Even after controlling for health status, French and Jones, (2004) find that medical expenses are very volatile and persistent. Thus we model the person-specific component of medical expenses, ψ_t, as

$$\psi_t = \zeta_t + \xi_t, \quad \xi_t \sim \mathcal{N}(0, \sigma_\xi^2) \tag{8}$$

$$\zeta_t = \rho_m \zeta_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2) \tag{9}$$

where ξ_t and ϵ_t are serially and mutually independent. ξ_t is the transitory component, while ζ_t is the persistent component, with autocorrelation ρ_m.

There are several different types of health insurance model. As a first step, it is useful to characterize an individual by his access to employer-provided health insurance (EPHI), which we denote by I_t. At the beginning of a period, the individual finds himself in one of three mutually exclusive states:

1. *retiree* health insurance that he can hold on to until his death.

2. *tied* health insurance that he will lose shortly after his current job terminates. If a worker with *tied* health insurance leaves his job, he can keep his health insurance coverage for that year. This is meant to proxy for the fact that most firms must provide "COBRA" health insurance to workers after they leave their job. After one year of *tied* coverage and not working, the individual's insurance ceases.[4]

3. *non-group* insurance, i.e., an individual is on his own. He has the choice between purchasing insurance on the *private* non-group market or being *uninsured*.

Accounting for the choices of those in the *non-group* category, there are four types of privately-provided health insurance: *retiree*, *tied*, *private*, and *uninsured*.

Workers move between these insurance categories according to the rules defined in appendix A.1

[4]Although there is some variability across states as to how long individuals are eligible for employer-provided health insurance coverage, by Federal law most individuals are covered for 18 months (Gruber and Madrian, 1996). Given a model period of one year, we approximate the 18-month period as one year. We do not model the option to take up COBRA, assuming that the take-up rate is 100%. The actual take-up rate is around $\frac{2}{3}$ (Gruber and Madrian, 1996). In French and Jones, (2011) we conducted a robustness test where we simulated the model assuming that the rate was 0%, so that individuals transitioned from *tied* to *non-group* as soon as they stopped working, and found very similar labor supply patterns.

(Table 10).

In addition to private coverage, individuals may receive Medicare and/or Medicaid benefits, according to the following rules:

1. *Medicare* insurance if he is either *disabled* or has reached the age of 65.[5]

2. An individual will qualify for *Medicaid* insurance if he is poor enough to receive Supplemental Security Income and he is eithers *disabled* or has reached the age of 65.[6]

Both Medicare and Medicaid operate on top of the private coverage, although some combinations are impossible. We model the interaction of public and private health insurance as follows:

1. In actual practice the interaction of employer (*retiree* or *tied*) coverage with Medicare is complicated, depending on employment and firm size (Centers for Medicare and Medicaid Services, 2014). We assume instead that all individuals receiving both EPHI and *Medicare* share a "joint" plan that differs only by demographics (t and SP_t).

2. Many households purchase "Medigap" insurance to help pay for expenses not covered by Medicare. Our model abstracts away from this choice, and our empirical estimates will combine the two coverages.

3. *Medicaid* insurance is intended to be the payer of last resort, which is to say that *Medicaid* covers only the co-payments and deductibles left behind by other insurers.

4. While *Medicaid* covers *Medicare* premia, it does not cover the premia associated with private insurance. As Brown and Finkelstein, (2008) show, the latter provision can at times strongly discourage the purchase of private insurance.

5. The eligibility rules of the DI program require that the individual not work during the application period, although he may work later. As a result, an individual with *tied*

[5] Individuals who have paid into the Medicare system for at least 10 years become eligible at age 65. A more detailed description of the Medicare eligibility rules is available at `http://www.medicare.gov/`.

[6] Our definition of Medicaid is that of "categorically needy" recipients, who qualify because their income and wealth are low, regardless of their medical conditions. The provision of Medicaid through other mechanisms, the most important of which is the "medically needy" provision, is captured by the consumption floor.

coverage will lose this coverage if he transitions to DI and the associated Medicare and Medicaid coverage.

Let I_t^+ denote the health insurance coverage the household receives after it has (possibly) decided whether to purchase *private* non-group coverage and after i's Medicaid eligibility has been determined. The realized value of I_t^+ determines how the total health care cost Z_t translates into out-of-pocket expendituresM_t via the formula

$$
\begin{aligned}
M_t &= \mathit{premium}(I_t^+,\, t,\, P_t,\, \widehat{Z}_t,\, SP_t) + \mathit{copay}(I_t^+,\, Z_t), \\
\widehat{Z}_t &= \mathbb{E}[Z_t \,|\, t, H_t, \zeta_{t-1}].
\end{aligned} \tag{10}
$$

Here *premium*$(\cdot)$ is the health insurance premium, which can depend on expected medical expenditures, $\widehat{Z}_t$; the function *copay*(I_t^+, Z_t) determines how much of Z_t is assigned to the individual via co-payments and deductibles. We estimate both *premium*$(\cdot)$ and *copay*$(\cdot)$-function directly from the MEPS data. See the appendix for more details.

3.4 Marital Status and Spousal Income

Because spousal income can serve as insurance against medical shocks, and because marital status affects eligibility for Medicaid and other government programs, we include it in the model. We assume that when a spouse is present, spousal income ys_t takes on two values: (i) zero; or (ii) a positive value that varies with age. We assume the transition probabilities for marital status, and whether the spouse has positive income depend on its current current marital status and income, current health, and age: see appendix B.5 for details.

3.5 Wages

We assume that the logarithm of wages at time t, $\ln W_t$, is a function of health status (H_t), age (t), hours worked (N_t) and an autoregressive component, ω_t:

$$\ln W_t = W(H_t, t) + \alpha \ln N_t + \omega_t. \tag{11}$$

The inclusion of hours, N_t, in the wage determination equation captures the empirical regularity that, all else equal, part-time workers earn relatively lower wages than full time workers. The autoregressive component ω_t has the correlation coefficient ρ_W and the normally-distributed innovation η_t:

$$\omega_t = \rho_W \omega_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2). \tag{12}$$

3.6 Social Security, Disability Insurance, and Pensions

Because pensions and Social Security generate potentially important retirement incentives, we model the two programs in detail.

Individuals receive no Social Security benefits until they apply. Individuals can first apply for benefits at age 62. Upon applying the individual receives benefits until death. The individual's Social Security benefits depend on his Average Indexed Monthly Earnings ($AIME$), which is roughly his average income during his 35 highest earnings years in the labor market.

The Social Security System provides three major retirement incentives.[7] First, while income earned by workers with less than 35 years of earnings automatically increases their $AIME$, income earned by workers with more than 35 years of earnings increases their $AIME$ only if it exceeds earnings in some previous year of work. Because Social Security benefits increase in $AIME$, this causes work incentives to drop after 35 years in the labor market.

Second, the age at which the individual applies for Social Security affects the level of benefits. For every year before age 65 the individual applies for benefits, benefits are reduced by 6.67% of the age-65 level. This is roughly actuarially fair. But for every year after age 65 that benefit application is delayed, benefits rise by 5.5% up until age 70. This is less than actuarially fair, and encourages people to apply for benefits by age 65.

Third, the Social Security Earnings Test taxes labor income of beneficiaries at a high rate. For individuals aged 62-64, each dollar of labor income above the "test" threshold (of \$9,120

[7] A description of the Social Security rules can be found in recent editions of the *Green Book* (Committee On Ways And Means, various years). Some of the rules, such as the benefit adjustment formula, depend on an individual's year of birth. Because we fit our model to a group of individuals that on average were born in 1933, we use the benefit formula for that birth year.

in 1998) leads to a 1/2 dollar decrease in Social Security benefits, until all benefits have been taxed away. For individuals aged 65-69 in 1998, each dollar of labor income above a threshold of $14,500 leads to a 1/3 dollar decrease in Social Security benefits, until all benefits have been taxed away. Although benefits taxed away by the earnings test are credited to future benefits, in 1998 the Social Security Earnings Test effectively taxes the labor income of beneficiaries aged 65-69.[8] When combined with the aforementioned incentives to draw Social Security benefits by age 65, the Earnings Test discourages work after age 65. In 2000, the Social Security Earnings Test was abolished for those 65 and older. Because those born in 1933 (the average birth year in our sample) turned 67 in 2000, we assume that the earnings test was repealed at age 67. These incentives are incorporated in the calculation of ss_t, which is defined to be net of the earnings test.

Associated with Social Security program is Disability Insurance (DI). Individuals with $H_t = disabled$ receive Disability benefits if their income is below a threshold. The level of the benefits is a function of $AIME$. Individuals with low $AIME$ and low assets also receive top-up benefits through the Supplemental Security Income (SSI) program. DI benefits are labor-income tested: individuals who earn more than $12,840 in 2014 do not receive any benefits. We model this period-by-period conditional on $H_t = disabled$.

Poor individuals who are elderly or disabled ($H_t = disabled$ or $t \geq 65$) can qualify for Supplemental Security Income (SSI). Individuals with income below Y^{SSI} and assets below A^{SSI} receive a transfer of $(Y^{SSI} - Y_t)$. As described in Table 10, they also qualify for Medicaid.

Pension benefits, pb_t, are a function of the worker's age and pension wealth. Pension wealth (the present value of pension benefits) in turn depends on pension accruals. We assume that pension accruals are a function of a worker's age, labor income, and health insurance type, using a formula estimated from confidential HRS pension data. The data show that pension accrual rates differ greatly across health insurance categories; accounting for these differences is essential in isolating the effects of employer-provided health insurance. When finding an individual's decision rules, we assume further that the individual's existing pension wealth is a

[8] The credit rates are based on the benefit adjustment formula. If a year's worth of benefits are taxed away between ages 62 and 64, benefits in the future are increased by 6.67%. If a year's worth of benefits are taxed away between ages 65 and 66, benefits in the future are increased by 5.5%. See Olsen and Romig, 2013 for more details on the earnings test.

function of his Social Security wealth, age, and health insurance type. Details of our pension model are described in Section 6.6; also see French and Jones, (2011).

3.7 Recursive Formulation

In addition to choosing hours, consumption, and potentially private non-group insurance vs. self-insurance, eligible individuals decide whether to apply for Social Security benefits; let the indicator variable $B_t \in \{0,1\}$ equal one if an individual has applied. In recursive form, the individual's problem can be written as

$$
\begin{aligned}
V_t(X_t) = \max_{C_t, N_t, B_t, I_t^+} \Bigg\{ & \frac{1}{1-\nu} \left(C_t^{\gamma} \left(L - N_t - \phi_{Pt} P_t - \phi_{RE} RE_t - \phi_H(H_t) \right)^{1-\gamma} \right)^{1-\nu} \\
& + \beta(1 - s_{t+1}) b(A_{t+1}) \\
& + \beta s_{t+1} \int V_{t+1}\left(X_{t+1} \right) dF(X_{t+1} | X_t, t, C_t, N_t, B_t) \Bigg\}
\end{aligned} \tag{13}
$$

subject to equations (5) and (6). The vector $X_t = (A_t, B_{t-1}, H_t, AIME_t, I_t, P_{t-1}, \omega_t, \zeta_{t-1}, \Upsilon_t)$ contains the individual's state variables, while the function $F(\cdot|\cdot)$ gives the conditional distribution of these state variables, using equations (4) and (7)–(12). [9] The solution to the individual's problem consists of the consumption rules, work rules, insurance choice rules, and benefit application rules that solve equation (13). These decision rules are found numerically using value function iteration.

4 Modeling changes induced by the ACA

4.1 Medicaid expansion

After 2014 low-income people can get Medicaid through the categorically needy channel, regardless of asset levels. In particular the eligibility test changes from

$$
I_t^+ = \textit{Medicaid} \text{ if } \left\{ Y_t < Y^{\text{cat-needy}}(SP_t) \text{ and } A_t < A^{\text{cat-needy}}(SP_t) \right\}, \tag{14}
$$

[9] Because we impute pension benefits as a function of the other state variables (as in French and Jones 2011), pension wealth is not a state variable.

to

$$I_t^+ = Medicaid \text{ if } \quad Y_t < Y^{\text{cat-needy}}(SP_t) \quad . \tag{15}$$

In both cases the eligibility thresholds depend on marital status. As before the reform, individuals who fail the Medicaid eligibility tests but have catastrophic medical spending receive the minimum consumption level given by equation (6).

4.2 Health insurance exchanges, tax subsidies and penalties

The ACA will affect the $premium(\cdot)$ and $copay(\cdot)$-functions for the non-group market. First, the $premium(\cdot)$ function will no longer depend on expected medical expenses except those related to age, and $copay(\cdot)$ function will have to satisfy the actuarial value and out-of-pocket limits specified by the law. Second, qualifying households will receive premium credits and cost-sharing subsidies. In addition, those who *self-insure* will have to pay the shared responsibility penalty for not buying insurance.

5 Estimation

To estimate the model, we adopt a two-step strategy, similar to the one used by Gourinchas and Parker, (2002) and French, (2005). In the first step we estimate or calibrate parameters that can be cleanly identified without explicitly using our model. For example, we estimate mortality rates and health transitions straight from demographic data. In the second step, we estimate the preference parameters of the model, as well as the consumption floor, using the method of simulated moments (MSM).

5.1 Moment Conditions

The objective of MSM estimation is to find the preference vector that yields simulated life-cycle decision profiles that "best match" (as measured by a GMM criterion function) the profiles from the data. The moment conditions that comprise our estimator are:

1. Because an individual's ability to self-insure against medical expense shocks depends upon his asset level, we match 1/3rd and 2/3rd asset quantiles by age. We match these quantiles in each of T periods (ages), for a total of $2T$ moment conditions.

2. We match job exit rates by age for each health insurance category. With three health insurance categories (*non-group*, *retiree* and *tied*), this generates $3T$ moment conditions.

3. Because the value a worker places on employer-provided health insurance may depend on his wealth, we match labor force participation conditional on the combination of asset quantile and health insurance status. With 2 quantiles (generating 3 quantile-conditional means) and 3 health insurance types, this generates $9T$ moment conditions.

4. To help identify preference heterogeneity, we utilize a series of questions in the HRS that ask workers about their preferences for work. We combine the answers to these questions into a time-invariant index, $pref \in \{high, low, out\}$, which is described in greater detail in Section 6.7. Matching participation conditional on each value of this index generates another $3T$ moment conditions.

5. We match hours of work and participation conditional on our binary health indicator. This generates $4T$ moment conditions.

6. Whether it is more attractive to buy private non-group health insurance or to self-insure against medical expense risk primarily depends on a household's asset level. Conditional on neither having access to employer-provided health insurance nor being eligible for Medicare or Medicaid, we match the fraction of households purchasing private insurance. Since everybody becomes eligible for Medicare at age 65, this generates $3T_{65}$ moment conditions, where T_{65} denotes all ages included in the model before 65.

Combined, the five preceding items result in $21T + 3T_{65}$ moment conditions.

5.2 Initial Conditions and Preference Heterogeneity

A key part of our estimation strategy is to compare the behavior of individuals with different forms of employer-provided health insurance. If access to health insurance is an important factor in the retirement decision, we should find that individuals with *tied* coverage retire later than those with *retiree* coverage. In making such a comparison, however, we must account for the possibility that individuals with different health insurance options differ systematically along other dimensions as well. For example, individuals with retiree coverage tend to have higher wages and more generous pensions.

We control for this "initial conditions" problem in three ways. First, the initial distribution of simulated individuals is drawn directly from the data. Because households with retiree coverage are more likely to be wealthy in the data, households with retiree coverage are more likely to be wealthy in our initial distribution. Similarly, in our initial distribution households with high levels of education are more likely to have high values of the persistent wage shock ω_t.

Second, we model carefully the way in which pension and Social Security accrual varies across individuals and groups.

Finally, we control for unobservable differences across health insurance groups by introducing permanent preference heterogeneity, using the approach introduced by Heckman and Singer, (1984) and adapted by (among others) Keane and Wolpin, (1997) and van der Klaauw and Wolpin, (2008). Each ndividual is assumed to belong to one of a finite number of preference "types", with the probability of belonging to a particular type a logistic function of the individual's initial state vector: his age, wealth, initial wages, health status, health insurance type, medical expenditures, and preference index.[10] We estimate the type probability parameters jointly with the preference parameters and the consumption floor.

In our framework, correlations between preferences and health insurance emerge because people with different preferences systematically select jobs with different types of health insurance coverage. Workers in our data set are first observed in their fifties; by this age, all else equal, jobs that provide generous post-retirement health insurance are more likely to be held by workers that wish to retire early. One way to measure this self-selection is to structurally model the choice of health insurance at younger ages, and use the predictions of that model to infer the correlation between preferences and health insurance in the first wave of the HRS.

[10] These discrete type-based differences are the only preference heterogeneity in our model. For this reason many individuals in the data make decisions different from what the model would predict. Our MSM procedure circumvents this problem by using moment conditions that average across many individuals.

Because such an approach is computationally expensive, we instead model the correlation between preferences and health insurance in the initial conditions.

5.3 Wage Selection

We estimate a selection-adjusted wage profile using the procedure developed in French, (2005). First, we estimate a fixed effects wage profile from HRS data, using the wages observed for individuals who are working. The fixed-effects estimator is identified using wage growth for workers. If wage growth rates for workers and non-workers are the same, composition bias problems—the question of whether high wage individuals drop out of the labor market later than low wage individuals—are not a problem. However, if individuals leave the market because of a wage drop, such as from job loss, then wage growth rates for workers will be greater than wage growth for non-workers. This selection problem will bias estimated wage growth upward.

We control for selection bias by finding the wage profile that, when fed into our model, generates the same fixed effects profile as the HRS data. Because the simulated fixed effect profiles are computed using only the wages of those simulated agents that work, the profiles should be biased upwards for the same reasons they are in the data. We find this bias-adjusted wage profile using the iterative procedure described in French, (2005).

6 Data and Calibrations

6.1 HRS Data

We estimate the model using data from the Health and Retirement Survey (HRS) which is nationally representative sample of initially non-institutionalized individuals, and their spouses. We use data from everyone in the HRS who is at least age 51, which is the youngest age that core members of the sample are interviewed. With the exception of assets and medical expenses, which are measured at the household level, our data are for male household heads. The HRS surveys individuals every two years, so that we have 11 waves of data covering the period 1992-2012. The HRS also asks respondents retrospective questions about their work history that allow us to infer whether the individual worked in non-survey years.

As noted above, the Social Security rules depend on an individual's year of birth. To ensure that workers in our sample face a similar set of Social Security retirement rules, we fit our model to the data for the cohort of individuals born in the 1940s. However, when estimating the stochastic processes such as marital status, health and spousal income we use the full sample, including older individuals. With the exception of wages and spousal income, we do not adjust the data for cohort effects. Because our subsample of the HRS covers a fairly narrow age range, this omission should not generate much bias.

6.2 Health and mortality

We estimate health transitions and mortality rates simultaneously by fitting the transitions observed in the HRS to a multinomial logit model. We allow the transition probabilities to depend on age and current health status. We estimate annual transition rates: combining annual transition probabilities in consecutive years yields two-year transition rates we can fit to the HRS data. Appendix B.2 describes this process in detail.

We assign individuals to one of three health states: good, bad or disabled. First, we give individuals a health status of "good" if their self-reported health is excellent, very good or good, and a health status of "bad" if their self-reported health is fair or poor. We reclassify individuals as disabled if they are receiving Medicare and/or Medicaid benefits and are younger than 65, regardless of self reported health. We use this measure of disability because we wish to capture both the cash transfers, and even more importantly, the Medicare or Medicaid insurance received by the disabled. Because DI recipients are transferred to Social Security at age 65, and virtually everyone qualifies for Medicare at the same age, we are able to identify disability in our data only up to age 64. From age 65 forward, we collapse the space of health outcomes to back to {bad, good}. This requires us to estimate three health transition specifications: one for the three-state health measure; one for the two-state measure; and one for the transtion from three states to two between ages 64 and 65.[11] To simplify the structural model, we assume that people in the "disabled" and "bad" health categories share the same

[11]Because we can assign people to good or bad health at any age, the data we use to estimate the two-state models encompass a broader age range than is used in the structural model.

(total) medical expense distributions.

Ages 50 → 51: Three states → three states				
	Next Year			
Current Year	Disabled	Bad	Good	Deceased
Disabled	95.4	0.9	0.5	3.1
Bad	10.8	64.7	21.2	3.3
Good	0.3	4.3	94.9	0.6
Ages 60 → 61: Three states → three states				
	Next Year			
Current Year	Disabled	Bad	Good	Deceased
Disabled	92.8	1.4	0.7	5.0
Bad	3.9	72.5	20.1	3.5
Good	0.5	6.2	92.8	0.5
Ages 64 → 65: Three states → two states				
	Next Year			
Current Year	Bad	Good	Deceased	
Disabled	62.8	31.0	6.2	
Bad	78.7	17.5	3.8	
Good	5.8	93.2	1.0	
Ages 70 → 71: Two states → two states				
	Next Year			
Current Year	Bad	Good	Deceased	
Bad	77.1	15.3	7.6	
Good	8.6	90.0	1.4	
Ages 80 → 81: Two states → two states				
	Next Year			
Current Year	Bad	Good	Deceased	
Bad	73.4	12.9	13.6	
Good	12.5	83.5	4.0	

Table 1: Health Transition Probabilities

Table 1 shows transition probabilities for selected ages. As people age, good health becomes less persistent, and mortality rates rise. Disability is very persistent.

6.3 The MEPS dataset

An important limitation of the HRS data is that it only contains data on out-of-pocket medical spending and lacks information on other payors of medical care, such as Medicaid, Medicare and private health insurance. Although there there are some self-reported survey data on total billable medical expenditures in the HRS, these data are mostly imputed, and are considered to be of low quality. To circumvent this issue, we use data from the 1996-2012 waves of the Medical Expenditure Panel Survey (MEPS).

The MEPS is a nationally representative survey. Respondents are asked about health status, health insurance, and health care expenditures paid out-of-pocket, by Medicaid, by Medicare, private insurance and by other sources. The MEPS data are matched to information provided by providers. Although it does not capture certain types of medical expenditures, such as nursing home expenditures, it captures most of the sources of medical spending that are faced by individuals in their 50s and 60s. Sing, Banthin, Selden, Cowan, and Keehan, (2006) and Pashchenko and Porapakkarm, (forthcoming) provide extended comparisons between the MEPS data and the aggregate statistics.

MEPS respondents are interviewed up to 5 times over a 2 year period, forming short panels. We aggregate the data to an annual level. We use the same sample selection rules in the MEPS that we use for the HRS data. Specifically, we keep only men (although we also keep information on spouses of married men) ages 51 and older. drop those who were observed to be married over the sample period, work, or be younger than 72 in 1996, 74 in 1998, etc. As with the HRS data, we assign individuals a health status of “good” if self-reported health is excellent, very good or good, and are assigned a health status of “bad” if self-reported health is fair or poor.

6.4 Total Medical Spending

MEPS has data on total medical spending by all providers. We aggregate medical spending to the household level and model the mean of logged medical expenses modeled as a function of: a quartic in age, current health status, marital status, marital status interacted with age, health interacted with marital status, and health status interacted with age. We estimate these profiles using a fixed-effects estimator.

We use fixed-effects rather than OLS for two reasons. First, differential mortality causes the composition of our sample to vary with age, while we are interested in how medical expenses vary for the same individuals as they grow older. Second, cohort effects are likely

Table 2 does not include insurance premia. We describe insurance premia in section 6.5 below.

	Younger than 65			
	Individual		Household	
	Total	Out-of-pocket	Total	Out-of-pocket
Mean	5,590	990	10,310	1,860
Median	1,860	440	4,780	1,060
90^{th} percentile	12,670	2,420	24,030	4,370
95^{th} percentile	22,450	3,670	38,470	6,130
	65 and Older			
	Individual		Household	
	Total	Out-of-pocket	Total	Out-of-pocket
Mean	8,640	1,370	13,750	2,180
Median	3,690	720	6,900	1,310
90^{th} percentile	21,250	3,190	32,770	5,000
95^{th} percentile	34,440	4,620	48,660	7,000

Table 2: Distribution of Medical Spending, both Total and Out-of-Pocket, by Age, Individual versus Household

6.5 Health Insurance and Medical Expenses

We assign individuals to one of four mutually exclusive health insurance groups: *retiree*, *tied*, *private* and *uninsured*, as described in Section 3. In addition, they can have *Medicaid* or *Medicare* coverage if they are disabled. We allow for the fact that many people have *Medicaid* or *Medicare* coverage in addition to other coverage they might have. Both the HRS and MEPS have their own advantages for understanding the effect of health insurance. MEPS has better information on the copays and premia of different types of insurance. HRS has better information to understand the impact of health insurance on savings and labor supply. In both datasets individuals are asked similar questions, and we code the data similarly in both the HRS and MEPS.

Our interest is in understanding how insurance affects the male head of household within a family. However, many of these male heads are married. A head of household may be uninsured although his spouse may receive insurance from her own employer, for example. To address this issue we aggregate medical spending variables to the household level, so that we can focus on household medical expense risk, but use the head's insurance status. For this

	EPHI	EPHI - Medicare	EPHI- Medicaid	EPHI- Medicare- Medicaid	Private	Medicare	Medicaid	Medicare- Medicaid
				Under 65				
Constant	1,237	2,132	820	986	2,907	1,299	182	162
Married	1,210	826	1,109	260	2,069	503	63	60
Predicted medical spending					0.11			
				Over 65				
Constant		2,027		640		1,447		59
Married		1,478		812		1,415		314

Notes: EPHI includes both Retiree and Tied coverage. For those over age 65, we assume everyone is covered by Medicare, so the EPHI EPHI-Medicaid, Private, and Medicaid entries are empty.

Table 5: Insurance Premia

or complete community rating. Furthermore, Buchmueller and DiNardo, (2002) and Herring and Pauly, (2006) show that even in states that did not mandate community rating, higher expected medical spending only leads to modestly higher insurance premia. Furthermore, Herring and Pauly, (2006) show that this is not due to selection issues coming from people being denied coverage when facing high medical expenses. Perhaps unsurprisingly, we found that predicted medical expenses has little impact on insurance premia for those with employer or government provided insurance. For this reason, and for parsimony, we set these coefficients to zero.

To better understand the parameters in Tables 4 and 5, Figure 2 uses the parameters to show predicted out-of-pocket medical spending for married households with different different levels of medical spending, summing over both insurance premia and copays. Of special interest is the difference between the budget line for the uninsured versus the budget line for those with private non-group insurance. The budget line shows that for households with less than $17,000 in total medical spending, choosing to be uninsured is cheaper than choosing to be insured in the private non-group market. However, many households with lower total medical spending will select private insurance, since total medical spending is uncertain at the time of selection of insurance. Many of those whose expect medical spending above $17,000 ex ante will wind up with medical spending below this level ex post. Furthermore, risk aversion implies that those with expected medical spending below $17,000 may still purchase insurance to insure themselves against the risk of higher medical spending. Likewise, many households with total medical spending over $17,000 will select to be uninsured, for two reasons. First, they may have expected medical spending that is lower than realized. Second, if they have low assets and income, their copays will be covered by the consumption floor. Thus they will use the implicit insurance from the consumption floor as a substitute for private insurance.

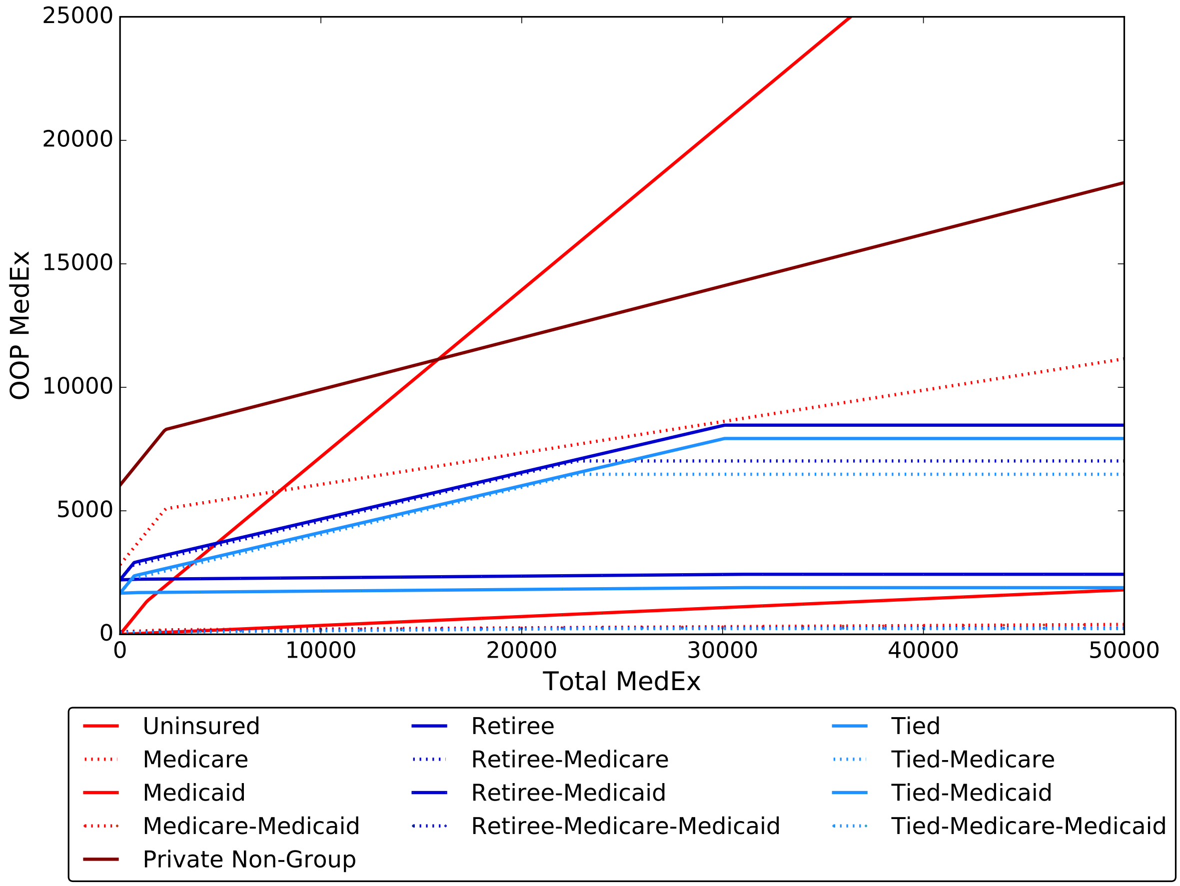

Figure 2: BUDGET SETS BY HEALTH INSURANCE TYPE, TOTAL EXPENSES UP TO $50,000

Idiosyncratic Shocks

The parameters for the idiosyncratic process ψ_t, $(\sigma_\xi^2, \sigma_\epsilon^2, \rho_m)$, are taken from French and Jones, (2004, "fitted" specification). Table 6 presents the parameters, which have been normalized so that the overall variance, σ_ψ^2, is one. Table 6 reveals that at any point in time, the transitory component generates almost 67% of the cross-sectional variance in medical expenses. The results in French and Jones reveal, however, that most of the variance in cumulative *lifetime*

Parameter	Variable	Estimate (Standard Errors)
ρ_m	autocorrelation of persistent component	0.925 (0.003)
σ_ϵ^2	innovation variance of persistent component	0.04811 (0.008)
σ_ξ^2	innovation variance of transitory component	0.6668 (0.014)

Table 6: Variance and Persistence of Innovations to Medical Expenses

medical expenses is generated by innovations to the persistent component. For this reason, estimates of the cross sectional distribution of medical expenses understate the lifetime risk of medical expenses. Given the autocorrelation coefficient ρ_m of 0.925, this is not surprising.

6.6 Pension Accrual

Our formula for pension accrual rates comes from French and Jones, (2011), who estimate them using confidential HRS pension data. Figure 3, taken from French and Jones, (2011), shows the average pension accrual rates generated by this formula when we simulate the model.

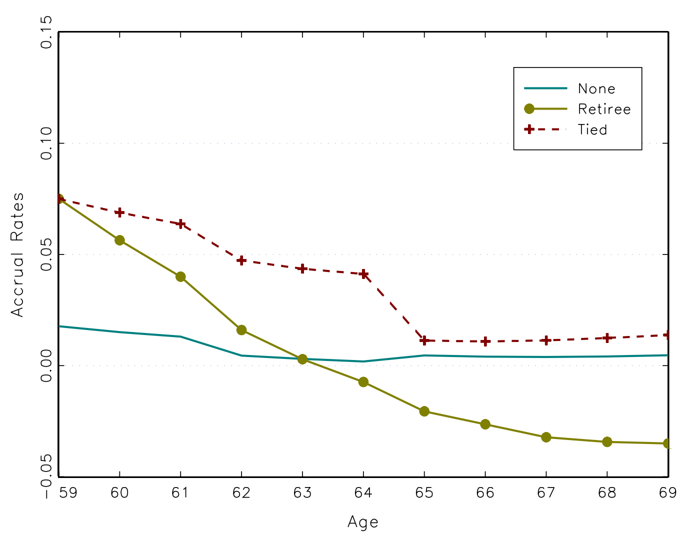

Figure 3: Average Pension Accrual Rates, by Age and Health Insurance Coverage

Figure 3 reveals that workers with *retiree* coverage face the sharpest drops in pension accrual after age 60.[13] While *retiree* coverage in and of itself provides an incentive for early retirement,

[13]Because Figure 3 is based on our estimation sample, it does not show accrual rates for earlier ages.

the pension plans associated with *retiree* coverage also provide the strongest incentives for early retirement. Failing to capture this link will lead the econometrician to overstate the effect of *retiree* coverage on retirement.

6.7 Preference Index

In order to better measure preference heterogeneity in the population (and how it is correlated with health insurance), we estimate a person's "willingness" to work using three questions from the first (1992) wave of the HRS. The first question asks the respondent the extent to which he agrees with the statement, "Even if I didn't need the money, I would probably keep on working." The second question asks the respondent, "When you think about the time when you will retire, are you looking forward to it, are you uneasy about it, or what?" The third question asks, "How much do you enjoy your job?"

To combine these three questions into a single index, we regress wave 5-7 (survey year 2000-2004) participation on the response to the three questions along with polynomials and interactions of all the state variables in the model: age, health status, wages, wealth, and AIME, medical expenses, and health insurance type. Multiplying the numerical responses to the three questions by their respective estimated coefficients and summing yields an index. We then discretize the index into three values: *high*, for the top 50% of the index for those working in wave 1; *low*, for the bottom 50% of the index for those working in wave 1; and *out* for those not working in wave 1.

6.8 Wages

Recall from equation (11) that $\ln W_t = \alpha \ln(N_t) + W(H_t, t) + \omega_t$. Following Aaronson and French, (2004), we set $\alpha = 0.415$, which implies that a 50% drop in work hours leads to a 25% drop in the offered hourly wage. This is in the middle of the range of estimates of the effect of hours worked on the offered hourly wage.

We estimate $W(H_t, t)$ using the methodology described in section 5.3.

Estimates that include the validation sample show, however, that those with retiree coverage have the highest pension accrual rates in their early and middle 50s.

The parameters for the idiosyncratic process ω_t, (σ_η^2, ρ_W) are estimated by French, (2005). The results indicate that the autocorrelation coefficient ρ_W is 0.977; wages are almost a random walk. The estimate of the innovation variance σ_η^2 is 0.0141; one standard deviation of an innovation in the wage is 12% of wages.

6.9 Remaining Calibrations

We set the interest rate r equal to 0.03. Spousal income depends upon an age polynomial and health status. Health status and mortality both depend on previous health status interacted with an age polynomial.

7 Data Profiles and Initial Conditions

7.1 Data Profiles

Figure 4 puts some of the labor market behavior that we seek to explain in relation to the health insurance status. By correctly estimating the structural parameters linking the two (and the broader environment), we will be able to predict the effects of the ACA on exit rates and participation. The top panel of Figure 4 shows empirical job exit rates conditional on the initially observed health insurance type. Recall that Medicare should provide the largest labor market incentives for workers that have *tied* health insurance. If these people place a high value on health insurance, they should either work until age 65, when they are eligible for Medicare, or they should work until age 63.5 and use COBRA coverage as a bridge to Medicare. The job exit profiles in the top panel provide some evidence that those with tied coverage do tend to work until age 65. While the age-65 job exit rate is similar for those whose health insurance type is *tied* (17%), *retiree* (17%), or *non-group* (14%), those with *retiree* coverage have higher exit rates at 62 (20%) than those with *tied* (15%) or *non-group* (13%). At all ages other than 65, those with *retiree* coverage have higher job exit rates than those with *tied* coverage, often much higher. These values for the 1940s cohort are very similar to those reported by French and Jones, (2011) for the 1931-1936 cohort.

The low job exit rates before age 65 and the relatively high job exit rates at age 65 for those

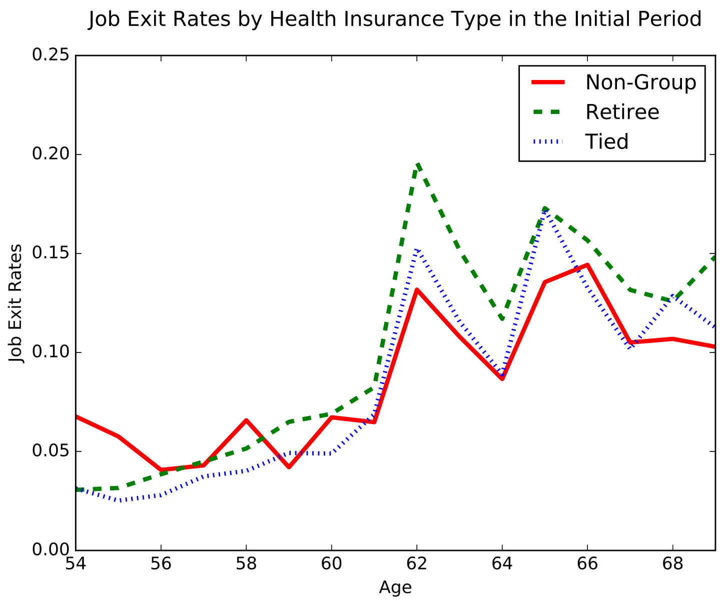

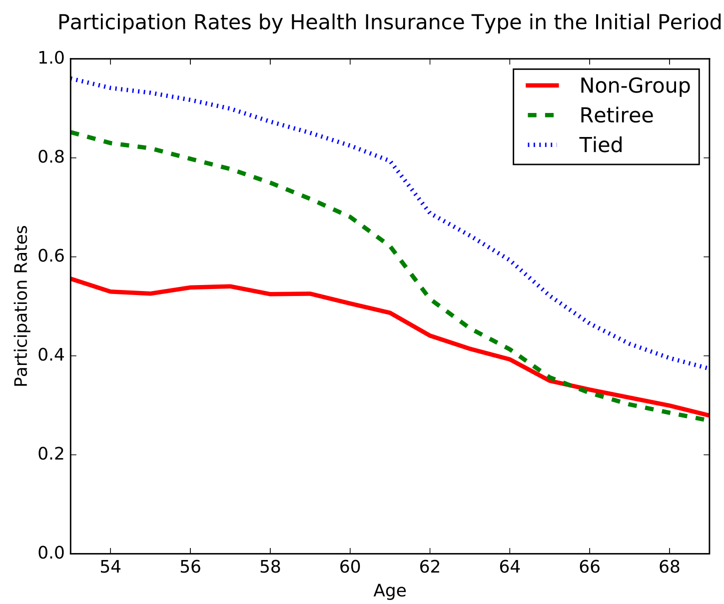

Figure 4: Job Exit and Participation (Employment) Rates, Data

Variable	Statistic	Private Non-Group	Uninsured	Medicare	Medicare-Medicaid
Age	Mean	53.5	53.3	53.7	54.1
	Median	53.0	53.0	54.0	54.0
	Std.Dev.	2.2	2.1	2.4	2.2
AIME / 1000	Mean	13.6	9.5	13.4	11.6
	Median	12.8	7.4	7.8	5.4
	Std.Dev.	7.5	7.0	18.3	26.7
Assets / 1000	Mean	768.7	144.5	136	24.1
	Median	375.6	35.4	41.2	1.4
	Std.Dev.	1341.8	347.5	322.2	170.1
Pension Wealth / 1000	Mean	50.1	21.9	21.1	4.0
	Median	0.3	0.0	0.0	0.0
	Std.Dev.	130.7	87.8	58.6	15.5
Works	Fraction	0.87	0.7	0.11	0.09
Wage if working	Mean	25.8	17.3	18.5	24.8
	Median	21	13.3	9.9	16.7
	Std.Dev.	16.5	12.5	15.9	18.6
Hours if working	Mean	2507.0	2240.0	1261.1	1505.7
	Median	2184.0	2080.0	1040.0	1742.0
	Std.Dev.	954.9	907.0	799.0	740.1
Total health care costs / 1000	Mean	24.93	4.93	24.37	132.58
	Median	10.30	.90	1.97	39.04
	Std.Dev.	29.36	21.32	51.94	335.89
OOP health care costs / 1000	Mean	6.74	3.60	4.41	1.37
	Median	3.94	0.90	1.97	0.25
	Std.Dev.	6.38	14.43	6.95	2.06
Good health	Fraction	0.87	0.69	0.0	0.17
Pref. index 0	Fraction	0.13	0.3	0.89	0.91
Pref. index 1	Fraction	0.75	0.57	0.08	0.09
Pref. index 2	Fraction	0.13	0.13	0.03	0
Married	Fraction	0.75	0.65	0.56	0.34
Not married	Fraction	0.25	0.35	0.44	0.66
Observations	Count	102	265	63	70

Table 8: Summary Statistics for the Initial Distribution of Individuals without Employer-Provided Health Insurance

Note: Source: HRS data. Individuals included in this table are those in the column called "No Employer Coverage" in Table 7. Total health care costs are imputed from out-of-pocket health care costs by inverting the budget sets described in 6.5, estimated off MEPS data.

8 Conclusion

The Affordable Care Act (ACA) is the most significant reform to the health care sector in since the 1960s. The ACA's provisions fall into four main categories: (1) an expansion of Medicaid; (2) an overhaul of private non-group insurance, including community rating, coverage standards, the introduction of exchanges, subsidies, and purchase mandates; (3) a mandate for large employers to offer health insurance coverage, and subsidies for smaller employers; (4) miscellaneous provisions including reforms to coverage standards, the tax code, and the management of Medicare.

In this paper, we consider the following two sets of provisions. First, the ACA expands Medicaid eligibility for low-income households younger than 65. Prior to the ACA, low-income households nearing retirement qualified for Medicaid only if they were disabled. Moreover, under the ACA Medicaid applicants no longer face an asset test, meaning that they can qualify for Medicaid even if they hold significant wealth. The ability to carry wealth into retirement should make Medicaid more attractive for older workers. Overall, the Medicaid expansion could either increase or reduce labor supply by the elderly. Perhaps most likely, fewer people will work, as they can now qualify for Medicaid if they retire.

The second set of provisions involves non-group insurance. The ACA establishes exchanges where households without group coverage can purchase insurance. The policies offered on these exchanges must meet coverage standards, and they must be community-rated, i.e., insurers cannot price-discriminate by health. The ACA also requires uninsured households ineligible for Medicaid to purchase insurance, provides tax subsidies for most purchases, and levies penalties on those not complying. These changes should significantly alter the customer base and actuarial costs in the non-group market. Although the subsidies will allow most households to purchase non-group insurance more cheaply, healthy and/or lightly subsidized individuals may see their premiums rise. Because many workers lose their employer-provided insurance after they leave their job (and the COBRA buy-in period expires), changes in the price of non-group insurance may change their retirement decisions. Because most people will be able to buy non-group health insurance more cheaply, early retirement will probably increase. Balancing against this, the subsidies provided under the ACA will allow uninsured low-income workers to purchase cheap insurance in the non-group market. Prior to the ACA these people may have used default on medical bills as a substitute for health insurance. However, default is a good substitute for insurance only when income and assets are low. Acquiring health insurance may encourage these workers to work and save more (Hsu, 2013).

Our goal is to assess the quantitative importance of these effects. To do this, we extend the structural labor supply and retirement model in French and Jones, (2011) to account for these reforms. We extend their model by adding in a much more detailed model of medical spending and insurance. We model explicitly how different types of health insurance plans affect the premiums and coinsurance rates that households face. We use data from the Health and Retirement Study (HRS) and the Medical Expenditure Panel Survey (MEPS) to estimate the structural model. We use the MEPS data to measure current medical expenditures, as well as who pays for these expenditures (out-of-pocket, private insurance, Medicaid, etc.). We use this information to estimate a dynamic programming model of labor supply and retirement behavior where individuals face realistic medical expense risk. Upon estimating the model, we conduct counterfactual experiments, where we modify the premia and co-insurance rates, net of subsidies and penalties, that households face.

We construct a retirement model that includes health insurance, uncertain medical costs, a savings decision, a non-negativity constraint on assets and a government-provided consumption floor.

We present evidence that those who cannot keep their employer-provided health insurance when they leave their job tend to remain on their job until age 65. Those who can maintain their insurance after they leave their job tend to exit the labor market earlier. This provides evidence that access to health insurance reduces labor supply. Interestingly, however, recent evidence on Medicaid expansions suggests small if any disemployment effect of Medicaid (Levy, Buchmueller, and Nikpay, 2015).

We show differences in both total and out-of-pocket medical spending prior to the enactment of the ACA. We show that average total medical spending in MEPS is high for all groups. Perhaps surprisingly, those with no health insurance do not spend much more out-of-pocket than those who private insurance. Those uninsured receive health care through a variety of sources such as worker's compensation and default on medical bills, which we refer to as a "consumption floor", which protects low income individuals against catastrophic medical spending. Those who appear to have the highest resources appear to be those who pay the most for health care, consistent with the view that those with low resources are covered by the consumption floor, whereas those with high resources face the most medical expense risk and might have the largest labor supply responses. We choose the consumption floor to match these, and other facts. Thus we model the ACA as a change in government insurance provisions rather than the provision of insurance where none existed before.

References

Aaronson, Daniel and Eric French (2004). "The Effect of Part-Time Work on Wages: Evidence from the Social Security Rules". In: *Journal of Labor Economics* 22.2, pp. 329–352.

Blau, David and Donna Gilleskie (2006). "Health Insurance and Retirement of Married Couples". In: *Journal of Applied Econometrics* 21.7, pp. 935–953.

Blau, David M. and Donna B. Gilleskie (2008). "The Role of Retiree Health Insurance in the Employment Behavior of Older Men". In: *International Economic Review* 49.2, pp. 475–514.

Brown, Jeff and Amy Finkelstein (2008). "The Interaction of Public and Private Insurance: Medicaid and the Long Term Care Insurance Market". In: *American Economic Review* 98.5, pp. 837–880.

Buchmueller, Thomas and John DiNardo (2002). "Did Community Rating Induce an Adverse Selection Death Spiral? Evidence from New York, Pennsylvania, and Connecticut". In: *American Economic Review* 92.1, pp. 280–294.

Center on Budget and Policy Priorities (2015). "Key Facts: Cost-Sharing Reductions". `http://www.healthreformbeyondthebasics.org/cost-sharing-charges-in-marketplace-health-insurance-plans-part-2/`.

Centers for Medicare and Medicaid Services (2014). "Medicare Secondary Payer". `https://www.cms.gov/Medicare/Coordination-of-Benefits-and-Recovery/Coordination-of-Benefits-and-Recovery-Overview/Medicare-Secondary-Payer/Medicare-Secondary-Payer.html`.

Cogan, John (1981). "Fixed Costs and Labor Supply". In: *Econometrica* 49.4, pp. 945–963.

Committee On Ways And Means, U.S.House of Representatives (various years). "Green Book". `http://greenbook.waysandmeans.house.gov/`.

Congressional Budget Office (2015). "Budgetary and Economic Effects of Repealing the Affordable Care Act". `https://www.cbo.gov/publication/50252`.

De Nardi, Mariacristina, Eric French, and John Bailey Jones (2010). "Why Do the Elderly Save? The Role of Medical Expenses". In: *Journal of Political Economy* 118.1, pp. 39–75.

De Nardi, Mariacristina, Eric French, John Bailey Jones, and Jeremy McCauley (forthcoming). "Medical Spending on the U.S. Elderly". In: *Fiscal Studies*.

De Nardi, Mariacristina (2004). "Wealth Inequality and Intergenerational Links". In: *Review of Economic Studies* 71.3, pp. 743–768.

Fernandez, Bernadette (2014). "Health Insurance Premium Credits in the Patient Protection and Affordable Care Act (ACA)". Congressional Research Service, Library of Congress R41137.

French, Eric (2005). "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour". In: *Review of Economic Studies* 72.2, pp. 395–427.

French, Eric and John Bailey Jones (2004). “On the Distribution and Dynamics of Health Care Costs”. In: *Journal of Applied Econometrics* 19.6, pp. 705–721.

— (2011). “The Effects of Health Insurance and Self-Insurance on Retirement Behavior”. In: *Econometrica* 79.3, pp. 693–732.

Gourinchas, Pierre-Olivier and Jonathan A. Parker (2002). “Consumption over the Life Cycle”. In: *Econometrica* 70.1, pp. 47–89.

Gruber, Jonathan and Brigette Madrian (1996). “Health Insurance and Early Retirement: Evidence from the Availability of Continuation Coverage”. In: *Advances in the Economics of Aging*. Ed. by David A. Wise. Chicago: University of Chicago Press, pp. 115–143.

Gustman, Alan and Thomas Steinmeier (2005). “The Social Security early entitlement age in a structural model of retirement and wealth”. In: *Journal of Public Economics* 89.

Harris, Edward and Shannon Mok (2015). “How CBO Estimates the Effects of the Affordable Care Act on the Labor Market”. Congressional Budget Office Working Paper 2015-09.

Heckman, James J. and Burton Singer (1984). “A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data”. In: *Econometrica* 52.2, pp. 271–320.

Herring, Bradley and Mark V. Pauly (2006). “The Effect of State Community Rating Regulations on Premiums and Coverage in the Individual Health Insurance Markets”. NBER Working Paper 12504.

Hsu, Minchung (2013). “Health Insurance and Precautionary Saving: A Structural Analysis”. In: *Review of Economic Dynamics* 16.3, pp. 511–526.

Hubbard, R. Glenn, Jonathan S. Skinner, and Stephen P. Zeldes (1994). “The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving”. In: *Carnegie-Rochester Conference Series on Public Policy* 40, pp. 59–125.

— (1995). “Precautionary Saving and Social Insurance”. In: *Journal of Political Economy* 103.2, pp. 360–399.

Jones, John and Yue Li (2016). “The Effects of Collecting Income Taxes on Social Security Benefits”.

Kahn, James (1988). “Social Security, liquidity, and early retirement”. In: *Journal of Public Economics* 35.1, pp. 97–117.

Keane, Michael P. and Kenneth I. Wolpin (1997). “The Career Decisions of Young Men”. In: *Journal of Political Economy* 105.3, pp. 473–522.

Levy, Helen, Thomas Buchmueller, and Sayeh Nikpay (2015). “The Effect of Health Reform on Retirement”. MRRC Working Paper.

Marken, Stephanie (2016). “U.S. Uninsured Rate at New Low of 10.9% in Third Quarter”. `http://www.gallup.com/poll/196193/uninsured-rate-new-low-third-quarter.aspx?version=print`.

Mulligan, Casey B. (2013). "Average Marginal Labor Income Tax Rates under the Affordable Care Act". NBER Working Paper 19365.

Olsen, Anya and Kathleen Romig (2013). "Modeling Behavioral Responses to Eliminating the Retirement Earnings Test". In: *Soc. Sec. Bull.* 73, p. 39.

Pashchenko, Svetlana and Ponpoje Porapakkarm (forthcoming). "Medical Spending in the US: Facts from the Medical Expenditure Panel Survey Dataset". In: *Fiscal Studies*.

Rust, John and Christopher Phelan (1997). "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets". In: *Econometrica* 65.4, pp. 781–831.

Scholz, John Karl and Ananth Seshadri (2013). "Health Insurance and Retirement Decisions". Michigan Retirement Research Center Research Paper 2013-292.

Sing, Merrile, Jessica S. Banthin, Thomas M. Selden, Cathy A. Cowan, and Sean P. Keehan (2006). "Reconciling Medical Expenditure Estimates from the MEPS and NHEA, 2002". In: *Health Care Financing Review* 28.1, pp. 25–40.

The Henry J. Kaiser Family Foundation (2013). "Summary of the Affordable Care Act". `http://kff.org/health-reform/fact-sheet/summary-of-the-affordable-care-act/`.

van der Klaauw, Wilbert and Kenneth I. Wolpin (2008). "Social Security and the Retirement and Savings Behavior of Low-Income Households". In: *Journal of Econometrics* 145.1-2, pp. 21–42.

A Cast of Characters

Preference Parameters		**Health-related Parameters**	
γ	consumption weight	H_t	health status
β	time discount factor	Z_t	total medical expenses
ν	coefficient of RRA, utility	I_t	employer-provided HI type
θ_B	bequest weight	$z(\cdot)$	mean shifter, logged medical expenses
κ	bequest shifter	$\sigma(\cdot)$	volatility shifter, logged medical expenses
C_{min}	consumption floor	ψ_t	idiosyncratic medical expense shock
L	leisure endowment	ζ_t	persistent medical expense shock
ϕ_H	leisure cost of bad health	ϵ_t	innovation, persistent shockk
ϕ_{Pt}	fixed cost of work	ρ_m	autocorrelation, persistent shock
ϕ_{P0}	fixed cost, intercept	σ_ϵ^2	innovation variance, persistent shock
ϕ_{P1}	fixed cost, time trend	ξ_t	transitory medical expense shock
ϕ_{RE}	re-entry cost	σ_ξ^2	variance, transitory shock
		M_t	out-of-pocket medical expenses
Decision Variables		**Wage-related Parameters**	
C_t	consumption	W_t	hourly wage
N_t	hours of work	$W(\cdot)$	mean shifter, logged wages
L_t	leisure	α	coefficient on hours, logged wages
P_t	participation	ω_t	idiosyncratic wage shock
A_t	assets	ρ_W	autocorrelation, wage shock
B_t	Social Security application	η_t	innovation, wage shock
I_t^+	health insurance type	σ_η^2	innovation variance, wage shock
Financial Variables		**Miscellaneous**	
$Y(\cdot)$	after-tax income	s_t	survival probability
τ	tax parameter vector	$pref$	discrete preference index
r	real interest rate	X_t	state vector, worker's problem
ys_t	spousal income	$\lambda(\cdot)$	compensating variation
Υ	spousal income indicator	SP	spouse indicator
ss_t	Social Security income	T	number of years in GMM criterion
$AIME_t$	Social Security wealth		
pb_t	pension benefits		

Table 9: Variable Definitions, Main Text

A.1 Health Insurance

Table 10 presents the health insurance state transitions that we allow for. These transitions depend on work status, age, and whether the individual is eligibile for the DI (Disability Insurance) or SSI (Supplemental Seccurity Income) programs. The last column of the table presents the payment sources and, if applicable, whether the individual has the choice between purchasing private non-group insurance and staying uninsured. We allow for the most common health insurance transitions observed in our data, and do not allow for transitions that are so uncommon empirically that we are unable to estimate the budget set parameters for these groups. Hence, there are a few things to note about Table 10:

- People will hold on to *Retiree* Health Insurance until their death, unless they become eligible for Medicaid by falling into the categorically needy category. Although in principle individuals could keep employer provided coverage while drawing Medicaid, our estimated insurance premia and co-insurance rates suggest that Medicaid alone is less expensive than Medicaid plus employer provided coverage since Medicaid pays for the same care private insurance pays for, but Medicaid does not pay the private insurance premia. Furthermore. eligibility requirements for Medicaid are so strict that these people would have to spend a very large fraction of their income on *Retiree* insurance premia. Reassuringly, very few people hold both empolyer provided and Medicaid insurance simulataneously.

- Closely related to the above case, individuals loose their *Tied* coverage if they become eligible for Medicaid. Again, very few people in the data simultaneously hold both employer provided insurance and Medicaid. Being eligible for Medicaid while working requires a very low wage or a very low number of hours, and individiuals with low hours or wages usually do not receive employer-provided coverage.

- Before age 65, the combination of *Tied* insurance and Medicare is impossible in the model because of individuals must be out of work for one year to be eligible for Disability Insurance.

- Everybody becomes eligible for Medicare at age 65. We do not model supplemental coverage. Hence, the "Uninsured" and "Private Non-Group" categories disappear after age 65. The same holds true for standalone Medicaid.

Table 10: Health Insurance State Transitions

I_{t-1}	$P_{t-1} = 1$	I_t	t	H_t = disabled	cat. needy Y_t, A_t	Payment sources
retiree	.	retiree	< 65	no	.	R
				yes	no	R + MC
			≥ 65	.	no	R + MC
		non-group	< 65	yes	yes	(MC +) MA
			≥ 65	.	yes	MC + MA
tied	yes	tied	< 65	no	.	T
			≥ 65	.	no	T + MC
		non-group	≥ 65	.	yes	MC + MA
	no	non-group	< 65	no	.	{U, P}
				yes	no	MC
					yes	(MC +) MA
			≥ 65	.	no	MC
					yes	MC + MA
non-group	.	non-group	< 65	no	.	{U, P}
				yes	no	MC
					yes	(MC +) MA
			≥ 65	.	no	MC
					yes	MC + MA

Legend for payment sources:

R	Employer's retiree plan	$P_t = 1$ if working
T	Employer's tied plan	I_t = insurance type
U	Uninsured	
P	Privately purchased insurance plan	
MC	Medicare	
MA	Medicaid	

B Key Changes to the Model for the Situation pre-ACA (relative to French and Jones, 2011)

B.1 Better Modeling of Medical Spending

We change from:

$$\ln M_t = m(H_t, I_t, t, P_t) + \sigma(H_t, I_t, t, P_t) \cdot \psi_t$$

to

$$\begin{aligned}
\ln Z_t &= \mu_z(H_t, SP_t, t) + \sigma_z(H_t, SP_t, t) \times \psi_t \\
M_t &= \mathit{premium}(I_t^+,\, t,\, P_t,\, \widehat{Z}_t,\, SP_t) + \mathit{copay}(I_t^+,\, Z_t), \\
\widehat{Z}_t &= \mathbb{E}[Z_t \,|\, t, H_t, \zeta_{t-1}].
\end{aligned}$$

where Z_t denotes **total** medical expenses, *premium*$(\cdot)$ is the health insurance premium, and the function *copay*$(\cdot)$ determines how much of Z_t is assigned to the individual via co-payments and deductibles.

We will estimate the parameters of these functions using MEPS.

B.2 Health States and their Transitions

Health can take on the following possible values: good, bad or disabled. Because we use both the HRS and MEPS, we exploit measures that exit in both datasets. We assign individuals a health status of "good" if self-reported health is excellent, very good or good; and we assign a health status of "bad" if self-reported health is fair or poor. "Disabled" is identified by an indicator equal to 1 if the individual is receiving Medicare and/or Medicaid benefits and is younger than 65, regardless of self reported health. We use this measure of disability because we wish to capture both the cash transfers, and even more importantly, the Medicare or Medicaid insurance received by the disabled. Unfortunately, however, this measure of disability status is missing for ages 65 and older since virtually everyone becomes Medicare eligible at age 65, and at the same age disability benefits are rolled into Social Security benefits. For this reason we assume that, conditional on age, those who are disabled or in bad health have the same distributions of medical spending, spousal income, and wages.

Let $H_t \in \{0,1,2,3\}$ denote death ($H_t = 0$) and the 3 mutually exclusive health states of the living (disabled = 1, bad = 2, good = 3, respectively). Let x be a vector that includes a constant, a quadratic in age, and indicators for previous health and previous health interacted with age. Our goal is to construct the likelihood function for the transition probabilities.

Prior to age 65, we allow the disabled to have different health transition probabilities than those in in bad or good health. Because we lack data on disability after age 65, we assume that at age 65 all disabled people become either dead, in bad or good health (with transition probabilities taken from the data), then after 65 nobody becomes disabled: the health states after 65 are dead, bad health and good health. Thus we must estimate three separate health transition probability models, for before 65, at age 65, and after 65. Although this causes jumps in the probability of being in either good and bad health at age 65, our estimates suggest there is no predicted jump in mortality rates around age 65.

Using a logit specification, we have, for $i \in \{1 \text{ or } ,3\}$, $j \in \{0, 1 \text{ or } 2, 3\}$,

$$\begin{aligned}
\pi_{ij,t} &= \Pr(h_{t+1} = j |\, h_t = i) \\
&= \gamma_{ij} \Big/ \sum_{k \in \{0,1,2,3\}} \gamma_{ik}, \\
\gamma_{i0} &\equiv 1, \quad \forall i, \\
\gamma_{1k} &= \exp\left(x\beta_k\right), \quad k \in \{1,2,3\}, \\
\gamma_{2k} &= \exp\left(x\beta_k\right), \quad k \in \{1,2,3\}, \\
\gamma_{3k} &= \exp\left(x\beta_k\right), \quad k \in \{1,2,3\},
\end{aligned}$$

Using a nested logit specification, we have, for $i \in \{1 or 2, 3\}$, $j \in \{0, 1 or 2, 3\}$,

$$\begin{aligned}
\pi_{ij,t} &= \Pr(h_{t+1} = j |\, h_t = i) \\
&= \gamma_{ij} \Big/ \sum_{k \in \{0,1,2,3\}} \gamma_{ik}, \\
\gamma_{i0} &\equiv 1, \quad \forall i, \\
\gamma_{1k} &= \exp\left(x\beta_k\right), \quad k \in \{1,2,3\}, \\
\gamma_{2k} &= \exp\left(x\beta_k\right), \quad k \in \{1,2,3\}, \\
\gamma_{3k} &= \exp\left(x\beta_k\right), \quad k \in \{1,2,3\},
\end{aligned}$$

where $\{\beta_k\}_{k=0}^{3}$ are sets of coefficient vectors and of course $\Pr(h_{t+1} = 0 | h_t = 0) = 1$.

The formulae above give 1-period-ahead transition probabilities, $\Pr(h_{t+1} = j | h_t = i)$. What we observe in the HRS dataset, however, are 2-period ahead probabilities, $\Pr(h_{t+2} = j | h_t = i)$. The two sets of probabilities are linked, however, by

$$\begin{aligned} \Pr(h_{t+2} = j | h_t = i) &= \sum_k \Pr(h_{t+2} = j | h_{t+1} = k) \Pr(h_{t+1} = k | h_t = i) \\ &= \sum_k \pi_{kj,t+1} \pi_{ik,t}. \end{aligned}$$

This allows us to estimate $\{\beta_k\}$ directly from the data using maximum likelihood.

B.3 Health Insurance Types

For health insurance type, we will have retiree/tied/private/self-insure/Medicaid/Medicare. We assume Medicaid+Medicare is available to everyone who is disabled. The vast majority of individuals in the 50-64 age range who are drawing Medicaid or Medicare benefits do so because of DI/SSI recipiency. We assume that all disabled people can draw DI benefits, and SSI benefits if they earn below a threshold level and their DI benefit would have been low in the absence of the benefit. Consistent with the facts (describe here), many people lose cash benefits because of work status, but do not lose their health insurance benefits.

B.4 Discrete Hours Choices

We assume that individuals can choose hours on a discrete grid. To be precise, $N_t \in \{500, 1000, 1500, 2000, 2500, 3000, 3500\}$. The main reason is that it is much easier to handle Medicaid/Medicare eligibility cutoffs in terms of earned income this way, compared to the alternative of interpolating between these values as in French and Jones, (2011).

B.5 Spousal Income

Because spousal income can serve as insurance against medical shocks, and because marital status affects eligibility for Medicaid, we include it in the model. We denote the presence of a spouse with the indicator SP_t, which equals 1 if the head is married and is 0 if he is single. For married $\rightarrow$ unmarried transitions, we do not distinguish between divorce and spousal death.

We assume that when a spouse is present, spousal income ys_t takes on two values: (i) zero; or (ii) a positive value that varies with age. With this assumption, we can collapse marital status and spousal earnings into a single variable, $\Upsilon_t \in$ {single, spouse with no income, spouse with positive income}. We assume the transition probabilities for Υ_t are logistic functions of its current value, the health of the household head, and age. We estimate the probabilities form the HRS, using the same approach to reconcile 1-year and 2-year transition probabilities that we used when estimating the health transition probabilities.

Ages 50 → 51			
	Next Year		
Current Year	Single	Spouse without income	Spouse with income
Single	97.2	0.8	0.2
Spouse without income	1.7	86.9	11.4
Spouse with income	1.6	8.7	89.7
Ages 60 → 61			
	Next Year		
Current Year	Single	Spouse without income	Spouse with income
Single	96.1	1.2	2.7
Spouse without income	1.2	83.0	15.8
Spouse with income	1.0	6.1	92.8
Ages 70 → 71			
	Next Year		
Current Year	Single	Spouse without income	Spouse with income
Single	97.5	0.5	2.0
Spouse without income	2.3	66.2	31.5
Spouse with income	1.3	2.5	96.2
Ages 80 → 81			
	Next Year		
Current Year	Single	Spouse without income	Spouse with income
Single	99.0	0.2	0.9
Spouse without income	5.4	48.4	46.2
Spouse with income	2.6	1.2	96.1

Table 11: Spousal Transition Probabilities: Household Head in Good Health

Table 11 shows transition probabilities for selected years when the household head is in good health. (The patterns are similar for all health states.) As households age, they are more likely to become single. In addition, as households age spouses without income are more likely to transition to having income. This likely reflects the initiation of Social Security or SSI benefits.

Next, we estimate mean spousal income, conditional on positive income, as a function of health status and age.