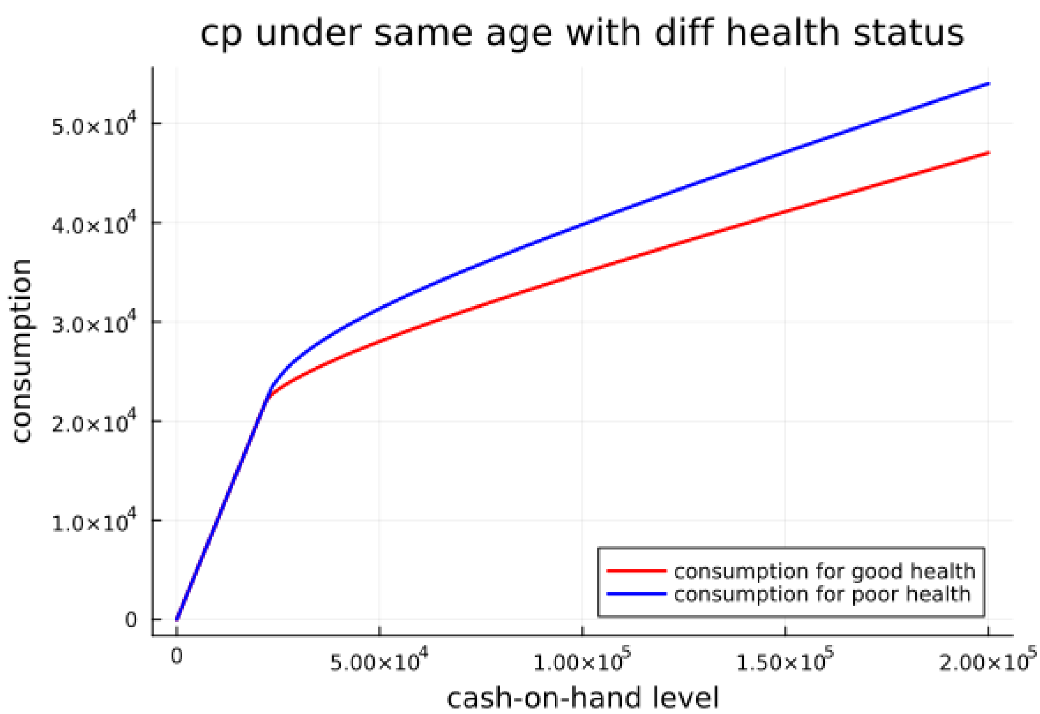

Figure B.3: Consumption Policy for Different Health Groups

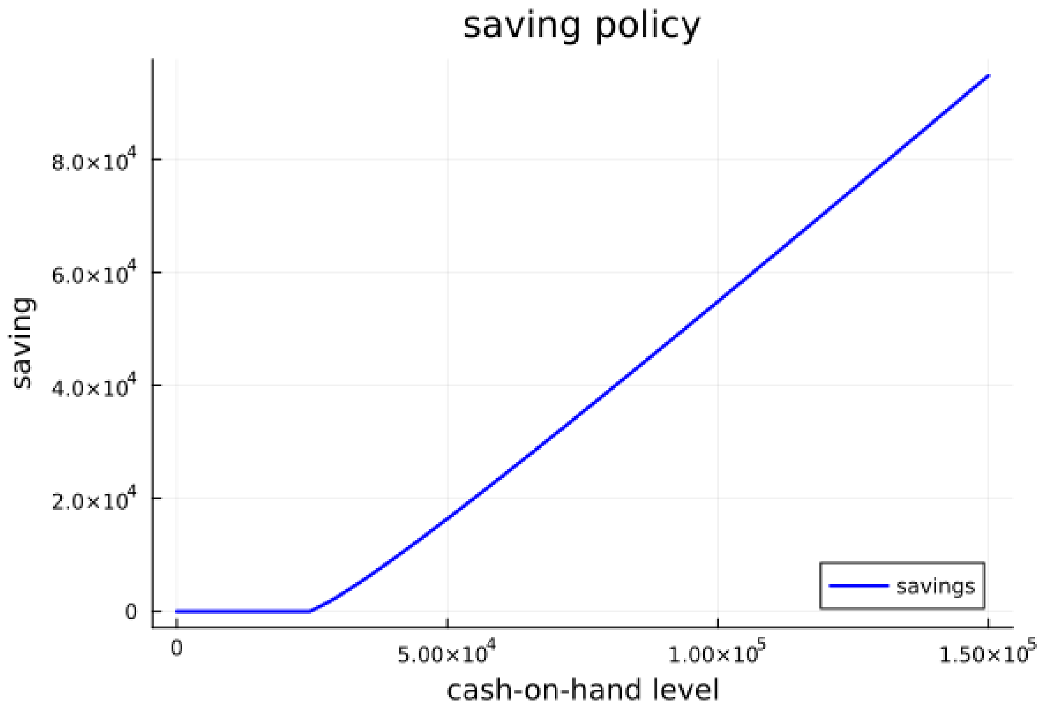

Figure B.4: Saving Policy

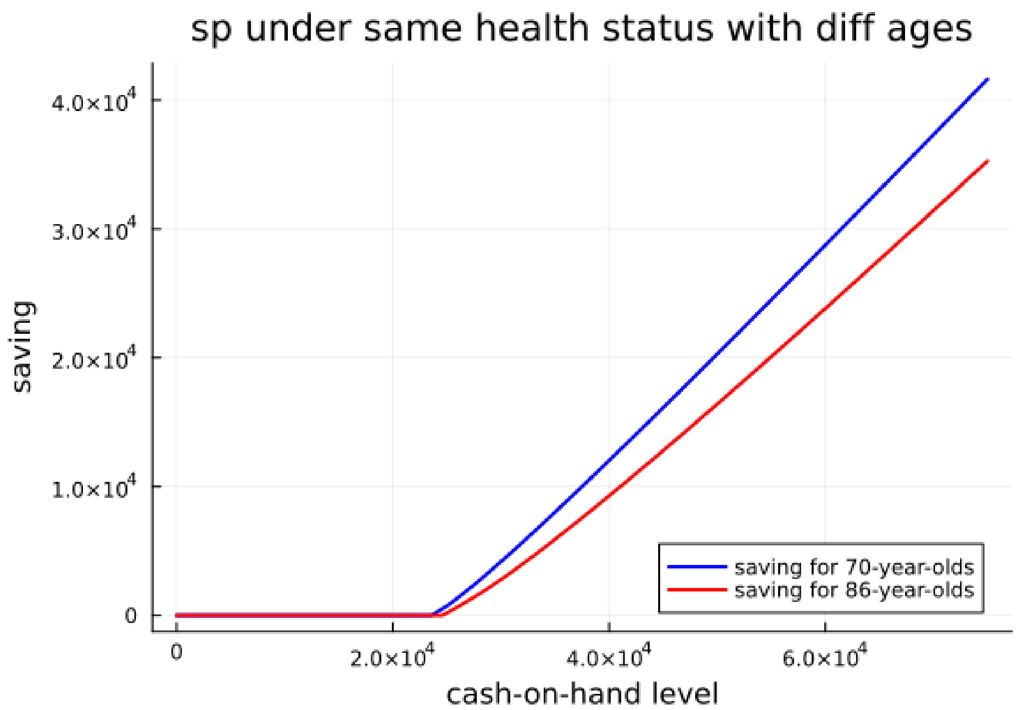

Figure B.5: Saving Policy for Different Ages

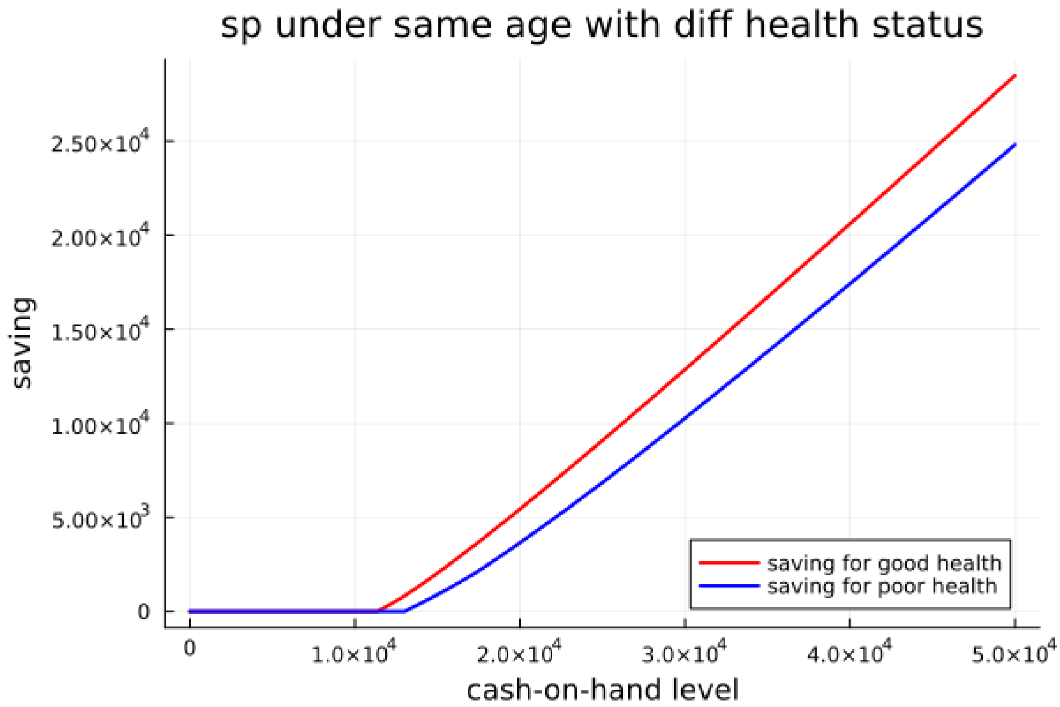

Figure B.6: Saving Policy for Different Health Groups

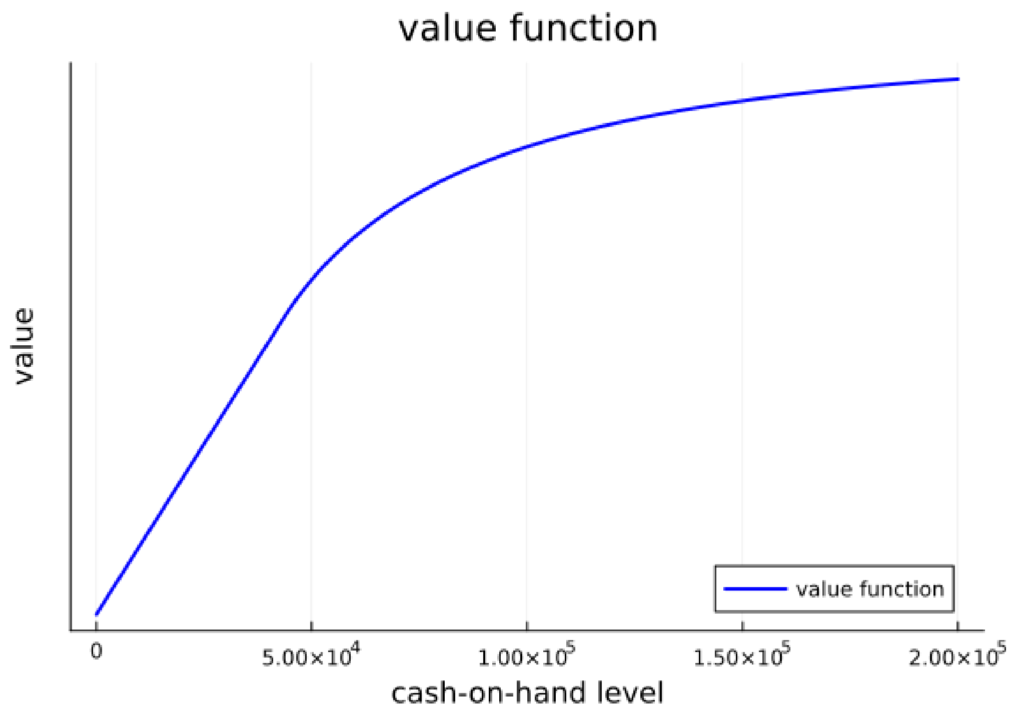

Figure B.7: Value Function

Appendix C

Derivation of the Variance of the Risky Asset

The risky return is modeled as:

$$R_{s,t} = \bar{R}_s \cdot \varepsilon_{s,t},$$

where $\log(\varepsilon_{s,t}) \sim \mathcal{N}\left(-\frac{\sigma_s^2}{2}, \sigma_s^2\right)$.

$$\mathbb{E}(\varepsilon_{s,t}) = \mathbb{E}(e^X), \quad X \sim \mathcal{N}(\mu, \sigma^2)$$

where $\mu = -\frac{\sigma_s^2}{2}$, $\sigma^2 = \sigma_s^2$

Use the moments generating function of normal distribution:

$$\mathbb{E}[e^{aX}] = e^{a\mu + \frac{1}{2}a^2\sigma^2}$$

Hence:

$$\mathbb{E}(\varepsilon_{s,t}) = e^{\mu + \frac{1}{2}\sigma^2}$$

$$\mathbb{E}(\varepsilon_{s,t}^2) = \mathbb{E}(e^{2X}) = e^{2\mu + 2\sigma^2}$$

The variance of $\varepsilon_{s,t}$:

$$\text{Var}(\varepsilon_{s,t}) = \mathbb{E}(\varepsilon_{s,t}^2) - (\mathbb{E}(\varepsilon_{s,t}))^2 = \left(e^{\sigma^2} - 1\right) e^{2\mu+\sigma^2}$$

Substituting μ and σ, we get:

$$\text{Var}(\varepsilon_{s,t}) = \left(e^{\sigma_{s,t}^2} - 1\right)$$

Then, the variance of $R_{s,t}$ is:

$$\text{Var}(R_{s,t}) = \bar{R}_s^2 \cdot \text{Var}(\varepsilon_{s,t}) = \bar{R}_s^2 \cdot \left(e^{\sigma_{s,t}^2} - 1\right)$$

Given $\bar{R}_s = 1.065$, and $\sigma_{s,t} = 0.18$, we can compute:

$$\text{Var}(R_{s,t}) \approx (1.065)^2 \cdot (1.03295 - 1) \approx 0.0374$$

Appendix D

AI disclosure

We used OpenAI's ChatGPT during the thesis writing process. It was primarily employed to debug code and to improve the clarity and precision of English phrasing. All modeling decisions, analysis, and core academic content were developed independently by the authors, and all output from the tool was used under careful consideration.