Late Career Job Loss and the Decision to Retire

Irina Merkurieva*

This version: April 12, 2016†

Abstract

This paper provides an empirical analysis of the effect of involuntary job loss on the lifetime income and labor supply of older workers. I develop and estimate a dynamic programming model of retirement and savings with costly job search and exogenous layoffs. The structural estimates from the Health and Retirement Study data show that older displaced workers lose up to one and a half years of pre-displacement earnings over the remaining lifetime. Most of this loss (80%) is due to the permanent wage penalty following displacement, while the rest is explained by search frictions. Involuntary job loss makes an average worker retire fifteen months earlier. However, workers who were approaching retirement at the onset of the Great Recession will increase their labor supply by approximately five months in response to the joint impact of changes in the value of household assets and the probabilities of losing and finding a job.

Keywords: retirement, life-cycle labor supply, layoff cost, saving, cyclical unemployment
JEL classification: J14, J26, J64

---

*School of Economics and Finance, University of St Andrews. E-mail: i.merkurieva@st-andrews.ac.uk. I am very grateful to my advisors John Kennan, J. Karl Scholz and James R. Walker for their guidance and support. I also thank J. David Brown, Charles M. Becker, Chao Fu, Rasmus Lentz, Rebecca Lessem, Christopher Taber, Giorgio Topa, Wilbert van der Klaauw and seminar participants at the University of Wisconsin and the Federal Reserve Bank of New York for their helpful comments. James R. Walker helped me to get access to the restricted part of the HRS data. My work with restricted data was supported by the Social Sciences Research Services at UW-Madison. The computations in this paper have been carried out with Flash High Performance Computing Cluster maintained by the Social Science Computing Cooperative at UW-Madison. All remaining errors are my own.

†The most recent version is available at http://www.st-andrews.ac.uk/%7Eim58/research/JobLoss.pdf.

1 Introduction

Every year over the last decade, well over a million US workers age 55 and older lost their jobs to layoffs. More than one fifth of the 2010 population between ages 60 and 65 have been displaced at least once since turning fifty. In the meantime, the labor force participation of workers older than 55 went up to 40%, a ten percentage points increase relative to 1990. As the fraction of people who remain employed well into their sixties keeps on growing, those workers increasingly rely on old age earnings. The exposure and vulnerability of an aging workforce to the consequences of job loss have manifested themselves most recently during the Great Recession, when unemployment rates among older persons have reached a historical maximum.[1]

In this paper I examine how job loss experienced by senior workers at different stages of the business cycle affects their labor force attachment, the take up of Social Security, and the lifetime income. I show that involuntary separation at older age results in substantial economic loss for affected individuals. On average job loss costs a sixty year old male worker thirty four thousand dollars, which is the amount required to keep him indifferent between the states of employment and unemployment due to displacement. This is approximately equivalent to an annual wage earned at this age, a large amount taking into account that the remaining work life for the majority of affected workers is relatively short. The main source of this cost is the reduction of post-displacement wages, which accounts for 80% of the total. The remaining 20% are due to search frictions, including the cost of job search and the loss of earnings over the spells of unemployment that follow a job loss.

The effect of involuntary separations on the timing of retirement is not straightforward. Laid off individuals face the cost of searching for a new job that is likely to yield lower earnings. Displacement may encourage some of these workers to leave the labor force sooner than they had planned and to take up Social Security at an earlier age. Others may prefer to work longer to replenish their retirement savings that have been depleted in the course of a post-displacement unemployment spell. These decisions will obviously depend on labor market conditions and vary over the business cycle. The problem is further exacerbated by fluctuations in the value of retirement assets that, in recent years, have coincided with cyclical movements in the labor market.

This paper is the first to evaluate how retirement decisions are affected simultaneously by both labor market and asset shocks. The majority of displaced workers retires earlier than they would had they not lost a job, by fifteen months on average. In the meantime, negative shocks

[1] Data sources: The number of workers displaced between 2001 and 2011 is computed from the Current Population Survey (CPS) Displaced Workers Supplement. The fraction of older workers affected by layoffs comes from the Health and Retirement Study (HRS) dataset. The labor force participation and unemployment rates are obtained from the CPS.

to assets delay retirement from the labor force. I estimate that the cohort of workers who were approaching retirement at the onset of the Great Recession would stay in the labor force approximately five months longer in response to the joint impact of changes in the value of household assets and the probabilities of losing and finding a job. If the same cohort had been affected by labor market shocks alone, their retirement would have happened approximately one and a half months sooner than in a benchmark scenario with the job loss, job finding and asset returns fixed at their long term averages.

The paper also makes an important methodological contribution. Introducing labor market frictions into a life cycle model of labor supply, I show that involuntary job loss is a major retirement channel. Without search frictions, the model overestimates the rates of employment among sixty five year old workers by ten percentage points, an understatement of retirement prevalence at this age of more than 20%. In a frictionless environment, most of this difference is absorbed by the fixed cost of work. The fixed cost of work generates sharp retirement in many models, including French and Jones (2011). Rogerson and Wallenius (2013) argue that its typical estimated values are unrealistically high and lack solid economic interpretation. With search frictions, I increase the share of retirements that can be explained without relying on fixed cost by approximately 11 percentage points.

My results are based on a dynamic programming model of optimal consumption and labor supply decisions with costly job search. The model includes uncertainty about survival, health status, medical expenses, asset returns, wages and availability of jobs. It accounts for Social Security rules, Medicare, employer provided health insurance, government transfers, unemployment insurance, taxes, and intentional bequests. The structural framework allows me to isolate the effect of layoffs, job finding and asset dynamics on the labor supply decisions of older workers. It also helps to define a relevant reference group for studying displaced workers. Varying the probabilities of job loss, job finding and asset returns, I account for the cyclical movements in the labor market along with the dynamics of the housing and stock prices. I estimate the model with the method of simulated moments using data from the Health and Retirement Study.

This paper is related to several strands in the literature. Most important, it builds upon earlier work on retirement from the labor force within the life cycle framework, including Blau and Gilleskie (2006, 2008); French and Jones (2011); Gustman and Steinmeier (1986); Haan and Prowse (2010, 2014); Rust and Phelan (1997); Scholz et al. (2006); van der Klaauw and Wolpin (2008). I extend the results in this literature by incorporating exogenous layoffs, search frictions and time-varying asset returns. None of these previous papers accounts for the possibility of involuntary job loss followed by endogenous costly search. Adding these features, I address an entirely new set of questions about retirement behavior.

Next, I contribute to the literature focused on the economic cost of job loss (Farber, 1993;

Stevens, 1997; Jacobson et al., 1993; Couch and Placzek, 2010; Davis and von Wachter, 2011). This literature does not accommodate endogenous retirement, and produces mixed conclusions on the relationship between age and the cost of displacement. I fill in this gap by providing a detailed account of the cost of involuntary job loss for older workforce. A small number of papers evaluate the impact of layoffs on the employment of older workers (Chan and Stevens, 2001, 2004; Elder, 2004). My estimated probabilities of post-displacement employment rates are consistent with Chan and Stevens (2001), however I evaluate a broader range of displacement consequences, most important the economic cost of job loss over the business cycle and its sources. I also model a number of features that were not addressed in these other papers, yet are essential to the retirement behavior, such as health and unemployment insurance, medical expenses, Social Security take up, and the dynamics of the labor and financial markets. While Chan and Stevens (2004) found that the pension-related incentives can only explain a small part of the observed changes in the timing of retirement associated with job loss, I show that post-displacement wage drop and search frictions can account for a large part of the response.

Finally, there is literature that analyzes the dynamics of retirement over the business cycle. Coile and Levine (2011a,b); Goda et al. (2011); McFall (2011) implemented a reduced form approach to study retirement during recessions. Gustman et al. (2010) estimated the impact of the stock market decline on retirement using a life cycle model; however, they abstracted from the role of the labor market shocks. In contrast to the existing work, my model quantifies indpendent and combined long term effects of wealth changes, layoffs and job finding on the labor market behavior of older workers.

The rest of the paper is divided into seven sections. Section two provides the essential data facts on job loss and retirement. In section three I develop the life cycle model with labor market frictions. Section four describes the dataset used in the structural estimation, explains the choice of the estimation sample and initial conditions. Section five contains the details of the method of simulated moments and its implementation, as well as the estimation methods used to specify exogenous probability transitions and policy environment. Section six discusses the estimates of the structural parameters. Section seven summarizes the results based on the counterfactuals and policy experiments. Section eight concludes.

2 Facts about job loss and retirement

I motivate the paper by highlighting the key facts on displacement and labor force attachment using the data from the Health and Retirement Study (HRS, http://hrsonline.isr.umich.edu/), a nationally representative panel of individuals over age fifty. I measure the prevalence of layoffs among older workers and explore associations between job loss and transitions from work to retirement. The data analysis is based on monthly employment histories of older males that

I construct using biennial labor force flows and employment changes reported between the survey waves.

The definition of retirement used in this paper corresponds to non-participation. I define a person as retired if he neither worked for pay nor looked for work in the last twelve months. Because the HRS does not contain information on job search between the survey waves, I base this definition on both the previous and current labor force status. By taking lagged job search reports, I can identify some of the long term searches that could otherwise be mistakenly taken as labor force exits. Overall I observe 4,798 exits from the labor force made by 4,166 workers after age fifty.

I compute the number of individuals displaced at older age from the data on labor force status and reported reasons of employment termination.[2] Involuntary job loss is represented by two of the separation reasons included in the HRS questionnaires, laid off/let go and business closed. All other separation reasons, including quit, health, family, new job, retirement or financial incentives, are classified as quits. The dataset contains information on 2,091 layoffs that represent 23% of recorded separations. Following these definitions, it is clear that involuntary job losses among older workers are not at all uncommon, with one in six retirees having been laid off at least once since age fifty.

This estimate, however, is a conservative lower bound because of a previously documented misreporting of the separation reasons. For example, Poterba and Summers (1984) find that 25% of the Current Population Study (CPS) quits have been reported as layoffs in the following survey month. Consistently with this pattern, I observe that the HRS quits are often preceded by adverse changes in employment situation. Roughly 12% of quits have been encouraged by co-workers or happened in anticipation of layoff, reduction in wages or hours, or similar circumstances. Once retired, as many as one third of the respondents claimed that they had been at least to some extent forced into their decision. Taking such instances into account, I conclude that a stunning 30% of the first time retirements from the labor force could be preceded by involuntary separations from employment.

The relationship between job loss and retirement can be characterized by post-displacement labor force transitions. Three quarters of the observed older laid off workers with completed unemployment spells do get reemployed, although it takes them almost twice longer than the national average to find a new job. The remaining one quarter retire, either immediately or give up on unsuccessful search. Under the definition of retirement adopted in this paper, the vast majority of layoff induced retirements (80%) happen right after displacement. These workers leave the labor force within the first year following a job loss. The other 20% stay and keep on searching, eventually retiring after an average unemployment spell of 27 months. Noticeably,

[2]The terms layoff, job loss and displacement are used interchangeably throughout the paper. The intended meaning is best described as involuntary separation that is initiated by the firm rather than the worker, as opposed to quit.

the fraction of immediate retirements rises rapidly with the age of job loss, starting at only 3% of completed post-displacement spells at age 50 and exceeding 50% by age 70.

Figure 1 helps understand post-displacement employment patterns. It follows the labor force status of the HRS male workers laid off at age sixty over a period of ten years after displacement. For comparison it also shows the percent of retired individuals by age conditional on being employed at sixty and having no recorded layoff experience. Twelve months after a layoff, around half of the workers displaced at sixty are reemployed at another job, 20% are unemployed and looking for work, and close to 30% are out of the labor force. Notice that at the same time merely 1% of non-displaced sixty year olds transitioned from work to retirement, and it would take them another three years to achieve at least a 20% retirement rate. As displaced workers keep on leaving the labor force, the gap in the retirement rates of the two groups carries on even after ten years. Eventually non-displaced workers spend one and a half years more in employment over the remaining lifetime.

An obvious concern is whether it is appropriate to use never displaced workers as a control group to estimate the cost of job loss. If people at higher risk of displacement have different tastes for work, the gap in the timing of retirement may be affected by selection. To explore this possibility, I check whether the two groups share the same retirement expectations. I infer expected time of retirement from a set of the HRS questions that ask about the year in which the respondent plans to retire or stop working, using the earliest value available for each individual. Comparing these expectations to the subsequently recorded actual age of retirement, I find that all workers regardless of displacement status tend to overstate their projected attachment to the labor force. Retirement on average happens five quarters earlier than expected. Remarkably, workers who would be laid off at sixty initially expect to retire more than a year later than the non-displaced. Eventually we confront a group of people who start with intended retirement age above the average, and yet end up among the first to leave the labor force after involuntary separation. This result indicates that the scope of early retirement among displaced workers may be substantially larger than the estimates based on the comparison to never displaced workers may suggest.

There are further issues that complicate the choice of the reference point. Because the labor market behavior of older workers varies substantially with age, comparison of individuals displaced at different ages may not be relevant. The choice between work and retirement after a layoff depends on labor market conditions. More people tend to stay in the labor force when search is fast and cheap, and human capital depreciation is low. Hence, comparison of workers who have been laid off at different stages of the business cycle may also be inappropriate. In some cases the data samples available for empirical analysis can be restrictive and fairly small. This is especially true for the demographic groups that so far have experienced fewer layoffs. Yet, because of the changes in occupational and industrial structure, knowing the impact of possible job loss for these groups is as important as it is for those who historically have been among the most affected.

These considerations motivate me to estimate the costs and consequences of late career job loss using a structural model of work and retirement with search frictions. I address selection by introducing two types of workers who are different in their exposure to the risk of layoff. In addition, I differentiate workers in the model by year of birth and use time varying probabilities of losing and finding a job to control for the labor market environment at different ages and stages of the business cycle. The structural setting enables me to resolve the problems of control group choice and limited sample sizes by using counterfactuals to simulate the impact of job loss on the labor market outcomes of the elderly.

3 Model of labor supply, search and retirement

I now develop a dynamic life cycle model of labor supply and retirement with search frictions and stochastic returns on assets. The model describes the decision problem of an older worker who faces uncertainty about available employment opportunities, wages, asset returns, health, medial expenses, and survival. In each period of life, individuals make decisions on labor supply, consumption and savings, job search, and the take up of Social Security. The model incorporates essential Social Security rules, taxes, unemployment insurance, government transfers, Medicare and intentional bequests. It accounts for individual heterogeneity in terms of the lifetime earnings, year of birth, the risk of layoff, and access to health insurance. The model provides a salient framework to analyze the impact of job loss and volatility of asset returns on the labor market behavior of older workers. I proceed with discussion of the key elements.

3.1 Preferences

In each period of life an individual aged $t = t_0, t_1 \ldots, t_T$ derives utility from consumption C_t and leisure L_t. Individuals belong to different types τ^i determined by the year of birth, lifetime earnings, the risk of being laid off, and access to employer provided health insurance. The within-period utility function is nonseparable in consumption and leisure as in French and Jones (2011),

$$U(C_t, L_t) = \frac{1}{1-\theta_2} \cdot \left(C_t^{\theta_1} L_t^{1-\theta_1} \right)^{1-\theta_2}. \tag{1}$$

Parameter $\theta_1 \in [0,1]$ is the weight placed on consumption, and $\theta_2 \geq 0$ determines the degree of risk aversion.

Individuals face exogenous mortality risk. An agent at age t survives to age $t+1$ with probability $\pi_t(H_t, \tau^i)$ that depends on age, health H_t and individual type τ^i. Everyone dies with probability one upon reaching the terminal age t_T. Individuals who die leave all remaining assets as bequest to their heirs. The value of bequest in the amount A_t is determined by a bequest function

$$b(A_t) = \frac{b_1}{1-\theta_2} \cdot \Big(b_2 + A_t\Big)^{\theta_1(1-\theta_2)}. \tag{2}$$

This formulation has been derived as a reduced form of an altruistic bequest motive in overlapping generations model by Abel and Warshawsky (1988), and later used to study saving decisions of the elderly (e.g. De Nardi, 2004; De Nardi et al., 2010). The coefficient $b_1 \geq 0$ captures the strength of the bequest motive, while $b_2 \geq 0$ characterizes the curvature of the bequest function and determines the extent to which bequests are a luxury good.

3.2 Time constraints

The quantity of consumed leisure depends on the amount of labor hours N_t supplied out of a fixed endowment $L > 0$, the fixed cost of work, health, and the cost of job search:

$$L_t = L - N_t - \big(\phi_0 + \phi_1 t + \psi \cdot \mathbb{1}\{H_t = 0\}\big) \cdot \mathbb{1}\{N_t > 0\} - c \cdot s_t. \tag{3}$$

The two indicator functions denote respectively bad health status ($H_t = 0$) and labor force participation ($N_t > 0$). The job search decision, $s_t = \{0, 1\}$, assigns the value one to actively searching individuals and zero to everyone else. The time constraints (3) say that the maximum amount of leisure equals to the entire time endowment which is further reduced depending on labor supply, search and health status.

First, the leisure of employed individuals is decreased by the number of hours worked and the fixed cost of work. The fixed cost of work is a linear function of age, $\phi(t) = \phi_0 + \phi_1 t$, with non-negative intercept and slope, $\phi_{0,}\phi_1 \geq 0$. Next, work requires more effort of individuals with serious health problems ($H_t = 0$), and hence those who stay employed in spite of illness lose additional ψ hours of leisure a month. Finally, unemployed individuals looking for work ($s_t = 1$) have to sacrifice c hours a month to search related activities.

The fixed cost of work implies that as people grow older, they have to give up increasing amounts of their leisure in order to stay employed. Using the fixed cost of work is one way to generate sharp retirement from the labor force. In its absence workers would retire by gradually decreasing the hours of labor supply, a pattern inconsistent with the prevalence of direct transitions from full time work to retirement in the data (Cogan, 1981; Rogerson and Wallenius, 2013). Notice however that this model can also generate sharp retirement through the labor market frictions, which distinguishes it from the rest of the literature.

3.3 Budget constraints

Individuals collect taxable income Y_t from wage earnings, receive Social Security, and gain returns on assets:

$$Y_t = N_t \cdot W_t \cdot \left(1 - d_1 \cdot \ell_t\right) \cdot \left(1 - d_2 \cdot \mathbb{1}\left\{N_t < 120\right\}\right) + SSB_t\left(T_t^a, \tau^i, \tau^{ss}\right) + r_{t,\tau^i} \cdot A_t. \tag{4}$$

Wage earnings are determined by the wage rate W_t and the hours of labor supply N_t. The binary indicator of layoff status $\ell_t = \{0, 1\}$ takes a value one if a worker has ever been laid off from a job and zero otherwise. The wages of displaced workers are permanently reduced by a factor $d_1 \in [0, 1]$. This reduction is interpreted as a loss of human capital or obsoleteness of skills that caused a layoff. If a worker chooses part-time employment, his hourly wage rate is decreased by a factor $d_2 \in [0, 1]$. This parameter accounts for an empirical pattern of part-time workers earning substantially less than comparable full-time workers after adjusting for the number of hours worked.

The size of Social Security benefits, $SSB_t(\cdot)$, is calculated based on the age at the time of take up T_t^a, individual type τ^i and parameters of the Social Security system τ^{SS}. Individuals who have applied draw a constant stream of benefits until death. Further particulars of the US Social Security incorporated into the model are discussed in detail in Section 5.3.

The assets A_t of an individual belonging to type τ^i are invested at the rate of return r_{t,τ^i} that varies over time. The rate of return depends on the type of worker and accounts for the structure of portfolios held by investors in different income brackets. For example, individuals with higher lifetime wealth invest more of their capital in stocks, while those with lower wealth invest primarily in housing. Because of such variations in the portfolio structure, any shocks to the housing or stock markets would impact individuals with various levels of wealth differently. No borrowing is allowed, $A_t \geq 0$. The budget constraint is

$$A_{t+1} = A_t + y(Y_t, \tau) + G_t + B_t - M_t(H_t, I_t, \tau^i) - C_t, \tag{5}$$

where function $y(Y_t, \tau)$ gives after tax income for a tax code τ. Non-taxable sources of income are the government transfers $G_t \geq 0$ and unemployment benefit $B_t \geq 0$. $M_t(\cdot)$ is out-of-pocket medical expenses that depend on health, access to health insurance I_t, and individual type. The budget constraints (5) state that the value of assets in the period $t+1$ is determined from the previously accumulated assets A_t augmented by the total value of the current period's after-tax and non-taxable income net of expenditures on consumption and healthcare.

In the data, assets are measured with an error ε_t^a that reconciles the model predictions with observed individual decisions:

$$\widehat{A}_t = A_t \cdot \exp(\varepsilon_t^a) \tag{6}$$
$$\varepsilon_t^a \sim N(0, \sigma^2_{\varepsilon^a}).$$

The government provides a minimum consumption level $\underline{C} > 0$ through transfers G_t according to the rule

$$G_t = \max\{0,\ \underline{C} - [A_t + y(Y_t, \tau) + B_t - M_t(H_t, I_t, \tau^i)]\}, \tag{7}$$

as in Hubbard et al. (1995). After tax income in this formulation reduces the amount of government transfers dollar for dollar. Individuals must totally deplete personal assets before gaining access to the consumption floor provided by the government, and the entire amount of received transfers must be consumed in the current period. Government transfers in the model serve as an aggregate approximation for all public transfer programs other than unemployment insurance and Medicare. Unlike the general government transfers that can be awarded to anyone with income below the threshold, unemployment benefits B_t are limited to actively searching unemployed. In the model, unemployment benefits are assigned conditional on search with probability $\lambda_B \in [0,1]$ to avoid keeping track of the unemployment duration that would be computationally unfeasible.

3.4 Recursive formulation

The vector of state variables S_t for an individual of type τ^i alive at age t includes the wage rate, medical expenses, health state, and assets. The state is also described by the employment and search decisions made in the previous period, the age of Social Security take up, job offer indicator $\omega_t = \{0,1\}$, layoff experience and the draw of unemployment benefits:

$$S_t(\tau^i) = (W_t, M_t, H_t, A_t, N_{t-1}, s_{t-1}, T^a_{t-1}, \omega_t, \ell_t, B_t). \tag{8}$$

Given the current labor force status, an individual makes decisions D_t about the levels of consumption and savings, labor supply, job search and take up of Social Security. The vector of decision variables is $D_t(S_t) = (C_t, N_t, s_t, T^a_t)$. The number of labor supply hours $N_t = 0$ for nonworkers, and $N_t \cdot s_t = 0$, which rules out on the job search.

Timing in the model is as follows. In the beginning of each period an individual can be working, unemployed and looking for work, or out of the labor force. All individuals regardless of their labor force status start the period with receiving two exogenous shocks that determine their health state and medical expenses. Because of the large spread of treatment costs for various conditions with similar overall impact on health, two persons in the same state of health may face very different medical expenses. Even in case of the same diagnosis, the cost of care may be different depending on the provider, for example primary physician compared to an emergency room. Hence the health shock determines the mean of medical expenses through the state of health, while a separate medical expenses shock generates variance that is high enough to match the data.

Employed individuals receive another type-specific shock that destroys their jobs with probability $\delta_{t,\tau^i} \in (0,1)$. Workers who did not lose jobs then obtain updated information on their wages. In the meantime, non-employed individuals get wage offers with probability

$$\lambda_{t,\tau^i}(s_t) = \begin{cases} \lambda^n_{t,\tau^i} \in (0,1) & \text{if } s_t = 0 \\ \lambda^u_{t,\tau^i} \in (0,1) & \text{if } s_t = 1 \end{cases} \tag{9}$$

where $\lambda^u_{t,\tau^i} > \lambda^n_{t,\tau^i}$ are respectively job finding probabilities for unemployed and non-participating. While the model in principle allows for getting an offer without search, the chances are much smaller than when an effort is put into searching for employment opportunities. In addition, only the unemployed who have been searching for jobs in the previous period may receive unemployment benefits with probability λ_B. Once the wages are revealed, workers who retained their jobs or received new offers make labor supply decisions.

Workers who just lost their jobs or quit, unemployed who did not receive acceptable offers, and workers who have been out of the labor force decide whether they want to invest time into looking for work. If they decide to search, they start the next period as active unemployed, have a higher chance to receive a job offer and may retain their unemployment benefits. Otherwise they enter the next period out of the labor force.

Next, all individuals make decisions on Social Security application, collect income, pay medical expenses and taxes, enjoy leisure, consume and save. Finally, they receive survival shock and either move to the next period or die and leave bequests.

In each period and state S_t, an individual chooses a decision rule D_t to maximize expected discounted lifetime utility subject to the exogenous processes for mortality, employment, health, medical expenses and wage determination, a set of time and budget constraints (3) and (5), government transfer rule (7), and policies for taxes, unemployment benefits and Social Security. The discount rate for future payoffs is β. The values of states with and without a wage offer, V^w_t and V^n_t, are defined recursively. The value of a state with a wage offer contains four additive terms: the current utility from consumption and leisure and the expected present values of unemployment in case of layoff, employment if not affected by job loss, and utility of bequest if the worker does not survive to the next period:

$$
\begin{aligned}
V_t^w(S_t, D_{t-1}, \theta) = \max_{D_t} \Bigg\{ & U_t(S_t, D_t, \theta) + \\
& \beta \cdot \Bigg[(1-\pi_t) \cdot \Bigg(\delta_{t,\tau^i} \int V_{t+1}^n\left(S_{t+1}, D_t, \theta\right) \cdot p(dS_{t+1}|S_t, D_t, \theta_p) + \\
& \left(1-\delta_{t,\tau^i}\right) \int V_{t+1}^w\left(S_{t+1}, D_t, \theta\right) \cdot p\left(dS_{t+1}|S_t, D_t, \theta_p\right) \Bigg) + \pi_t b(A_{t+1}) \Bigg] \Bigg\},
\end{aligned} \tag{10}
$$

where $\theta = \{\theta_s, \theta_p\}$ is a vector of model parameters that includes parameters of the state transition probability function θ_p and structural parameters θ_s. Similarly, the value of a state without a wage offer, V_t^n, is

$$
\begin{aligned}
V_t^n(S_t, D_{t-1}, \theta) = \max_{D_t} \Bigg\{ & U_t(S_t, D_t, \theta) + \\
& \beta \cdot \Bigg[(1-\pi_t) \cdot \Bigg(\lambda_{t,\tau^i}(s_t) \int V_{t+1}^w(S_{t+1}, D_t, \theta) \cdot p(dS_{t+1}|S_t, D_t, \theta_p) + \\
& (1-\lambda_{t,\tau^i}(s_t)) \int V_{t+1}^n(S_{t+1}, D_t, \theta) \cdot p(dS_{t+1}|S_t, D_t, \theta_p) \Bigg) + \pi_t b(A_{t+1}) \Bigg] \Bigg\}
\end{aligned} \tag{11}
$$

Because the model does not have a closed form solution, the decision rules it generates must be found numerically. I use backward induction to solve the value functions at monthly time intervals. In estimation, the terminal age is set to 100, the maximum working age to 75, the starting ages to the age of each type in 2000, and part-time work is under 120 hours per month.[3] I estimate the structural parameters of the model using the method of simulated moments. The details of the estimation procedure are discussed in Section 5.

4 Data

I estimate the model using the Health and Retirement Study (HRS), HRS RAND dataset (version M), and the restricted HRS data on earnings records from the Social Security Administration (SSA). I extract the variables from ten waves of the HRS that cover a period from 1992 to 2010 at biennial intervals. In this section I describe the sample and variables and explain how I use the data to draw initial conditions for structural estimation.

[3]Setting the maximum working age reduces dimensionality of the problem. Only 10% of the HRS estimation sample is employed after 75, and less than 3% holds full-time jobs.

4.1 Sample construction and variables

The estimation sample consists of white non-Hispanic males age fifty and older. I further restrict the sample to moderate the diversity of paths that respondents may take within the system of Social Security benefits. First, I drop individuals who started receiving Social Security before turning sixty two as well as the beneficiaries of the Supplemental Security Income and Social Security Disability Insurance (SSI/SSDI). These conditions aim to exclude workers with disabilities who may have quite different health transition processes, access to health insurance and non-labor income. Second, I eliminate anyone with total employment record under ten years or over ten years of work in the government. These respondents would either not qualify for Social Security, which requires ten years of contributions, or have access to the government pension schemes that are not modeled in the paper. The sample resulting from these restrictions contains information on 8,531 individuals, or 48,686 person-year observations.

To estimate the processes that govern the state transition probabilities, I use information on labor supply, wages and average lifetime earnings, assets, mortality, health, medical expenses, and health insurance. Except for the average lifetime earnings, all variables are taken from the HRS RAND file. The monetary values are converted into constant 2000 dollars. I use the SSA data to compute the average lifetime earnings. In addition, I construct a proxy of the SSA's average indexed monthly earnings (AIME) by indexing wages to year 2000 and taking an average of the top 35 values, replaced by zeroes if not enough years were reported. Indexing is done with the national wage index. I impute the missing values of medical expenses, average lifetime earnings and wages using a regression based procedure (David et al., 1986).[4]

Labor supply is computed as a product of the usual hours of work in a week and the average number of weeks in a month. Hourly wage rates are obtained from the data on earnings, periodicity of payments, and hours of labor supply. Search decisions are constructed from information on the current labor force status and questions about job search activity. Layoff incidents and the calendar months in which they happen are extracted from self-reported employment histories and separation reasons.

Assets are measured as the net value of financial and housing wealth of the household. They account for the value of housing, vehicles, businesses, ownerships in IRA, financial instruments and investment funds, and other savings and debt. An important concern is that this variable does not reflect the value of pensions, possibly underestimating assets for almost 80% of the respondents who receive or are eligible for at least one pension. In addition, omission of pension wealth may understate the impact of stock price volatility on the assets of individuals with defined contributions (DC) plans who comprise about a third of the sample. Although the

[4]The fractions of the imputed values of medical expenses and lifetime earnings that are used in the computations is 2.8% and 2.5%, respectively. Missing earnings data arise because only about three quarters of the respondents granted access to their administrative records.

number of pension holders is high, I expect that leaving out pensions data will have limited effect overall since the value of DC holdings is insubstantial in comparison with Social Security and housing wealth of the households. For example, Gustman et al. (2010) estimate that the average value of DC pensions for the HRS cohort born between 1948 and 1953 makes up only 8% of their household wealth. I expect that this fraction will be even smaller in my sample containing individuals born in earlier years.

Health status is a binary variable based on the question that asks respondents to evaluate their own health on a five-point scale as excellent, very good, good, fair, or poor. I define health status as good for the top three categories, and bad otherwise. Medical expenses are individual out-of-pocket payments for health care. These expenses include the cost of stays in the hospitals and nursing homes, visits to doctors and dentists, outpatient surgeries, prescription drugs, home health care, and special facilities paid by the respondent in a two-year reference period. They do not include expenses covered by health insurance, but do include insurance premiums paid by the respondent.

Access to health insurance is described by a categorical variable that takes three values: health insurance provided by the current or previous employer to respondents under 65, respondents older than 65 covered by Medicare, and respondents younger than 65 not insured by employer. The last category involves individuals without any health insurance (59%) as well as insured by privately purchased plans (41%). Both groups pay higher medical expenses on average, either as insurance premiums or as direct cost of health services. I exclude from the estimation of state transition probabilities for medical expenses individuals whose only insurance comes from their spouse's employer or government plans other than Medicare, such as Medicaid, CHAMPVA and TRICARE. These respondents together account for 7% of the estimation sample, with the largest share (5.9%) made up by insured through employers of the spouses.

4.2 Individual heterogeneity and the initial distribution of state variables

To estimate the structural parameters of the model, I only use six waves of the HRS data covering the period from 2000 to 2010. Earlier waves are excluded because of the changes in Retirement Earnings Test that took place in 2000 and substantially changed the policy environment for employment decisions of older workers. I draw initial conditions for the joint distribution of state variables from a sample of individuals born between 1938 and 1943. The model starts in 2000 when these respondents were between 57 and 62 years old. After accounting for missing data, these restrictions leave 422 persons with complete information available in 2000 who are eligible for initial state draws. A larger sample of 1,119 individuals with 5,302 available person-year observations is used to compute the moments matched in the estimation.

The model distinguishes between thirty six individual types. Each type is defined by the year of birth, position in the distribution of earnings, and the risk of being laid off. The six birth year types are essential in establishing the age of laid off workers at different stages of the business cycle. They form the basis of inference on the impact of cyclical fluctuations in the asset values, job loss and job finding probabilities on the labor supply of older workers as they approach retirement. Three earnings types are defined by the AIME tertiles. Two layoff risk types, low and high, are identified based on the HRS question that asks respondents to evaluate the likelihood of losing a job within a year. I predict the likelihood of layoff from a regression of self-reported probabilities on the individual characteristics. An individual is assigned to a high risk type if predicted probability exceeds 0.5, and to a low risk type otherwise. It has been shown in the literature that job loss expectations have a high predictive power for actual job loss (Stephens, 2004). Accordingly, my risk indicator appears to be a valid predictor of the actual layoff experiences. The monthly probability of a layoff in the estimation sample is 0.001 for low and 0.003 for the high risk type, a statistically significant difference. The fractions of workers who have experienced job loss in the two groups are 0.132 and 0.293, respectively.

The initial joint distribution of ages, assets, wage rates, AIME, health, medical expenses, health insurance, labor supply and layoff risk for 10,000 simulated workers is drawn from the 2000 HRS dataset using individual sampling weights. I assume that the initial state summarizes all relevant information from earlier ages, and that it does not reflect any anticipation of future changes to the policy environment, such as reforms of Social Security or tax code. Drawing initial conditions from the data accounts for some important empirical relationships between the variables of the model, as reflected in Table 1. The table summarizes descriptive statistics for the initial distribution. For example, the risk of layoff and the rates of unemployment are highest for the low income type, a correlation essential for the conclusions of the paper. Individuals with higher lifetime income understandably hold more assets and earn higher wages. The initial state also reflects a complicated connection between financial wellbeing and health. Workers of the highest income type are healthier and pay rather high out-of-pocket medical expenditures while being more likely to hold jobs with employer provided health insurance. Heterogeneity in layoff exposure accounts for an important fact that low skilled workers are more likely to lose jobs.

5 Estimation methods

I estimate the model parameters in two stages similar to Gourinchas and Parker (2002); De Nardi et al. (2010); French and Jones (2011). First, I obtain the values of parameters determining exogenous probabilities of transition between the points of state space that can be identified without using the entire model, θ_p. This estimation stage yields the transition processes for health, survival, medical expenses, wages, job loss and job finding, and asset returns. At this stage I also set up the policy rules for taxes, Social Security and unemployment insurance.

Second, I use the method of simulated moments (McFadden, 1989; Pakes and Pollard, 1989; Gouriéroux and Monfort, 1996) to estimate the structural parameters of the model, while taking the first stage estimates as given. The vector of $m = 14$ parameters obtained at the second stage, $\theta_s \in \Theta$, includes the coefficients of utility and bequest functions, the fixed costs of work, health and search cost, wage losses due to human capital depreciation and part time employment, government transfers, leisure endowment, the probability of getting unemployment benefits, and the variance of measurement error in assets:

$$\theta_s = \{\theta_1, \theta_2, b_1, b_2, L, \phi_0, \phi_1, \psi, c, d_1, d_2, \underline{C}, \lambda_B, \sigma^2_{\varepsilon^a}\}. \tag{12}$$

The parameter space $\Theta \subset \mathbb{R}^m$ is restricted to account for the lower and upper feasibility boundaries of the parameter values, as dictated by the model. Details of the methods applied at each stage of estimation are discussed in the rest of this section.

5.1 State transition probabilities

Individuals in the model face uncertainty about future survival, health, medical expenses, employment and income. I assume that beliefs about the likelihood of future events are rational, described by the Markov probability function $\Pr(S_{t+1}|S_t, D_t, \theta_p)$. In addition, I assume conditional independence that enables me to represent the state transition probabilities as products of the marginal conditional probabilities for individual state variables. Based on these assumptions, each individual component of the state transition probability function is estimated as follows.

Health and mortality

I estimate biennial survival and health transition probabilities conditional on age, lagged health, and average lifetime income from a set of logit models. Table 2 contains estimated marginal effects. Based on these estimates, I predict conditional probabilities of health and survival transitions that are further supplied as an input in the simulations.

Health insurance and medical expenses

I assume that the health insurance status of employed individuals remains unaffected when they change jobs. A worker insured by employer automatically retains access to health insurance at a new job, while uninsured worker can not improve his status by switching jobs. This assumption holds for a large fraction of respondents in the data: over three-fourths of workers who switched jobs between the HRS survey waves retained their health insurance status in a later wave. Loss of job leads to immediate termination of coverage for insured workers.[5] The consumption floor provided by the government accounts for the role of Medicaid. Employer provided health insurance becomes irrelevant after age sixty five, when all workers in the model become eligible for Medicare.

Out-of-pocket medical expenses M_t follow an error components process with autoregressive error term:

$$\begin{aligned} \log M_t &= M(t) + \zeta_t^m \\ \zeta_t^m &= \rho^m \zeta_{t-1}^m + \varepsilon_t^m \\ \varepsilon_t^m &\sim N(0, \sigma_{\varepsilon^m}^2), \end{aligned} \tag{13}$$

where ζ_t^m is persistent AR(1) component of medical expenses with autocorrelation ρ^m, and ε_t^m is white noise. The mean of log medical expenses $M(t)$ depends on age and health status. I estimate parameters of Equation (13) using conditional maximum likelihood for each category of health insurance: employer provided, Medicare, and not insured by employer. To implement estimation in logs, I drop reports of zero medical expenses.[6] The estimation results are shown in Table 3. When simulating the model, I approximate AR(1) process for the medical expense shock ζ^m with a first order three-node discrete Markov chain using Rouwenhorst method (Rouwenhorst, 1995; Kopecky and Suen, 2010).

Wage transition probabilities

The estimation of the wage process is similar to medical expenses, except that I first impute the unobserved wages of nonworking individuals. I then model the wage transitions as an error components process with AR(1) disturbances:

$$\begin{aligned} \log W_t &= W(t) + \zeta_t^w \\ \zeta_t^w &= \rho^w \zeta_{t-1}^w + \varepsilon_t^w, \\ \varepsilon_t^w &\sim N(0, \sigma_{\varepsilon^w}^2). \end{aligned} \tag{14}$$

The mean wage $W(t)$ depends on age quadratic, AIME and health. The estimates are reported in Table 4. In simulations the autoregressive component is discretized into five nodes discrete

[5]Post-displacement jobs, on average, are less likely to offer health insurance than previously held career jobs, so the model provides slightly better employment prospects for workers of insured type. I do not model COBRA or employer provided health insurance that can be retained after separation. Because individuals covered by COBRA typically pay higher premiums than insured by employer, the process estimated for the group without employer provided health insurance provides a reasonable representation of their medical expenses.

[6]Zero medical expenses are about 13% of the sample. There are no observable differences between respondents with zero expenses and the rest of the sample. I treat these observations as a result of measurement error, since very few people would literally spend zero on health and medications over a two year period.

Markov chain.

Job finding and layoff probabilities

I compute job finding probabilities from the gross flows of workers between three labor force states: employment, unemployment and out of the labor force. I obtain the sample-weighted monthly gross flows from the rotating part of the CPS microdata on males ages 50 and above.[7] The series are seasonally adjusted using the Census Bureau's X-13 procedure. The monthly transition probabilities are computed from the gross flows and corrected for time aggregation bias using Shimer's method (Shimer, 2012). An adjusted probability that a worker unemployed in month t will become employed within one month is then used as job finding probability for actively searching individuals. The probabilities of transition from non-participation to employment are used as job finding probabilities for those who do not search.

Layoff probabilities are computed from the HRS retrospective job histories. The CPS does not collect data on the separation reasons for workers who have left the labor force, making it unusable for estimation of involuntary separations among older workers who are highly likely to respond to a layoff by leaving the labor force. I adjust the monthly layoff probabilities by a factor of 1.3 for workers at high risk of layoff and by a factor of 0.5 for workers at low risk. Adjustments are based on the differences of job loss probabilities by layoff risk types in the data, so that the weighted average for the two types is consistent with the overall probability of involuntary job loss.

I use twelve month centered moving averages as job finding and layoff probabilities in the model for the months with available data, 2000-2014. Beyond this period, I converge the values to their sample averages over a period of three years. The mean values of the estimated series are $\lambda = 0.24$ and $\delta = 0.003$. Both probabilities for older workers are lower than their counterparts for all ages. This is consistent with the studies documenting that mobility between the labor force states declines with age (e.g. Menzio et al., 2015). With rational expectations, workers have full information on these probabilities but do not know exactly whether and when layoffs and job finding are going to affect their employment path.

5.2 Asset returns and discounting

Workers invest their wealth in housing, stocks and an additional composite asset. The share of wealth invested in each asset depends on worker's income type. Workers of higher income types tend to hold more wealth in stocks, while housing is the main asset for those with low

[7] I use the CPS instead of the HRS here because the former provides monthly information on labor force status and has larger sample size. Retrospective histories from the HRS seem to underestimate transitions between the labor force states and are very volatile. I use the data from January 2000 to October 2012. April and May 2004 data cannot be merged because of a change in the data structure, so I use an average of the adjacent months to get the numbers for May 2004.

lifetime income. I obtain the net values of housing, direct stock holdings, IRA and total assets from the HRS RAND data on the composition of assets. Stock holdings reflect both direct stock ownership and the share of IRA invested in stocks, and I take the share of stocks in the IRA from Gustman et al. (2010). Using these data, I compute the share of each asset in a typical portfolio held by workers of different income types.

The return on a portfolio held by a worker of given income type is a weighted average of returns on individual assets. The rate of return on stocks is a twelve months moving average of total real returns on large company stock taken from Ibbotson Associates (2013). Housing appreciation is computed from the FHFA all-transaction indexes. I take a twelve-month moving average of these series and adjust for inflation using the CPS "All items less shelter" index. Annualized rate of return on other assets is 4%, and the portfolio returns converge to the 4% rate over three years after the end of the observed series.

Workers expect that the annual rate of return on their assets is the long-term average of 4%, which is consistent with estimates by McGrattan and Prescott (2003). Yet the realized return values are computed so as to account for the housing bubble and dynamics of the stock market that are assumed to be entirely unanticipated by the workers. The annual discount rate for future payoffs $\beta = 0.96$.

5.3 Government

Social Security

The model incorporates three major features of the Social Security policy. First, awarded benefit is determined from the primary insurance amount (PIA), a transformation of the AIME over two bend points defined by the SSA. Following the SSA's approach, the PIA is calculated as a sum of 90% of the AIME under the first point, 32% of the AIME between the two points, and 15% of the AIME in excess of the second point. I use the SSA bend points that are set for each year of birth.

Second, the benefits are subject to the lower and upper limit. The minimum benefit is guaranteed to low wage earners who have contributed to the system for eleven years or longer. I set the minimal PIA to the SAA's year specific minimum for a worker with twenty years of contriburions. The maximum benefit received by a family is computed with three threshold values. The formula is 150% of the PIA under the first threshold, plus 272% of the PIA between thresholds one and two, plus 134% of the PIA between thresholds two and three, and plus 75% of the PIA above threshold three. I incorporate this rule into the model with the threshold values assigned to each year of birth.

Third, there is a penalty for Social Security take up before the normal retirement age. The earliest Social Security eligibility age is sixty two, and the normal retirement age of the simulated cohorts varies between sixty five and sixty six. The PIA determines the amount awarded to a worker who takes up the benefits at normal retirement age. This amount is reduced (or increased) proportionately to the time left until (or elapsed since) reaching the normal retirement age. Ignoring credit for delayed retirement, I only reduce the PIA by 8% for each year before the normal retirement age. Since the model does not offer incentives for delayed retirement, I assume that everybody applies for Social Security upon reaching the normal retirement age.

Taxes

Individuals pay federal, state and payroll taxes. I take the federal tax rates from the annual tax rate schedules, head of household tables.[8] The state tax rates are taken from 2001 Rhode Island tax rates schedule.[9] Earnings are subject to Social Security and Medicare payroll tax at rate of 7.65% up to a year specific maximum. Earnings above this maximum are taxed at rate 1.45% that only cover Medicare's Hospital Insurance.[10]

In addition, the Social Security benefits of early retirees may be taxed in accordance with Social Security Earnings Test (SSET) when wages exceed retirement earnings exempt amounts. I use the lower exempt amount that applies in years before reaching the normal retirement age, and ignore the kink in the last few months before reaching the normal retirement age. Because the sample period only covers the years after 2000, SSET does not apply after the normal retirement age. The benefits are taxed at a rate $1 for every $2 of earnings in excess of the lower exempt amount. I also ignore the refund of withheld benefits after the normal retirement age.

Unemployment insurance

The amount of unemployment benefits is set to replace 50% of pre-displacement earnings, with a maximum of 1,300 dollars per month. The maximum value roughly corresponds to the state average published by the US Department of Labor in 2001 "Comparison of State Unemployment Laws". According to the same source, the average duration of a person collecting unemployment insurance was 3.8 months in the period from January 2000 to December 2014. The estimated probability of unemployment benefits reward targets this average duration.

[8]Source: http://www.irs.gov/pub/irs-prior/i1040–2005.pdf

[9]Ideally, I would like to use state tax rates for the respondent's state of residence. Unfortunately, I do not have access to the HRS state identifiers, so I have to impose a uniform state tax. I chose Rhode Island because it was one of the last states to abandon piggyback taxes in 2001. It gives me a straightforward way to estimate tax liability based on federal taxes. French and Jones (2011) claim it to be a fairly representative state in terms of income tax rates.

[10]Source: http://www.ssa.gov/oact/progdata/taxRates.html

5.4 Estimation of structural parameters

The structural parameters of the model are estimated with the method of simulated moments (MSM) which minimizes the distance between the sample moments and their equivalents in the simulated dataset. I exploit five groups of moments for each age between 60 and 70. These are the rates of employment conditional on health, the fraction of part time workers among employed, the means of monthly labor supply, the rates of job search among non-participating, and the quartiles of the asset distribution. In addition, I match the variance of assets and the mean length of unemployment spell, in total yielding $l = 90$ moment conditions.

Most of the data moments, $\hat{\mathbf{M}}_{l\times 1}$, are computed from the HRS estimation sample. The two exceptions are the length of unemployment spell and the search rates. The former is determined from the US Department of Labor unemployment statistics, and the latter is obtained from the CPS sample of older workers. The HRS search moments are very volatile and therefore hard to match because of the smaller sample size.

Simulated moments, $\mathbf{M}_{l\times 1}(\theta, \hat{\theta}_p)$, are computed from a dataset that is generated by the model. Setting up the simulation, I first randomly draw the joint values of variables that form initial conditions from the estimation dataset. Each data point is selected into initial state with a probability inversely proportional to the individual HRS weight of an observation. Second, I generate a matrix of random shocks that determine the realization of exogenous processes. These two inputs remain fixed across the simulations. The value functions (10) and (11) can then be solved numerically by backward induction, yielding decision rules for any feasible choice of structural model parameters $\theta_s \in \Theta$. I discretize the state space and use linear interpolation to evaluate the value functions between the grid points. Decisions that are ruled out by the model constraints are assigned large negative values. Finally, I put together initial conditions, shocks and decision rules to simulate a dataset that contains a life path for each of the $n = 10,000$ modeled individuals.

For the vector of sample moment conditions $\bar{\mathbf{m}}_n(\hat{\mathbf{M}}, \theta_s, \hat{\theta}_p) = \hat{\mathbf{M}} - \mathbf{M}(\theta_s, \hat{\theta}_p)$, the estimate $\hat{\theta}_s \in \Theta$ minimizes the MSM criterion,

$$\hat{\theta}_s = \underset{\theta_s \in \Theta}{\operatorname{argmin}} \; \bar{\mathbf{m}}_n(\hat{\mathbf{M}}, \theta_s, \hat{\theta}_p)' \mathbf{W}_n \bar{\mathbf{m}}_n(\hat{\mathbf{M}}, \theta_s, \hat{\theta}_p), \tag{15}$$

where $\mathbf{W}_n$ is a symmetric positive definite weight matrix of size l. The optimal weight matrix that returns asymptotically efficient estimates is the inverse of the variance-covariance matrix of the population moments. Its method of moments estimator is the inverse of the variance-covariance matrix of the sample moments, which however may generate bias in small samples (Pischke, 1995). To avoid this problem, I use a diagonal weight matrix with non-zero elements given by the reciprocals of the sample moment variances, as suggested by Altonji and Segal (1996). This choice of weights implies that the moments estimated more precisely receive higher weight in the MSM criterion, while dependence among the moments is ruled out. In this way, I solve Eq. (15) numerically, using a simplex algorithm to search over parameter space on a 64-node cluster.

Pakes and Pollard (1989) show that under certain regularity conditions optimization estimator $\hat{\theta}_s$ is consistent and asymptotically normally distributed with variance-covariance matrix given by

$$\mathbf{V}_{\theta_s} = (\mathbf{J}'\mathbf{J})^{-1}\mathbf{J}'\Omega\mathbf{J}(\mathbf{J}'\mathbf{J})^{-1}, \tag{16}$$

where $\mathbf{J} = \mathbb{E}\frac{\partial}{\partial \theta_s'}\mathbf{m}_i(\theta_s)$ is Jacobian matrix of the moment conditions and $\Omega = \mathbb{E}(\mathbf{m}_i\mathbf{m}_i')$ is their variance-covariance matrix. I use method of moments estimators of these matrices to compute the standard errors. The core estimation results are presented in the next section.

6 Baseline estimates of the structural parameters

In this section I discuss the estimates and explain how each structural parameter is identified from the data. Furthermore, I show how the model fits the data and highlight the role of search frictions, a new modeling component that is central to the main contribution of this paper.

6.1 Estimates and identification

Table 5 contains the MSM estimates of the model parameters, standard errors and the value of chi-square statistic for over-identification test. I discuss below each estimate relative to the benchmarks available in the literature, and explain the identification strategy.

Individual preferences are characterized by the parameters of the CRRA utility and bequest functions. The consumption weight, $\hat{\theta}_1 = 0.58$, reflects the relative tastes for leisure and consumption, with higher values implying stronger preference for work. It is identified primarily from the hours of labor supply. The estimate is within the interval obtained by French and Jones (2011) for different values of willingness to work index, $\theta_1 \in [0.412, 0.967]$. The coefficient of relative risk aversion for a consumption-leisure bundle, $\hat{\theta}_2 = 4.8$, is consistent with the life-cycle retirement literature, e.g. French and Jones (2011) and French (2005). Under assumption of fixed labor supply, this estimate translates into the coefficient of relative risk aversion for consumption $R(C) = 1 - \hat{\theta}_1(1 - \hat{\theta}_2) = 3.2$, a conventional measure that is easy to compare to the broad labor supply literature. The coefficient of relative risk aversion is identified by asset quartiles that reflect the extent of precautionary savings accumulated to insure against adverse income and health shocks.

Parameters of the bequest function are identified from the upper quartiles of the asset distribution. In the data wealthier individuals keep on saving instead of divesting even after they have accumulated sufficient assets to insure against the risks incorporated into the model environment. The presence of an altruistic bequest motive generates this behavior in the model. In the terminal period $t = t_T$ the bequest motive becomes operational once the level of assets reaches the threshold $A_b = \hat{b}_2 \left(\frac{\hat{b}_1}{\hat{L}^{(1-\hat{\theta}_1)(1-\hat{\theta}_2)}} \right)^{\frac{1}{\hat{\theta}_1(1-\hat{\theta}_2)-1}}$, which after adjusting for the scaling of the coefficients gives a value of twenty three thousand dollars. This is the minimum amount of saving an agent needs in order to leave non-zero bequest to a heir instead of consuming the entire amount himself, low enough to suggest that at the terminal age bequest motive is operational for the majority of workers. The marginal propensity to consume out of the final period wealth is one for individuals whose assets are below this threshold, and $MPC = \frac{dC}{dA} = 0.02$ for those who hold more. Analogous relationships can be derived for other ages, hence maintained level of saving in the top of wealth distribution and the rate of saving among those whose assets are above the operational bequest threshold jointly identify the intensity of the bequest motive $\hat{b}_1 = 0.985$ and bequest shifter $\hat{b}_2 = 1.084$.

The monthly time endowment, L, is 375 hours, identified jointly by the employment and labor supply moments. The fixed cost of work for a sixty year old worker is 69 hours per month. The estimated slope of the linear fixed cost function suggests that over a year the monthly fixed cost of work increases by approximately 3 hours. Notice that while my model yields approximately the same monthly fixed cost as French and Jones (2011) for a sixty year old worker, estimated trend is almost 20 times smaller.[11] This is a major difference due to the introduction of labor market frictions in my model, and I discuss it further in Section 6.3. The fixed cost parameters are identified by declining trends in the participation rates and labor supply. The additional cost of work caused by ill health is 0.1 hours per month. This parameter is identified by the rates of labor force participation conditional on health status. The estimated cost of search parameter is 85 hours per month, implying that looking for work takes about the same time as half-time employment. Identification of the cost of search is based on the search intensity moments.

Together these estimated parameters of the utility function and the time constraint capture the way in which workers would intertemporally substitute their labor supply. For example, Frisch intertemporal elasticity of substitution for a healthy sixty year old worker employed $N_{60} = 160$ hours a month would be approximated by

$$\eta = \frac{\hat{L} - N_{60} - \hat{\phi}(60) - \hat{\psi}}{N_{60}} \cdot \frac{1 - \hat{\theta}_1(1 - \hat{\theta}_2)}{\hat{\theta}_2} = 0.594,$$

assuming that the borrowing constraint is not binding. This value falls just above the middle of the interval between 0.37 and 0.70 which according to Chetty et al. (2011) defines the range of the micro estimates for Frisch intensive margin elasticity.

[11]The annual fixed cost of work at age sixty in French and Jones (2011) is $\phi_{P0} = 826$, and the age trend is $\phi_{P1} = 54.7$ hours

Turning to the parameters of the budget constraint, the most important is the wage loss associated with displacement. It is estimated as 16% of the wage received by otherwise similar non-displaced worker. The cost of search and wage loss are identified by the share of non-working individuals who look for work at each age, and by the variation of job finding probabilities across time. The estimate is close to several reference points available in the displacement literature. For example, Couch and Placzek (2010) estimate the mid-term wage losses of displaced workers six years after separation in the range from 13 to 15%. Similarly, Davis and von Wachter (2011) find that displaced workers in the US lose 10 to 20% of their earnings in the long term depending on whether separation occurred during an expansion or a contraction.

The premium for full time employment, d_2, is 19% of the wage rate. This parameter is identified by the rate of part time employment, and the value is consistent with 25% part time employment penalty found by Aaronson and French (2004). The monthly government transfer $\underline{C} = \$375$ is identified by lower asset quantiles because guaranteed consumption minimum discourages saving among the poor. In annual equivalent, it is on the same scale with the consumption floor in Hubbard et al. (1995).

6.2 Model fit

Figure 2 shows the model fit for matched moments. I further provide some results to check whether the model captures data facts that were not directly used in the estimation. The model predicts well the number of workers affected by layoffs over the lifetime and the difference in the timing of labor force exit for displaced and non-displaced workers. Consistently with the data, the share of workers who retire after a layoff grows steadily until peaking at age 65. The model also reproduces the matrix of transitions between the labor force states, with an exception of exit rates from unemployment that are somewhat lower than in the data.

6.3 The role of search frictions

The life-cycle models of retirement from the labor force typically ignore search frictions, allowing no job destruction and costless access to positive draws from a wage distribution. Usually under these assumptions the transition from work to retirement is driven by changes in wage and non-wage income, health, government policies, or fixed cost of work. I show that search frictions represent an additional incentive to exit the labor force, accounting for up to 20% of retirements. Hence, omission of frictions results in nontrivial bias. To show how search affects the model predictions and estimates, I revert to a standard life-cycle retirement model by assuming that jobs are always available ($\lambda_{t,\tau} = 1$), never destroyed ($\delta_{t,\tau} = 0$), search is free ($c = 0$), and there is no unemployment insurance ($\lambda_B = 0$). I compare the original model and the frictionless one in two experiments.

In the first experiment, I simulate a dataset assuming frictionless labor markets and using parameter estimates of the original model. The employment rates generated by this experiment for different ages are five to twenty percentage points lower than in the data. For example, the labor force participation of 65 year old workers is underestimated by ten percentage points, which is over 20% of retirement prevalence at this age. These differences are large and due directly to the labor market frictions, indicating that without search a model of retirement would overload other channels in order to match the observed retirement dynamics.

In the second experiment, I estimate a frictionless model and identify parameters that turn out to be the most sensitive to the omission of search. I find that the main change occurs to the fixed cost function, increasing the fixed cost of work for a sixty year old worker by 80%. This result appeals to earlier discussions of the role of fixed costs. Rogerson and Wallenius (2013) estimate that the time fixed cost required to generate sharp retirement from the labor force for acceptable values of the intertemporal elasticity of substitution of labor supply exceeds the time required to hold a half-time job. This is implausibly high, and my experiment suggests that one of the reasons is that the fixed costs absorb the role of other retirement incentives that were omitted from the model. By introducing labor market frictions, I explain a portion of this excessive fixed cost, which is an important methodological contribution.

7 Discussion and counterfactuals

I now use the model to analyze the impact of involuntary job loss on the labor market outcomes of older workers. Taking full advantage of the structural approach, I construct counterfactual life paths of twin workers who are identical in all initial characteristics and received shocks, except that one of the twins is displaced at some point in his life while the other is not. Comparison of the outcomes of these twin workers over the life cycle yields a flexible tool helpful to evaluate a wide range of job loss consequences. I start with the analysis of labor force attachment following a layoff. Next, I estimate the cost of displacement and identify its main sources. Because the incidence and cost of displacement are related to the probabilities of losing and finding a job, I consider in a separate section how older workers respond to involuntary job loss at different stages of the business cycle. Finally, I show how the labor market conditions interact with unanticipated asset shocks affecting the retirement of displaced workers. To illustrate this point, I estimate the effect of the Great Recession on retirement decisions of workers who have reached their late fifties by the beginning of the economic downturn. In conclusion, I discuss several possible policies aimed at reducing the cost of job loss.

7.1 Labor force attachment after involuntary job loss

First, I show how displacement affects the labor force attachment and the timing of retirement. I take a sample of older workers and simulate their behavior under several different scenarios. In the first, baseline, scenario workers never lose their jobs even though they anticipate such possibility when making their decisions. In each of the other, displacement, scenarios workers are laid off at a specific age that ranges between 58 and 70. Apart from job loss, in all scenarios, workers receive the same shocks. Therefore, their simulated histories completely coincide with the baseline scenario up to the age of a layoff, and any differences after that age are solely ascribed to the job loss. Comparison of the levels of labor supply chosen by workers under the baseline and displacement scenarios shows how labor force attachment is affected by job loss.

On average for all displacement ages, almost one third of laid off workers drop out of the labor force following a job loss. This number consists of the 29% who leave the labor force without attempting to search for new work and another 4% who retire after unsuccessful job search. The response varies by the age of job loss. The fraction of workers who retire without trying to search first peaks at 26% among workers displaced at age 62, the early retirement age. It then declines slightly before rising again to 32% at the full retirement age, and eventually exceeds 40% among seventy year old workers. These results are broadly consistent with the data, as discussed in Section 2. The MSM estimator did not target directly any of the moments that characterize the labor supply response to job loss, and hence this is an important out-of-sample validation of the model.

Involuntary job loss on average decreases the lifetime labor supply by 2,050 hours, an equivalent of one year's full time work. Coming to fifteen months over the lifetime, the reduction of labor supply following displacement is even larger on the extensive margin. This indicates that displaced workers decrease their labor supply on the extensive as well as intensive margins. Indeed, two years after job loss, displaced workers are twice more likely to work part time and 24% less likely to hold a job than if they kept steady employment.

Finding that job loss is associated with early retirement is consistent with earlier studies (e.g., Chan and Stevens, 2001). However, looking beyond the averages reveals the existence of a sizable group of indivduals who respond to job loss by working more in the years following displacement. One in five laid off workers would increase the lifetime labor supply by an average of 1,500 hours, or eight months in full time equivalent. Workers who delay their retirement following displacement are systematically different in their observable characteristics. In particular, they earn substantially higher pre-displacement wages. Layoffs therefore have a differentiated impact: they are more likely to drive low income workers out of the labor market, while making higher wage earners stay longer.

Because responses to layoffs vary, displacement outcomes in the economy depend on the composition of affected workers. For example, if job loss only affected low wage earners in the bottom 10% of the wage distribution, then most of the displaced workers would retire at the same time or earlier and only 12% would delay retirement from the labor force. However, if affected workers were concentrated in the median 20% of the wage distribution, the rate of delayed retirement would rise to 16%. It would further increase to 25% if the top 20% of wage earners lost their jobs. These experiments are important because the compositions of workers affected by layoffs has been changing over time and is likely to change again in the future. For example, the Great Recession has pushed into unemployment an unprecedented number of white collar workers. This analysis is informative of the way in which the impact of job loss on the transition from work to retirement depends on the composition of the older workforce and possible concentration of layoffs in specific industries or occupations.

7.2 The cost of job loss

There is ample evidence that displaced workers suffer from a broad range of negative consequences of job loss, including earnings and employment instability, deteriorating health, increase in mortality, and adverse effects on family, such as higher divorce rates and lower educational achievement of the children (Davis and von Wachter, 2011). Here I measure the cost of job loss for older workers as a monetary transfer required to make displaced worker indifferent between his current state and otherwise equivalent state in which he was not laid off from a job.

Panel A of Figure 3 shows the estimated cost of job loss for each age of layoff between fifty eight and seventy when the probabilities of job loss and job finding are fixed at their long-term averages. On average, the lifetime cost of job loss for workers displaced at these ages amounts to thirty four thousand dollars, which is about one year of earnings for this age group. The cost of job loss declines with displacement age, and the Social Security rules appear essential to the understanding of this dynamics. Workers displaced in their late fifties bear the highest cost, losing around forty five thousand dollars over the remaining lifetime – one third more than the average. The cost of job loss falls with the approach of early retirement age, so that workers displaced at sixty two lose 20% less than if they lost their jobs four years earlier. The cost of job loss does not change much until the full retirement age when it sharply falls by an additional quarter relative to the cost of job loss for fifty eight year olds. It remains consistently high, over twenty thousand dollars, even for workers in their late sixties.

To understand the link between the cost of job loss and the Social Security retirement age, recall previously documented spikes in the rates of labor force exit at early and full retirement ages (e.g. Rust and Phelan, 1997). At any age, involuntary job loss is not going to affect the lifetime outcomes of a worker who intended to retire exactly at this age in any event. A mass of workers who plan to leave the labor force around early and full retirement ages will therefore lower the average cost of displacement at these ages. The expected remaining working time of those who reached retirement age is shorter, both because these workers are now older and because they have an option to take up Social Security. The shortening of expected employment time prevents the cost of job loss to increase again after passing the peaks of retirement. Of course, lower averages do not mean that specific individuals are better off being displaced around the normal retirement age, as in each case the cost of job loss would depend on the anticipated individual retirement date that may be well past the Social Security age.

To see what generates the cost of job loss, note that displaced workers are adversely affected via two channels. First, they suffer a permanent reduction in future wages. Second, they lose earnings and access to health insurance due to search frictions. The two channels represent respectively intensive and extensive margins of the cost associated with job loss. While a large part of losses through each channel can be traced to forgone earnings, the overall cost of job loss is substantially more complicated than merely the change in the lifetime earnings. The cost also accounts for any effects on the dynamics of assets, differences in medical expenses due to the loss of health insurance, or the effect of job loss on individual labor supply and participation decisions, including the timing of Social Security take up and the amount of benefits collected over the lifetime.

The two channels can be explored separately. The intensive margin can be shut down by eliminating displacement wage penalty ($d_1 = 0$), while the extensive margin is removed under the assumption of frictionless labor markets ($\lambda_{t,\tau} = 1$ and $c = 0$). Decomposing the total cost by channel, I find that in stable economic conditions most of the losses are due to the reduction of post-displacement wages (intensive margin), which accounts for 80% of the total. The remaining 20% is due to the loss of earnings over the periods of unemployment that follow a job loss (extensive margin). The share of the extensive margin depends on the age of displacement, as shown in Figure 3. Search frictions are relatively more important for workers laid off around the normal retirement age due to the increasing number of displaced workers who give up potential earnings by retiring early without attempting to search. The fraction of costs due to search peaks at displacement age sixty four where it accounts for almost a third of the total.

7.3 Job loss over the business cycle

The probabilities of job loss and job finding vary over the business cycle, and their dynamics affects both the incidence and the cost of job loss. In this section I evaluate the difference in the cost of job loss for workers displaced in a recession relative to the workers displaced in more favorable economic conditions. I assume that workers do not recognize cyclical fluctuations prior to job loss, and expect that the probabilities of job loss and job finding stay at their corresponding long-term average values. Upon losing a job, they realize the true state of the economy. If displacement happened during a recession, they update their expectations of job loss and job finding probabilities to the new (recessionary) levels. They believe a recession will last for one year, after which job finding and job loss probabilities will revert to the long-term averages. These assumptions on the structure of recession and revision of expectations are chosen to explain the main behavior patterns, however many other designs are straightforward to incorporate into the model if the estimates for specific economic conditions are required.

Using these assumptions, I repeat a set of experiments similar to those from sections 7.1 and 7.2 to estimate the impact of job loss in recession on the labor force attachment and the cost of displacement. I set the recessionary probabilities of job finding and job loss equal to the estimates obtained at the trough of the Great Recession ($\lambda_{rec} = 0.14$ and $\delta_{rec} = 0.005$), so that displaced workers now face lower probability of job finding and higher chances of repeated displacement. For workers displaced at different stages of the business cycle, both the differences in the labor force attachment and the cost of job loss depend primarily on their reemployment prospects, captured by job finding probabilities. While there is little difference in the fraction of displaced workers who retire immediately after job loss, the rates of discouragement in a recession increase by a factor of 1.4. Jointly, additional discouragement and immediate retirement decrease the rate of reemployment after job loss in recessions by three percentage points.

Business cycle fluctuations exacerbate the negative impact of displacement on labor force attachment. The lifetime labor supply of older workers displaced in a recession decreases by additional 400 hours, over two months of full time work. Accounting for the extended duration of search, their total time in the labor force is three months lower than that of the identical workers displaced in more favorable economic conditions. Recessions further polarize the response by workers who adjust their hours of work in the opposite directions. Those who work less over the lifetime now supply 10% hours less, and those who work more increase the number of hours by additional 6%. In monetary terms, these changes in the labor force attachment translate into an additional loss of seven thousand dollars over the lifetime, generated almost exclusively by search frictions which in a recession account for 35% of the total loss (panel B of Figure 3).

7.4 Retirement during the Great Recession

Not only do displaced workers lose more if laid off during a recession, there are more displacements that occur in a weak economy. What is the overall impact of an economic downturn, like the Great Recession of 2007-2009, on the labor force attachment of workers who would have planned to retire around this time? To answer this question, I consider an artificial cohort of workers whose age is between 57 and 62 at the end of 2004, just over three years past the trough of the previous recession in November 2001. When the new recession ensues three years later in December 2007, these workers (aged 60-65 by that time) will be hit directly in the process of transition to retirement.

To estimate the impact of the recession on the retirement behavior of this cohort, I simulate two datasets. In the first simulation, the probabilities of layoff and job finding are fixed at their long term averages. The second simulation uses the actual job finding and layoff probabilities registered between the beginning of the simulation period in December 2004 and December 2010, a year and a half past the trough of the Great Recession. This experiment introduces the labor market cyclicality into the environment of older workers as they transition to retirement. Comparing labor force transitions in the two datasets, I can now discern how the rates of exit from the labor force vary at different stages of the business cycle.

The solid line in Figure 4 represents the percentage difference in the monthly new retirements observed in the simulations with the cyclical probabilities against the benchmark case that uses the long-term averages. The positive and negative values represent respectively higher and lower rates of retirement relative to the stable economy scenario. As the probabilities of job loss and job finding deviate from their average values, the new retirements respond by substantial variation with the range of changes of about 13%. The number of new retirements steadily increases throughout the recession, with the peak difference of 6.9% reached in June 2009. With the economy improving, the number of new retirements declines to the low difference point of -6% towards the end of 2010. Overall, the entire cohort retires one and a half months earlier in the recession.

Variations in job finding and layoff probabilities therefore matter, with more retirements occuring in weak labor markets. The picture however would be incomplete without accounting for another major feature of the Great Recession, namely the negative shock to the returns on assets. The retirement wealth of the simulated cohort was depleted as the stock markets shrank by more than 50% over a year and a half, accompanied by a comparable decline in the value of the housing equity. It is commonly acknowledged that while layoffs accelerate, asset losses postpone retirement. For example, Coile and Levine (2010) document how the media view of the Great Recession impact on retirement evolved over time. Initially envisioning delayed retirement due to financial markets crash, it later increasingly turned to predicting early retirement caused by unfavorable labor market conditions.

The two forces indeed work in the opposite direction in the simulations. The dotted line in Figure 4 shows percentage change in new retirements during the recession once the variation in asset returns has been taken into account. The early retirement peak shifts to the left, so that workers who were not yet fully affected by the negative asset shocks self-select to retire at the very beginning of the recession. Later on however more workers delay retirement from the labor force in an attempt to make up for wealth losses. With both asset and labor market dynamics taken into consideration, I find that the workers retiring during the Great Recession would on average postpone retirement by approximately five months. Intuitively, asset losses dominate the final outcome for two reasons. First, the two forces do not entirely cancel each other out as job loss itself leads to delayed retirement for some of the displaced workers. Second, fewer people are affected by layoffs than by changes in asset returns, even when the impact of a layoff is bigger in magnitude.

7.5 Policy implications

The main factor that drives the cost of job loss is the dip in post-displacement earnings. Wage penalty therefore would be the first target parameter for a policymaker who attempted to alleviate the consequences of involuntary job loss. The model predicts that a fifty percent reduction of the wage penalty from 0.16 to 0.08 will decrease the cost of displacement by 28%. It will also promote stronger labor force attachment among displaced workers, cutting down the total reduction of the labor supply hours in the aftermath of displacement by one fourth. A compensation mechanism attended to achieve this result could be designed through tax credits for reemployed displaced workers, similar to the EITC.

Compared to the reduction of the wage penalty, improvement of the job finding prospects has limited effect on the the cost of job loss. It is also arguably harder to achieve. For example, doubling the probability of job finding decreases the cost of job loss by just 14%. However, the same change generates a substantial, 20%, increase of labor supply over the lifetime. This result is much closer to the one obtained for the compensation of wage penalty. While improving the probability of job finding on its own is not as effective as reducing the wage penalty, a combination of the two measures yields a powerful tool to combat the cost of job loss. Indeed, a reduction of wage penalty by one half accompanied by doubling of the job finding probabilities leads to a 46% reduction of the cost of job loss. The total effect is 10% higher than the additive impact of the two interventions considered separately. This synergy implies that the impact of tax credits to reemployed displaced workers will be boosted in expansion. Likewise, a tax credit would work best when implemented alongside programs aimed at facilitating reemployment of older displaced workers.

The cost of job loss may not be a major concern in the design of Social Security reforms, yet alterations in the main parameters of Social Security generate serious side effects for displaced workers. The following examples are indicative of the consequences associated with the measures that are frequently proposed in order to maintain the solvency of Social Security. The consequences of increasing the full retirement age to 67 years are grave, amounting to a 30% increase in the cost of job loss. Eliminating early retirement opportunity would increase the cost of job loss by 12%. On the contrary, a change of annual penalty for early Social Security withdrawal down from 0.08 to 0.04 leads to a 10% decline of the cost of job loss, and more than twice as much for workers displaced at the age of sixty two.

8 Conclusion

At least one in five older workers in the United States is affected by involuntary job loss shortly before retirement from the labor force. As the probabilities of job finding decrease with age, displaced older workers experience longer unemployment spells and substantial reduction in earnings over the remaining lifetime. Faced with the difficulty of reemployment, many stop searching and permanently exit the labor force. Involuntary job loss and labor market frictions therefore potentially represent an important retirement incentive. To estimate the cost of job loss for older workers and understand its effect on the transition from work to retirement, I construct a dynamic programming model of retirement with search frictions.

I estimate the structural parameters of the model using the method of simulated moments and the data from the Health and Retirement Study. The model confirms that search frictions and involuntary job loss are essential to the understanding of retirement, as they account for 11% of the retirement trend. In a frictionless environment, the model predicts that the fraction of retirements otherwise explained by search is predominantly absorbed by the fixed cost of work. This result explains why the fixed costs of work are excessively high in the standard lifecycle models of labor supply and retirement.

The average cost of involuntary job loss measured by a compensation required to keep a worker indifferent between the states of employment and unemployment due to displacement is equivalent to one year of full-time earnings. Eighty percent of this cost is attributed to the post-displacement wage reduction, the rest is due to search frictions. The cost of losing a job in a contraction with job loss and job finding probabilities equal to those observed at the trough of the Great Recession is approximately 20% higher. This difference is mainly generated by search frictions that account for substantially larger portion of the total cost during economic downturns.

Displaced workers on average retire about one year earlier than they would have done without a job loss, however 18% of the simulated workers postpone their retirement in response to a layoff. The fraction of displaced workers who retire early increases during recessions when the labor markets are weak. Yet, in the two most recent recessions, the negative impact of job loss on the lifetime labor supply was partially offset by retirement delay caused by devaluation of savings as they were hit by the plunge of the stock markets and the housing bubble. I estimate that the simultaneous shocks to the values of assets, layoff and job finding probabilities would lead the cohort of workers who approached retirement age at the onset of the recession to postpone retirement by approximately five months.

References

Aaronson, D. and E. French (2004). The effect of part-time work on wages: Evidence from the Social Security rules. *Journal of Labor Economics 22*, 329–352.

Abel, A. B. and M. Warshawsky (1988). Specification of the joy of giving: Insights from altruism. *The Review of Economics and Statistics 70*, 145–149.

Altonji, J. G. and L. M. Segal (1996). Small-sample bias in GMM estimation of covariance structures. *Journal of Business & Economic Statistics 14*, 353–366.

Blau, D. M. and D. B. Gilleskie (2006). Health insurance and retirement of married couples. *Journal of Applied Econometrics 21*, 935–953.

Blau, D. M. and D. B. Gilleskie (2008). The role of retiree health insurance in the employment behavior of older men. *International Economic Review 49*, 475–514.

Chan, S. and A. H. Stevens (2001). Job loss and employment patterns of older workers. *Journal of Labor Economics 19*, 484–521.

Chan, S. and A. H. Stevens (2004). How does job loss affect the timing of retirement? *Contributions in Economic Analysis & Policy 3*, –.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011). Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins. *The American Economic Review 101*, 471–75.

Cogan, J. F. (1981). Fixed costs and labor supply. *Econometrica 49*, 945–963.

Coile, C. C. and P. B. Levine (2011a). The market crash and mass layoffs: How the current economic crisis may affect retirement. *The B.E. Journal of Economic Analysis & Policy 11*, –.

Coile, C. C. and P. B. Levine (2011b). Recessions, retirement and social security. *American Economic Review: Papers & Proceedings 101*, 23–28.

Couch, K. A. and D. W. Placzek (2010). Earnings losses of displaced workers revisited. *The American Economic Review 100*(1), 572–589.

David, M., R. J. A. Little, M. E. Samuhel, and R. K. Triest (1986). Alternative methods for CPS income imputation. *Journal of the American Statistical Association 81*, 29–41.

Davis, S. J. and T. von Wachter (2011). Recessions and the costs of job loss. Brookings Papers on Economic Activity.

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

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.

Elder, T. E. (2004). Reemployment patterns of displaced older workers. Manuscript.

Farber, H. S. (1993). The incidence and costs of job loss: 1982-91. *Brookings Papers on Economic Activity*, 73–119.

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

French, E. B. and J. B. Jones (2011). The effects of health insurance and self-insurance on retirement behavior. *Econometrica 79*, 693–732.

Goda, G. S., J. B. Shoven, and S. N. Slavov (2011). What explains changes in retirement plans during the Great Recession. *American Economic Review: Papers & Proceedings 101*, 29–34.

Gouriéroux, C. and A. Monfort (1996). *Simulation-Based Econometric Methods*. New York: Oxford University Press.

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

Gustman, A. L. and T. L. Steinmeier (1986). A structural retirement model. *Econometrica 54*, 555–584.

Gustman, A. L., T. L. Steinmeier, and N. Tabatabai (2010). What the stock market decline means for the financial security and retirement choices of the near-retirement population. *Journal of Economic Perspectives 24*, 161–182.

Haan, P. and V. Prowse (2010). A structural approach to estimating the effect of taxation on the labour market dynamics of older workers. *Econometrics Journal 13*, S99–S125.

Haan, P. and V. Prowse (2014). Longevity, life-cycle behavior and pension reform. *Journal of Econometrics 178*, 582–601.

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1995). Precautionary saving and social insurance. *Journal of Political Economy 103*, 360–399.

Ibbotson Associates (2013). *Stocks, Bonds, Bills, and Inflation Yearbook*. Chicago: Morningstar, Inc.

Jacobson, L. S., R. J. LaLonde, and D. G. Sullivan (1993). Earnings losses of displaced workers. *The American Economic Review 83*(4), 685–709.

Kopecky, K. A. and R. M. Suen (2010). Finite state Markov-chain approximations to highly persistent processes. *Review of Economic Dynamics 13*, 701–714.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica 57*, 995–1026.

McFall, B. H. (2011). Crash and wait? The impact of the Great Recession on the retirement plans of older Americans. *American Economic Review: Papers & Proceedings 101*, 40–44.

McGrattan, E. R. and E. C. Prescott (2003). Average debt and equity returns: Puzzling? *The American Economic Review 93*, 392–397.

Menzio, G., I. A. Telyukova, and L. Visschers (2015). Directed search over the life cycle. *Review of Economic Dynamics*.

Pakes, A. and D. Pollard (1989). Simulation and the asymptotics of optimization estimators. *Econometrica 57*, 1027–1057.

Pischke, J.-S. (1995). Measurement error and earnings dynamics: Some estimates from the PSID validation study. *Journal of Business & Economic Statistics 13*, 305–314.

Poterba, J. M. and L. H. Summers (1984). Response variation in the CPS: caveats for the unemployment analyst. *Monthly Labor Review, March*, 37–43.

Rogerson, R. and J. Wallenius (2013). Nonconvexities, retirement and the elasticity of labor supply. *American Economic Review 103*, 1445–62.

Rouwenhorst, K. G. (1995). Asset pricing implications of equilibrium business cycle models. In T. F. Cooley (Ed.), *Frontiers of Business Cycle Research*, pp. 294–330. Princeton, New Jersey: Princeton University Press.

Rust, J. and C. Phelan (1997). How social security and Medicare affect retirement behavior in a world of incomplete markets. *Econometrica 65*, 781–831.

Scholz, J. K., A. Seshadri, and S. Khitatrakun (2006). Are Americans saving "optimally" for retirement? *Journal of Political Economy 114*, 607–643.

Stephens, M. (2004). Job loss expectations, realizations, and household consumption behavior. *The Review of Economics and Statistics 86*(1), 253–269.

Stevens, A. H. (1997). Persistent effects of job displacement: The importance of multiple job losses. *Journal of Labor Economics 15*(1), 165–188.

van der Klaauw, W. and K. I. Wolpin (2008). Social security and the retirement and savings behavior of low-income households. *Journal of Econometrics 145*, 21–42.

Figures and Tables

Figure 1: Labor force status following layoff at sixty

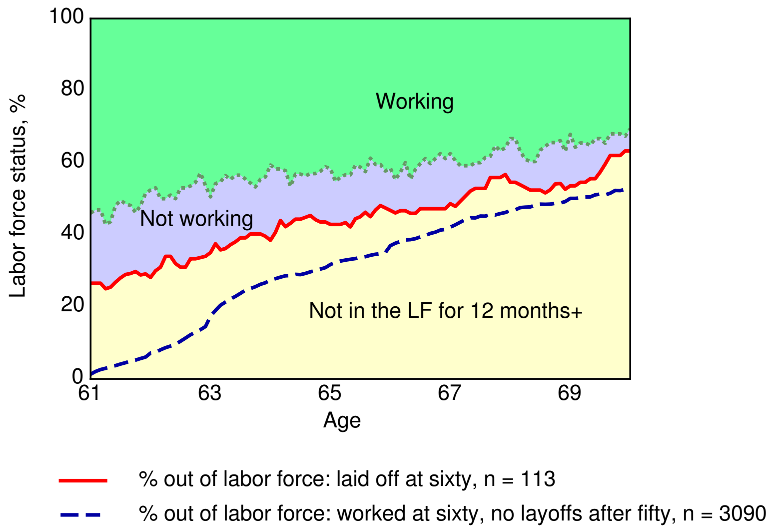

Notes: Monthly employment histories based on the HRS data, males fifty and older. The colored areas show fractions of employed, unemployed and out of the labor force in a group of workers laid off at age sixty.

Figure 2: Model fit: matched moments

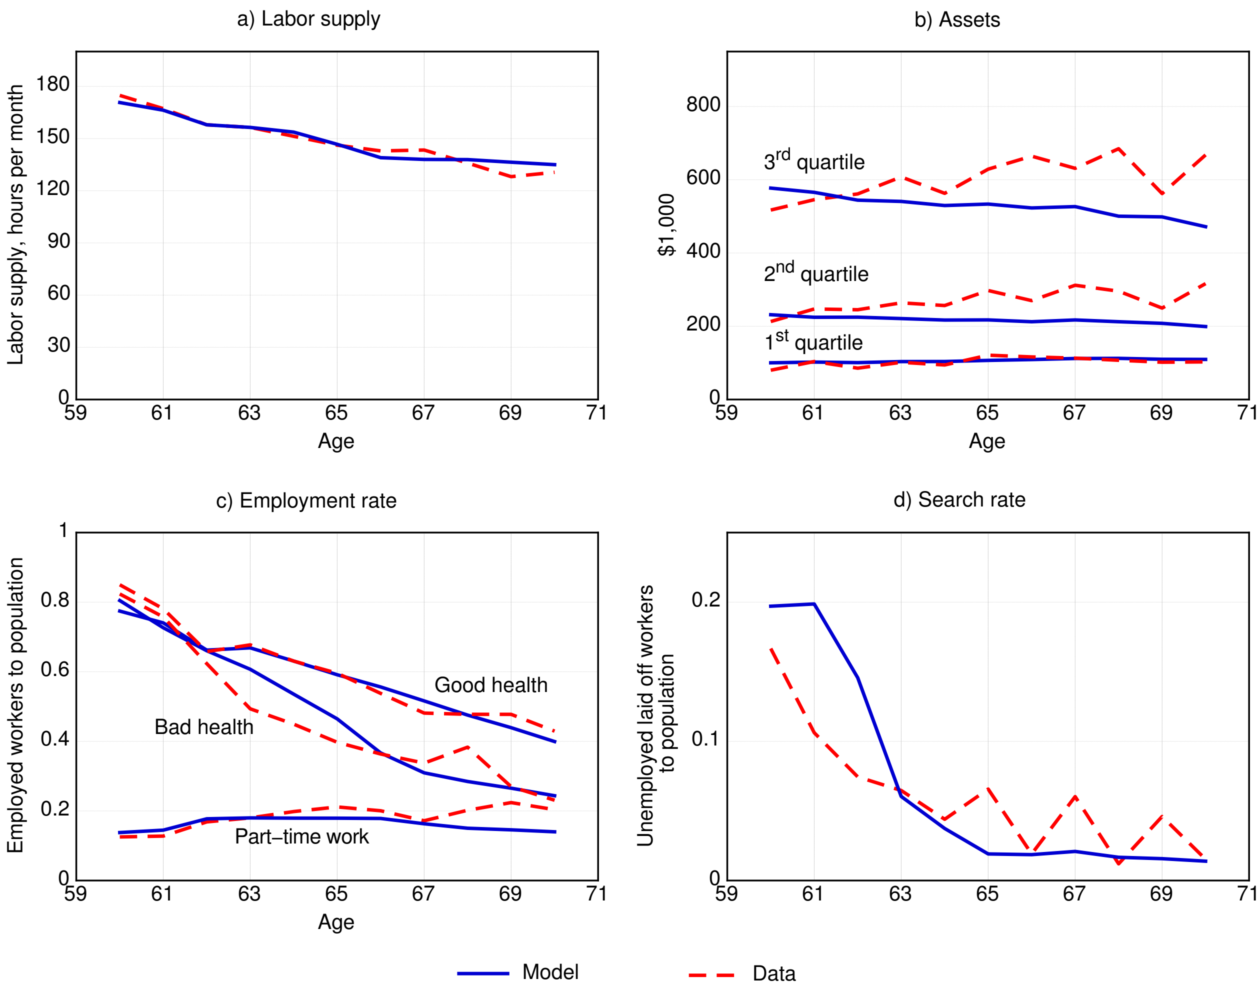

Figure 3: The cost of layoff by age of job loss

Notes: The cost of layoff is computed as a monetary compensation required to keep worker indifferent between the states of displacement and non-displacement. The solid line shows the total cost of layoff by age of job loss. The height of dashed area corresponds to the cost of job loss due to the reduction of the wage rate. The height of shaded area shows the cost of job loss due to search frictions.

Figure 4: The response of new retirements to the variations in job finding and layoff probabilities

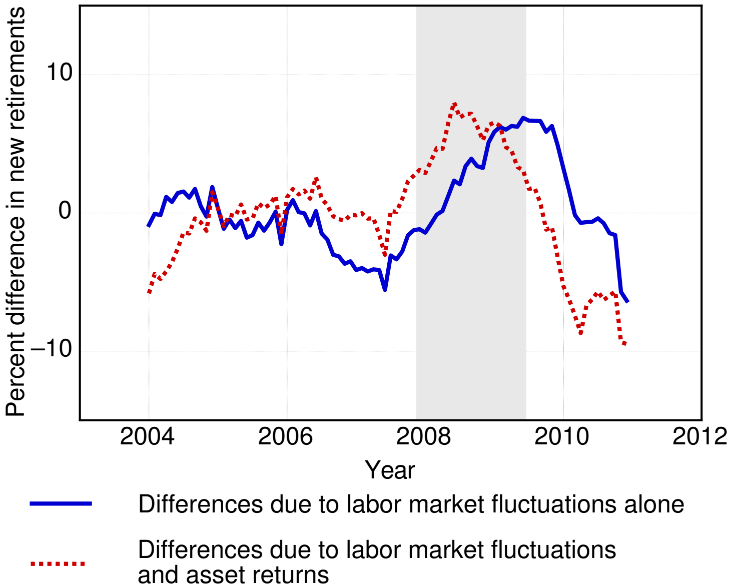

Notes: The solid blue line shows percentage difference in new retirements in simulations with actual probabilities of layoff and job finding against benchmark case where both probabilities are set to their long-term averages. Filter twelve months moving average of the simulated data. The dotted red line in addition accounts for percentage difference in new retirements due to the variations in asset returns. Grey areas correspond to the NBER contraction dates.

Table 1: Descriptive statistics for the initial state distribution

Variable	Income Type 1	Income Type 2	Income Type 3
Fraction in good health	.81	.87	.95
Fraction working	.84	.81	.87
Fraction unemployed	.03	.02	.01
Fraction at high risk of layoff	.49	.53	.39
Fraction with employer provided health insurance	.58	.73	.83
Labor supply (conditional on employment), in hours per month			
Mean	180	180	188
Median	160	160	180
Standard deviation	55.1	69.8	70.0
Wage rate (conditional on employment), in 2000 dollars			
Mean	21.4	34.0	38.2
Median	11.5	15.9	15.0
Standard deviation	18.3	35.4	36.6
AIME, in 2000 dollars			
Mean	1,789	3,464	4,878
Median	1,877	3,452	4,793
Standard deviation	684	379	541
Assets, in thousands of 2000 dollars			
Mean	309	428	704
Median	143	215	394
Standard deviation	493	520	838
Monthly medical expenses, in 2000 dollars			
Mean	55.3	82.5	63.9
Median	27.2	27.1	32.0
Standard deviation	87.1	204.9	128.2
Mean age, years	59.3	59.9	59.6
Sample observations	140	132	150

Notes: HRS white non-Hispanic male workers born in 1938-1943 with at least ten years of non-government employment. Excludes early recipients of Social Security (before 62) and recipients of SSI/SSDI. Income types 1-3 correspond to the tertiles of the AIME distribution.

Table 2: Logit estimates of survival and health transition probabilities

Explanatory variable	Income Type 1	Income Type 2	Income Type 3
		Survival probability:	
Age	-.007	-.006	-.005
	(.000)	(.000)	(.000)
Good health	.080	.098	.094
	(.009)	(.011)	(.014)
Number of observations	10,188	12,194	13,688
χ^2_2	989	919	596
		Probability of being in good health:	
Age	-.006	-.010	-.012
	(.001)	(.001)	(.001)
Good health	.537	.585	.694
	(.010)	(.014)	(.019)
Number of observations	8,810	10,929	12,855
χ^2_2	1,333	1,457	1,277

Notes: Estimation sample is HRS white non-Hispanic males 50 and older with at least ten years of non-government employment. Excludes early recipients of Social Security (before 62) and recipients of SSI/SSDI. The values reported are marginal effects of age and health on biennial transition probabilities from logit maximum likelihood estimation, computed at age 65 and lag health equal to 0. Standard errors are given in parentheses. Income types 1-3 correspond to the tertiles of the AIME distribution.

Table 3: Estimates of the medical expenses process

Variable	Health insurance status: Insured by employer	No insurance	Medicare (Age 65+)
Mean of log medical expenses, estimated from biennial data			
Age	.018	.039	.010
	(.004)	(.008)	(.002)
Good health	-.565	-.466	-.347
	(.045)	(.077)	(.024)
Constant	2.952	1.734	3.261
	(.189)	(.460)	(.146)
Autocorrelation of AR(1) disturbances	.333	.333	.453
	(.013)	(.035)	(.011)
Innovation variance of AR(1) disturbances	1.885	2.119	1.875
	(.031)	(.062)	(.025)
Number of observations	12,115	2,801	22,224
Parameters of monthly AR(1) error process			
Autocorrelation, ρ^m	.955	.955	.968
Innovation variance, $\sigma^2_{\varepsilon m}$	.165	.186	.120

Notes: Estimation sample is HRS white non-Hispanic males 50 and older with at least ten years of non-government employment. Excludes early recipients of Social Security (before 62) and recipients of SSI/SSDI. Conditional maximum likelihood estimates of equation (13). Standard errors are given in parentheses. Zero medical expenses were omitted from estimation.

Table 4: Estimates of the wage transition process

Variable	Estimate	Standard error
Mean of log wage rate, estimated from biennial data		
Age	.019	.022
$0.01 \times \text{Age}^2$	-.022	.018
AIME, log	.339	.024
Good health indicator	.075	.018
Constant	-.383	.702
Autocorrelation of AR(1) disturbances	.677	.028
Innovation variance of AR(1) disturbances	.529	.027
Number of observations	16,819	
Parameters of monthly AR(1) error process		
Autocorrelation, ρ^w	.984	
Innovation variance, $\sigma^2_{\varepsilon w}$	.031	

Notes: Estimation sample is HRS white non-Hispanic males 50 and older with at least ten years of non-government employment. Excludes early recipients of Social Security (before 62) and recipients of SSI/SSDI. Conditional maximum likelihood estimates of equation (14). Estimation sample excludes top and bottom 1% of wage rate values.

Table 5: Estimates of structural parameters

Parameter	Estimate	Standard error
θ_1: Consumption weight	0.580	0.019
θ_2: Coefficient of relative risk aversion	4.80	0.137
b_1: Intensity of bequest motive	0.985	0.086
b_2: Bequest shifter	1.084	0.228
L: Monthly leisure endowment, hours	375	13.4
ϕ_0: Fixed cost of work at 60, hours	69.3	3.56
ϕ_1: Fixed cost of work, slope	0.230	0.025
ψ: Cost of bad health, hours	0.100	0.022
c: Cost of search, hours	85.4	14.8
d_1: Wage loss due to displacement	0.163	0.024
d_2: Wage loss due to part time employment	0.191	0.016
$\underline{C}$: Government transfers, USD	375	11.3
λ_B: Probability of getting unemployment benefits	0.881	0.074
$\sigma^2_{\varepsilon^a}$: Variance of the measurement error in assets	0.237	0.010
χ^2 statistic $= 429$, degrees of freedom $= 76$		

Notes: Method of simulated moments estimates of the structural parameters, conditional on exogenous estimates of the state transition probabilities. Initial values of state variables are drawn from the 2000 HRS sample of white non-Hispanic males 50 and older with at least ten years of non-government employment, born in years 1938-1943. Excludes early recipients of Social Security (before 62) and recipients of SSI/SSDI. Estimation uses diagonal weighting matrix.