Essays in the Macroeconomics of Health Care

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Vegard Mokleiv Nygaard

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Manuel Amador, Adviser

Timothy J. Kehoe, Adviser

June, 2019

Acknowledgements

First and foremost, I am indebted to my advisors, Manuel Amador and Timothy J. Kehoe, for your guidance and support throughout my years at the University of Minnesota. Every sentence of this dissertation has been influenced and improved by your countless comments and suggestions. No one has taught me more about economics than the two of you.

Many people have been there for me during my years at the University of Minnesota. I have benefited greatly from comments from Juan Carlos Conesa, Mariacristina De Nardi, Doireann Fitzgerald, Fatih Guvenen, Hannes Malmberg, Gajendran Raveendranathan, Jack Rossbach, Kei-Mu Yi, participants of the International Trade Workshop, and my fellow students at the University of Minnesota.

Catherine Bach, Barbara Drucker, and Wendy Williamson have made my life as a PhD student infinitely easier by constantly helping me not only with the endless paperwork but also with priceless moral support.

To my Italian family, Mario Falcettoni and Nevia Turco, thank you for your encouragement and for the most delicious meals I have ever had in my life. And to Carol Critchley and Kurt Cederquist, thank you for the wonderful time I always have with you and for having become such an important part of my life.

This dissertation would not have been possible without my parents, Sølvi Mokleiv and Nils Nygård. You have always believed in me. Without your love and constant support, I never would have become the person I am today. And to my brothers, Håvard Mokleiv Nygård and Sondre Mokleiv Nygård, thank you for being the best brothers anyone could ever ask for.

Finally, I want to thank Elena Falcettoni. You have stood by me through all my stress and annoyance these years. Even though you had your own dissertation to worry about, you always listened patiently and helped me every step of the way. I would not have been able to do this without you.

Dedication

To my mother, Sølvi Mokleiv, and my father, Nils Nygård.

Abstract

This dissertation consists of three essays that study the macroeconomics of health care. The first essay studies how policies can be designed to reduce differences in life expectancy across income groups in the United States and examines what the consequences of these policies are for welfare and the macroeconomy. Using a calibrated structural life cycle model with incomplete markets, heterogeneous agents, and endogenous health, I find that a universal health insurance reform leads to higher life expectancy, lower life expectancy inequality, lower health care spending, higher GDP per capita, and higher welfare, even after controlling for the increased tax burden needed to finance the reform. The second essay develops a structural life cycle model with incomplete markets and heterogeneous agents to study how the ability to file for medical bankruptcy affects incentives to purchase health insurance. I find that the ability to file for medical bankruptcy crowds out private health insurance coverage. The majority of the population, however, is better off in the economy with medical bankruptcy because of the implicit insurance provided by this option. Finally, motivated by the considerable heterogeneity in GDP per capita across the states of the US, the third essay develops a model to quantify the welfare differences across the states as measured by the expected lifetime utility of being born in a particular state. Using a calibrated version of the model that allows for state-specific variation in mortality risk, consumption uncertainty, and educational attainment, I document large and persistent heterogeneity in welfare across the states of the US.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

This dissertation consists of an introductory chapter and three essays that study the macroeconomics of health care. This chapter provides a brief summary of the motivation, methods, and findings of each essay.

The richest 25 percent of Americans can expect to live about 7 years longer than the poorest 25 percent. Chapter 2 of my dissertation studies how policies can be designed to reduce life expectancy inequality and examines what the consequences of these policies are for welfare and the macroeconomy. To do this, I develop a structural life cycle model with incomplete markets and heterogeneous agents. In the model, agents' health evolves endogenously depending on the healthiness of their consumption basket. I calibrate the model to match several facts about the population health distribution that I document by constructing an objective measure of health called a frailty index. I show that there exists considerable heterogeneity in frailty in the population. To illustrate, 20-year-olds in the fourth frailty quartile are more frail than 50-year-olds in the first frailty quartile, even though the former group is 30 years younger than the latter. These differences in frailty have large implications for mortality risk, medical expenditure risk, and labor market outcomes. I then use the model to study the implications of health insurance and income tax reforms for life expectancy, life expectancy inequality, the macroeconomy, and welfare. I find that universal health insurance leads to higher life expectancy, lower life expectancy inequality, lower health care spending, higher GDP per capita, and higher welfare, even after controlling for the increased tax burden needed to finance the reform. Similarly, I find that the government can reduce life expectancy inequality by expanding Medicaid or by increasing redistribution through income tax reforms but that these reforms can have adverse implications for the macroeconomy.

Chapter 7 of the US bankruptcy code allows individuals to discharge their medical debt by filing for bankruptcy. Chapter 3 of my dissertation studies how the ability to file for medical bankruptcy affects incentives to purchase private health insurance. To do this, I develop a structural life cycle model with incomplete markets and heterogeneous agents that can account for salient features of the US health insurance system and Chapter 7 of the US bankruptcy code. I find that the ability to file for medical bankruptcy lowers the private insurance take-up rate by 2.0 percentage points. This means that 8 percent of the uninsured working-age population would take up private health insurance coverage if they did not have the option to file for medical bankruptcy. I then compute a transition from an economy with to an economy without medical bankruptcy to quantify the welfare implications of this law. I find that aggregate welfare is higher and that 79.6 percent of the population is better off in the economy with medical bankruptcy because of the implicit insurance provided by this option.

GDP per capita ranges by more than a factor of 2 across the states of the US. Chapter 4 of my dissertation develops a model to quantify the welfare differences across the states of the US as measured by the expected lifetime utility of being born in a particular state. The model allows for state-specific variation in mortality risk, consumption uncertainty, and educational attainment. I show that there exists considerable heterogeneity in welfare across the states. To illustrate, I find that consumption must be scaled down by 47 percent in all ages in the state with the highest welfare level, Massachusetts, to make a hypothetical agent indifferent between living her entire life in Massachusetts and the state with the lowest welfare level, Mississippi. Moreover, these differences are very persistent. Accordingly, although the lower ranked states in 2000 have generally experienced a higher welfare growth than the higher ranked states in 2000, there does not appear to be evidence of rapid convergence toward similar welfare levels.

Chapter 2

Causes and consequences of life expectancy inequality

2.1 Introduction

Recent research by Chetty et al. (2016) shows that, at age 40, the richest quartile of Americans can expect to live about 7 years longer than the poorest quartile. Moreover, this difference in life expectancy, commonly referred to as life expectancy inequality, has increased over time. Since 2001, men in the top earnings quartile have experienced a 2.8 year increase in their life expectancy, compared to only a 1.1 year increase for men in the bottom earnings quartile. This chapter studies how policies can be designed to reduce the disparities in life expectancy between high- and low-income individuals. To do this, I develop a structural life cycle model with incomplete markets, heterogeneous agents, and endogenous health that can account for the observed health distribution by age and income in the United States. I then use the model to study how policies that redistribute more resources from the rich to the poor such as an expansion of public health insurance programs or an increase in income tax progressivity affect life expectancy, life expectancy inequality, the macroeconomy, and welfare.

Since the objective of the paper is to understand how policies can be designed to reduce life expectancy inequality and to examine what the consequences of these policies are for welfare and the macroeconomy, it is important that the model closely matches both the determinants of health and the implications of health over the life cycle. I therefore start by documenting how health affects mortality risk, medical expenditure risk, and labor market outcomes, and show how health varies in the population with age and socioeconomic status. To do this, I use data from the Medical Expenditure Panel Survey (MEPS), a longitudinal study that collects detailed annual records on health indicators, medical spending, insurance coverage, income, and demographics for a representative sample of non-institutionalized individuals in the United States. To get a continuous and objective measure of an individual's health, I follow conventions in the gerontology literature (see e.g. Searle et al. 2008) and approximate an individual's health by her *frailty index*. This index counts how many health deficits the individual has. Health deficits are defined as disabilities and diseases, some examples of which include functional limitations, cognitive impairments, and health conditions such as cancer and diabetes.

Using this measure, I document the following facts. First, there is considerable heterogeneity in frailty in the population. To illustrate, 20-year-olds in the fourth frailty quartile (that is, the sickest 25 percent) are more frail than 50-year-olds in the first frailty quartile (that is, the healthiest 25 percent), even though the former group is 30 years younger than the latter. Next, there is a strong relationship between frailty and socioeconomic status as measured by educational attainment, income, and wealth. Across all age groups, educated individuals are less frail than non-educated individuals, low-income individuals are more frail than high-income individuals, and wealth-rich individuals are less frail than wealth-poor individuals. Moreover, these differences are economically significant. As an example, the healthiest 25 percent of 65-year-olds have 4 times higher median net wealth than the sickest 25 percent.

These differences in frailty have considerable implications for mortality risk, medical expenditure risk, and labor market outcomes. An increase in frailty leads to a large reduction in age-specific survival probabilities. To illustrate, the sickest 25 percent of 85-year-olds have a 3 times higher mortality probability than the healthiest 25 percent. Higher frailty also leads to increased medical expenditure risk. I find that average annual medical spending for 20-64 year-olds is about $\$1,200$ for individuals in the lowest frailty quartile. In contrast, 20-64 year-olds in the fourth frailty quartile spend an average of $\$9,600$ on health care. Similarly, the healthiest and sickest 25 percent of elderly spend an average of $\$3,400$ and $\$17,200$ per year on health care, respectively. Lastly, I examine how frailty affects labor market outcomes. I find a negative relationship between frailty and labor productivity as measured by hourly wage rates. In particular, I find that, conditional on initial wage rates, higher frailty is associated with a reduction in wage rates in the following year. Similarly, I find that, conditional on initial hourly labor supply, higher frailty is associated with a reduction in the number of hours worked per week in the following year. More importantly, frailty has large implications for labor force participation. As an example, I find that the sickest 25 percent of 50-year-olds have a 34 percentage point lower labor force participation rate than the healthiest 25 percent.

I then develop a structural life cycle model with incomplete markets and heterogeneous agents that can account for these facts. In the model, agents are heterogeneous along the dimensions of age, frailty, assets, labor productivity, medical spending, education, and health insurance coverage. Consistent with the health insurance system in the United States, health insurance is available in the form of private insurance, employer-provided insurance, Medicare, and Medicaid. The model also features a Social Security program that provides pensions to the elderly. As in the data, frailty has implications for longevity, medical expenditure risk, and labor market outcomes. In particular, an increase in frailty leads to lower survival probabilities and higher expected medical costs. In addition, higher frailty also lowers the agent's labor productivity and increases the agent's disutility of working, thereby increasing the probability that the agent will leave the labor force. This latter part enables me to match the large negative relationship between frailty and labor force participation rates in the data.

In the model, agents derive utility from consumption of healthy and unhealthy goods, and disutility from working. Their frailty evolves endogenously over the life cycle depending on their age, education, and choice of consumption basket. In particular, agents can improve their frailty transition probabilities by choosing a healthier consumption basket. This mechanism that frailty transitions depend on the healthiness of the agent's consumption basket is motivated by a large literature in medicine that studies the health consequences of healthy and unhealthy eating. To illustrate, Danaei et al. (2009) show that poor diets are among the leading causes of obesity, diabetes, cardiovascular diseases, and diet-related cancers. Estimates by Murray et al. (2013) suggest that dietary factors account for about 25 percent of deaths in the United States and about 14 percent of all disability-adjusted life-years lost. Similar results are reported by Lim et al. (2012), who show that dietary risk factors and physical inactivity are among the leading drivers of disability and mortality, with the most prominent dietary risks being diets low in fruits and those high in sodium.

I assume that healthy and unhealthy goods are perfect substitutes, but that unhealthy goods have a lower relative price. This is motivated by evidence from the medical literature (see e.g. Drewnowski and Specter (2004) and Drewnowski (2010)) that show that healthier foods such as fruits and vegetables tend to have a higher price per calorie than less healthy foods such as sweets and processed meats. Agents in the model therefore face a tradeoff between higher current period consumption and lower next-period expected frailty. In the model, optimal expected life span is longer for the rich than for the poor, which can only be achieved by investing more in healthy consumption. High-income agents therefore consume healthier baskets over the life cycle than low-income agents, which in turn leads to increasing frailty disparities between the income groups. Consistent with this, Rehm et al. (2016) document large variations in diet quality across both income and educational groups. To illustrate, they show that more than 60 percent of low-income individuals have a diet that is classified as poor by the American Heart Association because of its negative implications for various health outcomes such as cardiovascular diseases. In contrast, less than 30 percent of middle-income individuals have a diet that is classified as poor.

I calibrate the model to match the distribution of frailty by age and income in the population. Recall that the law of motion for frailty depends on the healthiness of the agent's consumption basket in the model. Since I do not observe consumption in the MEPS, I cannot estimate frailty transition probabilities directly from the data. I therefore use an indirect inference approach to calibrate the coefficients of the frailty transition matrix. In particular, since income predicts next-period frailty in the data, I calibrate the parameters of the model such that the model matches this relationship between income and next-period frailty. To do this, I first estimate a relationship between frailty and past frailty, age, income, and education using data from the MEPS. I then run the same regression on simulated data from the model. Lastly, I iterate on the preference parameters and on the coefficients of the frailty transition matrix until these two regressions coincide. This means that the model generates a relationship between income and next-period frailty that is consistent with what we observe in the data. Finally, since the way I model frailty coincides exactly with the way I measure frailty in the data, I can estimate how frailty affects mortality risk, medical expenditure risk, and labor market outcomes directly from the data.

After verifying that the model can account for both the frailty facts documented in this chapter and the relationship between income and life expectancy documented by Chetty et al. (2016), I then use the model to study how policies can be designed to reduce life expectancy inequality. Recall that frailty affects income through its effect on labor productivity and labor force participation. In addition, income indirectly affects frailty by facilitating higher consumption of healthy goods. There is therefore a two-way relationship, or a feedback loop, between income and frailty in the model. In particular, high frailty leads to low income, which in turn leads to lower consumption of healthy goods, which in turn leads to higher next-period expected frailty. This mechanism implies that the government might be able to improve population health outcomes by expanding public health insurance programs or by increasing income tax progressivity, both of which facilitate higher investments in healthy consumption in times of high medical costs or low labor productivity. To test this hypothesis, I first study a counterfactual economy with universal health insurance. In particular, I consider an economy where the government covers 86.2 percent of all health care expenses, and finances these costs through higher income taxes. I find that the reform leads to a 0.10 year increase in average life expectancy and a 0.41 year reduction in longevity differences at age 40 by income quartile due to higher investments in healthy consumption, especially among the previously uninsured. The improvements in the population health distribution leads to a 0.18 percent decline in health care spending per capita. Moreover, I find that GDP per capita increases by 0.50 percent due to both higher capital accumulation and higher labor supply. Because of these results, I find that the universal health insurance reform leads to higher *ex ante* welfare, even after controlling for the increased tax burden needed to finance the reform. In particular, I find that consumption must increase by 1.82 percent in all periods and contingencies in the benchmark economy to make an unborn agent under the veil of ignorance indifferent between the benchmark economy and the economy with universal health insurance.

I then compare the results from the universal health insurance reform with the results from two Medicaid reforms designed to increase insurance coverage among low-income agents. In particular, motivated by the recent Medicaid expansion under the Affordable Care Act, I first examine the effects of expanding Medicaid to all agents with income no greater than 138 percent of the federal poverty level, and then examine the effects of expanding Medicaid to all agents with income net of medical expenses no greater than 138 percent of the federal poverty level. For brevity, I refer to the two reforms as the categorical and medically needy reform, respectively. Consistent with the universal health insurance reform, I find that both Medicaid reforms lead to higher average life expectancy and lower life expectancy inequality as measured by the difference in life expectancy at age 40 by income quartile. Unlike the universal heath insurance reform, however, both Medicaid reforms have negative implications for the macroeconomy. To illustrate, I find that the medically needy reform leads to a 2.89 percent decline in consumption per capita and a 2.38 percent reduction in GDP per capita. These findings are driven by the observation that Medicaid discourages saving and working since eligibility for the program is tied to income. That is, uninsured agents with high medical costs have an incentive to reduce their saving and labor supply to qualify for the program. In contrast, these distortions are not present in the economy with universal health insurance since eligibility for public health insurance is independent of the agent's income in that environment.

Lastly, I examine how life expectancy is affected by the progressivity of the income tax schedule, which in turn governs the level of redistribution between the rich and the poor in the economy. I do this by comparing the benchmark model results with the results derived using three alternative income tax schedules: a proportional tax schedule, a tax schedule characterized by higher maximum but lower average marginal rates, and a tax schedule characterized by higher maximum marginal rates but also higher deductions for low-income agents. I find a positive relationship between tax progressivity and average life expectancy, and a negative relationship between tax progressivity and life expectancy inequality. In particular, life expectancy inequality is 0.33 years higher in the economy with proportional taxes than in the benchmark model, but 0.80 years lower in the environment with higher tax deductions. Similarly, I find a negative relationship between tax progressivity and health care spending. To illustrate, I find that health care spending per capita is 0.43 percent higher in the environment with proportional taxes than in the benchmark model, but 1.03 percent lower in the environment with higher tax deductions. On the other hand, consistent with the optimal tax literature, I find that higher tax progressivity leads to lower GDP per capita due to increased disincentive effects to both save and work for agents at the high end of the income distribution. Therefore, these findings suggest that policymakers can reduce life expectancy disparities by increasing tax progressivity, but that the gains from such reforms must be weighed against their adverse implications for labor supply and capital accumulation.

Related literature: This chapter relates to the literature that studies the relationship between socioeconomic status and longevity. In an early study, Kitagawa and Hauser (1973) find that low socioeconomic status is correlated with higher mortality. Deaton and Paxson (2001) use data from the Current Population Survey and the National Longitudinal Mortality Study to examine the links between mortality and income, income inequality, and education. They show that both income and education are protective against mortality. Several studies find similar results. Lin et al. (2003) show that life expectancy is positively associated with income, education, martial status, and employment. Attanasio and Emmerson (2003) show that wealth is an important determinant of mortality, even after controlling for initial health. Pijoan-Mas and Ríos-Rull (2014) find that education, wealth, and income are health-protective, but have otherwise little implications for two-year mortality rates conditional on health. In a recent study, Chetty et al. (2016) use data from the Social Security Administration to study the relationship between income and longevity in the United States between 2001 and 2014. They find a large positive relationship between income and life expectancy, although the magnitude of the relationship differs across geographical areas. Moreover, they show that differences in life expectancy by income have increased over time. Lastly, Milligan and Schirle (2018) use administrative data from Canada for the period 1966 to 2015 to examine how life expectancy differences by income have evolved over time. In contrast to the United States, they find that differences in life expectancy by income have remained roughly constant in Canada. All of these papers study how life expectancy varies with socioeconomic status. They do not, however, study how policies can be designed to reduce life expectancy inequality, which is the objective of this chapter.

Next, the chapter builds on a growing literature that incorporates health or medical expenditure risk in structural life cycle models. Hubbard et al. (1994), Palumbo (1999), De Nardi et al. (2010), Kopecky and Koreshkova (2014), and Nakajima and Telyukova (2018) use structural life cycle models to examine how medical expenditure risk affect savings of the elderly. De Nardi et al. (1999) and Attanasio et al. (2010) develop general equilibrium overlapping generations models to study the effects of aging. De Nardi et al. (2016) study the distribution of Medicaid transfers and Medicaid valuations across single retirees, and Braun et al. (2017) examine the welfare effects of means-tested social insurance programs such as Medicaid and Supplemental Social Insurance. Similarly, Paschenko and Porapakkarm (2013) and Conesa et al. (2018) develop models to examine the macroeconomic and welfare effects of the Patient Protection and Affordable Care Act (ACA) and Medicare, respectively.

All the papers mentioned above assume that health or medical expenditure risk evolves exogenously over the life cycle. I contribute to a subset of this literature that endogenizes the evolution of health. Hall and Jones (2007) and Fonseca et al. (2013) use a model where the evolution of health depends on medical spending to study the determinants of increasing health care costs and increasing life expectancy. Similarly, Scholz and Seshadri (2016) study how Medicare affects mortality risk in a model where health depends on medical spending and time investments. Such time investments can be thought of as a subset of healthy consumption. I extend their mechanism by also focusing on unhealthy consumption such as unhealthy eating and smoking. Jung and Tran (2016) study the welfare and macroeconomic effects of the ACA within a model where agents' health depends on medical spending. In a related paper, Cole et al. (2018) develop a life cycle model where the evolution of health is endogenously determined by the agents' effort choice, and use the model to study the tradeoff between the provision of social insurance and the incentives to maintain good health. In both of these papers, health is assumed to affect medical expenditure risk and labor productivity, but not longevity. Consequently, these papers are unable to examine how policy reforms such as an expansion of health insurance coverage would affect both life expectancy and life expectancy inequality.

Within the endogenous health literature, the two most closely related papers to mine are Ozkan (2014) and Kotera (2018). Both papers study the determinants of the health disparities by income in the United States. Ozkan (2014) develops a model with two types of health capital: preventive and curative. In his model, investments in preventive capital govern the distribution of shocks to physical capital, which in turn governs the agent's survival probabilities. Kotera (2018), on the other hand, develops a model where health depends on both medical spending and smoking. In his model, agents tradeoff the utility gains from smoking with its negative health implications.

This chapter complements these papers along several dimensions. First, this chapter differs in the way it models the evolution of health. Whereas Ozkan (2014) and Kotera (2014) focus on preventive health care spending and smoking, respectively, This chapter focuses on consumption of healthy and unhealthy consumption goods. This is motivated by a large literature in medicine that studies the health consequences of healthy and unhealthy eating and by the large variations in diet quality across income and educational groups. Second, it differs in the way it measures health, both in the model and in the data. I use data from the MEPS to construct an objective health measure called a frailty index. In the model, agents' health coincides exactly with this measure. This allows me to estimate the implications of health for survival probabilities, medical expenditure risk, and labor market outcomes directly from the data. In contrast, Ozkan (2014) and Kotera (2018) calibrate these effects since health stocks are unobserved in the data. Third, unlike these papers, this chapter examine how differences in both educational attainment and income affects life expectancy inequality, and find that both factors are key determinants of life expectancy inequality. Lastly, it differs in the policy reforms it studies. Whereas all three studies examine health insurance reforms, this chapter also studies the implications of several income tax reforms for life expectancy inequality, the macroeconomy, and welfare.

Lastly, the focus on objective health measures rather than the more commonly used subjective or self-rated health measures is related to ongoing work by Hosseini et al. (2018). They develop a life cycle model where agents' health is measured by their frailty index, and use the model to study how differences in health outcomes affect lifetime earnings. It also relates to recent work by De Nardi et al. (2017), who show that the dynamics of health in the data are not consistent with the low-order Markov processes typically used in structural life cycle models. They develop a life cycle model with a richer health process that allows for both history-dependence and fixed ex-ante heterogeneity, and use the model to quantify the lifetime costs of bad health. These papers, however, do not endogenize the evolution of health, and are therefore unable to examine how policies affect longevity.

The rest of the chapter is organized as follows. The next section describes the data used in this chapter, including how I construct the frailty index and how frailty affects mortality risk, medical expenditure risk, and labor market outcomes. Section 3 lays out the environment of the economy and sets up a quantitative life cycle model. This section also presents the different types of health insurance that are available in the economy. Section 4 describes how I calibrate the model, and examines how well the model matches both targeted and non-targeted moments of the data. Section 5 studies how policies can be designed to reduce life expectancy inequality and studies the consequences of these policies for welfare and the macroeconomy. I focus on a universal health insurance reform, two Medicaid reforms, and three income tax reforms. Lastly, section 6 concludes and gives directions for future research. Additional details about the data, the regressions, and the mechanism are given in the appendix. The appendix also examines how income inequality and differences in educational attainment affect life expectancy inequality.

2.2 Data

This section describes the data used in this chapter. Most of the data is obtained from the Medical Expenditure Panel Survey (MEPS). This longitudinal survey, which consists of two-year overlapping panels for the period 1996 to 2016, collects detailed annual records on health indicators, medical spending, insurance, income, and demographics for families and individuals across the United States. The participants in the survey are drawn from a nationally representative subsample of households that participated in the prior year's National Health Interview Survey conducted by the National Center for Health Statistics.

I start by explaining how I construct the frailty index, which is an objective measure of health commonly used in the medical literature. The medical literature has largely focused on how frailty affects the probability of death and institutionalization, and how an increase in frailty affects future health transitions. My chapter extends this literature by documenting how frailty affects economic outcomes such as labor productivity and labor force participation, and show how frailty covaries with socioeconomic characteristics such as income and wealth.

2.2.1 Frailty index

I approximate an individual's health by her *frailty index*. Following the literature in gerontology (see for example Searle et al. 2008), I let an individual's frailty index be given by the sum of her health deficits. Health deficits are defined as disabilities and diseases, some examples of which include functional limitations, cognitive impairments, and health conditions such as cancer and diabetes. A list of all the variables used to create the frailty index is given in appendix Table A.1. I follow conventions in the gerontology literature to create this index. First, all binary variables such as whether or not the person has a cancer diagnosis are assigned a value of '0' if the person does not have this health deficit and a value of '1' if the person does have this health deficit. Second, all variables that include more than two responses are assigned values from 0 to 1 depending on the number of possible responses. To illustrate, suppose the respondent is asked to rate her health limitations climbing stairs on the following 3-point scale: *not limited*, *limited a little*, and *limited a lot*. Then a value of '0' is assigned if the person is not limited, '0.5' if the person is limited a little, and '1' if the person is limited a lot. Lastly, following conventions in the gerontology literature, I sum the variables and normalize the index from 0 to 1 by dividing the score by the number of variables included in the frailty index.

The left panel of Figure 2.1 plots the distribution of frailty by age. Average frailty increases gradually over the life cycle from 0.03 at age 20 to 0.35 at age 85. Similarly, the dispersion, as measured by the standard deviation of frailty, also increases over the life cycle from 0.06 to 0.19 between the ages of 20 and 85. The right panel shows how the frailty percentiles vary with age. For each age group, I rank individuals according to their frailty index and compute the 25th, 50th, and 75th frailty percentile. The graph shows that frailty varies considerably in the population. To illustrate, 20-year-olds in the fourth frailty quartile (that is, the sickest 25 percent) have a higher frailty than 50-year-olds in the first frailty quartile (that is, the healthiest 25 percent). Hence, although the former group is 30 years younger, they still have a worse objective health score than the latter group. Similarly, the sickest 25 percent of 60-year-olds have a worse objective health score than healthiest

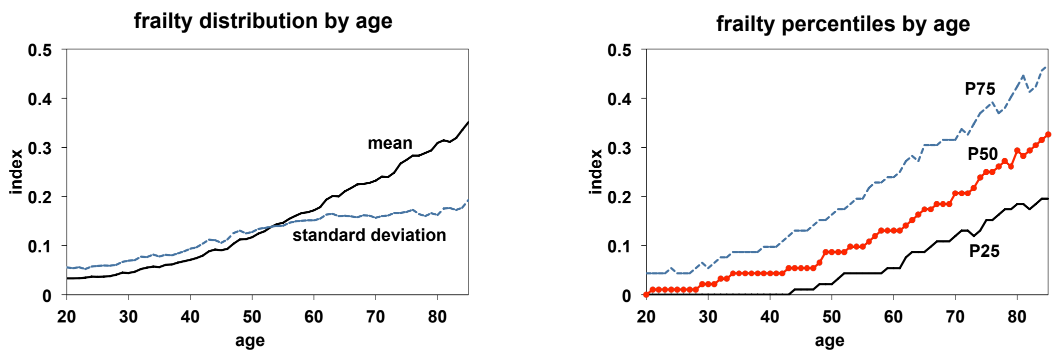

Figure 2.1: Frailty distribution by age

Notes: The left panel plots the average and the standard deviation of frailty by age. The right panel plots the 25th, 50th, and 75th percentile of frailty by age. Data source: MEPS.

25 percent of 85-year-olds, even though the former group is 25 years younger than the latter. As will be shown in the following subsections, these differences in frailty have large implications for mortality risk, medical expenditure risk, and labor market outcomes.

It is well known that health and income are positively correlated. To study the strength of this relationship, I first rank individuals in a given age group by their income and compute the 25th, 50th, and 75th income percentile. Next, I compute average frailty within each age and income group. This is illustrated in the left panel of Figure 2.2. The graph confirms that frailty and income are negatively correlated. Across all age groups, low-income individuals are consistently more frail than high-income individuals. Moreover, the difference in frailty between low and high-income individuals gradually increases over the course of the working-life. As a result, I find that 50-year-olds in the bottom income quartile are as frail as 65-year-olds in the top income quartile. It is also well known that health and educational attainment are positively correlated. The right panel of Figure 2.2 studies the strength of this relationship by plotting average frailty by age and educational attainment. Here, college refers to individuals with at least a four-year college degree. All other indi-

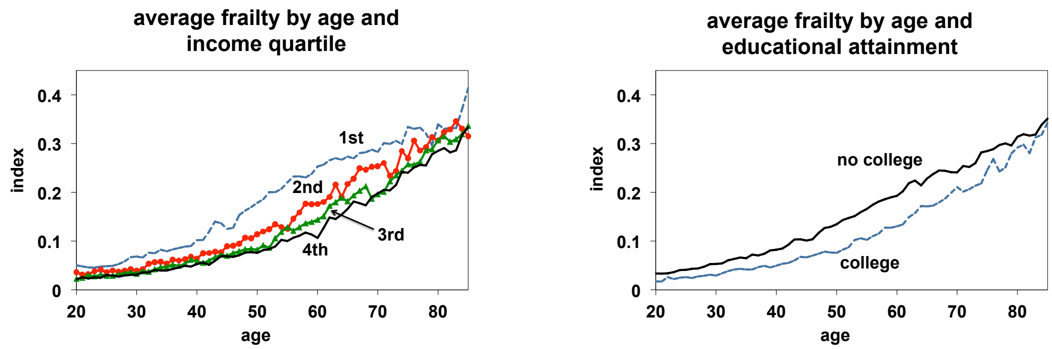

Figure 2.2: Average frailty by age, income quartiles, and educational attainment

Notes: The left panel plots average frailty by age and income quartile. The right panel plots average frailty by age and educational attainment. College refers to individuals with at least a 4-year college degree. Non-college refers to everyone else. Data source: MEPS.

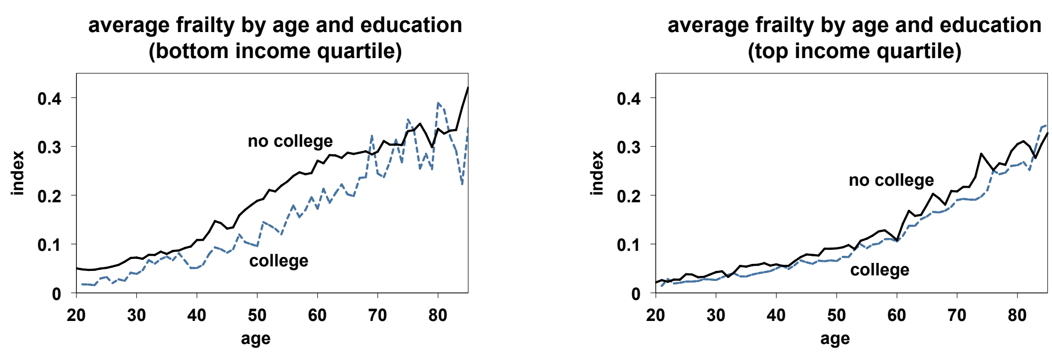

Figure 2.3: Average frailty by age and educational attainment for the bottom and top income quartiles

Notes: The left panel plots average frailty by age and educational attainment for individuals in the bottom income quartile. The right panel plots average frailty by age and educational attainment for individuals in the top income quartile. College refers to individuals with at least a 4-year college degree. Non-college refers to everyone else. Data source: MEPS.

viduals are classified as non-college. Across all age groups, I find that college educated individuals are consistently less frail than non-college educated individuals.

Since income and education are correlated, it is possible that the relationship between frailty and income documented above simply captures the relationship between frailty and education. To examine this, I split individuals by both their education and income. The left panel of Figure 2.3 plots average frailty by age and educational attainment for individuals in the bottom income quartile. The graph shows that, conditional on being in the bottom income quartile, college educated individuals are less frail than the non-college educated. The right panel of Figure 2.3 plots the same relationships for individuals in the top income quartile. I find that, conditional on being in the top income quartile, college and non-college educated individuals are almost equally frail on average. These observations show that frailty and education are correlated, but that the magnitude of the relationship is larger among low-income individuals. The relationship between income, educational attainment, and frailty is further examined in the appendix. I show in appendix Table A.3 that both income and education are predictive of individuals' next-period frailty, even after controlling for initial frailty, medical spending, county of residence, occupation, and a range of demographic variables that have been shown to covary with health such as marital status, sex, and gender.

Lastly, I examine how frailty covaries with wealth. Since the MEPS does not report wealth data, I use data from the Health and Retirement Study (HRS), which provides detailed longitudinal data on health, medical spending, income, wealth, and demographics for a sample of households whose head is at least 50 years of age. To maintain consistency with the MEPS, I construct individuals' frailty index in the HRS using the same approach as discussed above. A list of all the variables used to create individuals' frailty index in the HRS is given in appendix Table A.2. Next, for each age group, I rank individuals according to their frailty index and compute the 25th, 50th, and 75th frailty percentile. Lastly, I compute average and median net wealth by age and frailty quartile, where net wealth is given by the total value of housing and real estate, autos, money market accounts, savings accounts, T-bills, IRAs, Keoghs, stocks, farm or business assets, mutual funds, bonds, and other assets,

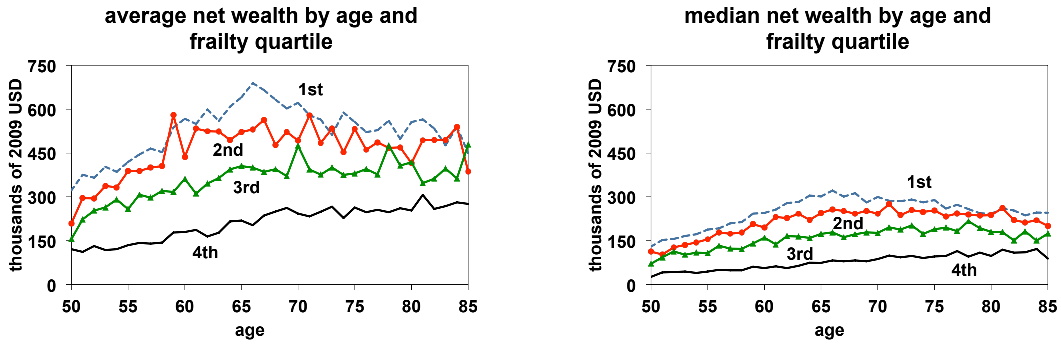

Figure 2.4: Average and median net wealth by age and frailty quartiles

Notes: The left panel plots average net wealth by age and frailty quartiles. The right panel plots median net wealth by age and frailty quartiles. Net wealth is given by the total value of housing and real estate, autos, money market accounts, savings accounts, T-bills, IRAs, Keoghs, stocks, farm or business assets, mutual funds, bonds, and other assets, net of mortgages and other debts. Numbers are in thousands of 2009 USD. Data source: HRS.

net of mortgages and other debts. These results are illustrated in the left and right panels of Figure 2.4. I find that, across all age groups, wealth-rich individuals are consistently less frail than wealth-poor individuals. Moreover, the difference in wealth is economically significant. As an example, the right panel of Figure 2.4 shows that the healthiest 25 percent of 65-year-olds have 4 times higher median net wealth than the sickest 25 percent.

So far, I have shown that frailty increases over the life cycle, that there exists considerable heterogeneity in frailty in the population, and that, across all age groups, educated individuals are less frail than non-educated individuals, low-income individuals are more frail than high-income individuals, and wealth-rich individuals are less frail than wealth-poor individuals. In what follows, I study the implications of these disparities in frailty for mortality risk, medical spending risk, and labor market outcomes.

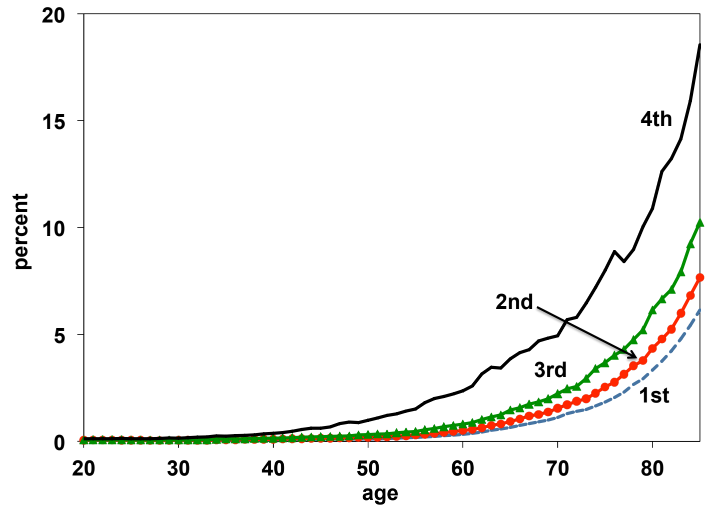

Figure 2.5: Mortality probability by age and frailty quartiles

Notes: The graph plots predicted mortality probabilities by age and frailty quartiles. Probabilities are derived from a logistic regression of mortality on age, frailty, high-order moments of age, and demographic controls. Data source: MEPS.

Relationship between frailty and survival I start by examining how frailty affects mortality risk. Using my MEPS sample, I first estimate a logistic regression of mortality on age, frailty, higher-order moments of age, and demographic controls. Age and frailty-specific survival probabilities can then be computed by applying the standard logistic formula. To illustrate how frailty affects mortality risk, I compute age-specific mortality probabilities by frailty quartiles. In particular, for each age group, I split the individuals in the MEPS into four frailty quartiles and compute the average frailty for individuals in each group. The derived mortality estimates are illustrated in Figure 2.5. I find that an increase in frailty leads to a large increase in mortality risk. To illustrate, I find that the sickest 25 percent of 85-year-olds have a 3 times higher mortality probability than the healthiest 25 percent.

Table 2.1: Medical spending by age groups and frailty percentiles

	Quartiles				Top	
Age group	1st	2nd	3rd	4th	90-95	95-100
20-64	1,200	1,900	3,400	9,600	11,000	18,300
65+	3,400	5,800	9,200	17,200	18,200	25,200

Notes: The table reports total average annual medical spending by age group and frailty percentiles. Total medical spending includes all costs covered by private insurance, public insurance, and out-of-pocket, but does not include spending on health insurance premia. Numbers are in constant 2009 dollars. Data source: MEPS.

Relationship between frailty and medical spending Next, I examine how frailty affects medical spending. Table 2.1 reports total average medical spending by age groups and frailty percentiles in the MEPS. Note that total medical spending includes all costs covered by private insurance, public insurance, and out-of-pocket, but does not include spending on health insurance premia. All current prices have been converted to 2009 dollars using personal consumption expenditure health (PCE-Health) price indices. I find that an increase in frailty leads to a large increase in medical expenditure risk. In particular, I find that average annual medical spending for 20-64 year-olds is about $\$1,200$ for individuals in the first frailty quartile. In contrast, 20-64 year-olds in the fourth frailty quartile spend an average of $\$9,600$ on health care per year. The last two columns report average annual medical spending for individuals at the top of the frailty distribution. As reported in the table, the sickest 5 percent of 20-64 year-olds spend almost $\$20,000$ per year on health care. Similarly, I find that health care spending is increasing in frailty for the elderly. To illustrate, the healthiest and sickest 25 percent of 65+ year-olds spend an average of $\$3,400$ and $\$17,200$ on health care per year, respectively, with the sickest 5 percent spending an average of more than $\$25,000$ per year.

Relationship between frailty and labor earnings Lastly, I examine how frailty affects labor market outcomes. Recall that the MEPS consists of two-year overlapping panels. Using this admittedly short panel dimension, I first estimate how frailty affects labor productivity by regressing the logarithm of hourly wages in year two on initial hourly wages, frailty, education, a quadratic in age, and demographic controls. I restrict the sample to 20-64 year-olds that work at least 10 hours per week. The estimates are reported in appendix Table A.5. The regression results show that the semi-elasticity of hourly wages with respect to frailty is -0.16. This finding suggests that frailty has a relatively modest effect on labor earnings since average frailty only increases by about 0.2 over the course of the working life. Such a conclusion, however, would understate the full effect of frailty on labor earnings. First, as reported in appendix Table A.6, I find that frailty has negative implications for hours worked conditional on working at least 10 hours per week. Second, and more importantly, I find that frailty has negative implications for labor force participation. Across all age groups, less frail individuals are more likely to participate in the labor force than more frail individuals. This can be seen in Figure 2.6, which plots labor force participation rates by age and frailty quartiles. Moreover, the magnitude of the effect is economically significant. As an example, I find that the sickest 25 percent of 50-year-olds have a 34 percentage point lower labor force participation rate than the healthiest 25 percent. This implies that frailty can have large implications for retirement decisions. In particular, because of the high persistence of frailty, more frail individuals are likely to opt for early retirement than less frail individuals. As a result, more frail individuals are not only likely to earn less and accumulate less wealth over the course of their working-life, they are also likely to claim lower Social Security benefits during retirement since the current Social Security system ties benefits to past contributions.

In what follows, I will develop a structural life cycle model with incomplete markets,

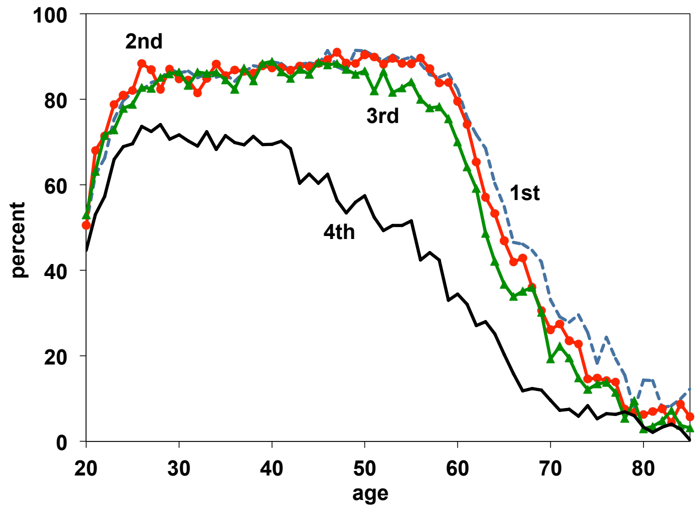

Figure 2.6: Labor force participation rate by age and frailty quartiles

Notes: The graph plots labor force participation rates by age and frailty quartiles. Data source: MEPS.

heterogeneous agents, and endogenous health that can account for the facts documented in this section. In particular, it can account for the considerable heterogeneity in frailty in the population, the implications of frailty for mortality risk, medical expenditure risk, and labor market outcomes, and the relationship between frailty and income, education, and wealth.

2.3 Model

This section develops a structural life cycle model that can account for the observations documented in the preceding section. The model is a discrete time general equilibrium life cycle model with incomplete markets, heterogeneous agents, and endogenous health.

2.3.1 Agents

The economy is populated by a continuum of ex-ante heterogeneous agents. Agents are indexed by type $s = (j, f, a, \xi, \eta, i, e)$, where j denotes age, f is frailty, a is assets, ξ denotes medical expenditures, η is labor productivity, i is the agent's private health insurance status, and e is educational attainment. The agent's educational attainment is assumed to be permanent over her life cycle and can take on one of two values: college or non-college. Throughout, I let $\Phi(s)$ denote the measure of agents of type s in the stationary distribution.

Agents are endowed with one unit of time that can be allocated to work, ℓ, or leisure. They choose consumption of healthy goods, c_h, and unhealthy goods, c_u, and how much to save, a'. Lastly, they choose whether or not to purchase private health insurance for the following period, i'. Starting at age j_r, all agents receive Social Security benefits that vary with their educational attainment, $SS(e)$, and health insurance from the government in the form of Medicare. Consistent with the data, agents are subject to mortality risk that varies with their age and frailty, $\psi(f, j)$. Agents that survive until age J are assumed to die with probability one. In the event of death, the agent's assets are uniformly distributed across the population by means of lump-sum transfers, B.

The law of motion for frailty is endogenous and depends on the agent's current age, frailty, education, and consumption basket. In particular, agents can affect their frailty transition probabilities, which follow finite-state Markov processes, through consumption of healthy goods:

$$\mathbb{P}\left(f'|f, j, e, c_h\right) = Prob\left(f' \in F \,:\, (f,\, j,\, e,\, c_h)\right). \tag{1}$$

This mechanism that frailty depends on the healthiness of the agent's consumption basket is motivated by a large literature in medicine that studies the health consequences of healthy and unhealthy eating. To illustrate, Danaei et al. (2009), Lim et al. (2012), and Murray et al. (2013) show that poor diets, especially diets low in fruits and those high in sodium, are among the leading causes of cardiovascular diseases, diabetes, diet-related cancers, and obesity. The assumption that frailty transitions depend on age and education enables me to capture both the observation that average frailty increases over the life cycle and the observation that, across all age groups, educated individuals are less frail than non-educated individuals. Disparities in frailty by educational attainment are partially driven by differences in smoking rates. In particular, data from the Centers for Disease Control and Prevention show that about 7 and 23 percent of educated and non-educated adults currently smoke in the United States, respectively. Therefore, letting frailty transitions vary with educational attainment enables me to indirectly account for smoking in the model, albeit in a reduced form way.

Next, building on the positive relationship between frailty and medical spending documented in Table 2.1, I assume that medical expenditure risk is increasing in frailty. In particular, I assume that the transition probabilities for medical expenses follow a finite-state Markov process that depends on the agent's current medical expenses, age, and frailty:

$$\mathbb{P}\left(\xi'|f,j,\xi\right) = Prob\left(\xi' \in X : (f,j,\xi)\right). \tag{2}$$

Similarly, labor productivity is given by a stationary finite-state Markov process:

$$\mathbb{P}\left(\eta'|\eta\right) = Prob\left(\eta' \in E : \eta\right). \tag{3}$$

I assume that healthy and unhealthy goods are perfect substitutes, but that unhealthy goods have a lower relative price, p. This is motivated by evidence from the medical literature that studies the relationship between energy density and energy cost. To illustrate, Drewnowski and Specter (2004) show that healthier foods such as fruits and vegetables tend to have a higher price per calorie than less healthy foods such as sweets and processed meats. Similar results are reported by Drewnowski (2010), who studies the relationship between energy density and energy cost for about 1400 different foods. Additional details are given in the appendix.

Building on Hall and Jones (2007), I let the utility function be given by

$$u(c_h, c_u, f, \ell) = b + \frac{\left((c_h + c_u)^{\gamma} (1 - \ell - \mu(f, j))^{1-\gamma}\right)^{1-\sigma}}{1-\sigma}, \tag{4}$$

where σ governs the relative risk aversion and γ governs the consumption share in intratemporal utility. The constant term in the utility function, b, governs the value of life in the model. Assuming $\sigma > 1$, a positive value of b is needed to ensure positive flow utility in the model. Note that these preferences imply that the marginal utility of consumption falls with consumption. As emphasized by Hall and Jones (2007), extending life does not run into the same kind of diminishing returns. High-income agents can thus avoid the diminishing returns to consumption by instead focusing on choices that extend their lifespan, thereby further increasing their lifetime utility. Consistent with this, Rehm et al. (2016) show that high-income individuals have a healthier diet than low-income individuals. In particular, using data from the National Health and Nutrition Examination Survey, they show that more than 60 percent of low-income individuals have a diet that is classified as poor by the American Heart Association because of its negative implications for various health outcomes such as cardiovascular diseases. In contrast, less than 30 percent of middle-income individuals have a diet that is classified as poor. Lastly, the disutility cost of working, $\mu(f, j)$, is assumed to be increasing in frailty to capture the negative relationship between frailty and labor force participation illustrated in Figure 2.6. It is also increasing in age to capture the drop in labor force participation at age 65. This is further discussed in Section 2.4.5.

2.3.2 Technology

Firms hire labor at wage w and rent capital at rate r from the agents to maximize profits. I assume that the aggregate technology can be represented by a constant returns to scale Cobb-Douglas production function:

$$Y = \theta K^{\alpha} N^{1-\alpha}, \tag{5}$$

where θ denotes total factor productivity, K is the aggregate capital stock, N denotes aggregate labor supply measured in efficiency units, and α is capital's share of income. Output is used for consumption of healthy goods, C_h, unhealthy goods, C_u, investment, $I = K' - (1-\delta)K$, government consumption, G, and to cover medical expenses, M:

$$C_h + pC_u + G + M + K' = \theta K^{\alpha} N^{1-\alpha} + (1-\delta)K, \tag{6}$$

where δ is the rate of depreciation.[1]

[1] For simplicity, instead of having two sectors in the economy, I assume that agents have access to a technology that transforms healthy to unhealthy goods at a fixed price, p. Since healthy and unhealthy goods are assumed to be perfect substitutes, but healthy consumption goods also facilitate better frailty transitions, unhealthy goods must be less expensive than healthy goods (that is, $p < 1$) to avoid a corner solution where all agents only consume healthy goods.

2.3.3 Health insurance

This section presents the different types of health insurance that are available in the economy. Health insurance is available in the form of private health insurance, employer provided health insurance, Medicare, and Medicaid. The agent's insurance status determines what fraction of her medical expenses must be paid out-of-pocket. Throughout, I let χ_P, χ_E, χ_{CARE}, and χ_{CAID} denote the copayment parameter for private health insurance, employer insurance, Medicare, and Medicaid, respectively.

Private health insurance Agents can purchase health insurance for the following period on the individual health insurance (IHI) market. I assume that insurance companies in the IHI market are allowed to price discriminate based on frailty, age, and current medical spending. That is, insurance pools are given by the triplet (f, j, ξ). Moreover, for computational simplicity, I assume that insurance premia are actuarially fair for each insurance pool. Lastly, I assume that Medicare is the primary payer for all agents aged j_r and older in the model. Hence, in the event that an agent has both Medicare and private health insurance, Medicare pays first. Insurance premia, $\pi(f, j, \xi)$, are then given by

$$\pi(f,j,\xi) = \begin{cases} \frac{\psi(f,j)(1-\chi_P)\int \xi' \mathbb{P}(\xi'|f,j,\xi)}{(1+r')} & \text{if } j < j_r - 1 \\ \\ \frac{\psi(f,j)(1-\chi_P)\chi_{CARE}\int \xi' \mathbb{P}(\xi'|f,j,\xi)}{(1+r')} & \text{if } j \geq j_r - 1, \end{cases} \tag{7}$$

where the second line follows from the assumption that Medicare is the primary payer for all the elderly.

Employer sponsored insurance I follow Conesa et al. (2018) and assume that a fraction of workers work for an employer that provides health insurance. The employer pools the medical expenses of all its employees and then splits these costs evenly between the workers. This gives the following expression for the employer-provided insurance premium:

$$\pi_E = \frac{(1-\chi_E)\int \xi \Phi(ds|i=i_E)}{\int \Phi(ds|i=i_E)}, \tag{8}$$

where by $\Phi(ds|i=i_E) = \Phi(dj \times dh \times da \times d\xi \times d\eta \times \{i_E\})$ it is understood that the integral is over all types s but restricted to agents with employer-provided health insurance, $i = i_E$. Following Conesa et al. (2018), I assume that agents cannot have both employer-provided and private health insurance, and that agents cannot opt out of employer-provided insurance.

Medicare and Medicaid The government runs two health insurance programs: Medicare and Medicaid. As noted earlier, Medicare covers a share $1-\chi_{CARE}$ of health expenses of all agents aged j_r and older. Medicaid, on the other hand, is a means-tested program that provides health insurance to the poor. Consistent with program rules, I model two ways to qualify for Medicaid. First, agents are eligible for Medicaid if the sum of their gross income and interest earnings is below a threshold y^{CAT}. Second, agents also qualify for Medicaid if the sum of their gross income and interest earnings net of out-of-pocket medical expenses is below a threshold y^{MN}. I refer to the two eligibility criteria as categorical eligibility and eligibility based on medical need, respectively.

2.3.4 Government

The government engages in three activities in the model. First, as already noted, it provides public health insurance in the form of Medicare and Medicaid. Second, it runs a Social Security program that provides public pensions for the elderly. In reality, Social Security payments are tied to an individual's earnings history. To account for this, I would have had to add individual earnings histories as an additional continuous state in the model. Therefore, for computational simplicity, I assume that benefits are independent of individual earnings histories. This means that the Social Security program will be more progressive in the model than it is in the United States. To partially correct for this, I let benefits vary with educational attainment, which is correlated with lifetime earnings both in the model and in the data. Let $d(e)$ denote the Social Security replacement rate conditional on educational attainment. Building on Conesa and Krueger (2006), I let Social Security benefits be given by

$$SS(e) = \frac{d(e)\,wN}{\int \Phi(ds|j < j_r)}, \tag{9}$$

where by $\Phi(ds|j < j_r) = \Phi(\{1,\ldots,j_r - 1\} \times dh \times da \times d\xi \times d\eta \times di)$ it is understood that the integral is over all types s but restricted to agents of age $j \leq j_r - 1$. Finally, the government runs a welfare program that guarantees minimum consumption. To qualify for this program in the model, agents have to forfeit all assets and earnings. In return, the government pays for all out-of-pocket medical expenses and guarantees a minimum consumption level, $\underline{c}$.

The government finances its expenditures by means of two taxes: a consumption tax, τ_c, and progressive income taxes, $T(y)$. Following Gouveia and Strauss (1994), I let the income tax schedule be given by

$$T(y) = a_0 \left(y - \left(y^{-a_1} + a_2 \right)^{-\frac{1}{a_1}} \right), \tag{10}$$

where y denotes income and (a_0, a_1, a_2) are parameters that determine the shape of the tax function. Here, a_0 determines the maximum marginal and average tax rate as $lim_{y\to\infty}(T(y)/y) = lim_{y\to\infty}T'(y) = a_0$. The second parameter, a_1, governs the progressivity of the tax function. To illustrate, for $a_1 = -1$ we obtain a constant tax independent of income, $T(y) = -a_0 a_2$, for $a_1 \to 0$ we obtain a proportional tax function, $T(y) = \alpha_0 y$, and for $a_1 > 0$ we obtain a progressive tax system where average and marginal tax rates are increasing with income. Lastly, the third parameter, a_2, affects the average income tax rate in the economy.

Let G_H denote total government spending on health care and welfare, and let G denote government consumption. For simplicity, I assume that government consumption is exogenous. Its only purpose is to equalize the size of the government sector in the model and the data, thereby ensuring that the tax burden in the model is consistent with the data. Throughout, I assume that the government balances its budget period-by-period. I then obtain the following expression for the government budget constraint:

$$G + G_H + \int SS(e)\,\Phi(ds|j \geq j_r) = \tau_c (C_h + pC_u) + \int T(y)\,\Phi(ds), \tag{11}$$

where by $\Phi(ds|j \geq j_r) = \Phi(\{j_r, \ldots, J\} \times dh \times da \times d\xi \times d\eta \times di)$ it is understood that the integral is over all types s but restricted to agents of age $j \geq j_r$.

2.3.5 Agent problem

The agent's choice set depends on her age. Throughout, I use the word *young* to denote agents less than age j_r and *old* to denote agents that are at least j_r years old.

Recall that an agent's type is given $s = (j, f, a, \xi, \eta, i, e)$, where j denotes age, f is frailty, a is assets, ξ denotes medical expenditures, η is labor productivity, i is the agent's private health insurance status, and e is educational attainment. Let $V^I(s)$ denote the value of young agents that do not work for an employer that provides health insurance. The value function is given by:

$$
\begin{aligned}
V^I(s) \quad = \quad & \max_{c_h, c_u, \ell, a', i'} u(c_h, c_u, f, \ell) + \beta \psi(f, j) \textstyle\sum_{\eta'} \sum_{\xi'} \sum_{f'} \mathbb{P}(\eta' | \eta) \mathbb{P}(\xi' | \xi, f, j) \mathbb{P}(f' | f, j, e, c_h) V(s') \\
& \text{s.t.} \quad c + pc_h + a' + m(\xi, i) + \mathbb{I}_{i'=i_P} \pi(f, j, \xi) = w\eta\varepsilon(f, j, e)\ell + (1+r)(a+B) \\
& \qquad\qquad -T(y) + TR(s) + \mathbb{I}_{Med}(s)(1 - \chi_{CAID}) m(\xi, i) \\
& \qquad y = w\eta\varepsilon(f, j, e)\ell + r(a+B) \\
& \qquad m(\xi, i) = \mathbb{I}_{i=i_P} \chi_P \xi + (1 - \mathbb{I}_{i=i_P}) \xi \\
& \qquad i' \in \{i_S, i_P\} \\
& \qquad c_h, c_u, \ell, a' \geq 0.
\end{aligned} \tag{12}
$$

$i = i_P$ means the agent has private health insurance and $i = i_S$ means the agent is self-insured. The indicator function, $\mathbb{I}_{Med}(s, \ell)$, equals one if the agent qualifies for Medicaid. Medicaid covers a share $1 - \chi_{CAID}$ of out-of-pocket medical expenses, $m(\xi, i)$, which is given by ξ for self-insured agents and $\chi_P \xi$ for agents with that purchased private health insurance in the preceding period. $TR(s)$ denotes transfers to agents whose resources net of medical expenditures are less than $\underline{c}$. Lastly, labor earnings depend on the agent's stochastic labor productivity, η, and deterministic life cycle productivity, $\varepsilon(f, j, e)$, the last of which varies with the agent's age, frailty, and education.

Similarly, let $V^E(s)$ denote the value of young agents that work for an employer that provides health insurance. The value function is given by:

$$\begin{aligned}
V^E(s) \quad &= \quad \max_{c_h, c_u, \ell, a'} u(c_h, c_u, f, \ell) + \beta \psi(f, j) \textstyle\sum_{\eta'} \sum_{\xi'} \sum_{f'} \mathbb{P}(\eta' | \eta) \mathbb{P}(\xi' | \xi, f, j) \mathbb{P}(f' | f, j, e, c_h) V(s') \\
&\text{s.t.} \quad c_h + pc_u + a' + m(\xi, i) + \pi_E = w\eta\varepsilon(f, j, e)\ell + (1+r)(a+B) \\
&\qquad\qquad -T(y) + TR(s) + \mathbb{I}_{Med}(s)(1 - \chi_{CAID}) m(\xi, i) \\
&\qquad y = w\eta\varepsilon(f, j, e)\ell + r(a+B) \\
&\qquad m(\xi, i) = \chi_E \xi \\
&\qquad i' = i_E \\
&\qquad c_h, c_u, \ell, a' \geq 0.
\end{aligned} \tag{13}$$

All agents with employer-provided insurance pays a health insurance premium π_E that is used to cover a fraction $1 - \chi_E$ of the workers' health care expenses.

Lastly, let $V^{O}(s)$ denote the value of old agents, which is given by

$$\begin{aligned}
V^{O}(s) \quad = \quad & \max_{c_h,c_u,\ell,a',i'} u(c_h,c_u,f,\ell)+\beta\psi(f,j)\textstyle\sum_{\eta'}\sum_{\xi'}\sum_{f'}\mathbb{P}(\eta'|\eta)\,\mathbb{P}(\xi'|\xi,f,j)\,\mathbb{P}(f'|f,j,e,c_h)\,V(s') \\
\text{s.t.} \quad & c_h+pc_u+a'+m(\xi,i)+\mathbb{I}_{i'=i_P}\pi(f,j,\xi)=SS(e)+w\eta\varepsilon(f,j,e)\ell+(1+r)(a+B) \\
& \quad -T(y)+TR(s)+\mathbb{I}_{Med}(s)(1-\chi_{CAID})\,m(\xi,i) \\
& y=SS(e)+w\eta\varepsilon(f,j,e)\ell+r(a+B) \\
& m(\xi,i)=\mathbb{I}_{i=i_P}\chi_P\chi_{CARE}\xi+(1-\mathbb{I}_{i=i_P})\,\chi_{CARE}\xi \\
& i'\in\{i_S,i_P,i_E\} \\
& c_h,c_u,\ell,a'\geq 0.
\end{aligned} \tag{14}$$

All agents start receiving Medicare and Social Security benefits at age j_r. Neither program is tied to retirement, and hence agents continue to receive both Medicare and Social Security benefits even if they choose to work in old age. Out-of-pocket medical expenses are given by $\chi_{CARE}\xi$ for agents that did not purchase private health insurance in the preceding period and $\chi_{CARE}\chi_P\xi$ for agents with private health insurance.

2.3.6 Definition of equilibrium

Given Social Security replacement rates $d(e)$, copayment parameters χ_P, $\chi_{CARE,}$ and χ_{CAID}, and initial conditions for capital K_1 and the measure of types Φ_1, an *equilibrium* in this model is a sequence of model variables such that:

1. Given prices, insurance premia, government policies, and accidental bequests, agents maximize utility subject to their constraints.

2. Factor prices satisfy marginal product pricing conditions.
3. Government policies satisfy the government budget constraint.
4. Goods, factor, and insurance market clearing conditions are met.
5. Aggregate law of motion for Φ is induced by the policy functions and the stochastic processes for idiosyncratic risk.

2.4 Calibration

This section describes how I map the model to the data. I start by discussing how I calibrate the preference and technology parameters, and how I compute health insurance coinsurance rates and Medicaid income limits. Next, I explain how I derive the frailty transition probabilities, the survival probabilities, and the medical expenditure probabilities. The following two subsections explain how I estimate the deterministic labor productivity profiles, how I calibrate both the stochastic labor productivity process and the disincentive effects of work, and how I parameterize the income tax schedule. Lastly, the final subsection compares how well the model matches both targeted and non-targeted moments.

2.4.1 Preference, technology, and health insurance parameters

Agents enter the model at age 20 and have a maximum life span of 100 years. Each model period corresponds to one year. The retirement age, j_r, is set to 46 such that agents start receiving Medicare and Social Security benefits at age 65. I set the annual population growth rate to 1.1 percent to match recent US population growth rates as reported by the Census. Following Castañeda et al. (2003), I set the depreciation rate, δ, to 0.059, and capital's share of income, α, to 0.360. The coefficient of risk aversion, σ, is set to 3, which is within the range of 2 to 4 commonly used in structural life cycle models. Lastly, I set the consumption share in intratemporal utility, γ, to 0.574 to match estimates in French (2005).

Following Conesa et al. (2018), I let the coinsurance rate on each insurance plan be given by the average share of expenses covered across individuals in the MEPS with that insurance plan as their primary insurance provider, where the primary insurance provider is defined as the insurer that pays for the largest share of the individual's expenses. This gives a coinsurance rate on private insurance and employer provided insurance of 22.9 percent, and a coinsurance rate on Medicare and Medicaid of 29.1 and 13.8 percent, respectively. Next, using estimates by the Kaiser Family Foundation, I set the Medicaid categorical income limit, y^{CAT}, equal to 90.2 percent of the federal poverty level (FPL), which was about $\$11,700$ in 2014, and set the Medicaid medically needy income limit, y^{MN}, equal to 41.9 percent of the FPL. A summary of the parameters determined outside the model is given in Table 2.2.

The final set of parameters is determined jointly in equilibrium. I normalize the total factor productivity parameter, θ, such that GDP per capita equals one in the benchmark model. I calibrate the discount factor, β, to match a capital to GDP ratio of 3.3. I calibrate the Social Security replacement rates, $d(e)$, to match average Social Security benefits by educational attainment. The constant term in the utility function, b, determines the value of life in the model. This parameter affects how much agents are willing to invest in healthy consumption. I calibrate this parameter to match average frailty at age 60. This gives a value of 6.5 for b, which is close to the value of 6.75 estimated by Ozkan (2014). Recall from equation (4) that healthy and unhealthy goods are assumed to be perfect substitutes, but that unhealthy goods have a lower relative price, p. This relative price governs the tradeoff between higher current period consumption and lower next-period expected frailty.

Table 2.2: Parameters determined outside the model

Parameter	Description	Source	Value
Preference and technology parameters			
J	Maximum life span		81
j_r	Agents receive SS and Medicare		46
α	Capital share of income	Castañeda et al. (2003)	0.360
δ	Depreciation rate	Castañeda et al. (2003)	0.059
γ	Consumption share in utility	French (2005)	0.574
σ	Risk aversion		3.000
	Population growth rate		0.011
	Percent of agents with college degree		0.320
Health insurance coinsurance rates			
χ_P	Private insurance coinsurance rate	MEPS	0.229
χ_E	Employer insurance coinsurance rate	MEPS	0.229
χ_{CARE}	Medicare coinsurance rate	MEPS	0.291
χ_{CAID}	Medicaid coinsurance rate	MEPS	0.138
Medicaid income limits			
y^{CAT}	Medicaid categorical income limit	Kaiser Family Foundation	0.197
y^{MN}	Medicaid medically needy income limit	Kaiser Family Foundation	0.092

Notes: The table lists the parameters that are determined outside the model.

I therefore calibrate this relative price to match life expectancy inequality by income, here defined as the difference in life expectancy at age 40 by income quartile. Lastly, I set the guaranteed consumption level, $\underline{c}$, equal to 0.070, which corresponds to a minimum consumption level of about $3,500 in 2014. This is close to the value used by De Nardi et al. (2017) and lies within the $1,000 to $7,000 range commonly used in the literature. A summary of the parameters determined outside the model is given in Table 2.3.

Table 2.3: Parameters determined jointly in equilibrium

Parameter	Description	Target	Value
Preference and technology parameters			
θ	Total factor productivity	GDP per capita = 1	0.647
β	Discount factor	Capital to output = 3.3	0.955
$d(c)$	SS college replacement rate	Avg. SS benefit college $\approx 14,200$	0.374
$d(nc)$	SS non-college replacement rate	Avg. SS benefit non-college $\approx 11,900$	0.313
b	Constant term in utility function	Average frailty at 60	6.504
p	Relative price unhealthy goods	Life expectancy diff. by income quartile at 40	0.923
$\underline{c}$	Guaranteed consumption	Minimum consumption $\approx 3,500$	0.070
	Eligible for emp. ins.	Perc. with pvt. or employer ins. = 0.508	0.484
	Scale for health care costs	health care spending to GDP = 0.165	1.803
Labor productivity process parameters			
σ_η	Variance	Labor earnings GINI = 0.670	3.941
η_{top}	Productivity at the top	Labor earnings top 1 percent = 0.148	25.021
π_{top}	Probability at the top	Laborn earnings top 10 percent = 0.435	0.004
ρ_η	Persistence	2-year persistence: Bottom 80 percent = 0.940	0.911
ρ_{top}	Persistence at the top	2-year persistence: Top 1 percent = 0.580	0.792

Notes: The table lists the parameters that are determined jointly in equilibrium.

2.4.2 Frailty transition probabilities

Recall that I normalize an individual's frailty in the data from 0 to 1. I follow the same approach in the model.[2] Next, I assume that agents draw their frailty from an ordered logistic distribution. The frailty transition probabilities are then given by the standard ordered logistic formula:

$$\mathbb{P}\left(f' = i|\mathbf{x}\right) = \frac{1}{1+\exp\left(-\kappa_i + \mathbf{x}\beta\right)} - \frac{1}{1+\exp\left(-\kappa_{i-1} + \mathbf{x}\beta\right)}, \tag{15}$$

where $\mathbf{x} = (f, e, c_h)$, β is a vector of parameters, and the κ's denote cutoffs. I use the following procedure to derive these probabilities. I start by estimating this ordered logistic function in the data by replacing healthy consumption with income. I then use the derived

[2] I use a uniformly spaced grid with 25 points for the frailty grid.

coefficient estimates as initial guesses for the coefficients of $\mathbb{P}\left(f'|f,j,e,c_h\right)$ in the model. Next, I estimate the same coefficients using simulated data from the model. Lastly, I iterate on the preference parameters and the coefficients of the ordered logistic function until the model and data regression estimates coincide.[3]

2.4.3 Survival probabilities

Using my MEPS sample, I first estimate age and frailty-specific survival probabilities by running a logistic regression of mortality on age, frailty, higher-order moments, and demographic controls. Let these probabilities be denoted by $\hat{\psi}(f,j)$. Next, I compute average frailty by age, $\Lambda(f,j)$, for 20-85 year-olds in the data. I then extrapolate using an exponential trend to obtain average frailty for 86-100 year-olds. This step is necessary since the MEPS top-codes age at 85. Lastly, for each age, I scale the survival probabilities estimated from the MEPS by a factor $\alpha(j)$ to match age-specific survival probabilities in 2014 as reported by the Social Security Administration, $\psi(j)^{SSA}$, where the scaling parameter solves the following equation:

$$\alpha(j) = \frac{\psi(j)^{SSA}}{\sum_f \Lambda(f,j)\,\hat{\psi}(f,j)},\ \forall j. \tag{16}$$

This step is necessary since the MEPS does not sample institutionalized individuals, which leads to upward biased survival estimates. Survival probabilities by age and frailty are then given by $\psi(f,j) = \alpha(j)\,\hat{\psi}(f,j)$. A similar approach is used by Attanasio et al. (2010) to derive health-specific survival probability estimates.

[3]The coefficients of the frailty transition matrix are calibrated jointly with the other parameters in the model to ensure that the model matches all targeted moments.

2.4.4 Medical expenditure transition probabilities

I use three points for the medical expenditure grid, ξ.[4] To estimate transition probabilities between these points, I split medical expenses into the following categories in the data: low, high, and very high, where the three categories refer to individuals in the MEPS with medical expenses less than $\$1,000$, between $\$1,000$ and $\$40,000$, and greater than $\$40,000$, respectively. I assume that these expenses are drawn from an ordered logistic distribution that depends on current medical spending, age, and frailty. Estimates of this regression are reported in appendix Table A.4. Transition probabilities between the medical expenditure states, $\mathbb{P}(\xi'|\xi, f, j)$, are then given by the standard ordered logistic formula.

I set the three values of the medical expenditure grid, $\xi \in \{\xi_1, \xi_2, \xi_3\}$, equal to the average value across individuals in the MEPS with low, high, and very high medical expenditures, respectively. These values are then scaled to match the ratio of health care spending to GDP in 2014. Again, this step is necessary since the MEPS does not sample institutionalized individuals, which account for a large share of total medical spending in the United States. Medical expenses derived from the MEPS must therefore be scaled up to match the aggregate statistics reported by the Center for Medicare & Medicaid Services.

2.4.5 Labor earnings

Agents in the model are endowed with one unit of time in every period that can be allocated to work and leisure. I let $\mu(f, j)$, which governs the disutility of work by frailty and age, be given by $\mu(f, j) = \mu_1 \exp(\mu_2 f) + \mathbb{I}_{j \geq j_r} \mu_3$. I calibrate μ_1 and μ_2 to match labor force participation by frailty, and set μ_3 to match labor force participation rates of the elderly.

[4] Because of the size of the state space, I restrict the size of the medical expenditure grid to reduce computational costs.

This last part is necessary to match the large drop in labor force participation at age 65. Although the model qualitatively matches the drop in labor force participation at age 65 due to eligibility for Social Security and Medicare, it falls short quantitatively. An additional disincentive effect, μ_3, is thus needed to match the data.

To reduce computational costs, I use a finite grid for labor supply. In particular, I assume that agents have the choice of working 0, 20, 40, or 60 hours per week. Having a continuous choice for labor supply is problematic in this environment since eligibility for Medicaid depends on current income. Infinitesimal changes in labor supply can thus lead to large discontinuous changes in the value function. That said, this finite grid for labor supply closely matches the observation in the MEPS that the majority of Americans work exactly 20, 40, or 60 hours per week.

I follow Castañeda et al. (2003) and let the stochastic labor productivity shocks be drawn from a right-skewed distribution whose moments are calibrated to match the observed earnings distribution in the United States. This process enriches the standard earnings process commonly used in structural life cycle models with a very high productivity at the top, whose level, probability, and persistence are calibrated to match the high end of the observed earnings distribution. A summary of the calibrated productivity process parameters is given at the end of Table 2.3. Lastly, recall that agents' life cycle labor productivity profiles vary with age, education, and frailty, $\varepsilon(e, f, j)$. I use the age and education-specific labor productivity profiles estimated by Conesa et al. (2018). These profiles are illustrated in appendix figure A1. I then combine these profiles with the estimates of the effect of frailty on hourly wages reported in appendix Table A.5.

2.4.6 Government

As noted in Section 2.3.4, the government finances its spending on consumption, Medicare, Medicaid, Social Security, and welfare by means of two taxes: a consumption tax, τ_c, and a progressive income tax, $T(y)$. Following Mendoza (1994), I set τ_c equal to 0.0567. I set a_0 equal to 0.258 and a_1 equal to 0.768 following estimates in Gouveia and Strauss (1994), and let a_2 adjust such that the government balances its budget period-by-period. Lastly, I set (unproductive) government consumption, G, such that aggregate government spending as a share of GDP equals 20.0 percent in the model. This value lies within the 17.8 and 21.3 percent range observed between 1989 and 2014 as reported by the Bureau of Economic Analysis.

2.4.7 Model fit

This section examines how well the model matches the data. Throughout this section, data and model results are illustrated using solid black and dotted blue lines, respectively. The left panel of Figure 2.7 plots average frailty by age. As shown in the graph, the model is able to match the gradual increase in average frailty over the life cycle. The model can also account for the high concentration of frailty observed in the data. This can be seen in the right panel of Figure 2.7, which compares the cumulative distribution function of frailty for 20 to 85-year-olds in the model and the data.

Recall that I calibrate the parameters of disutility function of work, $\mu(f, j)$, to match labor force participation rate by frailty quartile and the large drop in participation rates at age 65. To examine how well the model matches the data, Table 2.4 compares labor force participation rate by *age groups* in the model and the data, where the data refers to

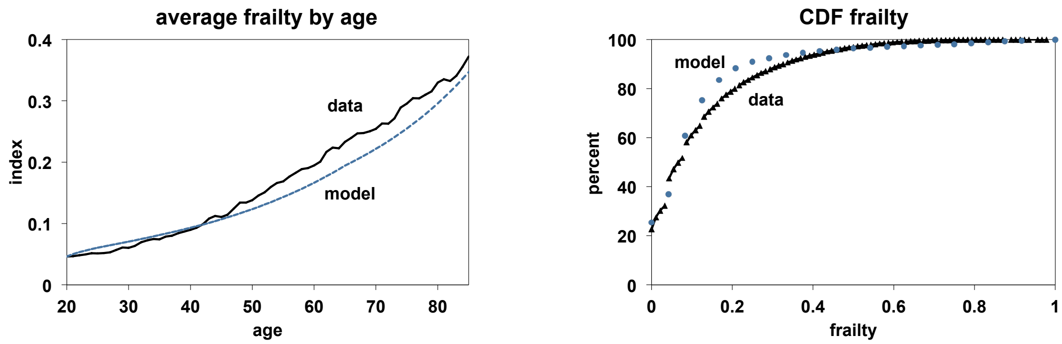

Figure 2.7: Frailty distribution by age: model versus data

Notes: The left panel plots average frailty by age in the model and the data. The right panel plots the cumulative distribution function of frailty for 20-85 year-olds in the model and the data. Data source: MEPS.

Table 2.4: Labor force participation rate by age group: model versus data

	Age group							
	20-24	25-34	35-44	45-54	55-64	65-69	70-74	75+
Data	0.75	0.83	0.84	0.82	0.64	0.29	0.17	0.06
Model	0.79	0.91	0.87	0.78	0.62	0.28	0.21	0.10

Notes: The table reports labor force participation rate by age group in the model and the data. The data refers to the 2006 civilian labor force participation rates by age groups as reported by the Bureau of Labor Statistics.

2006 civilian labor force participation rates by age groups reported by the Bureau of Labor Statistics. I find that the model slightly overestimates the participation rate of 20 to 44-year-olds and the 70+ year olds, but closely matches the participation rate of the rest of the population.

Next, I compare how well the model matches the relationship between wealth and frailty. Recall from Section 2.2.1 that the healthiest 25 percent of 50-65 year-olds have 3.1 times higher average net wealth than the sickest 25 percent. I find that the corresponding statistic in the model is 2.4. This shows that the model qualitatively matches the observed wealth-frailty gradient, but falls slightly short quantitatively.

Table 2.5: Labor earnings distribution (percent)

	Quintiles					Top			
	1st	2nd	3rd	4th	5th	90-95	95-99	99-100	Gini
Data	-0.40	3.19	12.49	23.33	61.39	12.38	16.37	14.76	0.63
Model	0.26	4.79	9.79	20.59	64.56	11.83	16.77	14.78	0.63

Notes: The table reports the labor earnings distribution in the model and the data. Cells denote shares of total. Data source: Kuhn and Ríos-Rull (2013).

Lastly, I examine how well the model matches the earnings and wealth distribution. As explained in Section 2.4.5, I calibrate the labor productivity process to match the empirical earnings distribution as documented by Kuhn and Ríos-Rull (2013). A comparison of the earnings distribution in the model and the data is given in Table 2.5. As shown in the first part of the table, the model closely matches each quintile's share of total earnings. Moreover, the model also matches the share of earnings held by the top quintile of the earnings distribution, and the overall concentration of earnings as measured by the GINI coefficient. Table 2.6 compares the wealth distribution in the model and the data. Although I do not calibrate the model to match the wealth distribution, the model is able to match the share of wealth held by each of the five wealth quintiles. It also matches the share of wealth held by the top 10 percent of the distribution, but fails to account for the very high wealth holdings of the top 1 percent. As a result, the model's GINI coefficient of wealth falls short of what we observe in the data.

2.5 Results

This section studies how policies can be designed to reduce life expectancy inequality and examines what the consequences of these policies are for welfare and the macroeconomy.

Table 2.6: Wealth distribution (percent)

	Quintiles					Top			
	1st	2nd	3rd	4th	5th	90-95	95-99	99-100	Gini
Data	-0.39	1.74	5.72	13.43	79.49	12.62	23.95	29.55	0.78
Model	0.22	0.48	4.93	18.27	76.10	15.98	25.79	15.15	0.75

Notes: The table reports the wealth distribution in the model and the data. Cells denote shares of total. Data source: Kuhn and Ríos-Rull (2013).

Recall that there is a two-way relationship between income and frailty in the model. High frailty leads to low income, which in turn leads to lower consumption of healthy goods and hence higher next-period expected frailty. This two-way relationship suggests that the government can improve population health outcomes by expanding insurance coverage or by increasing redistribution, both of which facilitate higher investments in healthy consumption in times of low net income. To test this hypothesis, I study the effects of three policy reforms. First, I examine the effects of a universal health insurance reform. Next, I compare the results from this reform with the results from two Medicaid reforms designed to increase insurance coverage among low-income agents. Lastly, I study the effects of three income tax reforms designed to alter the level of redistribution between high- and low-income agents in the economy.[5]

[5]The appendix provides additional details about the determinants of life expectancy inequality. In particular, I examine how income inequality and differences in educational attainment affect life expectancy inequality by studying two counterfactual environments. First, I study a model where frailty transitions are exogenous and hence independent of consumption. As reported in the appendix, I find that life expectancy inequality is 76 percent lower in this model than in the benchmark model. Next, I study a model where education affects income, but not frailty transitions. This model enables me to decompose the income effect of education from the direct health-protective effect of education. I find that life expectancy inequality is 22 percent lower in this model than in the benchmark model. These findings suggest that both income inequality and differences in educational attainment are key drivers of life expectancy inequality in the United States.

2.5.1 Universal health insurance

To understand how health insurance inequality affects life expectancy inequality, I compare the benchmark model with a counterfactual economy with universal health insurance. In particular, I consider an economy where the government covers 86.2 percent of all health care expenses, and finances these costs by increasing the average income tax rate in the economy (that is, by increasing a_2 in equation (11)). In effect, this means that the entire population becomes eligible for Medicaid.

The first part of Table 2.7 shows how life expectancy varies in the two models. I find a negative relationship between average life expectancy and health insurance inequality. In particular, I find that eliminating health insurance inequality leads to a small increase in life expectancy at age 20 from 78.90 to 79.00 years. Note, however, that this result is likely to provide a lower bound on the effect of health insurance inequality on life expectancy due to the assumption that health care spending does not improve health outcomes in the model. Although there is limited evidence of positive health returns to medical spending, there are papers that suggest that expanded insurance coverage can improve health outcomes. As an example, Card et al. (2008, 2009) find evidence that utilization of health care services increases when people become eligible for Medicare, and that this increased utilization lowers mortality rates.

Next, I examine how the reform affects life expectancy inequality. A comparison of the benchmark model and the model with universal health insurance shows that the richest 25 percent of 40-year-olds have the same life expectancy in the two environments. Life expectancy among agents in the first income quartile, on the other hand, increases by 0.41 years following the reform. Consequently, I find that life expectancy inequality declines from 7.01 to 6.60 years. That is, eliminating health insurance inequality through a uni-

Table 2.7: Comparative statics: Economy with and without universal health insurance

Variable	Benchmark	Universal insurance
Life expectancy (years)		
Life expectancy at 20	78.90	79.00
Life expectancy difference at 40 by income quartile	7.01	6.60
Macroeconomic aggregates (% change from bench.)		
GDP per capita	-	0.50
Capital per capita	-	0.68
Effective labor supply per capita	-	0.40
health care spending per capita	-	-0.18
Total consumption per capita	-	0.14
Government spending (percent)		
Government spending to GDP	20.00	24.56
Public health care spending to GDP	9.37	13.78
Social Security spending to GDP	4.77	4.82
Inequality		
GINI coefficient pre-tax earnings	0.67	0.67
GINI coefficient wealth	0.72	0.72

Notes: The table compares the benchmark model results with results from a counterfactual model with universal health insurance where the government covers 86.2 percent of all health care costs.

versal health insurance reform leads to a 6 percent decline in life expectancy inequality. Qualitatively similar results are obtained by Ozkan (2014) and Kotera (2018). To illustrate, Ozkan (2014) finds that universal health insurance leads to a 1.25 year increase in life expectancy at birth for individuals in the bottom income quintile.

So far, I have shown that the government can increase average life expectancy and reduce life expectancy inequality by expanding insurance coverage. To fully understand the consequences of an insurance expansion, I next examine the macroeconomic effects of the reform. The second part of Table 2.7 compares the key macroeconomic aggregates in the two model environments. I find that the reform leads to an increase in GDP per capita by

0.50 percent, owing to a 0.40 percent increase in labor supply per capita and a 0.68 percent increase in capital per capita. The increase in labor supply is driven by four factors. First, the increase in the size of the population caused by the rise in average life expectancy leads to a higher labor force. Second, the increase in wages brought about by the increased capital accumulation incentivizes agents to increase labor supply. Third, the presence of universal health insurance means that low-income and medically indebted agents no longer have an incentive to lower their labor supply to qualify for Medicaid. Lastly, the improvements in health outcomes caused by higher investments in healthy consumption increases labor productivity. I find that these effects more than offset the negative effect on labor supply brought about by the higher tax burden needed to finance the reform. Similarly, two factors contribute to the increase in capital accumulation. First, the increase in average life expectancy incentivizes agents to increase their savings in anticipation of a longer retirement period. In addition, the increase in wages and the reduction in average out-of-pocket medical spending brought about by the decline in the uninsurance rate contribute to an increase in disposable income, thereby facilitating higher saving, especially among the previously uninsured. My findings show that these effects more than offset the negative effect on saving brought about by both the higher tax burden and the lower medical expenditure risk, the last of which lowers the need for precautionary saving.

Recall that health care spending is exogenous in the model, but that the distribution of expenditure shocks depends on frailty. An expansion of insurance coverage therefore has two effects on health care spending in the model. On the one hand, increased insurance coverage lowers out-of-pocket medical expenditure risk, thereby reducing incentives to maintain low frailty by investing in healthy consumption. This is the standard moral hazard effect of health insurance. On the other hand, higher disposable income brought about by the reduction in out-of-pocket medical spending facilitates higher investments in healthy consumption, especially among the previously uninsured. This lowers the probability of experiencing high frailty shocks, thereby reducing expected health care spending. I find that the latter effect more than offsets the moral hazard effect of health insurance. Accordingly, the introduction of universal health insurance leads to a 0.18 percent decline in health care spending per capita. This result is qualitatively inconsistent with Ozkan (2014), who finds that universal health insurance leads to higher health care spending since the increase in preventive medical spending exceeds the reduction in curative medical spending. That said, the magnitude of the increase in health care spending in his model following the introduction of universal health insurance is very close to zero. Lastly, although the tax burden is higher in the model with universal health insurance, I find that consumption per capita increases by 0.14 percent following the reform due to the increase in wages and the reduction in out-of-pocket medical spending.

Next, I examine the implications for inequality, and find that both wealth and earnings inequality are unaffected by the universal health insurance reform. In terms of the fiscal implications of the reform, I find that the introduction of universal health insurance leads to a 4.56 percentage point increase in government spending to GDP. Most of the increase is driven by higher public health care spending, which increases from 9.37 to 13.78 percent of GDP. The increase in average life expectancy also generates an increase in the size of the elderly population, which in turn leads to a 0.05 percentage point increase in Social Security spending to GDP.

I end the subsection by studying the welfare implications of the reform. I quantify the welfare effects by means of consumption equivalent variations. In particular, I measure how much consumption must change in all periods and contingencies in the benchmark economy to make an unborn agent under the veil of ignorance indifferent between the benchmark economy and the economy with universal health insurance. This provides a measure of the *ex ante* welfare effect of the reform, that is, before the agent knows her type. I find that the universal health insurance reform leads to higher *ex ante* welfare. The results show that consumption must increase by 1.82 percent in all periods and contingencies in the benchmark model to equalize *ex ante* welfare in the two economies. A decomposition of this finding shows that about 20 percent of the welfare gain can be attributed to the improvements in average life expectancy. The remaining share is largely driven by the reduction in the dispersion of consumption (that is, by the improvements in consumption smoothing) brought about by the reduction in the uninsurance rate.

2.5.2 Expanding Medicaid

The previous subsection showed that universal health insurance would lead to higher life expectancy, lower life expectancy inequality, lower health care spending, higher GDP per capita, and generate large welfare gains, even after controlling for the increased tax burden needed to finance the reform. These results were driven by higher consumption of healthy goods, especially among the previously uninsured. This subsection studies the effects of two alternative health insurance reforms that are designed to increase insurance coverage among low-income agents. In particular, motivated by the recent expansion of Medicaid under the ACA, I study the effects of two Medicaid reforms. First, an expansion of Medicaid to all agents with income no greater than 138.0 percent of the FPL. To do this, I increase the Medicaid categorical income limit, y^{CAT}, from its benchmark value of 90.2 percent of the FPL to 138.0 percent of the FPL. Next, an expansion of Medicaid to all agents with income net of medical expenses no greater than 138.0 percent of the FPL. To do this, I increase the Medicaid medically needy income limit, y^{MN}, from its benchmark value of 41.9 percent of the FPL to 138.0 percent of FPL. For brevity, I will refer to the two reforms as the categorical reform and the medically needy reform, respectively. Both

Table 2.8: Comparative statics: Expanding Medicaid

Variable	Benchmark	y^{CAT}	y^{MN}
Life expectancy (years)			
Life expectancy at 20	78.90	78.92	78.94
Life expectancy difference at 40 by income quartile	7.01	6.80	6.49
Macroeconomic aggregates (% change from bench.)			
GDP per capita	-	-1.07	-2.38
Capital per capita	-	-1.83	-4.48
Effective labor supply per capita	-	-0.63	-1.18
health care spending per capita	-	-0.06	-0.12
Total consumption per capita	-	-1.04	-2.89
Government spending (percent)			
Government spending to GDP	20.00	20.75	22.69
Public health care spending to GDP	9.37	9.98	11.79
Social Security spending to GDP	4.77	4.80	4.81
Inequality			
GINI coefficient pre-tax earnings	0.67	0.68	0.68
GINI coefficient wealth	0.72	0.73	0.74

Notes: The table compares the benchmark model results with results from two counterfactual models: an expansion of Medicaid to all agents with income no greater than 138 percent of the federal poverty level ($y^{CAT} = 1.38FPL$), and an expansion of Medicaid to all agents with income net of medical expenses no greater than 138 percent of the federal poverty level ($y^{MN} = 1.38FPL$).

reforms are financed through an increase in average income tax rates, a_2. I then compare the results from these reforms with the results from the universal health insurance reform studied in the preceding section.

Consistent with the universal health insurance reform, I find that both Medicaid reforms lead to higher average life expectancy. In particular, as shown in Table 2.8, I find that the categorical and medically needy reform leads to an increase in average life expectancy from 78.90 to 78.92 and 78.94 years, respectively. Similarly, life expectancy inequality, as measured by the difference in life expectancy at age 40 by income quartile, declines from

7.01 to 6.80 years following the increase in the Medicaid categorical income limit, and to 6.49 years following the increase in the Medicaid medically needy threshold. Recall that life expectancy inequality declined by 0.41 years following the introduction of universal health insurance. This shows that expanding the Medicaid medically needy program, which leads to higher insurance coverage among low-income and medically indebted agents, leads to a larger reduction in life expectancy inequality than a universal health insurance reform.

The second part of Table 2.8 studies the implications of the reforms for the macroeconomy. Recall that the universal health insurance reform led to an increase in both capital accumulation and labor supply. As a result, GDP per capita was 0.50 percent higher in the model with universal health insurance than in the benchmark model. In contrast, I find that both Medicaid reforms lead to a reduction in capital accumulation and labor supply. As a result, GDP per capita declines by 1.07 and 2.38 percent following the categorical and medically needy reform, respectively. The reduction in capital per capita is partially driven by lower private insurance take-up rates. That is, the expansion of Medicaid crowds out private insurance enrollment. Since insurance premia enter the capital stock, this leads to a reduction in capital per capita. The reduction in precautionary saving brought about by lower out-of-pocket medical expenditure risk leads to further reductions in capital accumulation. The decline in labor supply can be attributed to the reduction in wages brought about by the reduction in capital accumulation, the increase in income tax rates needed to finance the reforms, and the Medicaid means tests. In particular, because of the Medicaid eligibility limits, uninsured agents with high medical costs have an incentive to reduce their saving and labor supply to qualify for the program. Since Medicaid crowds out private insurance enrollment, the share of privately insured agents in the economy declines following the expansion of Medicaid. As a result, more agents in the economy have an incentive to lower their saving and labor supply to qualify for Medicaid in the event of high medical costs. Note, however, that these adverse implications for the macroeconomy are partially offset by the higher saving and labor supply among agents at the low end of the income distribution. This follows from the observation that the increase in the Medicaid income limit allows these agents to both work and save more without losing eligibility for the program. Lastly, I find that the reduction in income brought about by the reforms leads to a 1.04 and 2.89 percent drop in total consumption per capita following the categorical and medically needy reform, respectively, even though health care spending per capita declines moderately due to the improvements in health outcomes.

The categorical and medically needy reform lead to a 0.75 and a 2.69 percentage point increase in government spending to GDP, respectively, almost all of which is driven by higher public health care spending. Consistent with the universal health insurance reform, I find that earnings inequality and wealth inequality are almost unaffected by the Medicaid reforms. In particular, I find that the GINI coefficient of pre-tax earnings increases from 0.67 to 0.68 under both reforms. Similarly, I find that the GINI coefficient of wealth increases from 0.72 to 0.73 under the Medicaid categorical eligibility reform, and to 0.74 under the Medicaid medically needy reform.

To summarize, the last two subsections have shown that expanding public insurance programs can increase average life expectancy and reduce life expectancy inequality, but that these reforms can have very different implications for the macroeconomy. In particular, whereas the universal health insurance reform led to higher consumption and higher GDP per capita, I found that both Medicaid reforms had negative implications for the macroeconomy because they crowded out private insurance enrollment and incentivized agents to reduce their saving and labor supply to qualify for the program. Therefore, a key advantage of the universal health insurance reform is that it eliminates these work and saving distortions since eligibility for public health insurance is no longer tied to income.

2.5.3 Income tax reform

This subsection examines how life expectancy inequality is affected by the progressivity of the income tax schedule. Recall from Section 2.3.4 that the shape of the income tax schedule is determined by a_0, a_1, and a_2. In what follows, I compare the benchmark model results with the results derived using three alternative tax schedules. First, a proportional tax schedule, where I set a_1 equal to zero and adjust a_0 to balance the government budget (a_2 is undefined when a_1 equals zero). Second, a tax schedule characterized by higher maximum but lower marginal tax rates, where I set a_0 equal to 0.400, a_1 equal to its benchmark value, and adjust a_2 to balance the government budget. Lastly, a tax schedule characterized by higher maximum marginal tax rates but also higher deductions, where I set a_0 equal to 0.400, a_2 equal to its benchmark value, and adjust a_1 to balance the government budget. For brevity, I refer to the three tax schedules as proportional, more progressive, and higher deductions, respectively. These tax schedules, together with the benchmark tax schedule used in the preceding sections, are illustrated in figure 8. The left and right panels plot average and marginal tax rates by income, respectively, where income has been normalized by GDP per capita.

The first part of Table 2.9 shows how life expectancy is affected by the progressivity of the tax schedule. I find a positive relationship between average life expectancy and the progressivity of the tax schedule. That is, higher tax progressivity is associated with higher average life expectancy. This can be seen by comparing the columns in Table 2.9, which shows that life expectancy at age 20 is 0.12 years lower in the economy with proportional taxes than in the benchmark model, but 0.18 and 0.31 years higher in the economy with more progressive taxes and higher deductions, respectively. Moreover, I find a negative relationship between tax progressivity and life expectancy inequality. To illustrate, life

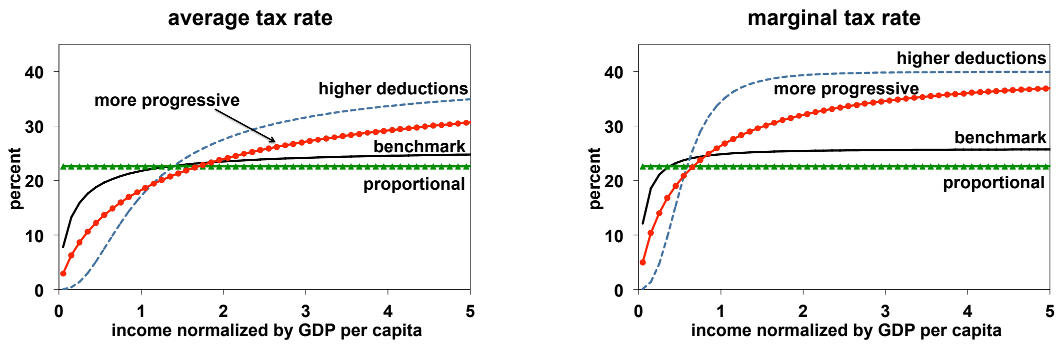

Figure 2.8: Average and marginal income tax rates in the different models

Notes: The left panel plots average tax rates by income, where income has been normalized by GDP per capita. The right panel plots marginal tax rates by income. The income tax function follows from Gouveia and Strauss (1994). See Section 2.3.4 for details.

expectancy inequality by income is 0.33 years higher in the economy with proportional taxes than in the benchmark model, but 0.80 years lower in the economy with higher tax deductions for low-income agents.

Consistent with the optimal tax literature (see for example Conesa and Krueger (2006)), I find that switching to a proportional income tax schedule leads to higher capital per capita and higher labor supply per capita. This reflects the reduced disincentive effects to both save and work for agents at the high end of the income distribution due to the reduced marginal income tax rates for this group. As a result, GDP per capita is 1.77 percent higher in the environment with proportional taxes than in the benchmark model. In contrast, GDP per capita is 1.52 percent lower in the model with more progressive income taxes, and 3.12 percent lower in the model with higher deductions for low-income agents. This follows from the increased marginal income tax rates, and hence increased disincentive effects to both save and work, for agents at the high end of the income distribution. Combining these findings with the results regarding the relationship between tax progressivity and life expectancy inequality discussed earlier suggests that there exists an equity-efficiency tradeoff

Table 2.9: Comparative statics: Economy with different tax codes

Variable	Bench.	Propor.	More prog.	High deduct.
Tax parameters				
α_0	0.258	0.226	0.400	0.400
α_1	0.768	0.000	0.768	2.540
α_2	3.161	N/A	0.598	3.161
Life expectancy (years)				
Life expectancy at 20	78.90	78.78	79.08	79.21
Life expectancy difference at 40 by income quartile	7.01	7.34	6.61	6.21
Macroeconomic agg. (perc. change from bench.)				
GDP per capita	-	1.77	-1.52	-3.12
Capital per capita	-	4.86	-4.24	-8.08
Effective labor supply per capita	-	0.08	0.04	-0.20
health care spending per capita	-	0.43	-0.59	-1.03
Total consumption per capita	-	0.13	-0.86	-2.56
Government spending (percent)				
Government spending to GDP	20.00	19.88	20.08	20.26
Public health care spending to GDP	9.37	9.30	9.34	9.38
Social Security spending to GDP	4.77	4.70	4.85	4.92
Inequality				
GINI coefficient pre-tax income	0.67	0.67	0.67	0.67
GINI coefficient wealth	0.72	0.73	0.70	0.68

Notes: The table compares the benchmark model results with the results derived using three alternative income tax schedules: a proportional tax schedule, a more progressive tax schedule characterized by higher maximum but lower marginal tax rates, and a tax schedule characterized by higher maximum marginal rates but also higher deductions. The income tax functions follow from Gouveia and Strauss (1994). See Section 2.3.4 for details.

in the economy, where equity is measured by life expectancy inequality and efficiency is measured by GDP per capita. That is, higher tax progressivity leads to lower life expectancy inequality but also lower GDP per capita, whereas lower tax progressivity leads to higher life expectancy inequality but also higher GDP per capita. Consistent with this, the United States has less progressive earnings taxes, but higher GDP per capita, than countries in Western Europe (see for example Guvenen et al. (2014)), but both lower average life expectancy and higher life expectancy inequality than these countries. Hence, although these countries differ along several other dimensions that are likely to affect both average life expectancy and life expectancy inequality such as the generosity and size of social insurance programs, the degree of income and wealth inequality, and the educational attainment of their citizens, these findings suggest that part of the difference might be driven by the higher earnings inequality and the lower tax progressivity in the United States.

Higher redistribution benefits low-income agents, who respond to the rise in disposable income by increasing their investments in healthy consumption. This lowers their probability of experiencing high frailty shocks, thereby reducing expected health care spending. High-income agents, on the other hand, respond to the increased tax burden by lowering their consumption of healthy goods, thereby increasing their medical expenditure risk. I find that the former effect is larger than the latter, and hence that increased tax progressivity leads to lower health care spending per capita. As shown in Table 2.9, health care spending per capita is 0.43 percent higher in the environment with proportional taxes than in the benchmark model, but 1.03 percent lower in the environment with higher tax-deductions. Again, consistent with this result, both tax progressivity and health care spending per capita is higher in the United States than in other developed countries.

Higher GDP per capita in the model with proportional taxes contribute to a reduction in government spending to GDP from 20.00 to 19.88 percent. In contrast, I find that government spending to GDP increases to 20.08 percent in the model with more progressive taxes, and to 20.25 percent in the model with higher tax deductions for low-income agents. A decomposition of the increase in government spending in these last two models shows that most of the increase is driven by higher Social Security spending, which follows from the increased old-age dependency ratio brought about by the rise in average life expectancy. Lastly, I examine the implications for inequality, and find that earnings inequality is unaffected by the tax reforms. This result is partially driven by the assumption that agents can only work 20, 40, or 60 hours per week. As a result, the model might underestimate the negative effect of more progressive taxes on hourly labor supply. It is thus plausible that earnings inequality would be moderately affected by the tax reforms if agents were allowed to choose their hourly labor supply from a continuous choice set. Wealth inequality, on the other hand, is affected by the reform. In particular, I find that switching to a proportional tax schedule leads to an increase in the GINI coefficient of wealth from 0.72 to 0.73. Switching to a tax code with higher deductions for low-income agents, on the other hand, leads to a reduction in the GINI coefficient of wealth from 0.72 to 0.68.

2.6 Conclusion

This chapter has developed a structural life cycle model with incomplete markets, heterogeneous agents, and endogenous health. I calibrated the model to match several facts about the population health distribution that I documented by constructing an objective measure of health called a frailty index. In particular, I showed that the model could account for the considerable heterogeneity in frailty in the population, the implications of frailty for mortality risk, medical expenditure risk, and labor market outcomes, and the relationship between frailty and income, education, and wealth.

I then used the model to study how policies could be designed to reduce life expectancy inequality, focusing on health insurance and income tax reforms. I started by studying the effects of a tax-financed universal health insurance reform. I found that the reform led to an increase in life expectancy at age 20 by 0.10 years and to a 0.41 year reduction in the difference in life expectancy at age 40 by income quartile. Moreover, I found that the improvements in the population health distribution brought about by the reform led to a 0.18 percent reduction in health care spending per capita. In addition, GDP per capita increased by 0.50 percent due to both higher capital accumulation and higher labor supply, which in turn contributed to a 0.14 percent increase in consumption per capita. Lastly, I found that the universal health insurance reform led to higher *ex ante* welfare. In particular, I showed that consumption had to increase by 1.82 percent in all periods and contingencies in the benchmark model to make an unborn agent under the veil of ignorance indifferent between the benchmark economy and the economy with universal health insurance.

I then compared the results from the universal health insurance reform with two Medicaid reforms designed to increase insurance coverage among low-income agents. Consistent with the universal health insurance reform, I found that the expansion of Medicaid led to higher average life expectancy and lower life expectancy inequality. On the other hand, I found that the expansion of Medicaid had negative implications for the macroeconomy since it incentivized more agents to reduce their saving and labor supply to qualify for the program. Lastly, I studied the implications of three income tax reforms. I found that the increased redistribution brought about by the higher tax progressivity led to higher life expectancy, lower life expectancy inequality, and lower health care spending per capita. Such increases in tax progressivity, however, had adverse implications for GDP per capita due to the increased disincentive effects to both save and work.

Chapter 3

Medical bankruptcy:
A rationale for uninsurance

3.1 Introduction

Individuals face considerable financial risk from medical expense shocks. Average annual medical expenses increase from $2,500 for 25-year-olds to well over $6,000 for people towards the end of their working-life. Yet, a quarter of the American working-age population, or about 48 million individuals, were uninsured in 2010. Understanding why so many Americans are uninsured has long been at the heart of the health economics literature. Although several theories have been considered, none of them can account for the high uninsurance rate observed in the data (Gruber, 2008).

Recent work by Mahoney (2015) suggests that the high uninsurance rate can be partially explained by the country's bankruptcy laws. The ability to file for Chapter 7 consumer bankruptcy distorts the insurance coverage decision by providing households with a form of high-deductible health insurance. Households thus respond by opting out of private health insurance. Using state-level variations in bankruptcy laws, Mahoney (2015) finds that bankruptcy does have an economically significant effect on the uninsurance rate. His estimates show that if the bankruptcy laws of the least debtor-friendly state were applied nationally, 8 percent of the uninsured population would take up coverage.

Despite the law's importance, very little is known about how much people value the ability to file for medical bankruptcy. For instance, how much would people lose if the law was eliminated? Little is also known about the interaction effects between Chapter 7 and public insurance programs such as Medicaid. In particular, does medical bankruptcy also crowd out Medicaid enrollment? To answer these questions, I develop a general equilibrium overlapping generations model that accounts for several features of the US health insurance system and Chapter 7 of the US bankruptcy code. Agents in the model face idiosyncratic risk to their labor productivity and medical expenses, the latter of which are partially insurable through private insurance and Medicaid. Agents also have the ability to file for medical bankruptcy. Consistent with program rules, filing for bankruptcy releases the debtor from all liabilities. In return, the filer must forfeit all non-exempt assets to the creditor. The filer then starts the following period with a bankruptcy flag on her credit report. This flag remains on the agent's credit report for several years, during which she cannot file for another Chapter 7 discharge.

I calibrate the model to match three sets of data moments: the percentage of Chapter 7 filers due to medical debt; the distribution of medical expenses by age and insurance provider; and both the persistence and distribution of labor earnings. Validation exercises show that the model also matches several non-targeted moments such as the concentration of wealth. I then use the calibrated model to study the effects of an unexpected elimination of Chapter 7. This policy experiment enables me to quantify the effects of medical bankruptcy on the private insurance take-up rate, Medicaid enrollment, and welfare.

Consistent with Mahoney (2015), I find that the ability to file for medical bankruptcy has an economically significant effect on the aggregate insurance rate in the economy. Chapter 7 lowers private insurance enrollment across the working-age population by 2.0 percentage points. With a working-age population of about 190 million, this translates into a reduction in private insurance take-up by 3.8 million individuals. In other words, 8 percent of the uninsured working-age population would take up private coverage if they did not have the option to file for bankruptcy. Medicaid enrollment, on the other hand, is positively affected by Chapter 7. I find that 0.2 percentage points of the working-age population would lose Medicaid eligibility if they did not have the option to file for bankruptcy. This follows from the observation that capital, and hence wages, are higher in the economy without Chapter 7, both of which lower the percentage of households that pass the Medicaid income and asset eligibility tests. As a result, I find that Chapter 7 lowers the insurance coverage rate by 1.8 percentage points.

I then compute a transition path between the steady state of the economy with and without Chapter 7. Private insurance enrollment increases by 1.7 percentage points, or about 3.2 million people, within the first two years following the elimination of Chapter 7. In other words, 85 percent of the increase in the private insurance take-up rate takes place within the first two years of the transition. Medicaid enrollment, on the other hand, declines steadily during the transition. As a result, I find that most of the increase in the insurance coverage rate would take place within the first years following the elimination of bankruptcy. In comparison, about 11.7 million people took up coverage during the first two years of the Affordable Care Act. This shows that reforming the bankruptcy laws can have considerable short run implications for the insurance coverage rate.

Lastly, I study the welfare implications of the ability to file for medical bankruptcy. The results show that 79.6 percent of the population is better off in the economy with medical bankruptcy because of the implicit insurance provided by this option. To quantify the welfare effects on those that benefit and lose from the reform, I compute the dollar value of how much wealth must change in the initial steady state to make the agents weakly better off in the economy with bankruptcy than in the transition to the economy without bankruptcy. I find that those that benefit from the reform experience an average welfare gain that is equivalent to receiving \$400 higher wealth in the steady state with bankruptcy. In contrast, those that lose from the reform experience an average welfare loss that is equivalent to a \$2,900 reduction in wealth. As a result, I find that aggregate welfare is higher in the economy with medical bankruptcy.

Related literature: This chapter is most closely related to Mahoney (2015), who also studies the interaction between bankruptcy and health insurance. Using data from the Medical Expenditure Panel Survey, the Survey of Income and Program Participation, and the Panel Study of Income Dynamics, he shows that households with a higher financial cost of filing for bankruptcy are more likely to be insured. This is consistent with the hypothesis that the ability to file for a Chapter 7 debt discharge crowds out conventional health insurance coverage by providing households with a form of high-deductible health insurance. My chapter differs from Mahoney's along several dimensions. First, I study the effect of bankruptcy on health insurance coverage within a general equilibrium life cycle model. Second, I use the model to study the interaction between Chapter 7 and Medicaid and the interaction between bankruptcy and the cost of health services. Lastly, I compute a transition from an economy with to an economy without Chapter 7. This enables me to quantify the welfare effects of the ability to file for medical bankruptcy.

The model I use follows closely the models used by Pashchenko and Porapakkarm (2013) and Jung and Tran (2016), who develop life cycle models that are consistent with the US health insurance system to quantify the effects of the Affordable Care Act on the insurance take-up rate and welfare. A similar model is used by Conesa et al. (2018) to study the general equilibrium effects of Medicare. I extend their models by introducing the possibility to default on medical debt by filing for bankruptcy. This enables me to quantify the general equilibrium effects of bankruptcy within a life cycle model that features both a rich set of institutional details of the US health insurance system and the key characteristics of the US bankruptcy code.

The chapter builds on the literature that studies consumer bankruptcy in incomplete market economies. Athreya (2002) develops a dynamic stochastic general equilibrium model with unsecured debt and default to study the welfare effects of bankruptcy. He finds large negative welfare effects of bankruptcy, brought about by the observation that bankruptcy lowers the interest rate on savings and increases the interest rate on loans, both of which adversely affect consumption smoothing. Livshits et al. (2007), on the other hand, using a heterogeneous agent life cycle model where households can file for bankruptcy, find that the US bankruptcy law may be welfare improving. Similarly, Chatterjee et al. (2007) study the welfare effects of alternative bankruptcy rules and find that a means test that discourages households with above-median income to file for bankruptcy increases welfare.

Lastly, I contribute to a long tradition in the literature that uses dynamic general equilibrium models with incomplete markets in the spirit of Bewley (1986), Hugget (1993), and Aiyagari (1994). The chapter belongs to the branch of the literature that augments the standard model with medical expenditure risk as in Hubbard et al. (1994), Palumbo (1999), and De Nardi et al. (2010).

The chapter is organized as follows. The next section presents the data used in this chapter. Section 3.4 and 3.5 lay out the environment of the model and present the recursive problem of the household. Next, I explain in Section 3.6 how I map the model to the data. Section 3.7 presents the quantitative analysis. Lastly, Section 3.8 concludes.

3.2 Data

Health expenditure data is obtained from the Medical Expenditure Panel Survey (MEPS), which provides detailed records on demographics, income, medical expenses, and insurance for a nationally representative sample of households. The survey consists of two-year overlapping panels for the period 1996 to 2013, from which I use data for individuals aged 20 to 85 years.[1]

Table 3.1 reports the insurance coverage rate by age in 2010. Here, private insurance

[1]Nominal series are converted to 2010 dollars by means of the GDP deflator.

Table 3.1: Insurance coverage rate by age (MEPS)

Age	18-24	29-29	30-34	35-54	55-64	18-64	65+
Private insurance	50.5	56.7	60.1	68.2	71.3	64.1	39.6
Uninsured	37.4	33.6	28.7	22.3	17.0	25.1	1.0

Notes: This table reports the insurance coverage rate by age groups. Data source: MEPS.

includes both employer-provided and privately purchased health insurance. As shown in the table, 74.9 percent of the working-age population was covered by a private or public health insurance plan. The remaining 25.1 percent of the population, or about 48 million people, was uninsured. The uninsurance rate was declining with age, from more than one-third for individuals in their twenties to 17.0 percent for people around the age of retirement. Most of this reduction could be attributed to private insurance enrollment, which increased from 50.5 to 71.3 percent over the course of an individual's working life. In fact, more than 80 percent of the insured non-elderly population was enrolled in a private or employer-based plan, while 17 percent was covered by Medicaid. In contrast, the insurance coverage rate among the elderly was near universal due to Medicare, which provided health insurance to 97 percent of the elderly population. More than 14 percent of the elderly were also covered by Medicaid. As a result, only 1 percent of the elderly population was uninsured in 2010.

3.3 Model

The following subsections lay out the environment of the economy. The recursive problem of the household and the definition of the equilibrium are given in Section 3.5.

3.3.1 Legal environment

Households can discharge their medical debt by filing for Chapter 7 bankruptcy. To qualify for a Chapter 7 discharge, agents first have to pass an eligibility test. In the model, households pass this test if their income net of allowable expenses is sufficiently low. Households also qualify for a Chapter 7 discharge if both their disposable income is low enough and annual payments of this amount over 5 years would pay back less than 25 percent of the agent's medical debt. More details about the allowable expense limits are given in Section 3.6.1 below. Agents that pass the eligibility test can file for bankruptcy to get their medical debt discharged. In return, they have to forfeit all non-exempt assets, $\underline{a}$, to the creditor (here, the hospital), and pay court filing fees, ψ. Current and future earnings, on the other hand, are retained by the filer. The agent then starts the following period with a bankruptcy flag on her credit record. I assume that households with a flag on their report are ineligible to file for bankruptcy. The flag also lowers their credit score, which increases their unit cost of consumption by ζ. This pecuniary cost of a bad credit record is motivated by the observation that the price of certain services such as auto insurance and mortgage payments are often linked to credit scores. Lastly, in every period, there is an exogenous probability λ that the bankruptcy flag will disappear from the agent's report. That is, there is a chance that the agent will have her record expunged in the following period. This is a computationally convenient way to model the fact that US law does not permit households to file for Chapter 7 more than once every eight years, and that bankruptcy flags only stay on an individual's credit record for a finite number of years.

3.3.2 Households

The economy is populated by a continuum of ex-ante heterogeneous agents. An agent's type is given by $s = (j, a, f, h, \eta, i)$, where j is age, a is assets, f is a bankruptcy flag, h is health status, η is labor productivity, and i is the household's health insurance status. Throughout, I will let the mass of agents of type s be denoted by $\Phi(s)$. The bankruptcy flag is a zero-one indicator, where zero means the household does not have a bankruptcy flag on her credit record. Health is stochastic and follows a stationary finite-state Markov process that depends on age and health. Labor productivity is governed by a deterministic and stochastic process. The deterministic process, ε_j, is a function of the household's age. The stochastic process is given by a finite-state Markov process with stationary transitions over time. Lastly, the household's health insurance status can take on one of three values: the agent is self-insured, has private health insurance, or has health insurance provided by her employer. Health insurance is used to cover non-discretionary medical services, m_{jh}, that vary with the agent's age and health status.

Households are endowed with one unit of time in every period that can be allocated to work or leisure. In addition to leisure, households also derive utility from consumption, c. The period-by-period return function is given by:

$$u(c, \ell) = \frac{\left[c^{\gamma}(1-\ell)^{1-\gamma}\right]^{1-\sigma}}{1-\sigma} \tag{1}$$

All households retire exogenously at age j_r, after which they receive Social Security benefits SS. They also receive government-provided health insurance in the form of Medicare. Lastly, households have a maximum lifespan of J years. They also face a positive probability of death, ψ_{jh}, in every period that depends on their age and health. In the event of death, the agent's assets are uniformly distributed across the population by means of lump sum transfers, B.

3.3.3 Hospital sector

Medical services, $m_{jh,}$ are produced by the hospital sector. The hospital sector takes in the composite consumption good as an input and transforms it one-for-one into medical services. In the model, as in the real world, some households might partially or fully default on their medical bills. Hospital revenues thus depend on both the mass of agents that repay their debt, and on how much the hospital receives in the event of a partial default. Let $d(s)$ denote the probability that a household of type s defaults on her medical bills. Hospital revenues are then given by:

$$\int \left[(1-d(s))\, m_{jh} + d(s)\, \nu(s)\, m_{jh}\right] \Phi(s) \tag{2}$$

where $\nu(s) \in [0,1)$ is the share of medical bills that an agent of type s repays in the event of a default. In order to ensure zero profits, hospitals charge a markup $H \geq 0$ over their marginal cost of production of medical services. H is thus set to balance the following equation:

$$\frac{1}{H}\int m_{jh}\Phi(ds) = \int \left[(1-d(s))\, m_{jh} + d(s)\, \nu(s)\, m_{jh}\right] \Phi(ds) \tag{3}$$

Note that this means that defaulters exert a negative externality on everyone else by increasing the list price of medical services.

3.3.4 Technology

Firms hire labor at wage w and rent capital at rate r from the households to minimize costs. The technology is given by a constant returns to scale Cobb-Douglas production function:

$$Y = AK^{\alpha}N^{1-\alpha} \tag{4}$$

where A denotes total factor productivity, K is the aggregate capital stock, N denotes aggregate labor supply in efficiency units, and α is capital's share of income. Output is used for consumption, investment, and medical expenses, M/H, where the markup H is consistent with zero profits in the hospital sector. This gives the following resource constraint:

$$C + \frac{M}{H} + K' = AK^{\alpha}N^{1-\alpha} + (1-\delta)K \tag{5}$$

where δ is the rate of depreciation.

3.3.5 Health insurance

Health insurance is available in the form of private insurance and public insurance, the latter of which is provided by the government through Medicare and Medicaid. Let χ_P, χ_{CARE}, and χ_{CAID} denote the copayment parameter on private insurance, Medicare, and Medicaid, respectively. These parameters govern the share of medical expenses that are covered by the different types of health insurance that are available in the economy.

Medicare provides health insurance to all elderly individuals. I assume that Medicare is the primary payer for all elderly households. That is, if an elderly agent has both Medicare and private insurance, Medicare pays first.[2] Medicaid, on the other hand, is a means-tested program that provides health insurance to low-income households. In order to qualify for Medicaid in the model, households must pass the categorical or medically needy eligibility test. Households pass the categorical test if the sum of their gross income and interest earnings is less than the categorical limit. Similarly, agents pass the medically needy test if the sum of their gross income and interest earnings net of out-of-pocket medical expenses is less than the medically needy income limit and their assets are less than the asset limit. I refer the reader to Section 3.6.2 for further details about the Medicaid means tests.

Lastly, households can purchase private insurance for the following period at an actuarially fair price that depends on the agent's age and health. This gives the following formula for the private health insurance premium:

$$\pi_{jh} = \begin{cases} \frac{\psi_{jh}(1-\chi_P)\int m_{j'h'} Q_j(h,dh')}{(1+r')} & \text{if } j < j_r - 1 \\ \\ \frac{\psi_{jh}(1-\chi_P)\chi_{CARE}\int m_{j'h'} Q_j(h,dh')}{(1+r')} & \text{if } j \geq j_r - 1 \end{cases} \tag{6}$$

where the second line captures the assumption that Medicare is the primary payer for all elderly agents.

[2] Elderly households have an incentive to purchase private insurance whenever $\chi_{CARE} > 0$, that is, whenever Medicare pays for less than 100 percent of the elderly's medical expenses.

3.3.6 Government

The government balances its budget period-by-period by means of payroll taxes, τ. Let b denote the Social Security replacement rate. Social Security benefits SS are then given by:

$$SS = \frac{bwN}{\int \Phi(\{1,\ldots,j_r-1\} \times dh \times d\eta \times da \times di)} \tag{7}$$

I assume that the proceeds from selling consumption goods at a price higher than marginal cost due to a low credit score goes to the government. That is, government revenues from taxing consumption is given by ζC_1, where C_1 denotes aggregate consumption across households with a bankruptcy flag on their credit report. The government also derives revenues from bankruptcy filing costs ψF_1, where F_1 is the measure of agents that files for a debt discharge. Lastly, let gov denote total government expenditure on Medicare, Medicaid, and food stamps, where the last guarantees a minimum consumption level, $\underline{c}$. Taxes on labor income then have to satisfy:

$$\tau = \frac{SS \int \Phi(\{j_r,\ldots,J\} \times dh \times d\eta \times da \times di) + gov - \zeta C_1 - \psi F_1}{wN} \tag{8}$$

3.4 Household problem

Before I present the household problem, let me go through the timing of events. At the beginning of the period, agents observe the realizations of their idiosyncratic health and labor productivity shock, the former of which determines the value of their medical expenses. Households that ended the preceding period with a bankruptcy flag on their credit record also observe whether or not their record have been expunged. Next, households receive transfers from accidental bequests and from private insurance companies if they purchased health insurance in the preceding period. Eligible households also receive transfers from the government in the form of Medicare, Medicaid, and Social Security. Then, households supply labor and capital to the firm, following which production takes place and households receive factor income. Next, agents pay out-of-pocket medical expenses and make their intertemporal decisions regarding saving and insurance. That is, households choose how much to allocate to consumption and saving, and whether or not to purchase private health insurance for the following period. Households that do not have sufficient resources to pay their medical bills forfeit their assets to the hospital and pay a fraction of their earnings or Social Security benefits net of health insurance expenses. Lastly, all agents that are eligible to file for bankruptcy compare the value of these choices with the value of filing for Chapter 7. If the latter is higher, the agent forfeits her non-exempt assets to the hospital. In return, her debt is discharged, and the agent starts the following period with a bankruptcy flag on her report.

3.4.1 Working-age households

Recall that an agent's type is given by $s = (j, a, f, h, \eta, i)$, where j is age, a is assets, f is a bankruptcy indicator, h is health status, η is labor productivity, and i is the household's health insurance status. Let $V^I(s)$ denote the value of a working-age household of type s that does not file for bankruptcy. Similarly, let $V^B(s)$ denote the value of filing for bankruptcy. Working-age households that qualify for a Chapter 7 debt discharge choose whether or not to file for bankruptcy:

$$V(s) = \max \left\{V^I(s), V^B(s)\right\} \tag{9}$$

where $V^I(s)$ is given by:

$$V^I(s) = \max_{c,a',\ell,i'} u(c+\mathbb{I}_{FS}(s,\ell)\underline{c},\ell) + \beta\psi_{jh}\sum_{f'}\lambda(f'|f)\iint V(s')Q(\eta,d\eta')Q_j(h,dh')$$

$$s.t. \quad (1+f\zeta)c + a' + m_{op} + \mathbb{I}_P(i')\pi_{jh} = w(1-\tau)\varepsilon_j\eta\ell$$
$$+(1+r)(a+B) + \mathbb{I}_{Med}(s,\ell)(1-\chi_{CAID})m_{op} \quad \text{if } x > 0$$

$$(1+f\zeta)c = (1-\kappa)\left[w(1-\tau)\varepsilon_j\eta\ell - \mathbb{I}_P(i')\pi_{jh}\right] \quad \text{if } x \leq 0$$

$$x = w(1-\tau)\varepsilon_j\eta\ell + (1+r)(a+B) + \mathbb{I}_{Med}(s,\ell)(1-\chi_{CAID})m_{op} - m_{op}$$

$$m_{op} = \mathbb{I}_P(i)\chi_P m_{jh} + (1-\mathbb{I}_P(i))m_{jh}$$

$$\mathbb{I}_{\mathbb{P}}(i) = \begin{cases} 1 & \text{if } i = i_P \\ 0 & \text{otherwise} \end{cases}$$

$$c,\ell,a' \geq 0$$

$$i' \in \{i_P, i_S\}$$

Here, $f=0$ means the agent does not have a bankruptcy flag on her report. Conversely, $f=1$ means the agent has a flag on her credit record due to a recent bankruptcy filing. The bankruptcy flag lowers the agent's credit score, which increases her unit cost of consumption by ζ. An insurance status of $i=i_P$ means the household has private insurance, while $i=i_S$ means the household did not purchase private insurance in the preceding period, and is hence self-insured today. Out-of-pocket medical expenses are denoted by m_{op}, which are given by m_{jh} for self-insured households and $\chi_P m_{jh}$ for agents with private health insurance. Households that are eligible for food stamps, $\mathbb{I}_{FS}(s,\ell)=1$, receive a transfer $\underline{c}$

from the government.[3] If the household qualifies for Medicaid, a fraction χ_{CAID} of her out-of-pocket medical expenses will be covered by the government. This is captured by the indicator function on the right-hand side of the first budget constraint, which equals one for all agents that qualify for Medicaid.

Next, x denotes the resources that are available after out-of-pocket medical expenses have been paid. $x \leq 0$ means the agent does not have sufficient resources to pay her medical debt, and is hence forced to partially default on her liabilities. Households that do not have the resources to pay their medical bills have to forfeit their assets to the hospital and pay a fraction κ of their earnings net of health insurance expenses. They are also not permitted to save. In that case, hospital revenues are given by:

$$\begin{aligned}
Hosp &= \kappa\left[w(1-\tau)\varepsilon_j\eta\ell - \mathbb{I}_P(i')\pi_{jh}\right] \\
&+ (1+r)(a+B) + \mathbb{I}_{Med}(s,\ell)(1-\chi_{CAID})m_{jh} < m_{jh} \text{ if } i = i_S \\
Hosp &= \kappa\left[w(1-\tau)\varepsilon_j\eta\ell - \mathbb{I}_P(i')\pi_{jh}\right] + (1+r)(a+B) \\
&+ (1-\chi_P)m_{jh} + \mathbb{I}_{Med}(s,\ell)(1-\chi_{CAID})\chi_P m_{jh} < m_{jh} \text{ if } i = i_P
\end{aligned} \tag{10}$$

where the inequalities follow from the assumption that households cannot default on their liabilities if they have sufficient resources to pay their debt. Lastly, as noted in Section 3.4.1, households that start the period with a bankruptcy flag on their report face a probability $\lambda(f'=0|f=1)$ that their record will be expunged in the following period. Households that start the period without a flag on their credit record, on the other hand, retain their current credit record with certainty, i.e., $\lambda(f'=0|f=0)=1$.

Alternatively, if the agent qualifies for a Chapter 7 debt discharge, that is, if she passes

[3] I abstract from the observation that individuals qualify for food stamps in the US. if their gross income is less than 130 percent of the federal poverty level, and assume that food stamps are only available to households that work zero hours and have insufficient assets to pay their medical debt (i.e., $(1+r)(a+B) \leq \mathbb{I}_{Med}(s,\ell)(1-\chi_{CAID})m_{op} - m_{op}$).

the eligibility test and does not currently have a bankruptcy flag on her report, she can file for bankruptcy:

$$
\begin{aligned}
V^{B}(s) \quad = \quad & \max_{c,\ell,i'} u\left(c+\mathbb{I}_{FS}(s,\ell)\underline{c},\ell\right)+\beta\psi_{hj}\iint V\left(s'\right)Q\left(\eta,d\eta'\right)Q_{j}\left(h,dh'\right) \\
s.t. \quad & c+\mathbb{I}_{P}(i')\pi_{jh}+\psi = w(1-\tau)\varepsilon_{j}\eta\ell+min\left\{(1+r)(a+B),\underline{a}\right\} \\
& f' = 1 \\
& a' = 0 \\
& \ell, c \geq 0 \\
& i' \in \{i_P, i_S\}
\end{aligned}
$$

The debt discharge releases the debtor from all liabilities. In return, the filer must forfeit all non-exempt assets, $\underline{a}$, to the hospital, and pay court filing fees, ψ. Current and future earnings, on the other hand, are retained by the agent. Hospital revenues are thus given by:

$$
\begin{aligned}
Hosp &= max\left\{(1+r)(a+B)-\underline{a},0\right\}+\mathbb{I}_{Med}(s,\ell)(1-\chi_{CAID})m_{jh} && \text{if } i=i_S \\
Hosp &= max\left\{(1+r)(a+B)-\underline{a},0\right\}+(1-\chi_P)m_{jh}+\mathbb{I}_{Med}(s,\ell)(1-\chi_{CAID})\chi_P m_{jh} && \text{if } i=i_P
\end{aligned} \tag{11}
$$

Remaining assets are used for consumption and purchase of health insurance. I assume that households are not permitted to save in the period they file for bankruptcy. This is motivated by the fact that US law does not permit those who invoke bankruptcy to simultaneously accumulate assets. Lastly, agents that file for bankruptcy start the following period with a bankruptcy flag on their report.

3.4.2 Retired households

Let $V^R(s)$ denote the value of a retired agent of type s that does not file for bankruptcy. Retired households that qualify for a Chapter 7 debt discharge choose whether or not to file for bankruptcy:

$$V(s) = \max \left\{V^R(s), V^B(s)\right\} \tag{12}$$

where $V^R(s)$ is given by:

$$V^R(s) = \max_{c,a',i'} u(c + \mathbb{I}_{FS}(s,0)\underline{c}, 0) + \beta\psi_{jh}\sum_{f'}\lambda\left(f'|f\right)\iint V\left(s'\right)Q\left(\eta, d\eta'\right)Q_j\left(h, dh'\right)$$

$$\begin{aligned}
s.t.\quad & (1+f\zeta)c + a' + m_{op} + \mathbb{I}_P(i')\pi_{jh} = SS \\
& \qquad + (1+r)(a+B) + \mathbb{I}_{Med}(s,\ell)(1-\chi_{CAID})m_{op} \quad \text{if } x > 0 \\
& (1+f\zeta)c = (1-\kappa)\left[SS - \mathbb{I}_P(i')\pi_{jh}\right] \quad \text{if } x \leq 0 \\
& x = SS + (1+r)(a+B) + \mathbb{I}_{Med}(s,\ell)(1-\chi_{CAID})m_{op} - m_{op} \\
& m_{op} = \mathbb{I}_P(i)\chi_P\chi_{CARE}m_{jh} + (1-\mathbb{I}_P(i))\chi_{CARE}m_{jh} \\
& \mathbb{I}_{\mathbb{P}}(i) = \begin{cases} 1 & \text{if } i = i_P \\ 0 & \text{otherwise} \end{cases} \\
& c, a' \geq 0 \\
& i' \in \{i_P, i_S\}
\end{aligned}$$

Retired households receive Social Security benefits and Medicare from the government. Medicare lowers the agents' out-of-pocket medical expenses, which are given by $\chi_{CARE}m_{jh}$

for households that did not purchase private insurance in the preceding period and $\chi_{CARE}\chi_P m_{jh}$ for agents that have private health insurance. In the event that an agent does not have sufficient resources to pay her medical debt, she forfeits her assets to the hospital and pays a fraction κ of her Social Security benefits net of health insurance expenses. The agent is also not permitted to save. Hospital revenues are then given by:

$$Hosp = max\left\{\kappa\left[SS - \mathbb{I}_P\left(i'\right)\pi_{jh}\right], 0\right\} + (1+r)(a+B) + (1-\chi_{CARE})\,m_{jh}$$
$$+\mathbb{I}_{Med}(s,\ell)(1-\chi_{CAID})\,\chi_{CARE}m_{jh} < m_{jh} \quad \text{if } i = i_S$$

$$Hosp = max\left\{\kappa\left[SS - \mathbb{I}_P\left(i'\right)\pi_{jh}\right], 0\right\} + (1+r)(a+B) + (1-\chi_{CARE})\,m_{jh} + (1-\chi_P)\,\chi_{CARE}m_{jh}$$
$$+\mathbb{I}_{Med}(s,\ell)(1-\chi_{CAID})\,\chi_{CARE}\chi_P m_{jh} < m_{jh} \ \text{ if } i = i_P \tag{13}$$

where the inequalities again follow from the assumption that households are not permitted to default on their debt if they have sufficient resources to pay the hospital.

Alternatively, as long as the agent passes the eligibility test and does not have a Chapter 7 bankruptcy flag on her credit record, she can file for bankruptcy:

$$V^B(s) \quad = \quad \max_{c,i'} \; u\left(c + \mathbb{I}_{FS}(s,0)\underline{c}, 0\right) + \beta\psi_{hj}\iint V\left(s'\right)Q\left(\eta, d\eta'\right)Q_j\left(h, dh'\right)$$

$$s.t. \quad c + \mathbb{I}_P(i')\pi_{jh} + \psi = SS + min\left\{(1+r)(a+B), \underline{a}\right\}$$

$$f' = 1$$

$$a' = 0$$

$$c \geq 0$$

$$i' \in \{i_P, i_S\}$$

As earlier, filers forfeit their non-exempt assets to the hospital. Current and future Social Security benefits, on the other hand, are retained by the agent. In return, the creditor releases the debtor from all liabilities. This gives the following expression for hospital revenues:

$$
\begin{aligned}
Hosp &= max\{(1+r)(a+B) - \underline{\mathrm{a}}, 0\} + (1-\chi_{CARE})\, m_{jh} \\
&\quad + \mathbb{I}_{Med}(s,\ell)(1-\chi_{CAID})\,\chi_{CARE} m_{jh} \quad \text{if } i = i_S \\
Hosp &= max\{(1+r)(a+B) - \underline{\mathrm{a}}, 0\} + (1-\chi_{CARE})\, m_{jh} + (1-\chi_P)\,\chi_{CARE} m_{jh} \\
&\quad + \mathbb{I}_{Med}(s,\ell)(1-\chi_{CAID})\,\chi_{CARE}\chi_P m_{jh} \ \text{ if } i = i_P
\end{aligned}
\tag{14}
$$

3.4.3 Definition of equilibrium

Given an asset exemption level, $\underline{\mathrm{a}}$, bankruptcy filing costs, ψ, a Social Security replacement rate, b, copayment parameters χ_P, χ_{CARE}, and χ_{CAID}, Medicaid eligibility thresholds, and initial conditions for capital K_1 and the measure of types Φ_1, an *equilibrium* in this model is a sequence of model variables such that:

1. Taking prices, insurance costs, government policies, and accidental bequests as given, households maximize utility subject to their constraints.

2. Hospital markup pricing is consistent with zero profits.

3. Factor prices satisfy marginal product pricing conditions.

4. Government policies satisfy the government budget constraint.

5. All market clearing conditions are met.

6. Aggregate law of motion for Φ is induced by the policy functions and the exogenous stochastic processes for idiosyncratic health and labor productivity.

3.5 Calibration

This section describes how I calibrate the model. The section starts by describing the calibration of the bankruptcy parameters. I then explain how I derive the insurance coinsurance rates and how the Medicaid income and asset limits are determined. Attention is also given to how these limits are affected when I introduce the Affordable Care Act in Section 3.7.2. Lastly, I discuss the calibration of the earnings process and the life cycle parameters.

3.5.1 Bankruptcy parameters

I set the probability λ that an agent with a bankruptcy flag on her report will have her record expunged in the following period to generate an average duration of bankruptcy spells of nine years. That is, filing for bankruptcy leaves a flag on the agent's report for the next nine years on average, during which the household cannot file for another Chapter 7 discharge. The bankruptcy flag also lowers the agent's credit score, which increases her cost of consumption. The probability is chosen to match the fact that US bankruptcy law does not permit households to file and receive a Chapter 7 debt discharge more than once every eight years and that bankruptcy flags remain on an agent's credit record for up to ten years.

Filers must pay court filing fees and legal fees.[4] These fees vary from state to state, but

[4]Throughout, I abstract from the observation that courts may waive the bankruptcy filing fees if the debtor's income is less than 150 percent of the federal poverty level.

typically range from $1,500 to $2,500. I pick an intermediate value and set the filing fee to $2,000. Filing for bankruptcy lowers the household's credit score and increases her unit cost of consumption by ζ. I calibrate this parameter to match the percentage of Chapter 7 filers due to medical debt. The literature has produced a wide range of estimates for this share, ranging from 16.4 percent (Chatterjee et al., 2007) to 62.1 percent (Himmelstein et al., 2009). Gross and Notowidigdo (2011) use cross-state variation in Medicaid expansions from 1992 to 2004 to quantify this and find that, among low-income households, 26 percent of consumer bankruptcies can be attributed to out-of-pocket medical expenses. A more recent study by Austin (2015), using responses from a nationwide survey of bankruptcy filers, finds that medical debt is the predominant causal factor in 18 to 26 percent of consumer bankruptcies. I pick a value between these estimates and assume that 25 percent of consumer bankruptcies can be attributed to medical debt. With a ratio of filers to total adult population of 0.50 percent, this gives a target share of 0.13 percent.[5]

The type and value of assets that can be seized in a Chapter 7 bankruptcy filing vary from state to state. As an example, Kansas allows households to exempt the full value of their property and claim $40,000 in vehicle exemptions. The corresponding exemption levels in Delaware, on the other hand, are zero for both home and vehicle equity. Some states also allow households to claim a wildcard exemption that can be used for all assets, ranging from $200 in Iowa to $60,000 in Texas. As the only asset in the model is saving, I set the exemption limit, $\underline{a}$, to match the $6,000 median wildcard exemption limit across the 33 states with positive wildcard exemptions.

Households that do not have the resources to pay their medical bills pay a fraction κ of their earnings or Social Security benefits net of health insurance expenses. This parameter

[5] According to data from the Administrative Office of the US Courts, 1.10 million people filed for Chapter 7 in 2010. This gives a ratio of filers to total adult population of 0.50 percent.

affects the households' precautionary saving motive. The higher the value, the lower the value of post-garnished resources available for consumption. This increases precautionary saving, which in turn lowers the share of households that have to partially default on their medical bills. I follow a similar calibration strategy as Chatterjee et al. (2007) and calibrate the wage garnishing parameter to match the percentage of individuals with medical debt.

Following the 2005 Bankruptcy Abuse Prevention and Consumer Protection Act, Chapter 7 is restricted to households that pass a means or repayment test. Agents pass the means test if their gross income is less than the state median income or if their disposable income is sufficiently low, where disposable income is defined as income net of allowable expenses for food, clothing, housing, health and child care, alimony, education, transportation, taxes, insurance, and secured debt payments. Each county and metropolitan region has different allowable expense amounts for these categories. Following Mahoney (2015), I set the disposable income limit to $\$1,320$ per year. Households pass the repayment test if both their disposable income is less than $\$2,190$ per year and annual payments of this amount over 5 years would pay back less than 25 percent of the agent's unsecured debt (Mahoney, 2015). The bankruptcy parameters are summarized in Table 3.2.

3.5.2 Health insurance and Affordable Care Act parameters

I follow Conesa et al. (2018) and let the coverage rate on each insurance plan be given by the average share of expenses covered across households in the MEPS with that insurance plan as their primary insurance provider, the last of which is defined as the insurer that pays for the largest share of the agent's expenses. This gives a coverage rate on private insurance, Medicare, and Medicaid of 77.1, 70.9, and 86.2 percent. The coinsurance rate on the three insurance plans are then given by the complement of these shares.

Table 3.2: Bankruptcy parameters determined jointly in equilibrium

Parameter	Description	Target or source	Value
Determined outside the model			
λ	Average duration of bankruptcy spell	Office of the US Courts	0.100
Ψ	Bankruptcy filing cost		0.043
$\underline{a}$	Asset exemption limit	Mahoney (2015)	0.113
Determined jointly in equilibrium			
ζ	Increase in cost of cons. after bankruptcy	Perc. of med. bankrupcty files = 0.1	0.164
κ	Wage garnishing parameter	Perc. with medical debt = 1.7	0.242
Chapter 7 eligibility limits			
	Disposable income test	Mahoney (2015)	0.025
	Repayment test	Mahoney (2015)	0.041

Notes: The table lists the bankruptcy parameters that are determined outside the model and those that are determined jointly in equilibrium.

The benchmark model is calibrated to match the key features of the US health insurance system prior to the Affordable Care Act (ACA). I therefore set the Medicaid income and asset thresholds to match the corresponding limits in the data that prevailed in the US at that time. Following Conesa et al. (2018), I set the income and asset eligibility thresholds to match the corresponding weighted average limits across US states, with weights given by each state's share of total health expenses. Using data from the Kaiser Family Foundation, I obtain a medically needy income limit of 41.9 percent of the federal poverty level (FPL), a medically needy asset threshold of $\$1,950$, and a categorical income limit of 90.2 percent of the FPL.[6]

Section 3.7.2 incorporates the key features of the ACA to study how the results depend on the specifics of the health insurance environment. The ACA expanded the Medicaid program to cover all households with income less than 133 percent of the FPL. In addition, all households that are not covered by Medicare or Medicaid must buy private insurance or pay a penalty, where the penalty is given by the maximum of \$695 and 2.5 percent of the

[6]The FPL is about 23 percent of GDP per capita.

Table 3.3: Health insurance and Affordable Care Act parameters

Parameter	Description	Source	Value
Insurance coinsurance parameters			
χ_P	Private insurance	MEPS	0.229
χ_{CARE}	Medicare	MEPS	0.291
χ_{CAID}	Medicaid	MEPS	0.138
Medicaid eligibility limits (pre ACA)			
	Categorical income	Kaiser Family Foundation	0.197
	Medically needy income	Kaiser Family Foundation	0.092
	Medically needy assets	Kaiser Family Foundation	0.041
ACA parameters			
	Income limit	Kaiser Family Foundation	0.291
$\underline{\tau}$	Minimum tax filing amount	Kaiser Family Foundation	0.189

Notes: The table lists the health insurance parameters that are determined outside the model.

household's income net of the minimum tax filing amount. I include the first feature of the law by replacing the three Medicaid eligibility tests by a single income test, whereby all households qualify for the program if their income is less than the corresponding limit in the data. Lastly, I set the minimum tax filing amount to $\$10,300$. A summary of the health insurance and ACA parameters are given in Table 3.3.

3.5.3 Life cycle parameters, medical expenditures, and the earnings process

The calibration of the life cycle and technology parameters follows closely that of Conesa et al. (2018). Each period in the model is one year. Households enter the economy at age 20, retire from the labor force at age 66, and have a maximum life span of 100 years. The population growth rate is set to 1.1 percent per year. I pick a depreciation rate of 0.059 and set capital's share of income to 0.360 to match commonly used values in the literature. Lastly, I pick σ to match an intertemporal elasticity of substitution of 0.5 and set the consumption share in intratemporal utility to match the median estimates in French (2005).

Recall that the stochastic process for health is given by a stationary finite-state Markov process that depends on the agent's current age and health. Using data from the MEPS, I estimate health transition and survival probabilities by running probit regressions of next period's health on current age, age squared, health, and interaction terms.

I calibrate the total factor productivity parameter to generate a GDP per capita of 1 in the steady state of the benchmark model. The discount factor is calibrated to match a capital-to-output ratio of 3. Next, I set the consumption floor to match average annual food stamps between 2006 and 2010 of about $\$1,300$, and set the Social Security replacement rate to match the $\$11,900$ average annual benefits of households without a college degree.

I follow Conesa et al. (2018) and split medical expenses for each age group into three categories: low, high, and catastrophic, where the three categories correspond to the average value of medical expenses between the 0-60th percentile, 60-99.9th percentile, and 99.9-100th percentile, respectively. That is, for each age group, I first pool all medical expenses in the MEPS, compute the 60th and 99.9th percentile, and then let the three expenditure states be given by the mean value of medical expenses between these data cutoffs. Lastly, I detrend the derived series using a log-linear trend, and scale the expenses to match the health-expenditure-to-GDP ratio observed in the data. A summary of the life cycle and technology parameters are given in Table 3.4.

The earnings process is calibrated to match the empirical earnings distribution in the US as reported by Díaz-Giménez et al. (1997). To match the top decile of this distribution, I follow Castañeda et al. (2003) and choose a right skewed labor productivity shock process characterized by two key features: a high labor productivity at the top and a low probability

Table 3.4: Life cycle and technology parameters

Parameter	Description	Target or source	Value
Determined outside the model			
	Maximum life span (100 years)		81
j_r	Retirement age (66 years)		47
	Population growth rate		0.011
α	Capital income share		0.360
δ	Depreciation rate		0.059
γ	Consumption share in utility	French (2005)	0.574
Risk aversion	IES=0.5	2.742	
Determined jointly in equilibrium			
A	Total factor productivity	GDP per capita = 1	0.764
β	Discount factor	Capital to output = 3	0.936
b	SS replacement rate	Avg. SS benefits non-college $\approx$ 11,900	0.329
$\underline{c}$	Consumption floor	Avg. food stamps $\approx$ 1,300	0.028
	Scale for health care costs	Health expend. to GDP = 0.165	1.902

Notes: The table lists the life cycle and technology parameters that are determined outside the model and those that are determined jointly in equilibrium. A period in the model is one year.

of transitioning to this state. This process has been shown to induce the large savings of wealth-rich and earnings-rich households that is needed to generate the concentration of wealth observed in the data.[7] As shown in the first part of Table 3.5, both the model with and without Chapter 7 successfully matches the earnings distribution observed in the data, where each data quintile is given by the ratio of group to sample average. The comparison of the wealth distribution in the second part of the table shows that both models also generate a concentration of wealth that is comparable to what we observe in the data.

3.6 Results

The following subsection studies the general equilibrium effects of eliminating Chapter

[7]Recent research by De Nardi et al. (2016) have shown that the process used by Castañeda et al. (2003) is inconsistent with the data. In particular, they find that the earnings risk faced by the top earners in the data is insufficient to generate the observed concentration of wealth. I leave it for future work to study the sensitivity of my results to alternative labor productivity processes.

Table 3.5: Labor earnings distribution (percent)

	Quintiles					Top			
	1st	2nd	3rd	4th	5th	90-95	95-99	99-100	Gini
Data	-0.40	3.19	12.49	23.33	61.39	12.38	16.37	14.76	0.63
Benchmark	0.00	4.13	11.05	19.84	64.98	12.10	16.47	14.78	0.63
w/o Chapter 7	0.00	4.14	11.05	19.84	64.97	12.11	16.47	14.78	0.63

Notes: The table reports the labor earnings distribution in the model and the data. Cells denote shares of total. Data source: Kuhn and Ríos-Rull (2013).

Table 3.6: Wealth distribution (percent)

	Quintiles					Top			
	1st	2nd	3rd	4th	5th	90-95	95-99	99-100	Gini
Data	-0.39	1.74	5.72	13.43	79.49	12.62	23.95	29.55	0.78
Model	0.23	0.41	3.18	16.91	79.26	16.98	27.18	15.92	0.77
w/o Chapter 7	0.23	0.41	3.29	17.03	79.04	16.92	27.08	15.86	0.77

Notes: The table reports the wealth distribution in the model and the data. Cells denote shares of total. Data source: Kuhn and Ríos-Rull (2013).

7. This experiment enables me to capture the effect of Chapter 7 on insurance coverage, welfare, macroeconomic aggregates, and the price of health services. I then extend the model to incorporate the key features of the Affordable Care Act to examine whether the results regarding the effects of Chapter 7 on the demand for health insurance depend on the specifics of the health insurance environment.

3.6.1 Chapter 7 consumer bankruptcy

This section studies the economic implications of an unexpected elimination of Chapter 7. The ability to file for bankruptcy leave households with a tradeoff: buy insurance today to lower expected medical costs tomorrow, or do not buy insurance today and file for a debt discharge tomorrow if hit by a severe medical shock. Eliminating the second option increases the cost of being uninsured. As a result, more households are likely to buy health insurance in an economy without Chapter 7. On the other hand, households are also likely to save more in an economy without bankruptcy. This raises capital and therefore wages in general equilibrium, which in turn lowers the likelihood that an agent will qualify for Medicaid, the latter of which follows from the Medicaid income and asset eligibility tests. The overall effect of Chapter 7 on the uninsurance rate thus depends on the relative effects of these two channels.

Table 3.6 compares the insurance take-up rate across working-age households in the steady state of the economy with and without Chapter 7. Eliminating the ability to file for bankruptcy increases private insurance enrollment across working-age households by 2.0 percentage points. With a working-age population of about 190 million, this translates into an increase in private insurance take-up by 3.8 million individuals. This result suggests that 8 percent of the uninsured working-age population would take up coverage if they did not have the option to file for medical bankruptcy. Medicaid enrollment, on the other hand, declines by 0.2 percentage points across working-age households, or 0.3 million people. As a result, I find that eliminating Chapter 7 leads to a 1.8 percentage point increase in the insurance coverage rate. These results, coupled with the fact that only 0.1 percent of households file for medical bankruptcy in the benchmark model, show that the ability to file for bankruptcy has an economically significant implication for the insurance coverage rate.

The reduction in Medicaid enrollment is driven by two forces. First, higher private insurance enrollment lowers out-of-pocket medical expenses. This increases income net of health expenses, and hence reduces the likelihood that an agent passes the Medicaid medically needy income test. Next, households respond to the removal of Chapter 7 by increasing saving. Higher saving affects Medicaid eligibility directly through the medically needy

Table 3.7: Insurance take-up rate across working-age agents in economy with and without Chapter 7

Variable	Without Chapter 7 (percentage change from benchmark)
Insurance coverage rate	1.8
Private insurance enrollment	2.0
Medicaid enrollment	-0.2

Notes: The table compares the steady state of the economy with and without Chapter 7 (see the text for details). The the first row reports the percentage point changes in the insurance coverage rate. The final two rows decompose the change into private insurance enrollment and Medicaid enrollment.

asset test. The additional saving also translates into higher wages, which in turn lowers the percentage of households that pass the Medicaid income tests.

I next solve for a transition path between the steady state of the economy with and without bankruptcy. This enables me to quantify how quickly the insurance take-up rate will adjust following the change in the bankruptcy law. Taking the costs of transitioning between the steady states into account also enables me to better quantify the welfare effects of the policy reform. I find that the private insurance take-up rate increases by 1.7 percentage points during the first two years following the elimination of Chapter 7. That is, 85 percent of the change in the private insurance take-up rate takes place within the first two years of the transition. Given the size of the working-age population, this translates into an increase in the private insurance take-up rate of 3.3 million people. In comparison, about 11.7 million people took up coverage during the first two years of the Affordable Care Act. This shows that reforming the bankruptcy laws can have considerable short run implications for the insurance coverage rate.

Lastly, I study the welfare effects of the ability to file for medical bankruptcy. I find that 79.6 percent of the population is better off in the economy with medical bankruptcy because of the implicit insurance provided by this option. To quantify the welfare effects on those that benefit and lose from the reform, I compute the dollar value of how much wealth must change in the initial steady state to make the agents weakly better off in the economy with bankruptcy than in the transition to the economy without bankruptcy. I find that those that benefit from the reform experience an average welfare gain that is equivalent to receiving $400 higher wealth in the steady state with bankruptcy. In contrast, those that lose from the reform experience an average welfare loss that is equivalent to a $2,900 reduction in wealth. Consequently, I find that aggregate welfare is higher in the economy with medical bankruptcy.

3.6.2 ACA

This subsection studies how the results derived earlier depend on the specifics of the health insurance environment. In particular, I examine whether the results regarding the effect of Chapter 7 on the demand for health insurance is robust to introducing the key features of the 2010 Patient Protection and Affordable Care Act in the model. The ACA introduced several mechanisms to improve health insurance coverage, quality, and affordability. First, health insurance providers were no longer permitted to refuse to cover or charge individuals more because of pre-existing conditions. Second, it introduced an insurance exchange marketplace where households could buy insurance at group-based premium rates with subsidies for households with income between 133 and 400 percent of the FPL. Third, it introduced penalties for not buying insurance. Lastly, it expanded the Medicaid program to cover all households with income less than 133 percent of the FPL.[8]

I build on work by Pashchenko and Porapakkarm (2013) and Jung and Tran (2016) and

[8]Following a recent Supreme Court ruling, states are free to opt out of the Medicaid expansion. As of 2016, 19 states have not expanded Medicaid.

incorporate the different features of the ACA by changing four aspects of the benchmark model. First, I change the price of the health insurance to π_j. That is, insurance companies are no longer permitted to condition the price of the insurance on current health. Next, I introduce a subsidy for purchasing insurance. To be eligible for a subsidy, the value of which varies with the agent's income and cost of insurance, an agent must have income between 133 and 400 percent of the FPL. Let $\tilde{y}(s,\ell)$ denote household income and $\tau_{sub}(s,\ell)$ denote the subsidy for purchasing insurance, where s denotes the agent's type and ℓ is her labor supply. The subsidy is then given by:

$$\tau_{sub}(s,\ell) = \begin{cases} max\left(0, \pi_j - 0.020\tilde{y}(s,\ell)\right) & if \quad \tilde{y}(s,\ell) < 1.33FPL \\ max\left(0, \pi_j - 0.030\tilde{y}(s,\ell)\right) & if \quad 1.33FPL \le \tilde{y}(s,\ell) < 1.50FPL \\ max\left(0, \pi_j - 0.040\tilde{y}(s,\ell)\right) & if \quad 1.50FPL \le \tilde{y}(s,\ell) < 2.00FPL \\ max\left(0, \pi_j - 0.063\tilde{y}(s,\ell)\right) & if \quad 2.00FPL \le \tilde{y}(s,\ell) < 2.50FPL \\ max\left(0, \pi_j - 0.081\tilde{y}(s,\ell)\right) & if \quad 2.50FPL \le \tilde{y}(s,\ell) < 3.00FPL \\ max\left(0, \pi_j - 0.095\tilde{y}(s,\ell)\right) & if \quad 3.00FPL \le \tilde{y}(s,\ell) \le 4.00FPL \end{cases} \tag{15}$$

Households that are not covered by Medicare or Medicaid must buy private health insurance or pay a penalty.[9] The penalty, $\bar{\pi}_P(s,\ell)$, varies with the agent's income and is given by the maximum of \$695 and 2.5 percent of the agent's income net of the minimum tax filing amount, $\underline{\tau}$. That is, $\bar{\pi}_P(s,\ell)$ is given by:

$$\bar{\pi}_P(s,\ell) = max\{\$695, 0.025\left(\tilde{y}(s,\ell) - \underline{\tau}\right)\} \tag{16}$$

[9]The penalty is waived if the price of health insurance exceeds 8 percent of the agent's income.

Lastly, I incorporate the Medicaid expansion by expanding Medicaid eligibility to all agents with income less than 133 percent of the FPL.

I find that the previous results regarding the effect of Chapter 7 on the demand for health insurance still applies in the post-ACA environment. In particular, I find that the ability to file for medical bankruptcy lowers private insurance enrollment by 1.5 percentage points across the working-age population. Similarly, Medicaid enrollment is 0.1 percentage points higher in the economy with Chapter 7. As a result, I find that eliminating the ability to file for medical bankruptcy would lower the uninsurance rate by 1.4 percentage points.

3.7 Conclusion

This chapter developed an equilibrium life cycle model with incomplete markets and heterogeneous agents to examine the general equilibrium effects of medical bankruptcy. I used the model to study the effects of an unexpected elimination of Chapter 7, which allowed me to quantify the effects of medical bankruptcy on the private insurance take-up rate, Medicaid enrollment, welfare, and macroeconomic aggregates. The ability to file for bankruptcy lowered private insurance enrollment across the working-age population by 2.0 percentage points and increased Medicaid enrollment by 0.2 percentage points. Eliminating Chapter 7 would thus lead to a 1.8 percentage point reduction in the uninsurance rate. Similar results were obtained in Section 3.7.2, where I extended the model to include the key features of the Affordable Care Act.

Next, I computed a transition path between the economy with and without Chapter 7, and found that 85 percent of the increase in the private insurance take-up rate took place within the first two years following the removal of bankruptcy. Medicaid enrollment, on the other hand, was found to gradually decline during the transition. As a result, I found that most of the increase in the coverage rate would take place within the first years of the transition. Lastly, I examined the welfare implications of Chapter 7. I found that 79.6 percent of the population was better off in the environment with bankruptcy because of the implicit insurance provided by this option.

Chapter 4

Comparing welfare across the states of the US

4.1 Introduction

A large literature has used economic models and various statistics to compare welfare across countries. Instead of comparing welfare across countries, this chapter compares welfare across the states of the US. The chapter is motivated by the considerable heterogeneity in GDP per capita across states, ranging from 32,000 dollars in Mississippi to 68,600 dollars in Alaska, compared to a national average of 50,600.[1] Moreover, total personal consumption per capita ranges by a factor of 1.7, and life expectancy at birth varies by almost 7 years. These differences are likely to have large implications for welfare as measured by the expected lifetime utility of being born in a given state. This chapter quantifies these welfare differences by computing how much consumption must change in all ages in a given state to make an agent indifferent between being born in this state and another state. I allow states to differ in longevity, consumption, and educational attainment, all of which have large implications for welfare.

My analysis shows that there exists considerable heterogeneity in welfare across the states. To illustrate, I find that consumption must be scaled down by 28 percent in all ages in the state with the highest welfare level, Massachusetts, to make a hypothetical agent indifferent between living her entire life in Massachusetts and the state with the median welfare level, Iowa, and has to be scaled down by an additional 19 percentage points to make her indifferent between living her entire life in Massachusetts and the state with the lowest welfare level, Mississippi. A decomposition of the welfare results shows that heterogeneity in consumption per capita accounts for the largest share of the variation in welfare levels. That said, variations in both life expectancy at birth and educational attainment also contribute considerably to the heterogeneity in welfare levels across the

[1] Numbers are in constant 2010 dollars.

states.[2]

Finally, I compare welfare across time by quantifying how much each state's welfare level has grown between 2000 and 2013, taking into account state-specific changes in consumption per capita, mortality risk, and educational attainment. I find that the variation in welfare across the states is very persistent. Accordingly, although the lower ranked states in 2000 generally experienced a higher welfare growth than the higher ranked states in 2000, there does not appear to be evidence of rapid convergence toward similar welfare levels.

The rest of the chapter is organized as follows. The next section provides summary statistics for all the states. In particular, I compare GDP per capita, consumption per capita, educational attainment, infant mortality rates, and life expectancy at birth across the states. Section 4.3 develops a model that can be used to quantify the welfare implications of these state-characteristics. Section 4.4 explains how I calibrate the model. In particular, it explains how I derive state-specific survival probabilities and state-specific consumption profiles, and how I calibrate the parameters of the utility function. The following section compares welfare levels across the states and over time. Finally, Section 4.6 concludes.

4.2 Data

I start by comparing income, consumption, educational attainment, and longevity across the states, all of which will play key roles in the following welfare analysis. Appendix Table B.1 provides summary statistics for all the states. Columns two and three report average GDP per capita and average consumption per capita for the period 2010 to 2017 across

[2] I split individuals into three educational groups: those without a high school degree, those with a high school degree but without a college degree, and those with a college degree, where a college degree refers to individuals with at least a bachelor's degree or a minimum of four years of college.

the states. Consumption refers to total personal consumption expenditures as reported by the Bureau of Economic Analysis (BEA). Population statistics are obtained from the intercensal population estimates reported by the Census. Both the GDP and consumption series have been deflated by means of the region-specific CPI numbers reported by the Bureau of Labor Statistics. The table shows that GDP per capita varies considerably across the states, from 32,000 dollars in Mississippi to 68,600 dollars in Alaska. Compared to an average GDP per capita of 50,600 dollars in the United States, this means that the richest state has a 36 percent higher GDP per capita relative to the national average, while the poorest state has a 37 percent lower GDP per capita relative to the national average. Similarly, consumption per capita ranges by nearly a factor of two between the states, from a low of 25,700 dollars in Mississippi to a high of 44,800 dollars in Massachusetts.

The fourth and fifth column of appendix Table B.1 compare the educational attainment of the 25+ year-old population across the states as reported by the Census. On average, 13.2 percent of 25+ year-olds did not have a high school degree in the United States between 2010 and 2017. The percentage without a high school degree varies from a low of 7.3 percent in Montana to a high of 18.1 percent in California. The variation in college attainment—measured by the share of the 25+ year-old population with at least a bachelor's degree or a minimum of four years of college—is even higher. Compared to a national average of 30.0 percent, only 19.2 percent of 25+ year-olds in West Virginia have a college degree. In contrast, 40.2 percent of 25+ year-olds in Massachusetts has at least a bachelor's degree. To illustrate, this means that West Virginia currently has the same college attainment as Massachusetts had during the 1980s.[3] As will be evident below, these differences in educational attainment have large implications for both consumption and longevity, and will hence have considerable implications for welfare.

[3]Historical educational attainment series by state and year are obtained from the Current Population Survey.

The final two columns of appendix Table B.1 report state-specific infant mortality rates and life expectancy at birth. Both statistics are derived from the Underlying Cause of Death database reported by the Centers for Disease Control and Prevention (CDC). It is well known that the United States has a high infant mortality rate compared to other developed countries. On average, 5.9 infants per 1000 live births died before the age of 1 in the United States between 2010 and 2017, about twice as high as in countries likes France and Germany, and about three times as high as in Japan. This national average, however, masks significant heterogeneity across the states, ranging from 4.1 in Massachusetts to 9.1 in Mississippi. To illustrate, this means that the current infant mortality rate in Mississippi is as high as the national infant mortality rate in the United States around 1990. Life expectancy by state is derived from state-specific mortality rates. It refers to the expected lifespan of a hypothetical individual who lives her entire life in the state she is born in.[4] Appendix Table C.1 shows that average life expectancy at birth between 2010 and 2017 varied by 6.9 years across the states, from a low of 74.7 years in Mississippi to a high of 81.6 years in Hawaii. Put differently, this means that the current life expectancy in Mississippi is as low as the average life expectancy in the United States 30 years ago.

So far, I have shown that states vary considerably in GDP per capita, consumption per capita, educational attainment, and longevity, all of which have important implications for welfare. Next, I look at the cross-state relationship between these variables. The left panel of Figure 4.1 shows that GDP per capita and longevity are positively correlated. That is, states with higher GDP per capita are more likely to have higher life expectancy at birth. Similarly, consumption per capita and longevity are positively correlated. This is illustrated

[4] The welfare analysis in Section 4.5 allows for the possibility of state-to-state migration. Consequently, the life expectancy of an individual who is born in a given state will not be equal to the life expectancy of a hypothetical individual who resides in that particular state her entire life. Note, however, that the life expectancy estimates that allow for migration and the estimates that do not allow for migration are very similar since, across all states, only 1-5 percent of residents migrate to another state on an annual basis. These numbers are derived from annual state-to-state migration statistics as reported by the Census.

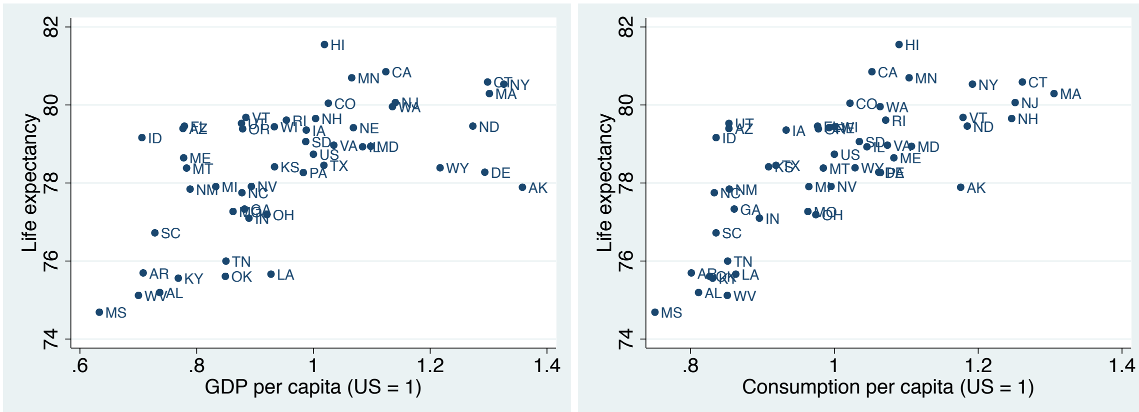

Figure 4.1: Relationship between life expectancy and GDP per capita, and relationship between life expectancy and consumption per capita

Notes: The left panel depicts the relationship between GDP per capita and life expectancy at birth. State-specific GDP per capita has been normalized by average GDP per capita in the United States. The right panel depicts the relationship between consumption per capita and life expectancy at birth. State-specific consumption per capita has been normalized by average consumption per capita in the United States. Consumption refers to total personal consumption expenditures. Life expectancy at birth is derived from the Underlying Cause of Death database reported by the Centers for Disease Control and Prevention. All numbers refer to average values for the period 2010 to 2017.

in the right panel of Figure 4.1, which shows that states with higher consumption per capita are more likely to have higher life expectancy at birth. This means that a comparison of welfare across states based solely on a comparison of GDP per capita or consumption per capita would understate the magnitude of the welfare differences between the states since it would fail to account for the fact that individuals in richer states are not only likely to consume more during each year of life but are also likely to benefit from a considerably higher life expectancy.

Figure 4.2 shows that states with lower educational attainment are more likely to have lower life expectancy at birth. The left panel of Figure 4.2 plots the relationship between the share of 25+ year-olds without a high school degree and life expectancy at birth. With some notable exceptions such as California, I find a negative relationship between these variables. Similarly, as depicted in the right panel of Figure 4.2, I find a positive relationship between

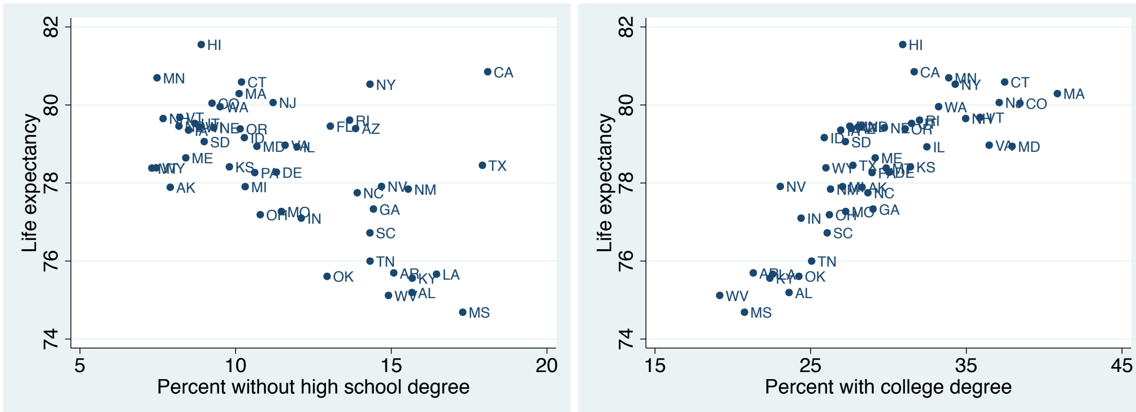

Figure 4.2: Relationship between life expectancy and educational attainment

Notes: The left panel depicts the relationship between the percentage of 25+ year-olds without a high school degree and life expectancy at birth. The right panel depicts the relationship between the percentage of 25+ year-olds with a college degree and life expectancy at birth, where a college degree refers to at least a bachelor's degree or a minimum of four years of college. Educational attainment by state is obtained from the Census. Life expectancy at birth is derived from the Underlying Cause of Death database reported by the Centers for Disease Control and Prevention. All numbers refer to average values for the period 2010 to 2017.

the share of 25+ year-olds with a college degree and life expectancy at birth.

To summarize, I have shown that consumption and life expectancy are positively correlated across states, and that states with lower shares of high school dropouts and higher shares of college graduates have higher life expectancy. These differences are likely to have large implications for welfare. The next section develops a model that can flexibly account for these differences. The model is an extension of the model used by Jones and Klenow (2016) to study welfare differences across countries. After mapping the model to the data in Section 4.4, I then use the model in Section 4.5 to quantify the magnitude of the welfare differences between the states.

4.3 Model

Let an agent's idiosyncratic state be given by her age, j, state of residence, s, and educational level, e. Agents derive utility from consumption, c, which follows a stochastic process that depends on the agent's age, education, and state. Let $\mathbb{E}_{jse}$ denote age-education-state-specific expected consumption. Agents are subject to mortality risk, $1-\psi_{jse}$, which varies with the agent's age, education, and state. Let $\Psi_{jse}=\prod_{k=0}^{j-1}\psi_{kse}$ denote the education-state-specific probability of surviving from age 0 to age j. I assume that agents that survive until age J die with probability one. For now, assume that agents live their entire life in the state they are born in. Moreover, assume that the educational level is revealed at birth and stays constant over the agent's lifespan. Both assumptions will be relaxed later in the chapter. Expected lifetime utility in state s is then given by

$$U(s)=\sum_{e}\pi(e|s)\left[\mathbb{E}_{0se}u(c)+\sum_{j=1}^{J}\beta^{j}\Psi_{jse}\mathbb{E}_{jse}u(c)\right], \quad (1)$$

where β denotes the discount rate and $\pi(e|s)$ is the state-specific probability of educational level e.

Let $U(s;\lambda)$ denote expected lifetime utility in state s if consumption is multiplied by a factor λ in all ages:

$$U(s;\lambda)=\sum_{e}\pi(e|s)\left[\mathbb{E}_{0se}u(\lambda c)+\sum_{j=1}^{J}\beta^{j}\Psi_{jse}\mathbb{E}_{jse}u(\lambda c)\right]. \quad (2)$$

I quantify the welfare difference between states s_i and s_{-i} by computing how much consumption must adjust in all ages in state s_i to make an agent indifferent between living her life in state s_i and s_{-i}. This corresponds to deriving the scaling factor, λ, that solves

$$U(s_i;\lambda)=U(s_{i-1};1). \quad (3)$$

The next section explains how I derive the inputs required for the welfare analysis.

4.4 Calibration

The following two subsections explain how I compute the age-education-state-specific survival probabilities and how I derive the stochastic process for consumption. The final subsection explains how I calibrate the parameters of the utility function.

4.4.1 Survival probabilities

This subsection explains how I derive age-education-state-specific survival probabilities, ψ_{jse}. I start by estimating age-state-specific survival probabilities, $\tilde{\psi}_{js}$. To do this, I first pool all death records for the period 2010 to 2017 from the Underlying Cause of Death database reported by the Centers for Disease Control and Prevention (CDC). The CDC reports each individual's age and state of legal residence at the time of death in the United States, with age top-coded at age 85. For 0–84 year-olds, I first compute raw age-state-specific survival probabilities directly from observed mortality rates. I then smooth the logarithm of the mortality rates by means of step-wise sixth-order polynomials in age. This helps ensure smooth mortality rates for smaller states like Vermont and Rhode Island. Unreported results show that the smoothed mortality rates fit the raw data almost perfectly. Beyond the age of 84, I approximate age-state-specific survival probabilities by means of Gompertz survival models. In a Gompertz model, the logarithm of the mortality rate is linear in age, $ln\left(m_{js}\right) = \alpha_s + \beta_s j$, where m_{js} is the mortality rate of j-year-olds in state s, and where α_s and β_s are state-specific coefficients. This log-linear approximation fits the CDC mortality rates for 40+ year-olds almost perfectly. I then use the estimated mortality regressions to predict age-state-specific survival probabilities for 85–100 year-olds. Note that all the results in Section 4.5 are robust to allowing for a higher maximum life span than

100 years.

Next, I compute age-education-specific survival probabilities, $\tilde{\psi}_{je}$. To do this, I first pool all death records for the period 2010–2017 from the National Vital Statistics System (NVSS) reported by the CDC. The NVSS reports each individual's age and educational attainment at the time of death in the United States, with age top-coded at age 85. For confidentiality purposes, the NVSS does not report individuals' state of residence or the state of occurrence at the time of death. I split individuals into three educational groups: less than high school, high school graduates but no college degree, and college graduates, where college graduates refer to individuals with at least a bachelor's degree or a minimum of 4 years of college. For brevity, I refer to the three groups as individuals without a high school degree, high school graduates, and college graduates. I then compute age-education-specific survival probabilities for 20–84 year-olds directly from observed mortality rates. Since the NVSS top-codes age at age 85, I assume that the difference in age-specific survival probabilities between individuals without a high school degree, with a high school degree, and with a college degree for 85+ year-olds is the same as that for 84-year-olds. Mortality probabilities by age and educational attainment are depicted in Figure 4.3. As shown in the graph, across all age groups, college-educated individuals are subject to a lower mortality risk than high school graduates, and high school graduates are subject to a lower mortality risk than individuals without a high school degree.

I then use data from the Current Population Survey (CPS) for the period 2010–2017 to compute the distribution of educational attainment by age and state, Λ_{jes}. That is, for each age and state, I compute the percentage of individuals without a high school degree, with a high school degree, and with a college degree. Since the CPS top-codes age at age 80, I assume that the distribution of educational attainment by age and state for 80+ year-olds is the same as that for 79-year-olds.

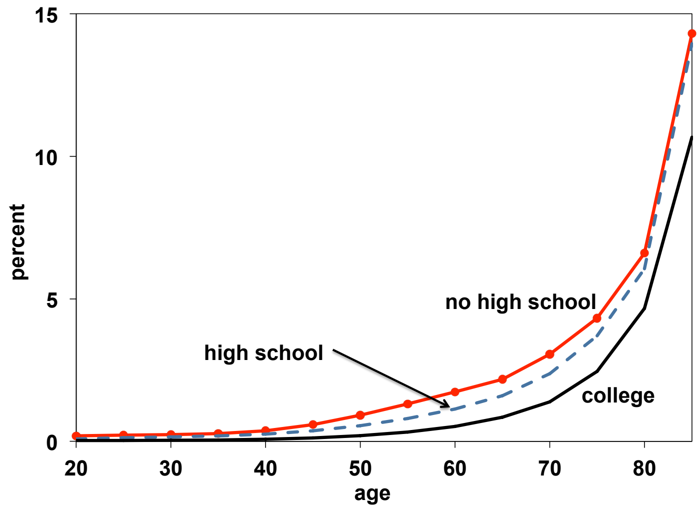

Figure 4.3: Mortality probability by age and educational attainment

Notes: The graph plots mortality probabilities by age and educational attainment. No high school refers to individuals without a high school degree, high school refers to individuals with a high school degree but less than a bachelor's degree, and college refers to individuals with at least a bachelor's degree. Data source: NVSS.

Lastly, given an initial guess, $\tilde{\psi}_{jes}$, I obtain age-education-state-specific survival probabilities, ψ_{jes}, by adjusting $\tilde{\psi}_{jes}$ to match the age-state-specific survival probabilities from the CDC, $\tilde{\psi}_{js}$, and the age-education-specific survival probabilities from the NVSS, $\tilde{\psi}_{je}$. This corresponds to deriving the scaling terms, a_{jes}, that solve the following system of equations:

$$\begin{aligned} \tilde{\psi}_{js} &= \textstyle\sum_e \Lambda_{jes} a_{jes} \tilde{\psi}_{jes} \\ \tilde{\psi}_{j2} - \tilde{\psi}_{j1} &= a_{j2s}\tilde{\psi}_{j2s} - a_{j1s}\tilde{\psi}_{j1s} \\ \tilde{\psi}_{j3} - \tilde{\psi}_{j2} &= a_{j3s}\tilde{\psi}_{j3s} - a_{j2s}\tilde{\psi}_{j2s}. \end{aligned} \tag{4}$$

The age-education-state-specific survival probabilities in the model are then given by $\psi_{jes} = a_{jes}\tilde{\psi}_{jes}$. Note that this approach relies on the assumption that the age-specific education survival premia are common across all states. That is, for each age, I assume that the mortality difference between individuals with and without a high school degree, and the mortality difference between college graduates and high school graduates, are common across all states.

4.4.2 Consumption

I assume that consumption is drawn from a lognormal distribution with age-education-state-specific mean, μ_{jes}, and variance, σ^2_{jes}. This subsection explains how I derive these parameters. I start by obtaining age-specific consumption profiles. To do this, I use data from the Consumer Expenditure Survey (CEX). The CEX provides detailed data on expenditures, income, and demographics for a representative sample of US households. Following Aguiar and Hurst (2013), I limit the analysis to nondurables excluding health care and education expenditures. I exclude the latter two categories because the utility from consuming these goods varies considerably over the life cycle. Each expenditure category is deflated by means of good/service-specific CPI series as in Aguiar and Hurst (2013). The CEX reports consumption at the household level. I follow Jones and Klenow (2016) and allocate consumption uniformly across household members. I test the sensitivity of the welfare results to this assumption by also considering alternative ways to allocate consumption across household members such as the OECD's equivalent scale. I adjust spending for cohort effects by splitting the individuals into 5-year cohort bins. I then compute life cycle consumption profiles for each cohort, following which I average the profiles across the cohorts. Lastly, I smooth the consumption profiles by means of fourth-order polynomials in age. Since I do not observe individuals older than 94 in the data, I extrapolate the consumption series for 95-100 year-olds. Let the logarithm of the derived consumption profile be denoted by $\tilde{c}_j$.

Following Storesletten et al. (2004), I compute age-education-specific consumption profiles by focusing on consumption of 25–85 year-old household heads in the CEX. That is, I use the age and educational attainment of the household head to derive education-specific consumption profiles. I continue to split individuals into three educational groups: less than high school, high school graduates but no college degree, and college graduates. Building on Aguiar and Hurst (2013), I define the household head as follows. If the household has more than one self-defined heads, I let the head be defined as the male one, the employed one, the oldest one, the married one, or the one with the highest educational attainment, in that order. I continue to adjust for cohort effects by splitting the household heads into 5-year cohort bins. Next, for each educational group, I compute life cycle consumption profiles for each cohort, average the profiles across the cohorts, and then smooth the profiles by means of fourth-order polynomials in age.

Let the logarithm of the derived consumption profile for those without a high school degree be defined as $\tilde{c}_{j1}$. Moreover, let $p_{j2} = \tilde{c}_{j2}/\tilde{c}_{j1}$ be defined as the age-specific ratio between consumption of high school graduates and high school dropouts, and $p_{j3} = \tilde{c}_{j3}/\tilde{c}_{j1}$ be defined as the age-specific ratio between consumption of college graduates and high school dropouts. For each j, let c_{j1} be defined as the solution to the following equation:

$$\tilde{c}_j = \Lambda_{j1} c_{j1} + \Lambda_{j2} p_{j2} c_{j1} + \Lambda_{j3} p_{j3} c_{j1} \tag{5}$$

where Λ_{je} is the age-education-specific distribution of individuals obtained using data from the CPS for the period 2010 to 2017. The adjustment of $\tilde{c}_{je}$ to c_{je} is necessary since the former refers to total consumption of household heads rather than individual consumption. Lastly, I scale c_{je} by a state-specific factor a_s to match total personal consumption expenditures per capita, C_s^{pc}, as reported by the BEA for the period 2010 to 2017:

$$C_s^{pc} = \sum_j \sum_e \Lambda_{jes} a_s \exp\left(c_{je}\right). \tag{6}$$

The age-education-state-specific mean of the lognormal consumption process is then given by $\mu_{jes} = a_s c_{je}$. For now, I let the variance of the process be independent of state and education, and let the variance be given by the average variance across age groups. I relax this assumption later in the chapter. This gives a value of 0.553 for the standard deviation of the consumption process, close to the value of 0.538 used by Jones and Klenow (2016).

4.4.3 Preferences

Building on Hall and Jones (2007), I let preferences be represented by a non-homothetic utility function:

$$u(c) = b + \frac{c^{1-\gamma}}{1-\gamma}, \tag{7}$$

where γ governs the relative risk aversion and the constant term in the utility function, b, governs the value of life in the model. I set γ equal to 2 for the benchmark analysis. Following Jones and Klenow (2016), I calibrate b such that a 40-year-old individual, facing the average cross-state age-specific consumption uncertainty, has a value of remaining life equal to \$7 million in 2010 prices. This gives a value of 4.82 for b, close to the value of 5.00 derived by Jones and Klenow (2016). The discount factor, β, is set to 0.99.

4.5 Results

This section reports the quantitative results. The next subsection quantifies the welfare differences across the states. Subsection 4.5.2 quantifies each state's welfare growth rate between 2000 and 2013.

4.5.1 Welfare across states

This section quantifies the expected lifetime utility of a hypothetical agent who lives her entire life in the state she is born in. I assume that the agent is subject to the age-education-state-specific mortality risk, consumption uncertainty, and educational uncertainty discussed in Section 4.4. Since educational attainment has increased considerably in the United States since the 1960s, I assume that agents draw their educational attainment from the current distribution of 25–29 year-olds.

Appendix Table B.2 summarize the results. Each cell reports the factor by which consumption must be scaled down in all ages in the state with the highest expected lifetime utility to make an agent indifferent between living her entire life in that state and any other state. The states are listed in descending order from highest to lowest welfare level, using the benchmark parameterization discussed in Section 4.4. I find that the states with the five highest welfare levels are, in descending order, Massachusetts, Connecticut, New York, New Jersey, and New Hampshire. In contrast, the five states with the lowest welfare level are, in ascending order, Mississippi, Alabama, Arkansas, Oklahoma, and West Virginia. Consistent with the high cross-state variation in GDP per capita reported in Section 4.2, I find that the magnitude of the welfare differences are substantial. To illustrate, appendix Table B.2 shows that consumption would have to be scaled down by 28 percent in all ages in

Massachusetts to make an agent indifferent between living her entire life in Massachusetts and the state with the median welfare level, Iowa, and would have to be scaled down by 47 percent to make her indifferent between living her entire life in Massachusetts and the state with the lowest welfare level, Mississippi.

To understand the determinants of the cross-state welfare differences, I compute welfare results under alternative parameterizations of the model. In particular, I study counterfactual cases where all states have the same survival probabilities, ψ_{jse}, average consumption, μ_{jse}, and college attainment, $\pi(e|s)$, as Massachusetts. Recall that most states have lower life expectancy than Massachusetts. Consequently, I find that the welfare differences are generally smaller in the model where all the states have the same survival probabilities as Massachusetts. To illustrate, I find that consumption has to be scaled down by 20 percent to make an agent indifferent between living her life in Massachusetts and Maryland under the benchmark parameterization of the model but by 17 percent if the two states had the same survival probabilities. Differences in mortality risk thus account for some of the variation in welfare across the states.

Column four of appendix Table B.2 reports the welfare results from the model when all the states have the same average consumption levels. A comparison of column two and four shows that variation in consumption accounts for most of the variation in welfare. As an example, I find that the welfare difference between Massachusetts and Kentucky would decrease by 23 percentage points if Kentucky had the same average consumption level as Massachusetts. That is, consumption would have to be scaled down by 18 percent to equalize welfare levels in Massachusetts and Kentucky if the two states had the same average consumption levels, compared to a 41 percent required reduction under the benchmark parameterization.

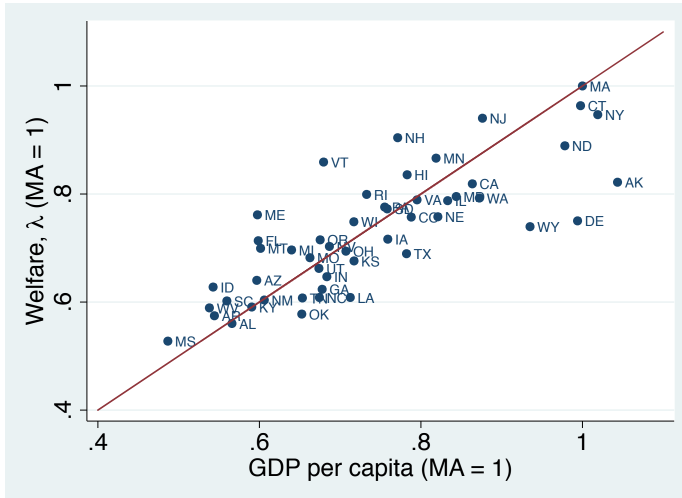

Figure 4.4: Relationship between welfare and GDP per capita

Notes: The graph plots the relationship between welfare and GDP per capita, where welfare refers to expected lifetime utility (see Section 4.3 for details). All calculations are based on data for the period 2010 to 2017. GDP per capita has been normalized by GDP per capita in Massachusetts and welfare has been normalized by welfare in Massachusetts. The red line is the 45 degree line. The correlation between welfare and GDP per capita is 0.80.

The final column of the table reports the results from the counterfactual case where all the states have the same educational attainment as Massachusetts. I find that differences in educational attainment account for some of the variation in welfare across the states. To illustrate, I find that the welfare difference between Massachusetts and Texas would decrease by 5 percentage points if Texas had the same educational attainment as Massachusetts.

I end this subsection by comparing two welfare measures. Figure 4.4 plots the relationship between welfare measure in this chapter and the more commonly used welfare measure, GDP per capita. Although there are important differences between the two measures, the graph shows that the two measures generally provide similar results. The correlation between the two welfare measures across all the states is 0.80.

4.5.2 Welfare across time

Instead of comparing welfare across states, this subsection compares welfare across time. In particular, for each state, I quantify the change in welfare between 2000 and 2013 by computing how much consumption must adjust in all ages in 2000 to make an agent indifferent between the 2000 and 2013 environment. The results are reported in appendix Table B.3. To lower the prevalence of business cycle fluctuations, I pool data for 7-year-periods. Accordingly, 2000 and 2013 refer to average values between 1997–2004 and 2010–2017.

The second column of appendix Table B.3 reports the state-specific annual growth rate in welfare, ordered from the state with the highest welfare growth rate over this time period, North Dakota, to the state with the lowest welfare growth rate, Arizona. There is considerable heterogeneity in welfare growth rates over this time period, ranging from 0.4 to 3.1 percent per year. To better understand the determinants of these growth rates, column two, three, and four of appendix Table B.3 report the annual growth rate in GDP per capita and consumption per capita, and the total change in life expectancy at birth over this time period.[5] I find large variations in the growth rate of consumption per capita across the states, ranging from an annual growth rate of 0.3 percent in Arizona to 3.0 percent in North Dakota. Hence, heterogeneous growth rates in consumption per capita account for most of the variation in the welfare growth rates. The final column shows large variations in life expectancy changes across the states. Compared to the average 1.8 year increase in life expectancy at birth in the United States between 2000 and 2013, life expectancy only increased by 0.2 years in West Virginia. In contrast, life expectancy at birth increased by more than 2.5 years in California, Maryland, New York, and Vermont. Accordingly, heterogeneity in life expectancy changes account for some of the variation in the welfare

[5] GDP per capita increased by 0.8 percent per year in the United States over this time period. The slow growth rate compared to the historical value of 2 percent per year is partially due to the Great Recession.

growth rates.

Figure 4.5 compares the ranking of welfare levels across the states over this time period. The horizontal axis gives the ranking in 2000 and the vertical axis gives the ranking in 2013, with a value of 1 referring to the lowest ranked state and a value of 50 referring to the highest ranked state. States that lie below the 45 degree line such as Colorado declined on the welfare level ranking between 2000 and 2013, while states that lie above the 45 degree line such as North Dakota increased on the welfare level ranking over this time period. The graph shows a high degree of persistence in welfare ranking over time, with most states clustered around the 45 degree line. The correlation between the welfare level ranking in 2000 and 2013 is 0.93. This enables me to draw two conclusions. First, as noted in Subsection 4.5.1, there exists considerable heterogeneity in expected lifetime utility levels across states. Second, these welfare differences are persistent. Hence, although appendix Table B.3 shows that the lower ranked states in 2000 generally experienced a higher welfare growth than the higher ranked states in 2000, there does not appear to be evidence of rapid convergence toward similar welfare levels.

4.6 Conclusion

This chapter developed a model to quantify the welfare differences across the states of the US. Consistent with the data, the model allowed the states to differ in longevity, consumption, and educational attainment. The analysis showed that there exists considerable heterogeneity in welfare across the states. In particular, I found that consumption had to be scaled down by 28 percent in all ages in the state with the highest welfare level, Massachusetts, to make a hypothetical agent indifferent between living her entire life in Massachusetts and the state with the median welfare level, Iowa, and had to be scaled down

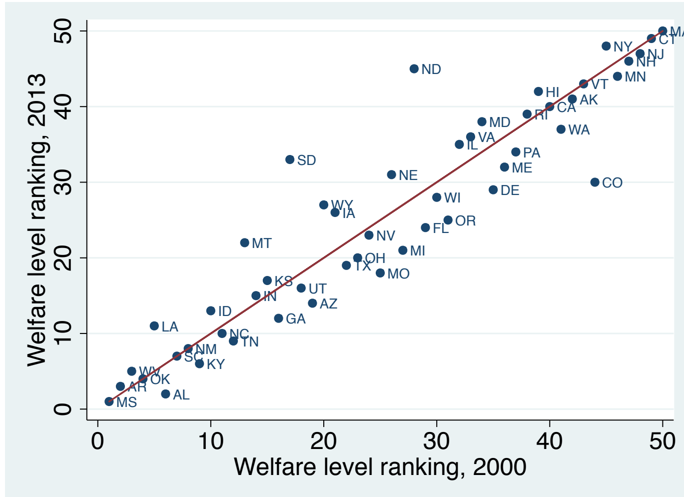

Figure 4.5: Comparing ranking of welfare levels across states in 2000 and 2013

Notes: The graph plots each state's welfare level ranking in 2000 and 2013. Numbers for 2000 are based on average values between 1997 and 2004, and numbers for 2013 are based on average values between 2010 and 2017. The correlation between the welfare level ranking in 2000 and 2013 is 0.93.

by 47 percent to make her indifferent between living her entire life in Massachusetts and the state with the lowest welfare level, Mississippi. A decomposition of the welfare results showed that heterogeneity in consumption per capita accounted for the largest share of the variation in welfare levels, with a smaller share accounted for by life expectancy at birth and educational attainment.

I also compared welfare across time by quantifying how much each state's welfare had grown between 2000 and 2013, taking into account state-specific changes in consumption per capita, mortality risk, and educational attainment. I found that the variation in welfare across the states is very persistent. Accordingly, although the lower ranked states in 2000 generally experienced a higher welfare growth than the higher ranked states in 2000, there does not appear to be evidence of rapid convergence toward similar welfare levels.

References

Aguiar, M., and E. Hurst (2013), “Deconstructing life cycle expenditures,” *Journal of Political Economy*, 121: 437–492.

Aiyagari, R. (1994), “Uninsured idiosyncratic risk and aggregate saving,” *Quarterly Journal of Economics*, 109: 659-684.

Athrey, K. B. (2002), “Welfare implications of the Bankruptcy Reform Act of 1999,” *Journal of Monetary Economics*, 49: 1567-1595.

Attanasio, O.P., and C. Emmerson (2003), “Mortality, Health Status, and Wealth,” *Journal of the European Economic Association*, 1(4): 821-850.

Attanasio, O.P., S. Kitao, and G.L. Violante (2010), “Financing medicare: a general equilibrium analysis,” in J.B. Shoven (Ed.), *Demography and the Economy*, University of Chicago Press.

Austin, D. A. (2015), “Medical Debt as a Cause of Consumer Bankruptcy,” *Maine Law Review*, 67: 1-23.

Bewley, T. (1986), “Stationary monetary equilibrium with a continuum of independently fluctuating consumers,” in W. Hildenbrand and A. Mas-Colell (eds.), *Contributions to Mathematical Economics in the Honor of Gerard Debreu*, North-Holland, Amsterdam, 79-102.

Braun, R.A., K.A. Kopecky, and T. Koreshkova (2017), “Old, Sick, Alone, and Poor: A Welfare Analysis of Old-Age Social Insurance Programmes,” *Review of Economic Studies*, 84: 580-612.

Card, D., C. Dobkin, and N. Maestas (2008), "The Impact of Nearly Universal Insurance Coverage on Health Care Utilization: Evidence from Medicare," *American Economic Review*, 98(5): 2242-2258.

Card, D., C. Dobkin, and N. Maestas (2009), "Does Medicare Save Lives?," *Quarterly Journal of Economics*, 124 (2): 597-636.

Castañeda, A., J. Díaz-Giménez, and J.-V. Ríos-Rull (2003), "Accounting for the U.S. Earnings and Wealth Inequality," *Journal of Political Economy*, 111: 818-857.

Chatterjee, S., D. Corbae, M. Nakajima, and J. Ríos-Rull (2007), "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," *Econometrica,* 75: 1525-1589.

Chetty, R., M. Stepner, S. Abraham, S. Lin, B. Scuderi, N. Turner, A. Bergeron, and D. Cutler (2016), "The Association Between Income and Life Expectancy in the United States, 2001-2014," *Journal of the American Medical Association*, 315(16): 1750-1766.

Cole, H.L., S. Kim, and D. Krueger (2017), "Analyzing the Effects of Insuring Health Risks: On the Trade-off between Short Run Insurance Benefits vs. Long Run Incentive Costs," forthcoming in *Review of Economic Studies*.

Conesa, J.C., D. Costa, P. Kamali, T. J. Kehoe, V.M. Nygard, G. Raveendranathan, and A. Saxena (2018), "The Macroeconomic Effects of Medicare," *Journal of the Economics of Ageing*, 11: 27-40.

Conesa, J.C., and D. Krueger (2006), "On the Optimal Progressivity of the Income Tax Code," *Journal of Monetary Economics*, 53: 1425-1450.

Danaei, G., Ding, E. L., Mozaffarian, D., Taylor, B., Rehm, J., Murray, C. J. L., and M. Ezzati (2009), "The Preventable Causes of Death in the United States: Comparative

Risk Assessment of Dietary, Lifestyle, and Metabolic Risk Factors," *PLoS Medicine*, 6(4): e1000058.

Deaton, A.S., and C. Paxson (2001), "Mortality, education, income, and inequality among American cohorts," in *Themes in the Economics of Aging*, University of Chicago Press, 129- 170.

De Nardi, M., G. Fella, and G. P. Pardo (2016), "The Implications of Richer Earnings Dynamics for Consumption, Wealth, and Welfare," NBER Working Paper 21917.

De Nardi, M., E. French, and J. B. Jones (2010), "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 118: 39-75.

De Nardi, M., E. French, and J. B. Jones (2016), "Medicaid Insurance in Old Age," *American Economic Review*, 106(11): 3480-3520.

De Nardi, M., S. İmrohoroğlu, and T. Sargent (1999), "Intergenerational Transfers and China's Social Security Reform," *Review of Economic Dynamics*, 2: 575-615.

De Nardi, M., S. Pashchenko, and P. Porapakkarm (2017), "The Lifetime Costs of Bad Health," NBER Working Paper No. 23963.

Díaz-Giménez, J., V. Quadrini, and J. Ríos-Rull (1997), "Dimensions of Inequality: Facts on the U.S. Distributions of Earnings, Income, and Wealth," *Federal Reserve Bank of Minneapolis Quarterly Review,* 21: 3-21.

Dobkin, C., A. Finkelstein, R. Kluender, and M. J. Notowidigdo (2016), "The Economic Consequences of Hospital Admissions," NBER Working Paper 22288.

Drewnowski, A., and S. E. Specter (2004), "Poverty and Obesity: The Role of Energy

Density and Energy Costs," *American Journal of Clinical Nutrition*, 79: 6-16.

Drewnowski, A. (2010), "The Cost of US Foods as Related to Their Nutritive Value," *American Journal of Clinical Nutrition*, 92(5): 1181-1188.

French, E. (2005), "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour," *Review of Economic Studies*, 72: 395-427.

Gross, T., and M. J. Notowidigdo (2011), "Health insurance and the consumer bankruptcy decision: Evidence from expansions of Medicaid," *Journal of Public Economics*, 95: 767-778.

Gruber, J. (2008), "Covering the Uninsured in the United States," *Journal of Economic Literature*, 46: 571-606.

Guvenen, F., G. Kaplan, J. Song, and J. Weidner (2017), "Lifetime Incomes in the United States over Six Decades," NBER Working Paper No. 23371.

Guvenen, F., B. Kuruscu, and S. Ozkan (2014), "Taxation of Human Capital and Wage Inequality: A Cross-Country Analysis," *Review of Economic Studies*, 81: 818-850.

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

Himmelstein, D. U., D. Thorne, E. Warren, and S. Woolhandler (2009), "Medical Bankruptcy in the United States, 2007: Results of a National Study," *American Journal of Medicine*, 122: 741-746.

Hosseini, R., K.A. Kopecky, and K. Zhao (2018), "How Important Is Health Inequality for Lifetime Earnings Inequality?," Mimeo.

Hubbard, R., J. Skinner, and S. P. Zeldes (1994), "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving," *Carnegie-Rochester Conference Series on Public Policy*, 40: 59-125.

Hugget, M. (1993), "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of Economic Dynamics and Control*, 17: 953-969.

Kitagawa., E.M., and P.M. Hauser (1973), "Differential Mortality in the United States: A Study in Socioeconomic Epidemiology," *Harvard University Press*.

Kopecky, K.A., and T. Koreshkova (2014), "The Impact of Medical and Nursing Home Expenses on Savings," *American Economic Journal: Macroeconomics*, 6(3): 29-72.

Kotera, T. (2018), "Understanding Life Expectancy Inequality," Mimeo.

Kuhn, M., and J.-V. Ríos-Rull (2013), "2013 Update on the U.S. Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomics," *Federal Reserve Bank of Minneapolis Quarterly Review*, 37 (1): 1-73.

Jones, C. I., and P. J. Klenow (2016), "Beyond GDP? Welfare across countries and time," *American Economic Review*, 106: 2426–2457.

Jung, J., and C. Tran (2016), "Market Inefficiency, Insurance Mandate and Welfare: U.S. Health Care Reform 2010," *Review of Economic Dynamics*, 20: 132-159.

Lim, S. S., et al. (2012),"A Comparative Risk Assessment of Burden of Disease and Injury Attributable to 67 Risk Factors and Risk Factor Clusters in 21 Regions, 1990-2010: A Systematic Analysis for the Global Burden of Disease Study 2010," *Lancet*, 380: 2224-2260.

Lin, C. C., Rogot, E., Johnson, N. J., Sorlie, P. D., and E. Arias (2003), “A further study of life expectancy by socioeconomic factors in the National Longitudinal Mortality Study,” *Ethnicity & Disease*, 13: 240-247.

Livshits, I., J. MacGee, and M. Tertilt (2007), “Consumer Bankruptcy: A Fresh Start,” *American Economic Review*, 97: 402-418.

Mahoney, N. (2015), “Bankruptcy as Implicit Health Insurance,” *American Economic Review*, 105: 710-746.

Milligan, K., and T. Schirle (2018), “The Evolution of Longevity: Evidence From Canada,” NBER Working Paper No. 24929.

Murray, J. L., et al. (2013),“The State of US Health, 1990-2010: Burden of Diseases, Injuries, and Risk Factors,” *Journal of the American Medical Association*, 310(6): 591-608.

Nakajima, and Telyukova (2018), “Medical Expenses and Saving in Retirement: The Case of U.S. and Sweden,” FRB Minneapolis, Opportunity and Inclusive Growth Institute Working Paper No. 8.

Ozkan, S. (2014), “Preventive vs. Curative Medicine: A Macroeconomic Analysis of Health Care over the Life Cycle,” Mimeo.

Palumbo, A. (1999), “Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle,” *Review of Economic Studies*, 66: 395-421.

Pashchenko, S., and P. Porapakkarm (2013), “Quantitative Analysis of Health Insurance Reform: Separating Regulation from Redistribution,” *Review of Economic Dynamics*, 16: 383-404.

Pijoan-Mas, J. and J.-V. Ríos-Rull (2014), “Heterogeneity in Expected Longevities,” *Demography*, 51: 2075-2102.

Rehm, C. D., Peñalvo, J. L., Afshin, A., and D. Mozaffarian (2016), “Dietary Intake Among US Adults, 1999-2012,” *Journal of the American Medical Association*, 315(23): 2542-2553.

Scholz, J.K., and A. Seshadri (2016), “Health and Wealth In a life cycle Model,” Mimeo.

Searle, S. D., Mitnitski, A., Gahbauer, E. A., Gill, T. M., and K. Rockwood (2008), “A standard procedure for creating a frailty index,” *BMC Geriatrics*, 8:24.

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

Appendix A: Appendix to Chapter 2

This section provides additional material. I start by listing the variables that are included in the construction of the frailty index both in the MEPS and in the HRS. Next, I provide additional information about the effects of medical spending, income, and education on frailty. The following subsection shows how age, education, and frailty affect labor market outcomes. Next, I show how age, education, and frailty affect medical expenditure risk. I then show that individuals with employer provided insurance tend to be richer, more educated, and less frail. The next subsection provides evidence that healthier consumption goods tend to be more expensive on average than unhealthy goods. Next, I show that richer and educated individuals lead healthier lives on average, and provide estimates of the magnitude of these relationships. Lastly, I use the benchmark model to study how income inequality and differences in educational attainment affect life expectancy inequality in the United States.

Frailty index

Table A.1 lists all the variables that are included in the construction of the frailty index in the MEPS. Following conventions in gerontology (see for example Searle et al., 2008), I include all disease diagnoses such a whether or not the individual has been diagnosed with for example cancer or diabetes. I also include functional limitations such as whether or not the individual needs help with activities of daily living such as eating, dressing, or bathing, cognitive limitations, and a measure of whether or not the individual needs help with instrumental activities of daily living such as taking medications or preparing meals. I follow the same approach to construct frailty indices in the HRS. A list of the variables used to construct the index is given in Table A.2.

Effects of medical spending, income, and education on frailty

This section provides additional details about how medical spending, income, and education affect frailty. To do this, I run an instrumental variable regression of frailty in year 2 on frailty, medical spending, income, and education in year 1, where I instrument medical spending by the density of physicians by county. I also control for smoking, exercising, county of residence, occupation, a cubic in age, county of residence, state plus cohort fixed effects, marital status, gender, and race. The results are reported in Table A.3. I find that, conditional on initial health, medical spending has no significant effect on frailty. In contrast, both education and higher income are predictive of lower next-period frailty. Note, however, that income is also likely to be endogenous, and it would therefore be incorrect to say that these results establish a causal effect of income on frailty, even though the regression controls for educational attainment, medical spending, and initial frailty.

Medical spending

This section provides additional details about how age and frailty affects medical spending risk. Recall from Section 2.4.4 that I split medical expenses into three categories in the MEPS: less than $\$1,000$, between $\$1,000$ and $\$40,000$, and greater than $\$40,000$. To compute transition probabilities between the medical expenditure states, I first run an ordered logistic regression of medical spending in year 2 on medical spending, age and frailty in year 1. I also controls for gender, race, and interaction terms between age, frailty, and medical spending. The results are reported in Table A.4. Transition probabilities between the medical expenditure states can then be derived by applying the standard ordered logistic formula.

Labor earnings

This section provides additional details about how age, education, and frailty affect labor market outcomes. I start by examining how labor productivity as measured by hourly wages is affected by frailty. Table A.5 reports the results from an ordinary least squares regression of the logarithm of hourly wages in year 2 on the logarithm of hourly wages, age, frailty, and education in year 1. The regression also controls for gender, race, and a quadratic in age. I restrict the MEPS sample to individuals aged 20 to 64 that work at least 10 hours per week. The regression shows that labor productivity is negatively affect by an increase in frailty. In particular, I find that the semi-elasticity of hourly wages with respect to frailty is -0.16.

Next, I examine how frailty affects hours worked, focusing on 20 to 64 year-olds that work at least 10 hours. To do this, I regress the logarithm of hours worked per week in year 2 on the logarithm of hours worked per week, age, frailty, and education in year 1. I also control for gender, race, and a quadratic in age. The results are reported in Table A.6. The regression shows that individuals respond to an increase in frailty by reducing the number of hours worked per week. In particular, I find that the semi-elasticity of hours worked per week with respect to frailty is -0.032.

Following Conesa et al. (2018), I estimate the agents' deterministic life cycle labor productivity by regressing the logarithm of hourly wages on age, education, higher-order moments of age, and interaction terms. Figure A.1 plots the derived labor productivity profiles by age and education, where I have normalized productivity at age 20 to 1.

Relative price of healthy and unhealthy goods

Drewnowski and Specter (2004) and Drewnowski (2010) show that low-energy dense goods tend to be more expensive than high-energy dense goods. Figure A.2 illustrates this relationship between cost and energy density for 9 consumption categories: vegetables; fruit; milk and milk products; eggs; dry beans, legumes, nuts, and seeds; meat, poultry, and fish; sugar, sweets, and beverages; grain products; and fats, oils, and salad dressings. Here, energy cost is measured by cost per 100 kcal and energy density is measured by kcal per 100 gram.

The prices of low energy dense goods have also increased more rapidly than the prices of high energy dense goods. Using consumer price index data from the BLS for the period 1980 to 2015, I find that the price for low-energy dense goods such as fruits and vegetables have increased faster than the price for high-energy dense goods such as sugar, sweets, and carbonated drinks. This is illustrated in Figure A.3.

Health behavior by income and education It is well known that the relationship between income and health, and the relationship between education and health, are partially driven by differences in health behavior. That is, richer and educated individuals lead healthier lives on average. This subsection provides estimates of the magnitude of these relationships. To do this, I use data from the National Health and Nutrition Examination Survey (NHANES), a program of studies designed to assess the health and nutritional status of adults and children. The NHANES provides detailed cross-sectional data on demographics, socioeconomic characteristics, dietary information, and health indicators for a nationally representative sample of Americans.

I start by testing for differences in healthy eating by income and educational attainment, focusing on the following measures: a self-rated measure of how healthy the respondent's diet is, a measure of whether the respondent can afford to eat balanced meals, and how often the respondent has fruits, dark green vegetables, and soft drinks available at home. To maintain consistency with Subsection 2.2.1, I continue to focus on college and non-college educated individuals, where college refers to individuals with at least a four-year college degree. All regressions control for the respondent's self-rated health, a quadratic in age, gender, and race. The results are reported in Table A.7. I find that educated and high-income individuals are more likely to have fruits and dark green vegetables available at home. Similarly, these groups are more likely to afford balanced meals. In addition, I find that educated individuals are less likely to have soft drinks available at home and more likely to rate their own diet as healthy.

Next, I test for differences in exercising, smoking, and drinking rates by income and educational attainment, focusing on whether or not the respondent engages in any vigorous- or moderate-intensity sports, fitness, or recreational activities, whether the respondent smokes or drinks alcoholic beverages, and how much the respondent smokes or drinks conditional on smoking and drinking, respectively. As earlier, all regressions control for the respondent's self-rated health, a quadratic in age, gender, and race. The results, which are reported in Table A.8, show that educated and high-income individuals are more likely to engage in both vigorous- and moderate-intensity recreational activities. Similarly, these groups are less likely to smoke. Moreover, I find that, conditional on smoking, educated individuals smoke fewer cigarettes per day than non-educated individuals. Lastly, I find that both educated and high-income individuals are *more* likely to drink alcoholic beverages. That said, I find that, conditional on drinking, non-educated and low-income individuals drink more than educated and high-income individuals. These findings confirm that educated and high-income individuals lead healthier lives, and hence that the relationships between income, education, and frailty documented in Subsection 2.2.1 are partially driven by differences in health behavior.

Determinants of life expectancy inequality This section provides additional details about the determinants of life expectancy inequality. In particular, I study how income inequality and differences in educational attainment affect life expectancy inequality as measured by the difference in life expectancy at age 40 by income quartile. I do this by comparing the benchmark model with two counterfactual model environments. First, I assume that frailty transitions are exogenous and hence independent of consumption. Next, I assume that education affects income but not frailty transitions. This latter model enables me to decompose the income effect of education from the direct health effect of education.

Recall from equation (1) that frailty transitions depend on current frailty, age, education, and consumption of healthy goods, $\mathbb{P}(f, j, e, c_h)$. To understand how income inequality affects life expectancy inequality, I solve a counterfactual model where frailty transitions are exogenous and hence independent of consumption. In practice, I set c_h in the agents' frailty transitions equal to a common constant, $\bar{c}_h$, and adjust its value such that average life expectancy at age 20 in this model is identical to the life expectancy in the benchmark model. That is, frailty transitions are assumed to be given by $\mathbb{P}(f, j, e, \bar{c}_h)$, where $\bar{c}_h$ is identical for all agents. Note that although income does not directly affect life expectancy in the environment with exogenous frailty transitions, income will still be positively correlated with longevity due to the negative effect of frailty on earnings. Lastly, since frailty transitions are independent of consumption, I assume that the relative price of unhealthy goods, p, is equal to one, and hence that agents are indifferent between the two consumption goods.

Table A.9 compares life expectancy and life expectancy inequality in the two models.

The first row reports average life expectancy at age 20. Due to the adjustment of $\bar{c}_h$, average life expectancy is identical in the benchmark model and the model with exogenous frailty transitions, as can be seen by comparing the columns labeled benchmark and exogenous. Next, I compare life expectancy inequality in the two models. To compute life expectancy by income, I first rank the 40-year-olds in the model by income and split the agents into four income quartiles. I then simulate life trajectories by using the policy functions derived from the model, the Markov processes for the shocks, and the age and frailty-specific survival probabilities. This enables me to compute how many of the agents will survive until a given age. To illustrate, conditional on being in the top income quartile at age 40, I can compute how many of these agents will survive until age 41, 42, etc. Life expectancy by income quartile at age 40 is then given by the number of years an agent in a given income quartile can expect to live. The second row of Table A.9 reports how much longer the richest 25 percent of 40-year-olds can expect to live compared to the poorest 25 percent. A comparison of the two models shows that life expectancy inequality is 7.01 years in the benchmark model, but only 1.72 years in the model with exogenous frailty transitions. That is, life expectancy inequality by income quartile is 5.30 years lower in the model with exogenous, consumption-independent, frailty transitions.

This finding that life expectancy inequality is 76 percent lower in the model with exogenous frailty transitions suggests that income inequality is a key driver of life expectancy inequality in the United States. Consistent with this finding, both life expectancy inequality and income inequality have increased considerably over time. Recent research by Chetty et al. (2016) shows that, between 2001 and 2014, life expectancy increased at an annual rate of 0.20 years for men in the top earnings quartile, but only 0.08 years for men in the bottom earnings quartile. Similarly, a number of papers have documented the increase in earnings inequality in the United States. As an example, using data from the Survey of

Consumer Finances, Kuhn and Ríos-Rull (2013) show that the GINI coefficient of labor earnings increased from 0.61 to 0.67 between 1989 and 2013. Recent research by Guvenen et al. (2017) also shows that *lifetime* earnings inequality has increased in the United States. Using individual earnings histories from the Social Security Administration, they document that, from the 1967 cohort to the 1983 cohort, only the top 10 percent of the distribution of men experienced a significant increase in lifetime earnings. More than 75 percent of the distribution of men, on the other hand, experienced either no increase, or even declining, gains in lifetime earnings.

I run two counterfactuals to examine how differences in educational attainment affect life expectancy inequality. First, I assume that all agents have the same frailty transitions as college educated agents. That is, I assume that frailty transition probabilities are given by $\mathbb{P}(f, j, e_c, c_h)$, where e_c denote college-educated. Next, I assume that all agents have the same frailty transitions as non-college educated agents, and hence that frailty transition probabilities are given by $\mathbb{P}(f, j, e_{nc}, c_h)$, where e_{nc} denotes non-college-educated. In both cases, I continue to assume that education affects income through its effect on labor productivity. These counterfactuals enable me to decompose the income effect of education from the direct health-protective effect of education.

The results are reported in the last two columns of Table A.9. A comparison of the benchmark column and the college column shows that life expectancy at 20 is 78.90 years in the benchmark model, but 79.69 years in the model where all agents have the same frailty transitions as college-educated agents. Note, however, that this 0.79 year increase in life expectancy is likely to provide a lower bound on the effect of education on longevity. This follows from the assumption that education only affects life expectancy through its effect on frailty transitions. This assumption is consistent with the finding in Pijoan-Mas and Ríos-Rull (2014) that socioeconomic characteristics such as education, income, and wealth are health-protecting, but have otherwise little impact on two-year mortality rates conditional on health. That said, given the large difference in life expectancy by education documented by a number of studies, it is possible that educated individuals have lower mortality rates even after controlling for other factors such as age, frailty, gender, race, and health behavior. I also find that education has a large effect on life expectancy inequality. As reported in Table A.9, the difference in life expectancy inequality is 1.54 years, or 22 percent, lower in the model where all agents are subject to college educated frailty transitions than in the benchmark model. Lastly, I compare the benchmark model environment with an economy where all agents have the same frailty transitions as non-college educated agents. As shown in Table A.9, life expectancy is 0.33 years lower in this model than in the benchmark model. Moreover, I find that life expectancy inequality is 0.41 years, or 6 percent, higher in the model where all agents are subject to non-college educated frailty transitions.

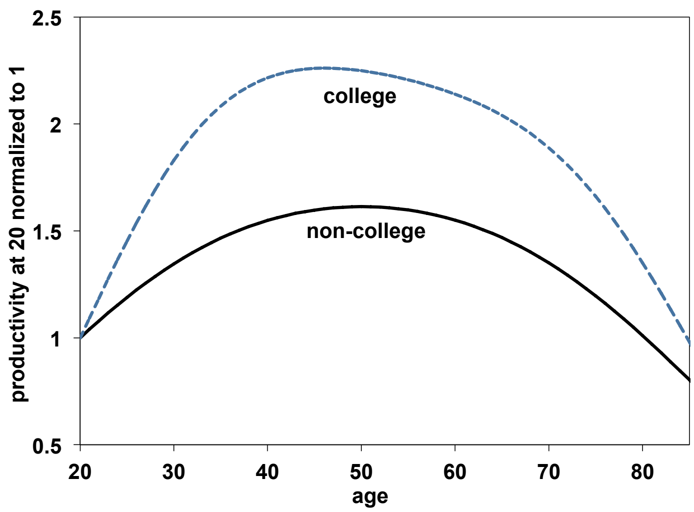

Figure A.1: Labor productivity profiles by age and education

Notes: The graph plots labor productivity profiles by age and education. Productivity profiles are derived from a regression of the logarithm of wages on age, education, higher-order moments of age, and interaction terms between age and education. Data source: MEPS.

Table A.1: Variables included in the frailty index (MEPS)

Variable	Cutoff points
Angina diagnosis	Yes = 1, No = 0
Arthritis diagnosis	Yes = 1, No = 0
Asthma diagnosis	Yes = 1, No = 0
Cancer diagnosis	Yes = 1, No = 0
Coronary heart disease diagnosis	Yes = 1, No = 0
High cholestorol diagnosis	Yes = 1, No = 0
Diabetes diagnosis	Yes = 1, No = 0
Emphysema diagnosis	Yes = 1, No = 0
High blood pressure diagnosis	Yes = 1, No = 0
Heart attack (MI) diagnosis	Yes = 1, No = 0
Other heart disease diagnosis	Yes = 1, No = 0
Stroke diagnosis	Yes = 1, No = 0
Need help with activities of daily living (ADLs)	Yes = 1, No = 0
Use assistive devices	Yes = 1, No = 0
Cognitive limitations	Yes = 1, No = 0
Need help with instrument activities of daily living (IADLs)	Yes = 1, No = 0
Limitation in physical functioning	Yes = 1, No = 0
Any limitations with work/housework/school	Yes = 1, No = 0
Health limitations climbing stairs	Limited a lot = 1, Limited a little = 0.5, Not limited = 0
Health limitations moderate activities	Limited a lot = 1, Limited a little = 0.5, Not limited = 0
Pain limit normal work	Extremely = 1, Quite a bit = 0.75, Moderately = 0.5, A little bit = 0.25, Not at all = 0
Accomplish less because of physical problems	All of the time = 1, Most of the time = 0.75, Some of the time = 0.5, Little of the time = 0.25, None of the time = 0
Work lim. because of phys. prob.	All of the time = 1, Most of the time = 0.75, Some of the time = 0.5, Little of the time = 0.25, None of the time = 0

Notes: The table lists the variables used to construct the frailty index in the MEPS. Cutoff points for the different variables are given in the second column. All binary variables such as whether or not the person has a cancer diagnosis are assigned a value of 0 if the person does not have this health deficit and a value of 1 if the person does have this health deficit. All variables that include more than two responses are assigned values from 0 to 1 depending on the number of possible responses. See Section 2.2.1 for details.

Table A.2: Variables included in the frailty index (HRS)

Dependent variable: Variable	Cutoff points
Ever had high blood pressure	Yes = 1, No = 0
Ever had diabetes	Yes = 1, No = 0
Ever had cancer	Yes = 1, No = 0
Ever had lung disease	Yes = 1, No = 0
Ever had heart problems	Yes = 1, No = 0
Ever had stroke	Yes = 1, No = 0
Ever had arthritis	Yes = 1, No = 0
Difficulty walking across room	Can't do = 1, Yes = 0.5, No = 0
Difficulty dressing	Can't do = 1, Yes = 0.5, No = 0
Difficulty bathing or showering	Can't do = 1, Yes = 0.5, No = 0
Difficulty eating	Can't do = 1, Yes = 0.5, No = 0
Difficulty getting out of bed	Can't do = 1, Yes = 0.5, No = 0
Difficulty using the toilet	Can't do = 1, Yes = 0.5, No = 0
Difficulty using a map	Can't do = 1, Yes = 0.5, No = 0
Difficulty using a phone	Can't do = 1, Yes = 0.5, No = 0
Difficulty managing money	Can't do = 1, Yes = 0.5, No = 0
Difficulty taking medication	Can't do = 1, Yes = 0.5, No = 0
Difficulty shopping for groceries	Can't do = 1, Yes = 0.5, No = 0
Difficulty preparing meals	Can't do = 1, Yes = 0.5, No = 0
Difficulty walking several blocks	Can't do = 1, Yes = 0.5, No = 0
Difficulty jogging a mile	Can't do = 1, Yes = 0.5, No = 0
Difficulty walking one block	Can't do = 1, Yes = 0.5, No = 0
Difficulty sitting for two hours	Can't do = 1, Yes = 0.5, No = 0
Difficulty getting up from chair	Can't do = 1, Yes = 0.5, No = 0
Difficulty climbing several flights of stairs	Can't do = 1, Yes = 0.5, No = 0
Difficulty climbing one flight of stairs	Can't do = 1, Yes = 0.5, No = 0
Difficulty stooping/kneeling/crouching	Can't do = 1, Yes = 0.5, No = 0
Difficulty lifting/carrying 10lbs	Can't do = 1, Yes = 0.5, No = 0
Difficulty picking up a dime	Can't do = 1, Yes = 0.5, No = 0
Difficulty reaching/extending arms up	Can't do = 1, Yes = 0.5, No = 0
Difficulty pushing/pulling large objects	Can't do = 1, Yes = 0.5, No = 0

Notes: The table lists the variables used to construct the frailty index in the HRS. Cutoff points for the different variables are given in the second column. All binary variables such as whether or not the person ever had cancer are assigned a value of 0 if the person does not have this health deficit and a value of 1 if the person does have this health deficit. All variables that include more than two responses are assigned values from 0 to 1 depending on the number of possible responses. See Section 2.2.1 for details.

Table A.3: Frailty determinants: IV regression results

Dependent variable: Frailty in year two	
Frailty	0.905
	(0.004)
Age	0.001
	(0.0005)
Logarithm of medical spending	-0.005
	(0.008)
Logarithm of income	-0.004
	(0.0004)
Education	-0.004
	(0.0007)
Constant	0.0005
	(0.008)
Centered R^2	0.850
Number of observations	56933

Notes: The table reports results from an instrumental variable regression of frailty in year 2 on frailty, age, medical spending, income, and education in year 1, where medical spending has been instrumented by the density of physicians by county (see ongoing work by Bairoliya et al. (2018) for further details). I also control for gender, race, a cubic in age, county of residence, and state plus cohort fixed effects. Education is split into two categories: college and non-college, where college refers to individuals with at least a 4-year college degree. Non-college refers to everyone else. The sample includes individuals aged 20 to 85. Numbers in parentheses denote clustered standard errors. Sample weights are used to account for complex survey design. Data source: MEPS.

Table A.4: Medical spending: Ordered logistic regression results

Dependent variable: Medical spending category in year two	
Medical spending	1.192
	(0.072)
Age	0.012
	(0.002)
Frailty	7.271
	(0.347)
Cutoff 1	3.081
	(0.114)
Cutoff 2	8.456
	(0.116)
R^2	0.261
Number of observations	77640

Notes: The table reports results from an ordered logistic regression of medical spending in year 2 on medical spending, age, and frailty in year 1. The regression also controls for gender, race, and interaction terms between age, frailty, and medical spending. Medical spending is split into three categories: less than $\$1,000$, between $\$1,000$ and $\$40,000$, and greater than $\$40,000$. Education is split into two categories: college and non-college, where college refers to individuals with at least a 4-year college degree. Non-college refers to everyone else. The cutoffs reflect the predicted cumulative probabilities at covariate values of zero. The sample includes individuals aged 20 to 85. Numbers in parentheses denote robust standard errors. Sample weights are used to account for complex survey design. Data source: MEPS.

Table A.5: Labor productivity: Ordinary least squares regression results

Dependent variable: Logarithm of hourly wages in year 2	
Logarithm of hourly wage	0.594
	(0.008)
Age	0.014
	(0.002)
Frailty	-0.155
	(0.043)
Education	0.167
	(0.007)
Constant	0.850
	(0.041)
R^2	0.468
Number of observations	32401

Notes: The table reports results from an ordinary least squares regression of the logarithm of hourly wages in year 2 on the logarithm of hourly wages, age, frailty, and education in year 1. The regression also controls for gender, race, and a quadratic in age. Education is split into two categories: college and non-college, where college refers to individuals with at least a 4-year college degree. Non-college refers to everyone else. The sample includes individuals aged 20 to 64 that work at least 10 hours per week. Numbers in parentheses denote robust standard errors. Sample weights are used to account for complex survey design. Data source: MEPS.

Table A.6: Hours worked: Ordinary least squares regression results

Dependent variable: Logarithm of hours worked per week in year 2	
Logarithm of hours worked per week	0.863
	(0.007)
Age	0.001
	(0.001)
Frailty	-0.032
	(0.011)
Education	0.003
	(0.002)
Constant	0.480
	(0.029)
R^2	0.788
Number of observations	40885

Notes: The table reports results from an ordinary least squares regression of the logarithm of hours worked per week in year 2 on the logarithm of hours worked per week, age, frailty, and education in year 1. The regression also controls for gender, race, and a quadratic in age. Education is split into two categories: college and non-college, where college refers to individuals with at least a 4-year college degree. Non-college refers to everyone else. The sample includes individuals aged 20 to 64 that work at least 10 hours per week. Numbers in parentheses denote robust standard errors. Sample weights are used to account for complex survey design. Data source: MEPS.

Table A.7: Healthy eating: Ordinary least squares regression results

	Dependent variable				
Regressor	Fruits	Vegetables	Soft drinks	Healthy diet	Balanced meals
Education	0.154	0.085	-0.413	0.239	0.027
	(0.025)	(0.032)	(0.051)	(0.019)	(0.006)
Income	0.053	0.051	-0.003	0.005	0.081
	(0.007)	(0.009)	(0.013)	(0.005)	(0.002)
R^2	0.057	0.032	0.033	0.199	0.133
Number of observations	9632	9631	9632	22907	25590

Notes: The table reports results from 5 ordinary least squares regressions: "How often do you have fruits available at home," "How often do you have dark green vegetables available at home," "How often do you have soft drinks available at home," "How healthy is your overall diet," and "How often could you not afford to eat balanced meals." For questions 1, 2, 3, and 5, higher values means *more often*. For question 4, higher values means *more healthy*. All regressions control for self-rated health, a quadratic in age, gender, and race. Income refers to the ratio of family income to the federal poverty level, which is topcoded at 5. Education is split into two categories: college and non-college, where college refers to individuals with at least a 4-year college degree. Non-college refers to everyone else. The samples include individuals aged 20 to 85. Numbers in parentheses denote robust standard errors. Sample weights are used to account for complex survey design. Data source: NHANES.

Table A.8: Smoking, drinking, and exercising: Ordinary least squares regression results

	Dependent variable					
Regressor	Vigorous act.	Moderate act.	Smoking	Smok. inten.	Alcohol	Alc. inten.
Education	0.145	0.133	-0.335	-0.342	0.033	-0.201
	(0.010)	(0.012)	(0.027)	(0.054)	(0.008)	(0.013)
Income	0.021	0.030	-0.092	-0.015	0.022	-0.030
	(0.003)	(0.003)	(0.007)	(0.010)	(0.002)	(0.004)
R^2	0.180	0.084	0.197	0.149	0.093	0.236
Number of observations	18741	18741	12623	5753	26909	17826

Notes: The table reports results from 6 ordinary least squares regressions: “Do you do any vigorous-intensity sports, fitness, or recreational activities,” “Do you do any moderate-intensity sports, fitness, or recreational activities,” “Do you currently smoke,” “What is the average number of cigarettes per day on days that you smoke,” “Have you had at least 12 alcoholic drinks in a year,” and “How many drinks did you have on days that you drank.” Questions 1, 2, 3, and 5 are binary questions, where a value of 0 means no. All regressions control for self-rated health, a quadratic in age, gender, and race. Income refers to the ratio of family income to the federal poverty level, which is topcoded at 5. Education is split into two categories: college and non-college, where college refers to individuals with at least a 4-year college degree. Non-college refers to everyone else. The samples include individuals aged 20 to 85. Numbers in parentheses denote robust standard errors. Sample weights are used to account for complex survey design. Data source: NHANES.

Table A.9: Life expectancy and life expectancy inequality: Counterfactual model environments

Variable	Benchmark	Adjust frailty transitions: Exogenous	Adjust frailty transitions: College	Adjust frailty transitions: Non-college
Life expectancy at 20	78.90	78.90	79.69	78.57
Life expectancy difference at 40 by income quartile	7.01	1.72	5.47	7.42

Notes: The column labeled exogenous corresponds to the model where frailty transitions are exogenous and given by $\mathbb{P}(f, j, e, \bar{c}_h)$, where $\bar{c}_h$ is a common constant across all agents (see equation (1) for details). The column labeled college corresponds to the model where all agents are assumed to have the same frailty transitions as college educated agents, but where an agent's education still affects income. Similarly, the column labeled non-college corresponds to the model where all agents are assumed to have the same frailty transitions as non-college educated agents, but where an agent's education still affects income.

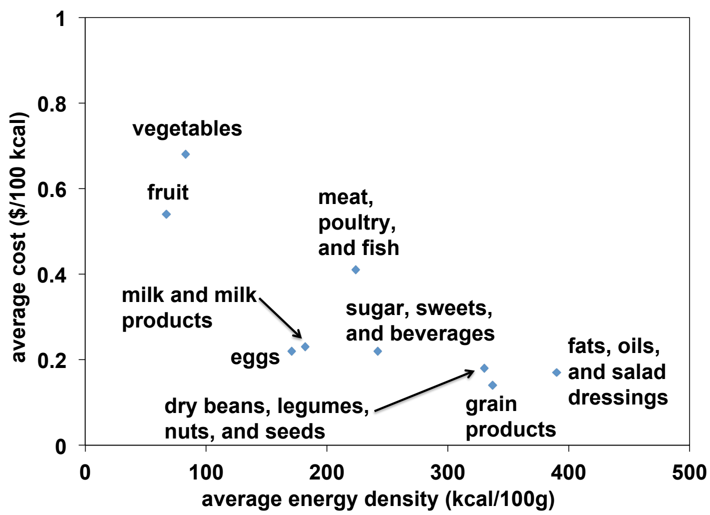

Figure A.2: Relationship between price and energy density for selected consumption goods

Notes: The graph plots the relationship between average energy cost, as measured by the cost per 100 kcal, and average energy density, as measured by kcal per 100 gram, for selected consumption goods. Source: Drewnowski (2010).

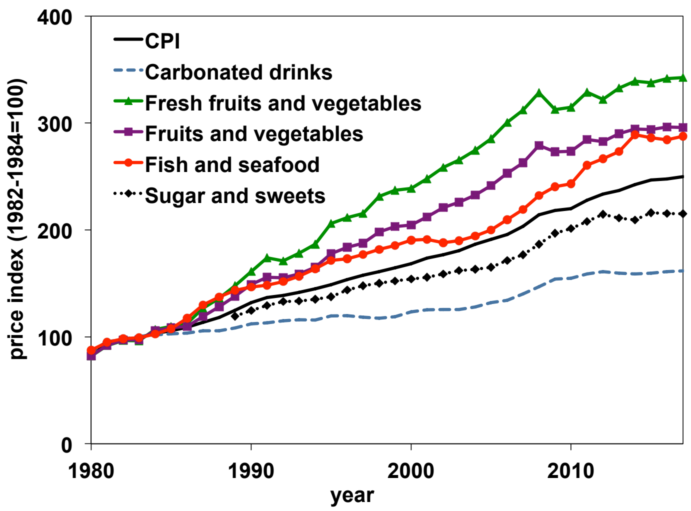

Figure A.3: Trends in price indices for selected consumption goods

Notes: The left panel shows how the price index for selected consumption goods have evolved since 1982. Each index has been normalized to 100 between 1982 and 1984. Data source: BLS.

Appendix B: Appendix to Chapter 4

Table B.1: Descriptive statistics

State	Income and consumption		Educational attainment		Longevity	
	GDP/capita (1000s of 2010 dollars)	Cons./capita (1000s of 2010 dollars)	No HS (% of 25+)	Bachelor's+ (% of 25+)	Infant mortality (deaths per 1000 births)	Life expec. (years)
United States	50.6	34.3	13.2	30.0	5.9	78.7
Alabama	37.2	27.8	15.7	23.6	8.5	75.2
Alaska	68.6	40.3	7.9	28.3	5.4	77.9
Arizona	39.3	29.3	13.9	27.6	5.7	79.4
Arkansas	35.8	27.5	15.1	21.3	7.6	75.7
California	56.8	36.1	18.1	31.7	4.4	80.9
Colorado	51.8	35.1	9.2	38.4	4.9	80.0
Connecticut	65.6	43.3	10.2	37.5	5.0	80.6
Delaware	65.4	36.4	11.3	30.0	7.6	78.3
Florida	39.4	33.5	13.0	27.5	6.3	79.5
Georgia	44.6	29.5	14.4	29.0	7.0	77.3
Hawaii	51.5	37.4	8.9	30.9	5.6	81.6
Idaho	35.7	28.7	10.3	25.9	5.2	79.2
Illinois	54.8	35.9	12.0	32.5	6.4	78.9
Indiana	45.0	30.7	12.1	24.4	7.3	77.1
Iowa	49.9	32.0	8.5	26.9	5.0	79.4
Kansas	47.2	31.2	9.8	31.4	6.1	78.4
Kentucky	38.9	28.5	15.7	22.4	6.8	75.6
Louisiana	46.9	29.6	16.5	22.6	7.9	75.7
Maine	39.3	37.2	8.4	29.1	6.3	78.6
Maryland	55.5	38.0	10.7	37.9	6.6	78.9
Massachusetts	65.8	44.8	10.1	40.8	4.1	80.3
Michigan	42.1	33.1	10.3	27.0	6.8	77.9
Minnesota	53.9	37.9	7.5	33.9	4.9	80.7
Mississippi	32.0	25.7	17.3	20.7	9.1	74.7
Missouri	43.6	33.0	11.5	27.2	6.5	77.3
Montana	39.6	33.8	7.3	29.9	5.7	78.4
Nebraska	54.0	34.1	9.3	29.7	5.4	79.4
Nevada	45.2	34.2	14.7	23.1	5.4	77.9
New Hampshire	50.7	42.8	7.7	35.0	4.3	79.7
New Jersey	57.6	42.9	11.2	37.1	4.6	80.1
New Mexico	39.9	29.3	15.5	26.3	5.6	77.8
New York	67.0	40.9	14.3	34.3	4.8	80.5
North Carolina	44.4	28.6	13.9	28.7	7.1	77.7
North Dakota	64.4	40.7	8.2	28.3	6.2	79.5
Ohio	46.5	33.4	10.8	26.2	7.4	77.2
Oklahoma	42.9	28.4	12.9	24.2	7.5	75.6
Oregon	44.4	33.5	10.1	31.1	4.9	79.4
Pennsylvania	49.7	36.5	10.6	28.9	6.5	78.3
Rhode Island	48.2	36.8	13.7	32.0	6.0	79.6
South Carolina	36.8	28.7	14.3	26.1	7.0	76.7
South Dakota	49.9	35.5	9.0	27.2	6.8	79.1
Tennessee	43.0	29.2	14.3	25.1	7.3	76.0
Texas	51.5	31.5	17.9	27.7	5.8	78.5
Utah	44.3	29.3	8.7	31.5	5.2	79.5
Vermont	44.7	40.5	8.2	35.9	4.4	79.7
Virginia	52.3	36.8	11.6	36.5	6.2	79.0
Washington	57.4	36.5	9.5	33.2	4.5	80.0
West Virginia	35.4	29.2	14.9	19.2	7.2	75.1
Wisconsin	47.2	34.3	8.8	28.1	6.0	79.4
Wyoming	61.5	35.3	7.5	26.0	5.5	78.4

Notes: The table provides summary statistics for each age. All numbers refer to average values for the period 2010 to 2017. See the text for details.

Table B.2: Welfare across states

State	Benchmark	Same survival	Same consumption	Same education
Massachusetts	1.00	1.00	1.00	1.00
Connecticut	0.96	0.95	0.99	0.99
New York	0.95	0.92	1.01	0.97
New Jersey	0.94	0.94	0.97	0.97
New Hampshire	0.90	0.92	0.95	0.93
North Dakota	0.89	0.90	0.94	0.94
Minnesota	0.87	0.85	1.00	0.89
Vermont	0.86	0.87	0.95	0.89
Hawaii	0.84	0.80	0.98	0.90
Alaska	0.82	0.87	0.87	0.88
California	0.82	0.79	0.97	0.88
Rhode Island	0.80	0.80	0.94	0.84
Maryland	0.80	0.83	0.92	0.82
Washington	0.79	0.79	0.94	0.84
Virginia	0.79	0.81	0.93	0.82
Illinois	0.79	0.81	0.94	0.81
Pennsylvania	0.78	0.81	0.91	0.81
South Dakota	0.77	0.79	0.92	0.82
Maine	0.76	0.78	0.90	0.81
Nebraska	0.76	0.76	0.95	0.79
Colorado	0.76	0.76	0.95	0.78
Delaware	0.75	0.78	0.88	0.80
Wisconsin	0.75	0.75	0.94	0.79
Wyoming	0.74	0.76	0.88	0.80
Iowa	0.72	0.72	0.95	0.75
Oregon	0.72	0.72	0.92	0.76
Florida	0.71	0.72	0.92	0.76
Nevada	0.70	0.73	0.86	0.77
Montana	0.70	0.73	0.89	0.74
Michigan	0.70	0.73	0.89	0.73
Ohio	0.69	0.73	0.86	0.73
Texas	0.69	0.70	0.89	0.74
Missouri	0.68	0.72	0.86	0.72
Kansas	0.68	0.70	0.91	0.70
Utah	0.66	0.67	0.92	0.70
Indiana	0.65	0.68	0.86	0.68
Arizona	0.64	0.64	0.91	0.68
Idaho	0.63	0.63	0.91	0.67
Georgia	0.62	0.65	0.87	0.66
Louisiana	0.61	0.66	0.82	0.64
North Carolina	0.61	0.63	0.88	0.64
Tennessee	0.61	0.65	0.84	0.63
New Mexico	0.60	0.63	0.84	0.65
South Carolina	0.60	0.63	0.86	0.63
Kentucky	0.59	0.63	0.82	0.62
West Virginia	0.59	0.64	0.81	0.62
Oklahoma	0.58	0.62	0.81	0.61
Arkansas	0.57	0.61	0.82	0.61
Alabama	0.56	0.60	0.80	0.60
Mississippi	0.53	0.57	0.78	0.56

Notes: Each cell reports the factor by which consumption must be scaled down in all ages in the state with the highest expected lifetime utility, Massachusetts, to make an agent indifferent between living her entire life in that state and any other state. *Same survival*, *Same consumption*, and *Same education* refer to the counterfactual cases where all states have the same survival probabilities, ψ_{jse}, same average consumption, μ_{jse}, and same educational attainment as Massachusetts.

Table B.3: Growth rates in welfare, GDP/capita, consumption/capita, and change in life expectancy between 2000 and 2013

State	Welfare growth	GDP/capita growth	Cons./capita growth	Life expec. change
North Dakota	3.1	4.5	3.0	0.8
South Dakota	2.3	2.1	2.2	1.2
Montana	1.9	1.8	1.8	1.1
Wyoming	1.8	2.4	1.7	1.5
Vermont	1.7	1.0	1.5	2.8
Hawaii	1.7	1.2	1.6	1.8
West Virginia	1.6	1.2	1.6	0.2
Nebraska	1.6	2.0	1.5	1.3
New York	1.6	1.3	1.4	2.5
Louisiana	1.5	1.6	1.4	1.6
Alaska	1.5	1.6	1.4	1.3
Iowa	1.4	1.7	1.4	0.8
Virginia	1.4	0.7	1.4	0.6
Wisconsin	1.4	1.0	1.3	1.3
Oklahoma	1.4	1.8	1.3	0.5
Rhode Island	1.4	0.9	1.3	1.6
Mississippi	1.4	0.6	1.3	1.1
Maryland	1.3	1.4	1.1	2.5
Kansas	1.3	1.4	1.2	1.0
Illinois	1.3	1.0	1.1	2.1
New Hampshire	1.3	0.6	1.2	1.1
Ohio	1.3	0.7	1.2	0.9
Arkansas	1.3	0.6	1.2	0.7
California	1.3	1.1	1.1	2.6
Maine	1.2	0.4	1.1	1.1
Pennsylvania	1.2	1.1	1.1	1.5
New Mexico	1.2	0.2	1.1	0.9
New Jersey	1.2	0.3	1.0	2.4
Nevada	1.2	-0.4	1.0	2.1
Washington	1.1	1.1	1.0	1.8
Idaho	1.1	0.2	1.0	1.1
Massachusetts	1.1	1.0	0.9	2.0
Missouri	1.1	0.5	1.0	1.2
Delaware	1.1	-0.3	0.9	1.8
Indiana	1.1	0.7	1.0	0.9
Minnesota	1.0	1.0	0.9	1.5
South Carolina	1.0	0.1	0.9	1.8
Connecticut	1.0	0.4	0.8	1.9
Utah	1.0	1.0	0.9	1.1
Michigan	1.0	-0.1	0.9	1.3
Texas	1.0	1.2	0.8	1.8
Florida	1.0	0.1	0.8	2.0
Kentucky	1.0	0.5	0.9	0.5
Oregon	0.9	0.4	0.8	1.7
Alabama	0.9	0.6	0.8	0.9
North Carolina	0.8	0.2	0.6	1.9
Tennessee	0.7	0.5	0.6	1.2
Georgia	0.7	-0.3	0.5	2.2
Colorado	0.5	0.2	0.4	1.8
Arizona	0.4	-0.2	0.3	2.0

Notes: The second column reports the state-specific growth rate in welfare between 2000 and 2013 as measured by expected lifetime utility (see the text for details). Column three and four report the annualized percentage growth in GDP/capita and consumption/capita. The final column reports the total change in life expectancy at birth over this time period (in years).