

ON THE DISTRIBUTION AND DYNAMICS OF HEALTH CARE COSTS

Eric French

Federal Reserve Bank of Chicago

230 South LaSalle Street

Chicago, IL 60604-1413

efrench@frbchi.org

John Bailey Jones*

Department of Economics, BA-110

University at Albany

State University of New York

Albany, New York 12222

jbjones@albany.edu

November 4, 2003

Abstract

Using data from the Health and Retirement Survey and the Assets and Health Dynamics of the Oldest Old survey, we estimate the stochastic process that determines both the distribution and dynamics of health care costs. We find that the data generating process for log health costs is well represented as the sum of a white noise process and a highly persistent AR(1) process. We also find that the innovations to this process can be modelled with a normal distribution that has been adjusted to capture the risk of catastrophic health care costs. Simulating this model, we find that in any given year 0.1% of households receive a health cost shock with a present value of at least \$125,000.

*We thank John Rust for encouraging us to write this paper. We also thank Meredith Crowley, Hamp Lankford, Jonathan Skinner, Dan Sullivan, and seminar participants at the Conference on Social Insurance and Pensions Research in Aarhus, Denmark, the Econometric Society, the Chicago Fed, and the NBER summer institute for useful comments. Three anonymous referees and especially editor Bent Jesper Christensen provided extremely detailed and helpful comments. Diwakar Choubey, Kate Godwin, Ken Housinger, Kirti Kamboj, Tina Lam, and Santadarshan Sadhu provided excellent research assistance. The research reported herein was funded in part by a grant from the U.S. Social Security Administration (SSA) to the Center for Retirement Research at Boston College. The views of the authors do not necessarily reflect those of the Federal Reserve System or the SSA.

1 Introduction

Despite nearly universal health insurance coverage, America's elderly still face the risk of catastrophic health care expenses (Crystal *et al.*, 2000). In response to this problem, there have been many proposals to expand the coverage provided by Medicare, America's public health insurance system for the elderly.¹ Hoping to clarify the debate on Medicare reform, numerous researchers have documented the health cost risk that older Americans face in any given year. However, this literature typically does not consider the distribution of health care expenditures over extended periods.² Using panel data from the Health and Retirement Survey (HRS) and the Assets and Health Dynamics of the Oldest Old (AHEAD) survey, we analyze both the distribution and dynamics of health care expenditures. This allows us to consider not only the risk of catastrophic expenses in a single year, but also the risk of moderate but persistent expenses that accumulate into a catastrophic lifetime cost.

We begin our analysis by examining the time series properties of health costs. To our knowledge, the only time series model of the health cost process from a nationally representative sample is that of Feenberg and Skinner (1994).³ We re-evaluate Feenberg and Skinner's results, using data that are 20 years more recent and a sample that is, in several ways, more comprehensive. We find that the autocorrelation structure of log health costs is reasonably well represented by the sum of an AR(1) component and a white noise component. The AR(1) component is quite persistent, so that health cost shocks can have a large impact on lifetime wealth. Because this sum can be rewritten as an ARMA(1,1) process, our results comport with Feenberg and Skinner's findings.

We then consider the cross-sectional distribution of health costs. Feenberg and Skinner assume that the cross section follows a lognormal distribution. Rust and Phelan (1997) argue that the right tail of the health cost distribution is better represented by a Pareto distribution. When these two distributions have the same variance, the Pareto distribution has a fatter

¹For example, Medicare does not pay for drugs or nursing home or hospital stays over 150 days.

²The persistence of health care costs has been considered more carefully in the literature investigating whether the elderly maintain high asset levels (Hubbard *et al.*, 1994; Palumbo, 1999; Dynan *et al.*, 2002) or delay retirement (Rust and Phelan, 1997; Blau and Gilleskie, 2001; French and Jones, 2003) in order to "buffer" themselves against uncertain medical expenses.

³We are aware of two other papers that consider the persistence of health costs. Eichner *et al.* (1998) non-parametrically model health costs of insured individuals from a single firm. Palumbo (1999) models the persistence of health costs coming from persistence in health status. Both models are useful contributions, but neither can be used to infer the lifetime incidence of health costs for a representative sample.

right tail than the lognormal, implying a higher probability of catastrophic health care costs. Neither set of authors, however, formally tests between the two distributions.

Using the likelihood ratio test developed by Vuong (1989), we conclude that the (truncated) lognormal and Pareto specifications fit the top decile of the health cost distribution equally well. In contrast, we find that the lognormal does a much better job of fitting the entire distribution. The lognormal distribution that best fits the overall cross section (in log-likelihood terms), however, understates the right tail of the distribution, and thus understates the risk of a catastrophic health cost shock. To address this problem, we construct the “fitted” lognormal distribution, which matches exactly the mean and the 99.5th percentile of the empirical distribution. This model fits the upper, catastrophic, portion of the health cost distribution much more closely than the standard lognormal model.

To complete our model of the stochastic process for health care costs, we need the distribution of log health cost innovations. Because the *sum* of an AR(1) and a white noise process is not a Markov process, we cannot use our time series model to back out an empirical distribution of log health cost innovations. We can, however, infer the innovation distribution from the cross-sectional distribution: if the innovations in our time series model follow a normal distribution, the cross-sectional distribution will follow a normal distribution as well. This allows us to estimate a complete model, by finding the Gaussian time series model that best matches the autocorrelation structure of log health costs and the mean and 99.5th percentile of health costs themselves. Using this model to simulate health cost histories, we find that in any given year 0.1% of households receive a health cost shock that exceeds \$125,000 in present value. This is considerably more risk than is generated by the models of Feenberg and Skinner (1994) and Hubbard *et al.* (1994).

The rest of the paper is organized as follows. In Section 2, we describe the health cost data contained in the HRS and AHEAD surveys. In Section 3, we examine the correlation of health care costs across time, and in Section 4, we examine the cross-sectional distribution. In Section 5, we estimate lifetime health cost risk by simulating our preferred health cost model. We conclude in Section 6. Additional results and discussion, including technical appendices, can be found in a companion paper that is available on line.⁴

⁴The companion paper can be found at <http://www.chicagofed.org/economists/EricFrench.cfm/> and <http://www.albany.edu/~jbjones/papers.htm>.

2 Data

We use data from the HRS and AHEAD surveys. Because the HRS and the AHEAD data are collected by the same researchers at the University of Michigan, the two data sets have similar sample designs, allowing us to merge them together. The HRS is a sample of non-institutionalized individuals, aged 51-61 in 1992, and their spouses. A total of 12,652 individuals in 7,608 households were interviewed in 1992. These individuals were interviewed again in 1994, 1996, 1998, and 2000. The AHEAD is a sample of non-institutionalized individuals, aged 70 or older in 1993. A total of 8,222 individuals in 6,047 households were interviewed for the AHEAD survey in 1993. These individuals were interviewed again in 1995, 1998, and 2000. Because the health insurance and health cost data are incomplete in wave 1 of both datasets, we use waves 2 through 5 in the analyses below.

Table 1 presents means and standard deviations of variables that measure health care costs, health insurance coverage, health care utilization, and demographic features. All of these variables are measured at the household level. To annualize the data, we divide the health cost and health care utilization measures by the number of years since the individual was last interviewed, which is, on average, two.

Virtually all Americans aged 65 and older are eligible for insurance through the government’s Medicare program. We therefore split the sample between those younger than 65 and those older than 65. Note that after age 65 nearly everyone has some form of insurance, although the fraction of individuals with employer-provided coverage falls, as many workers lose their employer-provided coverage when they leave their job. The other major form of publicly-provided health insurance is Medicaid, which is available to individuals with low income and very few assets. Those who report not having any insurance are assigned to the “none” category.

The variable of interest in this study is the total amount of health care costs paid by the household. Health care costs are the sum of what the household spends on insurance premia, drug costs, and costs for hospital, nursing home care, doctor visits, dental visits, and outpatient care.⁵ For our sample, mean household health care costs for those younger than 65 are \$2,365 and mean costs for those aged 65 and older are \$2,805. This compares to the US average of \$2,832 per capita for households headed by a non-institutionalized individual

⁵See French and Kamboj (2003) for a more detailed description of the data.

Variable	Age < 65		Age ≥ 65	
	Mean	Std. Dev.	Mean	Std. Dev.
Annual health care costs (in 1998 dollars)	2,365	(4,271)	2,805	(6,072)
Male head of household	0.64	(0.48)	0.51	(0.50)
Married	0.48	(0.50)	0.37	(0.48)
Age	58.5	(3.6)	76.9	(8.1)
No insurance (none)	0.15	(0.36)	0.01	(0.11)
Employer-provided insurance	0.61	(0.49)	0.28	(0.45)
Privately-purchased insurance	0.10	(0.31)	0.25	(0.43)
Medicaid	0.09	(0.28)	0.15	(0.36)
Medicare	0.05	(0.22)	0.31	(0.46)
Income (in 000s of 1998 dollars)	49.8	(100.5)	29.6	(45.5)
Assets (in 000s of 1998 dollars)	249.8	(756.9)	298.0	(1,113)
Annual doctor visits	6.2	(10.3)	7.2	(9.9)
Annual nursing home nights	0.7	(14.9)	8.0	(48.6)
Annual hospital nights	1.2	(4.9)	2.1	(6.8)
	$N = 15,990$		$N = 18,903$	

Table 1: SAMPLE STATISTICS

aged 65 or older (Federal Interagency Forum on Aging-Related Statistics, 2000). Note that even though most individuals are insured, there is a great deal of variation in health care costs. The standard deviation of health care costs is \$4,271 for those younger than 65 and \$6,072 for those older than 65. Although this figure is large, it is not surprising, because most health insurance plans have deductible and/or co-pay provisions. Moreover, many insurance plans (such as Medicare) do not cover prescription drugs.

One important reason why average health care costs in the HRS/AHEAD data are below the national average is that individuals in the HRS/AHEAD spend far fewer nights in a nursing home.⁶ In our sample, individuals aged 65 or older spent 8.0 nights per year in a nursing home, as opposed to the national average of 15.8 nights (National Center for Health Statistics, 1999). Because the HRS/AHEAD sample was initially drawn from the non-institutionalized population, which excludes individuals in nursing homes, this difference is not surprising. HRS/AHEAD members who enter a nursing home after the initial interview, however, are retained in the sample, and re-interviewed. In wave 5, individuals aged 65 or older spent 12.7 nights in a nursing home, which is much closer to the national average.

⁶Selden *et al.* (2001) find that 9% of total aggregate health costs and 13% of costs paid out-of-pocket arise from nursing home visits. Because of the skewness of nights spent in a nursing home, Palumbo (1999) argues that nursing homes are a significant source of health cost uncertainty for the elderly.

3 The Persistence of Health Care Costs

We begin our analysis by estimating the autocorrelation structure of log health costs. Feenberg and Skinner (1994) find that the autocorrelation structure is well represented by an ARMA(1,1) process. To re-examine their findings, we evaluate several time series models with a commonly-used error components methodology.⁷ This approach works well with short panels and it requires no distributional assumptions.

We estimate the following error components model:

$$\ln hc_{it} = X'_{it}\beta + R_{it}, \quad (1)$$

$$R_{it} = f_i + a_{it} + u_{it}, \quad (2)$$

$$a_{it} = \rho a_{it-1} + \epsilon_{it}, \quad (3)$$

$$u_{it} = \psi_{it} + \phi\psi_{it-1}, \quad (4)$$

where $X'_{it}\beta$ is the expectation of health costs conditional on the vector X_{it} , and R_{it} is the residual, which can be decomposed into: f_i , a permanent person-specific component; a_{it} , an autoregressive component; and u_{it} , a moving average component. Note that t denotes a two-year period.

The estimation procedure has two stages. In the first stage, we estimate the parameter vector β in equation (1) by regressing log health costs on the demographic and health insurance variables that households can use to forecast future health costs.⁸ Table 2 presents the parameter estimates. We use the OLS estimates throughout the paper, but the GLS estimates are very similar.⁹

Of particular interest is the coefficient on log income of 0.179. Households have some control over the quality of care they receive; ideally, the variation induced by this choice should be omitted from our measure of health cost risk.¹⁰ Most of the remaining parameter

⁷See Abowd and Card (1989) and Pischke (1995) for similar approaches.

⁸In all the analyses that follow, health care costs below \$250 (including reports of no expenditures) were recoded to \$250. One alternative bottom-coding scheme is employed by Hubbard *et al.* (1994), who drop zero-cost observations, but recode none of the others. Applying their rule would lead us to drop about 10% of our observations.

⁹The GLS estimates account for the correlation of observations using the empirical covariance matrix in Table 3. Both OLS and GLS standard errors account for the correlation of observations.

¹⁰Conversely, to the extent that income or health insurance coverage are driven by health costs, our estimates of β and health cost risk could be inconsistent.

estimates are of the expected sign. The one surprising finding is that those with no health insurance have lower health care costs than those with employer-provided insurance, even though our estimates exclude employer expenditures, which average over \$2,700 per employee (Employee Benefit Research Institute, 1999). As French and Kamboj (2002) show, one reason for this is that those with employer-provided insurance are more likely to obtain health care services. Households receiving Medicaid spend significantly less. Given that the government provides Medicaid for free to those with low income and assets, this is hardly surprising.

Variable	OLS Estimates		GLS Estimates	
	Coefficient	(S.E.)	Coefficient	(S.E.)
Male	-0.121	(0.020)	-0.106	(0.019)
Married	0.729	(0.020)	0.718	(0.019)
Age	0.0451	(0.009)	0.062	(0.008)
Age ²	-0.00020	(0.00006)	-0.00032	(0.00005)
Employer-provided × (age < 65)	0.227	(0.026)	0.269	(0.023)
Privately-purchased × (age < 65)	1.34	(0.035)	1.18	(0.031)
Medicaid × (age < 65)	-0.309	(0.039)	-0.281	(0.034)
None or Medicare × (age ≥ 65)	-0.019	(0.033)	-0.043	(0.029)
Employer-provided × (age ≥ 65)	0.186	(0.033)	0.201	(0.030)
Privately-purchased × (age ≥ 65)	0.972	(0.034)	0.844	(0.031)
Medicaid × (age ≥ 65)	-0.481	(0.037)	-0.438	(0.033)
Log income	0.179	(0.009)	0.155	(0.008)
Wave dummies included				
$N = 34,893$	$R^2 = 0.30, \sigma = 1.05$		$R^2 = 0.30, \sigma = 1.05$	

Table 2: LEAST SQUARES REGRESSIONS OF LOG HEALTH COSTS

In the second stage of the estimation procedure, we estimate the covariance matrix of the residuals from the first step regression, and fit to it the model described in equations (2)–(4). For tractability, we assume that a_{it} is a stationary process ($|\rho| < 1$ and ϵ_{it} is homoskedastic) and that the components in equations (1)–(4) are mutually orthogonal. We also assume that the person-specific effect f_i is unchanging over time, so that:

$$\text{Var}(f_i) = \sigma_f^2; \quad \text{Var}(a_{it}) = \sigma_a^2; \quad \text{Var}(\epsilon_{it}) = \sigma_\epsilon^2 = \sigma_a^2(1 - \rho^2). \quad (5)$$

We allow for heteroskedasticity, however, in the innovation to the MA(1) component:

$$\text{Var}(\psi_{it}) = \sigma_{\psi t}^2. \quad (6)$$

The variances and autocovariances implied by the model are:

$$\text{Var}(R_{it}) = \sigma_f^2 + \sigma_a^2 + \sigma_{\psi t}^2 + \phi^2 \sigma_{\psi t-1}^2, \quad \text{for all } t, \quad (7)$$

$$\text{Cov}(R_{it}, R_{i,t+1}) = \sigma_f^2 + \rho \sigma_a^2 + \phi \sigma_{\psi t}^2, \quad \text{for all } t, \quad (8)$$

$$\text{Cov}(R_{it}, R_{i,t+k}) = \sigma_f^2 + \rho^k \sigma_a^2, \quad \text{for all } k > 1, t. \quad (9)$$

Table 3 shows the empirical covariance matrix. Autocovariances appear below the diagonal of this matrix, variances appear along the diagonal, and autocorrelations appear above it. Because the data are unbalanced, Table 3 also shows the number of observations in each cell (in brackets). This covariance matrix gives us 10 moment conditions to match.

	Wave 2	Wave 3	Wave 4	Wave 5
Wave 2	1.1257 (0.0185) [7, 935]	0.3854	0.3221	0.3082
Wave 3	0.4386 (0.0220) [2, 804]	1.1504 (0.0286) [3, 938]	0.4113	0.3281
Wave 4	0.3552 (0.0171) [5, 358]	0.4585 (0.0214) [3, 037]	1.0806 (0.0168) [10, 826]	0.4139
Wave 5	0.3391 (0.0160) [5, 809]	0.3649 (0.0196) [3, 164]	0.4462 (0.0138) [8, 489]	1.0754 (0.0157) [12, 194]
Covariances lie below the diagonal, correlations above Standard errors in parentheses Sample sizes in brackets				

Table 3: EMPIRICAL COVARIANCE MATRIX

Using a minimum distance estimator, we fit several variants of the error components model to this covariance matrix. Details of the estimation procedure are in the companion paper. Table 4 shows parameter estimates and values of the overidentification test statistic.¹¹ When the model is true, this statistic will converge to a χ^2 distribution, with degrees of freedom equal to the number of moment conditions less the number of parameters.

¹¹We used optimal weighting throughout the paper. The results were similar when using the diagonal weighting matrix suggested by Pischke (1995).

The first column of Table 4 shows the results for a simple stationary AR(1) model, where $\sigma_{\psi t}^2 = \sigma_f^2 = 0$. This model is overwhelmingly rejected by the data; the overidentification test statistic is 394.9, with a p-value of 0. The reason for this failure can be seen in Table 3, which shows that while there is a large decline from the variance to the first autocovariance, the decline between the first and second (and between the second and third) autocovariances is much smaller. An AR(1) model, having a geometrically declining series of autocovariances, cannot replicate this progression.

Parameter	Model					
	1	2	3	4	5	6
σ_a^2	1.046 (0.010)	0.522 (0.019)	0.574 (0.453)	0.395 (0.043)	0.519 (0.020)	0.399 (0.044)
σ_ϵ^2	0.825 (0.014)	0.145 (0.047)	0.552 (0.513)	0.039 (0.115)	0.141 (0.049)	0.041 (0.114)
ρ	0.459 (0.009)	0.849 (0.018)	0.197 (0.186)	0.949 (0.043)	0.854 (0.018)	0.948 (0.042)
σ_{ut}^2		0.575 (0.019)	0.189 (0.470)	0.702 (0.059)	0.589 (0.027)	0.672 (0.070)
$\sigma_{\psi t}^2$		0.575 (0.019)	0.189 (0.470)	0.694 (0.039)	0.589 (0.027)	0.659 (0.049)
ϕ				0.104 (0.031)		0.102 (0.031)
σ_f^2			0.334 (0.022)			
χ^2 -statistic	394.9	18.1	10.9	10.9	7.3	0.2
Degrees of freedom	8	7	6	6	4	2
p-value	0.000	0.012	0.091	0.091	0.120	0.909
Standard errors in parentheses σ_a^2 = variance of autoregressive component of log health costs σ_ϵ^2 = innovation variance of the autoregressive component ρ = autoregressive coefficient of log health costs σ_{ut}^2 = variance of moving average component at wave t $\sigma_{\psi t}^2$ = innovation variance of the moving average component In models 5 and 6, σ_{ut}^2 and $\sigma_{\psi t}^2$ report averages across waves ϕ = moving average coefficient of log health costs σ_f^2 = variance of permanent person-specific component						

Table 4: PARAMETER ESTIMATES OF ERROR COMPONENTS MODELS

The above reasoning suggests adding a moving average component to the AR(1) model. We begin with the simplest case, setting $\phi = 0$, so that the moving average component is white noise, and assuming that this white noise component is homoskedastic across waves: $\sigma_{ut}^2 = \sigma_{\psi t}^2 = \sigma_\psi^2$. The sum of these two processes can be rewritten as the ARMA(1,1) process

studied by Feenberg and Skinner.¹² Estimates for this model are reported in the second column of Table 4. The overidentification test statistic is 18.1, implying a considerably better fit than the AR(1). Given that we have only 7 degrees of freedom, however, the model is still rejected, with a p-value of 0.012.

A common error components model of wages (see Abowd and Card, 1989, for example) includes a permanent person-specific effect, f_i , and allows the moving average component of wages to follow an MA(1) process instead of white noise. Columns 3 and 4 of Table 4 show the effects of these two changes. These models fit the data better than the AR(1) with white noise,¹³ although they are still rejected at the 10% level.

Allowing for heteroskedasticity in ψ_{it} across waves also improves goodness of fit. Such heteroskedasticity could reflect the wave-to-wave variation that exists in the survey questions used to generate the health cost measure. Results from this model are shown in column 5 of Table 4.¹⁴ Given that the empirical variance of health costs changes significantly from wave to wave, allowing for heteroskedasticity significantly improves the fit; the χ^2 statistic falls to 7.3.¹⁵ This model is not rejected at the 10% level. One last attempt to improve the fit of the model allows the moving average component of health costs to be an MA(1) with heteroskedastic innovations. Estimates are in column 6. The model does fit the data better, but introduces two additional parameters, leaving us with only two degrees of freedom.

For the exercises in Section 5 below, we use the AR(1)-plus-homoskedastic white noise time series model. Although the heteroskedastic models provide better fits, much of this heteroskedasticity likely reflects wave-specific differences in the wording of questions, rather than changing health cost risk over time. The homoskedastic model is also more parsimonious, and is more easily compared to other studies. Fortunately, the parameters ρ , σ_a^2 , and σ_{ut}^2 seem reasonably stable across models 2, 4, 5 and 6, so that all of the models have similar time series implications.

¹²See Hamilton (1994, p. 393) for a derivation.

¹³In both cases, moving from the AR(1)-plus-homoskedastic white noise model (model 2) to the more general model (model 3 or 4) reduces the χ^2 statistic by 7.2. Under the null that model 2 is correct, these decreases in χ^2 statistics are both distributed $\chi^2(1)$ (as models 3 and 4 both have one more parameter than model 2), so that the observed decrease of 7.2 has a p-value of 0.007.

¹⁴The reported estimates and standard errors for $\sigma_{\psi t}^2$ and σ_{ut}^2 are averages across waves.

¹⁵Moving from model 2 to model 5 adds three parameters and reduces the χ^2 statistic by 10.8. With a $\chi^2(3)$ distribution, this decrease has a p-value of 0.013.

4 Cross-Sectional Distribution

For the risk-averse, the possibility of catastrophic health care costs may be a matter of great concern. This means that when modelling the cross-sectional distribution of health care costs, special attention must be given to fitting the far right tail. Moreover, even if one prefers a nonparametric approach, the scarce data of the upper tail might require employing a parametric model. We thus proceed in two steps, considering first the upper tail, and then the entire distribution.

4.1 The Upper Tail

Previous studies have identified two statistical models for the upper tail of the health cost distribution. Feenberg and Skinner (1994) use the lognormal distribution. This implies that the conditional density function for large health costs, $f(\cdot)$, is

$$f(\ln hc | \ln hc \geq \ln hc_L) = \frac{1}{1 - \Phi([\ln hc_L - \mu]/\sigma)} \phi([\ln hc - \mu]/\sigma) \frac{1}{\sigma}, \quad (10)$$

where Φ and ϕ are the standard normal cdf and pdf, respectively; μ and σ are the mean and standard deviation of the untruncated distribution; and hc_L is the truncation point used to define the upper tail. Rust and Phelan (1997) use the Pareto distribution, which has the density

$$\tilde{g}(hc | hc \geq hc_L) = \gamma hc_L^\gamma hc^{-(1+\gamma)}. \quad (11)$$

A change of variables shows that if hc has a Pareto distribution, its logarithm has an exponential distribution:

$$g(\ln hc | \ln hc \geq \ln hc_L) = \gamma e^{-\gamma[\ln hc - \ln hc_L]}. \quad (12)$$

The two models can be compared formally with the likelihood ratio test developed by Vuong (1989) and extended by Rivers and Vuong (2002). Consider a sample of log health costs of size N . Let $L_N(\hat{\mu}_N, \hat{\sigma}_N^2)$ and $L_N(\hat{\gamma}_N)$ denote the maximized sample log-likelihoods for the truncated normal and exponential models, respectively. Suppose that $\frac{1}{N}\hat{\omega}_N^2$ consistently estimates the variance of $\frac{1}{N}[L_N(\hat{\mu}_N, \hat{\sigma}_N^2) - L_N(\hat{\gamma}_N)]$, the mean log-likelihood difference. Since the two models in question are strictly non-nested, it follows from Vuong (Theorem 5.1) and Rivers and Vuong (Theorems 1 and 3) that the adjusted statistic

$$D_N \equiv N^{-1/2} \left(\frac{1}{\hat{\omega}_N} [L_N(\hat{\mu}_N, \hat{\sigma}_N^2) - L_N(\hat{\gamma}_N)] - 1 \right) \quad (13)$$

will converge in distribution to a standard normal variable if the two models are equivalent. On the other hand, if the truncated normal model better represents the data generating process for log health costs, D_N will converge to infinity and the estimated p-value will converge to 0; if the exponential model is better, D_N will converge to negative infinity and the p-value to 1.

To perform the Vuong test, we treat our panel of health cost data as a single cross section.¹⁶ To account for the effects of age, gender, marital status, income, wave, and health insurance type, we repeat the linear regression shown in Table 2, compute the residuals, and add back the mean.¹⁷ The first column of Table 5 presents parameter estimates, log-likelihoods, and p-values of the Vuong statistic D_N for the top decile of this modified cross section. Recalling the discussion above, the p-value of 0.266 shown on the bottom line of Table 5 suggests that the two models are roughly equivalent.

Item	Entire Sample	Wave 5
90th percentile of log health costs ($\ln hc_{0.9}$)	8.41	8.47
Number of observations in top decile (N)	3,490	1,220
Truncated Normal		
$\hat{\mu}_N$	-24.13 (0.2224)	-24.29 (0.1567)
$\hat{\sigma}_N^2$	21.70 (0.3117)	22.31 (0.1275)
Log-likelihood	-1,938.2	-707.8
Exponential		
$\hat{\gamma}_N$	1.56 (0.0264)	1.52 (0.0436)
Log-likelihood	-1,938.7	-707.4
p-value of the Vuong test statistic D_N	0.2664	0.7925
Standard errors in parentheses		

Table 5: PARAMETER ESTIMATES AND LOG-LIKELIHOOD VALUES FOR THE TOP DECILE

¹⁶Although a household's health care costs are correlated across waves, we can calculate the likelihood values as if the observations were independent—the Vuong test is valid even if both of the competing models are misspecified. The variance estimate $\hat{\omega}_N^2$ must be calculated, however, in a way that captures this correlation; see the companion paper for details.

¹⁷We also controlled for the conditioning variables by breaking the data into cells by age, marital status, and health insurance type and repeating the analysis of this section for each cell. The companion paper contains these detailed results, which are qualitatively similar to the results presented here.

Recall that the wave 5 data may do a better job of capturing nursing home costs, which could skew the health cost distribution to the right. We therefore repeat the conditioning regression (with wave variables omitted) and the estimation with the wave 5 data alone. Perhaps not surprisingly, the second column of Table 5 shows that the fatter-tailed exponential model better fits the wave 5 data, although the difference is not significant at standard levels.

4.2 The Entire Cross Section

Although the Pareto and lognormal models fit the upper tail of the empirical health cost distribution equally well, the overall cross section allows us to discriminate between the two. Figure 1 shows the cross-sectional distribution for the entire sample. Once the effects of the conditioning variables have been removed, the empirical distribution is fairly close to lognormal.

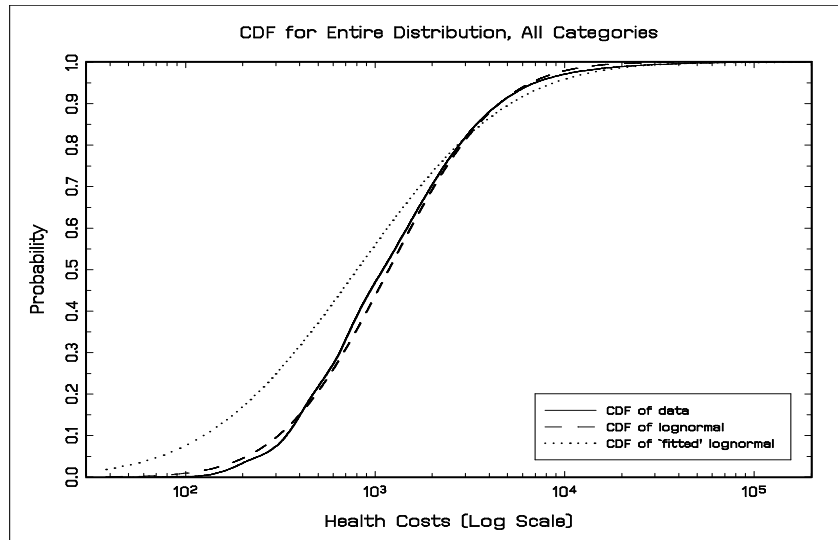


Figure 1: DISTRIBUTION OF HEALTH CARE COSTS

This conclusion is reinforced by the first four columns of Table 6, which show that the “standard” Pareto distribution, the one estimated on the entire cross section (with hc_L set to the sample minimum), is markedly inferior to the “standard” lognormal model in every dimension. Even if we could somehow extend a Pareto model of the top decile to the entire

health cost distribution, we would have difficulty incorporating it into the time series models estimated in section 3.¹⁸ While a stationary ARMA process with normal innovations (that is common across households) will generate a normally-distributed cross section, to our knowledge there is no closed-form innovation distribution for log health costs that generates an exponential cross section. These concerns lead us to abandon the Pareto as a model of the overall cross section.

But even though the lognormal model fits the overall distribution fairly well, as shown in Figure 1, it does not fit the right tail very well. As the graph of the top decile displayed in Figure 2 shows, the standard lognormal model significantly understates the right tail, and thus understates the possibility of catastrophic health costs. In contrast, the truncated lognormal and truncated Pareto models fit the upper tail almost perfectly.




Figure 2: CONDITIONAL DISTRIBUTION OF THE TOP DECILE OF HEALTH CARE COSTS

This concern leads us to an alternative estimator. We find the mean and variance of the lognormal distribution that matches both the mean and 99.5th percentile of health costs. In

¹⁸Although Rust and Phelan (1997) fit a Pareto distribution to the upper tail of their data, they use a discrete approximation to model the rest of the distribution. Because they assume health care costs are independent across time, this bifurcated model does not restrict their structural estimation.

particular, we pick values $\tilde{\mu}$ and $\tilde{\sigma}^2$ such that

$$e^{\tilde{\mu} + \tilde{\sigma}^2/2} = \widehat{E}(hc), \quad (14)$$

$$\Phi\left(\frac{\ln \widehat{hc}_{0.995} - \tilde{\mu}}{\tilde{\sigma}}\right) = 0.995, \quad (15)$$

where $\widehat{E}(hc)$ and $\widehat{hc}_{0.995}$ are the mean and the 99.5th percentile of the empirical cross section.

Table 6 shows the parameters of this “fitted” distribution, along with the parameters of the standard lognormal distribution, and standard errors for both sets of estimates.

Item	Standard Lognormal		Standard Pareto		Fitted Lognormal	
	Entire Sample	Wave 5	Entire Sample	Wave 5	Entire Sample	Wave 5
Log-likelihood: top decile	-2,224.1	-836.3	-4,956.6	-1,738.2	-1,997.0	-737.4
Log-likelihood: all deciles	-51,125	-17,722	-77,933	-27,234	-55,385	-19,788
Estimated value of μ	7.07	7.18	N.A.	N.A.	6.69	6.69
	(0.0056)	(0.0094)	N.A.	N.A.	(0.032)	(0.069)
Estimated value of σ^2	1.10	1.07	N.A.	N.A.	2.11	2.32
	(0.0083)	(0.0137)	N.A.	N.A.	(0.072)	(0.157)
Estimated value of γ	N.A.	N.A.	0.291	0.291	N.A.	N.A.
	N.A.	N.A.	(0.0016)	(0.0026)	N.A.	N.A.
Implied mean	2,036	2,250	$+\infty$	$+\infty$	2,300	2,558
Implied 99.5th percentile (in 000s)	17.5	18.9	3.02×10^6	3.37×10^6	33.8	40.5
Number of Observations (N)	34,893	12,194	34,893	12,194	34,893	12,194
Log-likelihoods calculated with log health costs						
Standard errors in parentheses						

Table 6: COMPARISON OF MODELS FOR THE ENTIRE DISTRIBUTION

Figure 2 shows that in addition to replicating average health costs, this fitted specification fits the far upper tail of the data distribution fairly well. As Figure 1 shows, one weakness of this specification is that it can provide a poor fit of the distribution’s lower tail. In practical terms, this is a relatively minor cost; at the lower tail of the distribution, large differences in logged health costs lead to relatively small changes in health costs themselves. Still more comparisons can be found in Table 6. The first row of Table 6 shows that the fitted distribution better fits the upper decile than the standard lognormal. The sixth and seventh rows compare the two models’ predictions of mean health care costs and the 99.5th percentile. While the fitted model matches these two statistics by construction, the standard lognormal

model often misses by a large margin.¹⁹ For the full sample, the standard lognormal implies a 99.5th percentile that is half of what is seen in the data.²⁰

A good model of the health cost distribution should be able to accurately measure the welfare losses associated with health cost uncertainty. We therefore conduct a simple numerical experiment, where we estimate the welfare loss that a household with certain health costs would experience if its health costs became uncertain. Assuming that lifetime utility is given by $V(A - hc)$, where $V(\cdot)$ is a value function and A is assets, we compute welfare losses by comparing $E(V(A - hc))$ to $V(A - E(hc))$, using three different health cost distributions: the empirical cross section; the standard lognormal; and the fitted lognormal. The results of the experiment, which we describe in some detail in the companion paper, show that while the fitted lognormal model and the empirical distribution generate similar welfare losses, the welfare losses generated by the standard lognormal are much smaller. By understating the risk of a catastrophic shock, the standard lognormal model understates the welfare cost of health care uncertainty.

Finally, it is useful to compare the parameter estimates in Table 6 to the lognormal estimates for the top decile shown in Table 5. The lognormal parameters estimated for the top decile are quite different from the lognormal parameters estimated for the full distribution, and are unlikely to fit the overall distribution very well. For example, the lognormal parameters shown on the first column of Table 5 imply that mean health care costs are less than one cent. All of these factors suggest that in terms of simultaneously matching both the overall distribution and its upper tail, the fitted lognormal provides the best approximation.

¹⁹The mean implied by the fitted distribution, \$2,300, is below the raw data mean of \$2,600, because the fitted distribution is estimated with data that have been purged of income, wave and demographic effects. Because this filtering reduced the variance of log health costs without changing their mean, health costs themselves are lower.

²⁰When the data are bottom-coded with Hubbard *et al.*'s rules (see footnote 8), the standard lognormal model fits the upper tail much better, and, moreover, is very close to the alternative model we develop here. Our alternative model, however, is estimated mostly from the upper tail, and does not rely on bottom-coding decisions.

5 Lifetime Health Cost Risk

5.1 The Annual Stochastic Process

The stochastic process for log health costs can be found by combining the time series model estimated in Section 3 with a model for the distribution of log health cost innovations. Unfortunately, the error components model that we estimate does not allow us to back out an empirical distribution of log health cost innovations, because the sum of an AR(1) and an MA process is not a Markov process. On the other hand, the lognormal model in Section 4 implies that the innovations are normal: if log health costs are normally distributed in the cross-section and follow a stationary ARMA process, then the innovations to that process must be normally distributed.²¹ Therefore, our preferred model of the health cost process combines the AR(1)-plus-homoskedastic white noise time series model (discussed in Section 3) with the fitted lognormal approximation of the cross section (discussed in Section 4).

An important limitation of this model is that the health cost data which it fits consist of two-year averages. In order to make our results comparable to other papers, we fit an annual model of log health costs to the data, using the Method of Simulated Moments. By simulating a large number of health cost histories at a one-year frequency and aggregating them into two-year data, we can find the summary statistics for two-year data implied by any set of one-year parameters. We estimate the model by finding the parameter values that come closest to replicating the mean and 99.5th percentile of health costs (as in Section 4), and the first three autocorrelations of the log health cost residuals, found in the HRS/AHEAD data. Details of our approach are in the companion paper. This “fitted lognormal” model is analogous to the fitted lognormal model derived above. We also estimate a one-year analog to the standard lognormal model, where we match mean log health costs and the variance and first three autocovariances of the log health cost residuals. Comparing the two models gives a sense of how failing to match the upper tail may lead to an understatement of health cost risk.

Table 7 presents estimates of the annual health cost process for the entire data set. As

²¹More generally, if households share a common stationary and ergodic health cost process, the unconditional health cost distribution for an individual household must equal the cross-sectional distribution across the population. A rigorous discussion of the necessary conditions for a stationary cross-sectional distribution can be found in Stokey and Lucas (1989). (Although the discussion there is couched in terms of Markov processes, one can derive stationary cross-sectional distributions for each component of our model, and then consider the sum.)

with the two-year data, the fitted estimates contain a much higher level of variance than the standard estimates. The last column of Table 7 shows overidentification test statistics, along with p-values.²²

	σ_a^2	σ_u^2	ρ	μ	χ_{10}^2 -statistic
HRS/AHEAD: Standard Lognormal	0.524	1.039	0.922	6.852	18.54
	(0.0195)	(0.0281)	(0.0100)	(0.0613)	[0.0466]
HRS/AHEAD: Fitted Lognormal	0.909	1.819	0.925	6.366	24.05
	(0.0485)	(0.0746)	(0.0034)	(0.0709)	[0.0075]
Feenberg and Skinner	0.269	0.100	0.896	N.A.	N.A.
Hubbard <i>et al.</i>	0.930	0.220	0.901	N.A.	N.A.
Standard errors in parentheses, p-values in brackets					

Table 7: PARAMETER ESTIMATES FOR ONE-YEAR HEALTH COST PROCESSES

Table 7 also includes estimates from Feenberg and Skinner (1994) and Hubbard *et al.* (1994); we know of no other studies that present estimates of these parameters. Feenberg and Skinner analyze health care costs of tax filers who deducted medical expenses, and obtain estimates of both σ_a^2 and σ_u^2 that are much smaller than ours.²³ There are several potential reasons for this difference. First, their data are from 1968-1973, when medical spending was lower and potentially less volatile. Second, they use a balanced panel in their analysis whereas we use an unbalanced panel. Given that a major reason for attrition is death, and those who die likely have higher medical expenses, Feenberg and Skinner likely underestimate the variance of health care costs. In contrast, the HRS/AHEAD survey utilizes follow-up interviews of the deceased’s survivors to obtain information on those who die.²⁴ Third, their sample consists only of individuals whose health care costs are high enough to be itemized on their income tax returns. The adjustments they make for truncation might not recover all of the underlying variance. Alternatively, the tax data could be less noisy than standard survey data. Note that their estimate of σ_u^2 is much smaller than ours, and if measurement error is transitory, datasets with less measurement error will have lower values of σ_u^2 .

²²The overidentification statistics for the fitted and standard lognormal models should be compared to each other and to the Table 4 statistics with some care, as they are calculated with different moment conditions.

²³We decompose Feenberg and Skinner’s ARMA(1,1) process into AR(1) and white noise components by utilizing the discussion in Hamilton (1994, p. 393). Neither Feenberg and Skinner or Hubbard *et al.* calculate the analog to μ . In the simulations below, we set μ so that the latter two models generate the same mean health costs as our fitted model.

²⁴When restricting our sample to those who have non-missing health care costs in all waves, the variance of health costs drops by 10% and the 99.5th percentile drops by 6%.

Hubbard *et al.* use cross-sectional data from the 1977 National Health Care Expenditures Survey and the 1977 National Nursing Home Survey to estimate the total cross-sectional variance of health care costs. Their estimated total, $\sigma_a^2 + \sigma_u^2$, is smaller than either of our estimates, perhaps because their data are 20 years older than ours and lack the health costs of those who died.²⁵ Because Hubbard *et al.* allocate total variance between σ_a^2 and σ_u^2 largely on the basis of Feenberg and Skinner’s estimates, they also attribute much more of the cross-sectional variance to the autoregressive component, σ_a^2 , than we do.

5.2 Lifetime Health Cost Risk

Using the stochastic processes described in Table 7, we can estimate the lifetime health cost risk that households face. In particular, we simulate 30-year health cost sequences for 1 million households. Each household begins at age 64 with a draw of a_{i64} from its invariant distribution and then realizes a 30-year sequence of innovations, $\{\epsilon_{it}, u_{it}\}_{t=65}^{94}$. Adding μ to these sequences of shocks and exponentiating yields a health cost history for each individual.²⁶ To measure lifetime health care costs, we discount this sequence back to age 65, using an annual interest rate of 3% and age- (but not health- or health cost-) specific mortality adjustments. Holding all other variables fixed, we then recompute the sequence with one or both of the age-65 innovations, $(\epsilon_{i65}, u_{i65})$, set to zero. The differences between the various discounted sequences give the lifetime effects of the age-65 innovations.

Table 8 shows the effects of the age-65 innovations on age-65 and lifetime health care costs. The first column of Table 8 shows results for the standard lognormal model, while the second column shows results for the fitted lognormal model. When the AR(1) and white noise innovations are considered together, the fitted lognormal implies that the lifetime cost variation induced by the age-65 shocks has a standard deviation of \$12,870. This is considerably larger than the standard deviation of the age-65 variation, \$7,990, indicating that persistence in health care costs is important. Moreover, the variation induced by the AR(1) innovation, ϵ_{i65} , has a lifetime standard deviation of \$10,440 and an age-65 standard

²⁵Although Hubbard *et al.* do not explicitly match extreme health cost events when estimating their variance, their bottom-coding decisions largely attenuate this problem. When we use Hubbard *et al.*’s bottom-coding rule (see footnote 8), the cross-sectional variance of the standard lognormal model increases from 1.56 to 2.35.

²⁶To restore the age effects that have been removed from the stochastic processes in Table 7, we let μ vary by age, using the coefficients given in Table 2. We have not attempted to account for other differences within or across individuals.

deviation of \$2,840. Although transitory shocks generate most of the cross-sectional and short-term variance, it is the persistent shocks (reflecting chronic conditions) that generate most of the lifetime health cost risk. Turning to catastrophic shocks, we find that under our fitted lognormal model, 1% of the population will receive an age-65 shock to lifetime health costs of at least \$43,500, and 0.1% will receive a shock of at least \$124,700. This is much more risk than is implied by the standard lognormal model.

	HRS/AHEAD: Standard	HRS/AHEAD: Fitted	Feenberg- Skinner	Hubbard <i>et al.</i>
Standard Deviation of Age-65 Health Care Costs (in \$000s)				
Due to ϵ_{i65}	1.19	2.84	0.61	1.59
Due to $\epsilon_{i65} + u_{i65}$	3.63	7.99	1.01	2.29
Standard Deviation of Lifetime Health Care Costs (in \$000s)				
Due to ϵ_{i65}	5.58	10.44	3.74	8.70
Due to $\epsilon_{i65} + u_{i65}$	6.57	12.87	3.82	8.86
Change in Lifetime Health Care Costs Due to $\epsilon_{i65} + u_{i65}$ (in \$000s)				
99th percentile	23.9	43.5	11.8	31.7
99.9th percentile	54.7	124.7	19.9	71.1
Total Increase in Lifetime Costs Given a \$1 Increase in Age-65 Costs				
Median ratio	\$1.55	\$1.61	\$3.01	\$3.82

Table 8: EFFECTS OF AGE-65 SHOCKS ON LIFETIME HEALTH CARE COSTS

The amount of health cost risk implied by our estimates is considerably higher than that found by Feenberg and Skinner. Redoing the simulations with Feenberg and Skinner’s parameter values, we find that the lifetime cost effects have a standard deviation of \$3,820, of which \$3,740 is attributable to the AR(1) innovation. When Hubbard *et al.*’s parameter values are used, the standard deviations rise to \$8,860 and \$8,700.

Lastly, we compute the total increase in lifetime health care costs associated with a \$1 increase in health care costs at age 65. For each simulated household we divide the change in lifetime costs generated by $\epsilon_{i65} + u_{i65}$ by the change in age-65 costs caused by the same two shocks. Taking the median of this ratio, we find that a \$1 shock to current health care costs leads to somewhere between \$1.55 and \$1.61 of total lifetime health care costs. Using Feenberg and Skinner’s parameter values and our methodology, we find that a \$1 health cost shock today leads to \$3.01 of lifetime health costs.²⁷ Using Hubbard *et al.*’s parameter values, the corresponding figure is \$3.82. Given that our estimates attribute a much smaller fraction

²⁷When mortality risk is omitted from the simulations, the lifetime effect rises from \$3.01 to \$3.62. This is very close to Feenberg and Skinner’s reported value of \$3.65.

of the variance to the autoregressive component, it is not surprising that health cost shocks have less persistent effects in our model.

6 Conclusion

Using data from the Health and Retirement Survey and the Assets and Health Dynamics of the Oldest Old survey, this paper presents estimates of the stochastic process that determines the distribution and dynamics of health care costs. We find that the data generating process for log health costs is well represented as the sum of an AR(1) and a white noise process. We also find that the innovations to the log health cost process can be modelled with a normal distribution. However, the variance of this innovation distribution and the mean for the overall process should be adjusted so that the model matches the mean and the 99.5th percentile of the empirical health cost distribution. This fitted lognormal distribution matches the right tail of the health cost distribution much better than the standard lognormal model, which understates the probability of catastrophic health costs. Simulating this fitted distribution reveals significant catastrophic health cost risk: in any given year 0.1% of households suffer a shock that costs at least \$125,000 over their lifetimes. The risk implied by our model is considerably more than is implied by previous estimates.

We conclude by pointing out six caveats to our analysis. First, to the extent that our data suffer from classical measurement error, our estimates will overstate the transitory variation in health care costs. Second, because the initial sample excluded those who were in nursing homes, we may be understating health care costs from this source, leading us to underestimate both the level and variability of health care costs. The third problem is that the quantity of health care services consumed is, to some extent, a choice. This means that households can reduce their health care costs by reducing the amount of medical services they consume. Fourth, low income, low wealth households have access to Medicaid, making health care services very inexpensive. While we have conditioned our estimates on several factors, including income and health insurance type, we might not have completely removed these two effects. Fifth, those with high health care costs often die shortly after their health cost shock. Because they die so soon, people who suffer from massive health cost shocks face less risk of being financially destitute (Pauly, 1990). Finally, we have assumed that health costs have no effect on income. However, it is likely to be the case that many of the shocks

that affect a household's health care costs also affect its members' ability and/or willingness to work.

References

- [1] **Abowd, J., and David Card** (1989), "On the Covariance Structure of Earnings and Hours Changes", *Econometrica*, 57(2), 411-445.
- [2] **Blau, D., and D. Gilleskie** (2000), "A Dynamic Structural Model of Health Insurance and Retirement", manuscript.
- [3] **Crystal, S., R. Johnson, J. Harman, U. Sambamoorthi, and R. Kumar** (2000), "Out-of-Pocket Health Care Costs Among Older Americans", *Journal of Gerontology Series B: Psychological Sciences and Social Sciences*, 55(1), S51-62.
- [4] **Dynan, K., J. Skinner, and S. Zeldes** (2002), "The Importance of Bequests and Life-Cycle Saving in Capital Accumulation: A New Answer", *American Economic Review*, 92(2), 274-278.
- [5] **Eichner, M., M. McClellan, and D. Wise** (1998), "Insurance or Self-Insurance? Variation, Persistence, and Individual Health Accounts", in D. Wise, ed., *Inquiries in the Economics of Aging*, 19-45.
- [6] **Employee Benefit Research Institute** (1999), *Health Benefits Databook*, EBRI-ERF.
- [7] **Feenberg, D., and J. Skinner** (1994), "The Risk and Duration of Catastrophic Health Care Expenditures", *The Review of Economics and Statistics*, 76(4), 633-647.
- [8] **Federal Interagency Forum on Aging-Related Statistics** (2000), *Older Americans 2000: Key Indicators of Well-Being*, U.S. Government Printing Office, available at <http://www.agingstats.gov/chartbook2000/default.htm>.
- [9] **French, E., and J. Jones** (2003), "The Effects of Health Insurance and Self-Insurance on Retirement Behavior", manuscript.
- [10] **French, E., and K. Kamboj** (2003), "The Effect of Health Insurance on Health Costs and Health Care Utilization of the Elderly", *Economic Perspectives*, 26(3), 60-72.
- [11] **Hamilton, J.** (1994), *Time Series Analysis*, Princeton University Press.
- [12] **Hubbard, R., J. Skinner, and S. Zeldes** (1994), "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving", *Carnegie-Rochester Series on Public Policy*, 59-125.
- [13] **National Center for Health Statistics** (1999), *Health and Aging Chartbook*, Hyattsville, MD.
- [14] **Palumbo, M.** (1999), "Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle", *Review of Economic Studies*, 66(2), 395-421.
- [15] **Pauly, M.** (1990), "The Rational Non-Purchase of Long Term-Care Insurance", *Journal of Political Economy*, 98(1), 153-168.

- [16] **Pischke, J-S** (1995), “Measurement Error and Earnings Dynamics: Some Estimates From the PSID Validation Study”, *Journal of Business & Economics Statistics*, 13(3), 305-314.
- [17] **Rivers, D., and Q. Vuong** (1997), “Model Selection Tests for Nonlinear Dynamic Models”, *Econometrics Journal*, 5, 1-39.
- [18] **Rust, J. and C. Phelan** (1997), “How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets”, *Econometrica*, 65(4), 781-831.
- [19] **Selden, T., K. Levit, J. Cohen, S. Zuvekas, J. Moeller, D. McKusick, and R. Arnett** (2001), “Reconciling Medical Expenditure Estimates from the MEPS and NHA, 1996”, *Health Care Financing Review*, 23(1), 161-178.
- [20] **Stokey, N., and R. Lucas** (1989), *Recursive Methods in Economic Dynamics*, Harvard University Press.
- [21] **Vuong, Q.** (1989), “Likelihood Ratio Tests for Model Selection and Non-nested Hypotheses”, *Econometrica*, 57(2), 307-333.