Uncertain Medical Expenses and Precautionary Saving
Near the End of the Life Cycle

by

Michael G. Palumbo

Department of Economics
University of Houston
Houston, TX 77204-5882
(713) 743-3824
internet: palumbo@uh.edu

May 1998

This paper introduces a dynamic, structural model of household consumption decisions in which elderly families consider the effects of uncertain future medical expenses when deciding current levels of consumption. The model with uncertain medical expenses implies a potentially important role for precautionary saving incentives to explain slow rates of dissaving among elderly Americans during retirement. Rather than just simulating the stochastic dynamic model, preference parameters are estimated using panel data on health, wealth and expenditures for retired families. The health uncertainty model predicts consumption levels closer to observed expenditures than a life cycle model with uncertain longevity. However, elderly families typically dissave their financial assets more slowly than even the baseline health uncertainty model predicts is optimal.

JEL Classifications: D91, J14, C61.
Keywords: Precautionary Saving, Life Cycle Model.

On this project, I received valuable advice from Jonathan Skinner and Steven Stern. Orazio Attanasio, Guglielmo Weber and Daniel Slesnick generously provided unpublished parameter estimates to me. Manuel Arellano and two anonymous referees made excellent suggestions to improve the paper. I appreciate receiving comments on earlier drafts from Alok Bhargava, Steven Craig, Eric Engen, Chinhui Juhn, Dan Levin, Jim Smith, David Papell and Peter Thompson. The Alfred P. Sloan Foundation and the Office of Sponsored Programs at the University of Houston helped to finance this research.

I. Introduction

The life cycle model of consumption, originally proposed by Modigliani and Brumberg (1954) and Ando and Modigliani (1963), is the dominant framework in economics for analyzing consumption, saving, and wealth accumulation. One implication of the simple life cycle hypothesis is that a retired household should divide its lifetime wealth by the number of years it expects to live and spend that amount each year.[1] Empirical research, however, finds saving behavior that is inconsistent with this simple life cycle implication. According to early studies, the elderly engage in no dissaving, but instead continue to amass wealth as they grow older (White, 1977; Mirer, 1979; Danziger, et al, 1982-1983). More recent articles report less dramatic conclusions: on average, wealth increases during the first few years of retirement and then decreases with age, although too slowly to be consistent with the simple life cycle model (King and Dicks-Mireaux, 1982; Diamond and Hausman, 1984; Hamermesh, 1984; Hurd, 1987). This underspending puzzle can be explained, at least in part, by modelling sources of uncertainty confronting the elderly.

In this paper I propose a dynamic, structural model of household consumption decisions in which elderly families consider the impact of uncertain future medical expenses when choosing their current levels of consumption. Uncertainty regarding future health care expenses, including those incurred during possible nursing home residences, effectively introduces random shocks to the pension incomes of retired households in the model. These random shocks provide the incentive for the elderly to engage in precautionary behavior with respect to current consumption. That is, households optimally maintain additional financial wealth to offset potentially large future out-of-pocket medical expenses in the health uncertainty model.[2,3] By ignoring sources of uncertainty, the standard life cycle model misses

---

[1] This implication holds strictly only when the household discounts the future at a rate equal to the real interest rate and each agent knows its horizon with certainty. If the rate of time preference exceeds the real interest rate, the family's consumption should decrease with age.

[2] It is well known that if "prudent" households do not know their future incomes with certainty, then they will reduce their current consumption levels to self-insure against possible low draws. Prudence is the term used by Kimball (1990) to describe utility functions that have a positive third derivative. Building on work by Leland (1968) and Sandmo (1970), Kimball shows that such a condition on the utility function is sufficient to imply precautionary behavior by optimizing agents who face income uncertainty.

[3] The most important feature of the model investigated in this paper is uncertainty about future out-of-

precautionary motives for household saving and, therefore, overpredicts household consumption expenditures if such motives are important.

Even in the presence of Medicare and private health insurance, out-of-pocket expenditures for health care represent an important source of risk to the elderly's wealth holdings. My calculations from the National Medical Care Expenditure Survey indicate that nearly ten percent of elderly households spend a fifth or more of their incomes on out-of-pocket medical expenses. Furthermore, these figures neglect what perhaps is the most important contributor to health care costs -- nursing home expenses. The likelihood that a typical sixty-five year old person enters a nursing home during her lifetime is forty-three percent. Once admitted, the average stay in a long term care facility exceeds one year. Because nursing home costs are virtually uninsured, admission to a long term care facility can quickly deplete one's financial wealth. In their examination of IRS tax files, Feenberg and Skinner (1994) find that two or three percent of elderly families incur medical expenses exceeding forty percent of their adjusted gross incomes.

Previous researchers have examined the potential for uncertain longevity to explain the slow rate of dissaving evidenced among elderly Americans. Davies (1981), Skinner (1985), and Engen (1993) report that modelling lifespan uncertainty helps to explain households' consumption decisions. My health uncertainty model, therefore, also includes uncertain longevity. However, Davies (1981) concludes that elderly Americans spend their financial assets much more slowly than is optimal under a life cycle model that includes only uncertain lifespans.[4] Indeed, the results I present in this paper confirm his conclusion and reveal that health uncertainty provides another important precautionary motive for consumption and saving decisions.

My analysis builds primarily on previous research by Kotlikoff (1988) and Hubbard, Skinner and Zeldes (1994, 1995), in which aspects of health uncertainty are incorporated into life cycle consumption

---

pocket medical expenses, but I use the term "health uncertainty model" throughout the paper for brevity. The primary role health status plays is to forecast future out-of-pocket medical expenses, but I also investigate utility functions that depend on health status explicitly.

[4] Neither Skinner (1985) nor Engen (1993) examines dissaving during retirement.

models. Kotlikoff (1988) examines the implications of different financing mechanisms for random health expenditures for macroeconomic saving rates with his model, largely abstracting the distribution of medical expenses observed among families in the U.S. In their simulation papers, Hubbard, Skinner and Zeldes build a life cycle model that incorporates uncertainty regarding annual earnings, medical expenses and longevity to explain the distribution of wealth holdings in the U.S. (their 1994 paper) and to study the consequences of a resource-tested Medicaid program for saving decisions by low- and middle-income families (their 1995 paper). These authors simulate their models for a given specification of household preferences, whereas I actually estimate preference parameters under alternative model specifications and compare predictions from the estimated models to actual household consumption data. My methodology, therefore, naturally allows alternative models to be compared using statistical criteria to gauge the accuracy of their predictions relative to the data.

Ample econometric literature exists about estimating structural dynamic models with discrete choice variables.[5] However, along with contemporaneous work by Lillard and Weiss (1996) and Gourinchas and Parker (1997), this paper joins Hurd (1989) by estimating a continuous control, stochastic dynamic programming model. My statistical method is to estimate the coefficient of relative risk aversion by choosing values for which the actual levels of consumption observed for a sample of elderly retired families are closest to the consumption levels predicted by the health uncertainty model. In addition, I estimate preferences under a life cycle model with uncertain longevity. I use three years of data for two separate samples (elderly couples and elderly persons living alone) from the Panel Study of Income Dynamics (PSID) to estimate many alternative model specifications.

From a policy standpoint, it is important to understand why retirees dissave slowly. If the elderly save for precautionary reasons due to medical expense risk, then changes in government health insurance programs influence their saving decisions by increasing or decreasing exposures to such risk. Identifying a model better capable of explaining the consumption decisions of elderly Americans and estimating parameters for that segment of the population allows a better understanding of the effects of reforming

---

[5] Eckstein and Wolpin (1989) provide a thorough survey of recent work.

many policies, such as Medicare and Medicaid (which finances much long term care in the U.S.).

The paper is organized as follows. Section II provides detail about how the health uncertainty model is specified. Section III briefly describes numerous sources of data used to parameterize the probability distributions for random variables in the model (details appear in the Appendix). In Section IV, I derive a procedure to estimate the preference parameters for two samples of elderly families from the PSID and discuss several estimation issues. Estimation results are presented in Section V and Section VI summarizes the paper's contributions.

II. A Life Cycle Model with Uncertain Out-of-Pocket Medical Expenses

The consumption model proposed in this paper differs from the standard life cycle model in two primary respects. First, households incur out-of-pocket medical expenses randomly each year. Second, each household is uncertain about its longevity (it does not know its date of death exactly). I estimate the probability distributions for these random variables using publicly available data sources and incorporate them in the health uncertainty model. Thus, the model allows the role of precautionary saving during retirement to be investigated under parameterizations of medical expense and lifespan risks consistent with available micro data.

In the health uncertainty model, each household finds itself in good, fair, or poor health each year. Denote by h_t = 1, 2, or 3, the discrete health status of the household in year t. The household knows h_t at the beginning of period t in this model, but future health outcomes are uncertain and household health status follows a Markov process. Transition probabilities from this year's health status to next year's depend on the family's current health status and other household characteristics.[6] Each family is assumed to know the probability distribution for next year's health outcome conditional on this year's health status.

Each period a family incurs out-of-pocket medical expenses, m_t, drawn from a distribution whose

---

[6] Section 3 of Appendix B describes how I use health data for a sample of elderly couples and single persons from the PSID to estimate parameters of the Markov health model. Among couples, h_t measures the household's health status, not just the health of the household head.

mean depends on current health status, h_t, the age of the head of the family, i_t, and other household characteristics, X_t. Health care expenses are not another consumption good subject to household choice in this model. Rather, I follow Kotlikoff (1988) and Hubbard, Skinner and Zeldes (1994, 1995) by assuming that medical expenses cover required health care and are incurred exogenously each year. Realized out-of-pocket medical expenses (in logs) are the sum of a deterministic function of observable characteristics (h_t, i_t and X_t) and an unpredictable element, denoted by ε_t. The disturbance term allows households with identical observable characteristics to incur different medical expenses. This model contains the intuitive implication that even if a family correctly guesses its health status next year, it cannot predict its out-of-pocket medical expenses with certainty.[7] Household data on out-of-pocket expenses show substantial variability even after controlling for observable household characteristics. As Section III and Appendix B describe, I use several data sources to parameterize the distribution of out-of-pocket medical expenses facing elderly Americans.

In an important simulation paper, Kotlikoff (1988) examines the consequences of health uncertainty for saving decisions under different financing schemes for health care. However, in his paper, random variables follow highly stylized probability distributions that are not based on empirical observation. Perhaps most importantly, Kotlikoff imposes the simplifying assumption that medical expenses are incurred independently over time. In concluding his paper, Kotlikoff (1988) suggests future researchers incorporate distribution functions for medical expenses estimated from micro data into the life cycle framework, which I do explicitly in this paper. My model incorporates temporally dependent health outcomes and distributions of family medical expenses based on econometric analysis of household data.

According to the health uncertainty model, families know the values of the nonstochastic variables, i_t and X_t, in all future periods. Each household knows m_t when deciding about c_t, but it does not know its

---

[7] Kotlikoff (1988) and Lillard and Weiss (1996), for example, assume that each health status outcome is associated with only one level of medical expenses. However, a household in good health overall may be involved in an accident or suffer some other injury that results in large medical expenses. On the other hand, a household seemingly in poor health may happen to be lucky, in terms of out-of-pocket expenses actually incurred this year.

future out-of-pocket medical expenses. The outcomes of the random variables, h_t and ε_t, are known at the beginning of period t, as well. Finally, the family is assumed to be able to evaluate distribution functions for m_{t+1}, m_{t+2}, ... m_L, and, therefore, can rationally account for uncertainty about future medical expenses when deciding on current consumption expenditures.

A household's problem in period t, upon observing h_t, is to choose consumption in all periods, c_τ $\{\tau = t, t+1, \ldots, L\}$, to

$$\max \ \{\nu(h_t, c_t) + E_t[\sum_{\tau=t+1}^{L} \beta^{(\tau-t)} (\prod_{j=t+1}^{\tau} s_{j-1}^{j})\, \nu(h_\tau, c_\tau)]\}, \tag{1}$$

where β is the discount factor, L is the maximum attainable age (assumed to be 100 years) and s_j^{j+1} denotes the household's probability of surviving to age j+1, conditional on having lived to period j. As detailed below, utility each period, $\nu(h_t,c_t)$, depends on health status and consumption expenditures. If b denotes the time preference rate, then β equals 1/(1+b). Because the model assumes the household is uncertain about its lifespan, utility k periods in the future is effectively discounted by the product of β raised to the power k and the product of the next k conditional survival probabilities.[8]

Utility each period is assumed to depend on current health status and consumption excluding out-of-pocket medical expenses, as follows:

$$\nu(h_t,c_t) = \delta(h_t)\, u(c_t) \ , \tag{2}$$

where $\delta(h_t)$ is a function that decreases as health becomes more poor and takes values between zero and one; and $u(c_t)$ is a utility function exhibiting constant relative risk aversion (denoted γ). I follow Kotlikoff (1988) and Viscuzi and Evans (1990) by modelling the effect of health status on the utility of consumption as a multiplicative constant. $\delta(h_t)$ is a vector of parameters measuring the "disutility" of fair and poor health outcomes, relative to good health. I experimented with several specifications of $\delta(h_t)$, however,

---

[8] Because households, not just individuals, are the decision units in my health uncertainty model, I compute appropriate survival probabilities for elderly couples using published age- and gender-specific mortality rates for individuals. This implies survival probabilities for couples exceed those for singles, at all ages, which means effective discount rates vary by family size, also. Naturally, effective discount rates vary by age because actuarial survival rates do.

incorporating health-state-dependent utility turns out not to improve the health uncertainty model's ability to fit PSID consumption data, nor does it significantly affect estimates of γ. Therefore, most of the results reported in this paper ignore health-state-dependent utility by setting all three elements (good, fair and poor health) of the disutility parameter equal to one.

In a given health state, household utility of consumption, $u(c_t)$, is represented by constant relative risk aversion, where γ denotes the coefficient of relative risk aversion. This functional form is popular in the literature because its positive third derivative implies precautionary behavior on the part of rational agents who face uncertain future incomes. Papers in the life cycle simulation literature typically define c_t as total family consumption, while most Euler equation estimation papers control for family size explicitly or model consumption per adult equivalent. I consider both definitions in the health uncertainty and life cycle models estimated here, using adult equivalence scales estimated by Slesnick (1993).

In equation (1), expectations are taken with respect to the density function of h_t and the density function for medical expenses, m_t, conditional on health. That is, expected utility during any future period is the integral over the three possible health states and all possible medical expense draws for each health outcome.

Consumption expenditures and medical expenses are financed out of financial wealth each period. Each year wealth is augmented by accrued interest, at rate r, and nonstochastic income (social security and pension payments). Financial wealth (W_t) cannot be negative in the model and consumption is bounded (slightly) above zero in the model. If a family's medical expenses in period t are so large that consumption above the minimum level cannot be attained, the family receives $\underline{c}-(W_t-m_t)$ from the government.[9] Naturally, in this case, the family owns no assets to carry into period t+1. Imposing such a stylized Medicaid program for medically-needy simply rules out the possibility of a negative asset position after medical expenses are incurred.

Conceptually, my health uncertainty model is the same as the retirement phase in the Hubbard,

---

[9] Most states set the maximum allowable resources for Medicaid eligibility at $2,000. I use this value for $\underline{c}$ in the model. Estimation results are not very sensitive to this parameter value.

Skinner and Zeldes (1995) model. Both models incorporate household longevity and out-of-pocket medical expenses as exogenous, stochastic events and neither includes explicit bequest motives. Our papers focus on the contribution of precautionary saving motives to explain observed consumption behavior. Hubbard, Skinner and Zeldes carefully model the pre-retirement phase of the life cycle, while I focus on the retirement period exclusively. Finally, my goal is to estimate the preferences under the health uncertainty model for a sample of elderly households, while Hubbard, Skinner and Zeldes simulate their model's implications for a given set of parameter values. I explicitly compare the ability of the estimated health uncertainty model to fit PSID consumption data relative to alternative life cycle model specifications using statistical criteria, rather than presume the health uncertainty model works better.

No closed-form solution exists for the optimal consumption plan $\{c_t, t = 1, 2, ... , L\}$ under the health uncertainty model (see Zeldes (1989) and Deaton (1992) for excellent discussions). Optimal consumption (the control variable) instead is calculated for a large number of combinations of the state variables: the age of the householder (or the number of periods away from L), the health status of the family, and the household's wealth minus medical expenses (discretionary financial wealth). The approximated consumption function is a matrix in which the optimal value of c_t is computed for each possible combination of i_t, h_t and W_t-m_t. Therefore, the procedure is feasible only for discrete state variables. By definition, age and the health status of the household are discrete variables, but discretionary wealth is continuous. I make W_t-m_t discrete by assuming arbitrary maximum and minimum values over which it is allowed to range and then dividing the range into a finite number of values. I settle on the grid specification through experimentation.

Appendix A describes numerical techniques to exploit the recursive nature of the dynamic programming problem to approximate optimal consumption plans. The intuition underlying the solution procedure is that a household surviving to period L has an easy decision to make: consume all remaining wealth. Knowing the period-L decision allows one to numerically integrate the expected value (at time L-1) of surviving to period L and, therefore, allows one to find the consumption decision that maximizes the period-L-1 value function. The same logic is carried back until the optimal consumption choice in the first period is computed for a large number of combinations of wealth and health. Taylor and Uhlig (1990) refer to the numerical solution algorithm as the Euler equation grid method, which is described further by Baxter, Crucini and Rouwenhorst (1990), as well as Appendix A.

III. Characterizing the Distributions of the Random Variables in the Health Uncertainty Model

Before the health uncertainty model can be solved for optimal consumption, I must parameterize the distributions for all random variables in the model: the household's future health status, its future out-of-pocket medical expenses, and the number of years it will remain alive. To accomplish this, I employ several sources of household and individual data. Appendix B describes the exact procedures and data used to estimate probability distributions needed to completely specify the health uncertainty model. Here I simply introduce the data and methods used.

I construct a longitudinal database from the PSID to estimate a Markov model for household health transitions over time. By incorporating an estimated dynamic model of health outcomes (described in Section 3 of Appendix B), my health uncertainty model captures two important characteristics of the data: the persistence of poor and good health states over time and the general deterioration of health that occurs as people age.

Section 1 of Appendix B explains how I estimate the distribution of out-of-pocket expenses for health care provided to elderly members living at home using data from the National Medical Care Expenditure Survey of 1977 (NMCES). The NMCES carefully catalogs annual expenses on health care (these are recorded as they occur during the survey year) and distinguishes between out-of-pocket expenditures and those paid for by Medicare and other third parties. Thus, NMCES data is most appropriate for my application. Several variables are used to explain the distribution of out-of-pocket medical expenses incurred for health care received by elderly persons living in the community: income, age, family size, retirement status, race and health status. A histogram of the least squares residuals is used to represent the distribution of the stochastic element of out-of-pocket medical expenses.

Because the NMCES data omits expenses incurred for health care received during periods of institutionalization, the regressions just mentioned ignore the impact of potential nursing home admissions on the distribution of out-of-pocket medical expenses. Expenses incurred during residences in intermediate care and skilled nursing facilities, however, represent a large, virtually uninsured risk to the wealth of elderly Americans. As Section 2 of Appendix B describes, results from two empirical studies allow me to parameterize the distribution of nursing home expenses across the elderly population. Cohen, Tell, and Wallack (1986) report age-specific probabilities of nursing home admissions for elderly persons. Liu and Manton (1984) estimate distribution of lengths of stay in long term care facilities for admitted patients. After correcting the Liu and Manton results for censoring bias due to nursing home deaths adjusting entry rates across poor and good health states, I combine results from the two papers to describe the probability density function for out-of-pocket medical expenses incurred for long term care that elderly Americans face.

Including potential nursing home expenses in the distribution of out-of-pocket medical expenses is important, even though (by definition) none of the elderly PSID heads of household for whom the health uncertainty model is being estimated reside in a nursing home. As Kotlikoff (1988) explains, prudent forward-looking families might respond to the possibility of a nursing home admission in the future by reducing expenditures today. Every elderly person faces the possibility of requiring an expensive nursing home admission during their lives, though most will not actually experience one. Explicitly incorporating this risk implies a precautionary motive potentially capable of explaining some of the empirical reluctance to dissave financial wealth relative to standard life cycle predictions.

Lifespan uncertainty is incorporated into the health uncertainty model by estimating the probability that one or both members of a typical elderly couple survive to the following year, conditional on having survived to the current year.[10] This is done by "aging" a hypothetical cohort of couples through their retirement years and allowing the proper number of males and females to leave the cohort each year. The proper exit frequencies come from age- and gender-specific estimates of mortality tables for individuals published by the U.S. Department of Health and Human Services (1985). Implementing standard nonparametric survival analysis to the hypothetical cohort gives the correct age-specific mortality

---

[10] Appropriately, I create new life tables for couples using published mortality rates for individuals.

probabilities for elderly households. Published life tables provide relevant survival probabilities for the sample of single elderly persons.[11]

IV. Estimating the Health Uncertainty Model by Maximum Likelihood

This section describes a novel approach to estimating the coefficient of relative risk aversion for the health uncertainty and life cycle models of consumption. I detail the econometric procedure, provide some motivation for the approach to estimation and relate the paper's methods to some other recent contributions to the literature.

My estimation procedure requires first using numerical methods to solve the health uncertainty model for a given value of the coefficient of relative risk aversion (γ) by computing optimal consumption as a function of the model's state variables -- household age, health and wealth (or, "cash on hand", in Deaton's (1992) terminology). I observe values of the state variables (age, health and wealth) for two samples of retired families in the Panel Study of Income Dynamics (PSID), so I can use the solution matrix from the health uncertainty model to predict a consumption level for each household during each year of the panel data.[12] Denote by $c^p(a_{it}, h_{it}, W_{it}; \gamma)$ the level of consumption predicted for household i during year t according to the optimal decision under the health uncertainty model parameterized with a value of γ given the family's observed age, health and wealth. Predicted consumption is a function of the coefficient of relative risk aversion (γ) because the approximated consumption function depends on household preferences according to the health uncertainty model.

In principle, effective household discount factors, which could depend on perceived mortality rates, subjective time preference, consumption equivalent scales or anticipated interest rate movements, also

---

[11] Because I am unaware of appropriate estimates, I do not adjust the published actuarial mortality rates for differences in marital status or wealth.

[12] I solve the health uncertainty and life cycle models for three different income groups (nonasset income is the appropriate definition) and use the appropriate solutions to predict consumption levels for the sample households. It might be preferable to solve the models using each family's actual income process, but that increases the computational burden substantially.

determine optimal consumption expenditures by age, health and wealth. I only have access to three years of panel data, which, in practice, turns out to be insufficient for accurately estimating effective discount factors in this context. Consequently, as described below, in this paper I examine the sensitivity of estimated risk aversion to alternative maintained assumptions about household discounting under both the health uncertainty and life cycle models of consumption.

A sample log-likelihood function for γ can be derived based on deviations between the natural log of consumption actually observed for a family, $\ln(c^a_{it})$, and the natural log of its predicted expenditure according to the health uncertainty model and given its observed state vector. Three years of data are available for each family (t = 1, 2, 3) and the difference between actual and predicted log-consumption levels can be represented by a serially correlated, normally distributed error term:

$$\begin{aligned} \ln(c^a_{it}) &= \ln(c^p(a_{it}, h_{it}, W_{it}; \gamma)) + \eta_{it}, \qquad t = 1, 2, 3, \\ \eta_i &= (\eta_{i1}, \eta_{i2}, \eta_{i3}) \sim \textit{iid}\ N((0,0,0), \Omega), \end{aligned} \tag{3}$$

where Ω is a three-by-three covariance matrix to be estimated. For household i, the trivariate normal density function describes the likelihood of observing a vector of errors over the panel, $\eta_i = (\eta_{i1}, \eta_{i2}, \eta_{i3})$. Under the standard regularity conditions and as long as the error vector η_i is uncorrelated with the function $\ln(c^p(.))$, maximizing the sample log-likelihood function with respect to the parameter γ produces consistent and asymptotically normal estimates.[13]

---

[13] A number of numerical techniques might have been employed to maximize the sample log-likelihood function with respect to the coefficient of relative risk aversion. Recall, the health uncertainty model does not yield a closed-form solution for predicted consumption as a function of either the observable state variables or the model parameters. This makes it difficult to use analytical derivatives in an optimization algorithm for the likelihood function. Furthermore, preliminary results indicated the possibility of multiple local maxima for some particular model specifications. Therefore, the preference parameter, γ, is estimated using a grid search. Experimentation using a (Nelder and Mead) simplex algorithm under GQOPT closely replicated my estimation results, but took significantly longer than my grid search algorithm to converge. Finally, I truncate grid searches at γ=25.00, which binds the estimate under some life cycle model specifications. Larger values of γ simply do not generate very different predictions for household expenditures because the numerical algorithm produces essentially the same approximation to the optimal expenditure function at such extreme parameter values.

Thus, the statistical criterion upon which estimates of γ are based in this paper come from adding a disturbance term to model-predicted consumption levels based on the optimal solution algorithm and observed information about age, assets and health for each family in the sample. At least two justifications exist for "adding an error term" to form the statistical criterion for estimation. First, as shown by MaCurdy (1985) and Altonji (1986), the existence of multiplicative preference shocks, unobserved by the econometrician, to a utility function exhibiting constant relative risk aversion results in an additive error term for optimally chosen log-consumption levels. In the context of the health uncertainty model or the life cycle model with uncertain longevity, multiplicative preference shocks could result from, for example, differences between perceived mortality rates among retirees and actuarial rates published in life tables; differences between actual adult equivalence scales for family consumption and estimated functions based on previous economic research (I use Slesnick, 1993); or, anticipated movements in real interest rates during the survey period. Second, total consumption levels are not directly observed in the PSID data I study in this paper. As described below and in detail in appendix D, I impute total consumption expenditures from the set of information directly available from the PSID survey. Thus, random errors from the imputation procedure would result in additive disturbances to log consumption in the sample. To summarize, both of these justifications for random deviations between model-predicted consumption and that observed in the micro data reasonably supports the assumption that the errors to be uncorrelated with the optimal consumption function. In these cases, γ should be estimated consistently by this paper's methods.

While the statistical formulation for estimating γ in (3) allows a general time-series correlation structure in consumption-errors for each family under the health uncertainty or life cycle models, consistent estimation requires independence among errors across families. Thus, consistent estimates do not necessarily follow in the presence of aggregate shocks to consumption expenditures, such as those operating, for example, through unanticipated macroeconomic shocks to real interest rates during the sample period. Simultaneous consideration of random income realizations and random real interest rates poses too complex a problem for numerical analysis given current computational constraints (consult, for example, Ludvigson and Paxson, 1997). Thus, a limitation in this and virtually all other papers in the precautionary saving literature follows from abstracting from the effects of unanticipated movements in real interest rates on family consumption decisions in the presence of income risk.

The sample log-likelihood value provides a metric for measuring the distance between actual and predicted consumption levels in the data -- the determinant of the covariance matrix of prediction errors.[14] Holding covariances constant, larger estimates for the error variance terms result in smaller log-likelihood values. Thus, models which predict consumption closer to observed levels on a household-by-household (and year-to-year) basis, yield smaller estimated error variances and obtain larger sample log-likelihood values. In comparing predictions from different model specifications in Section V below, I report both sample log-likelihood functions and estimated variances for the log-prediction errors under different model specifications.

Nearly all previous papers in the large body of recent microeconomic research on consumption and saving fall into one of two disjoint sets. The first set comprises econometric papers in which the objective is to estimate structural preference parameters by directly examining Euler equations consistent with intertemporal optimization and considering their ability to explain variation in consumption growth rates among samples of families.[15] The second set comprises simulation papers in which the objective is to contrast behavioral implications from specific models of intertemporal optimization to benchmark models, taking preference parameters as given, rather than estimating them directly.[16] Thus, the former set of papers achieves its econometric goals without explicitly modeling the stochastic environment families face

---

[14] The estimator does not depend critically on the prediction errors being normally distributed and, thus, could be relabeled quasi-maximum likelihood. Because maximizing the sample log-likelihood function is equivalent to minimizing the determinant of the covariance matrix, Ω, my estimation procedure is the same as Hurd's (1989) nonlinear GLS.

[15] Browning and Lusardi (1996) identify 25 relevant papers in their survey of the literature. Deaton (1992) and Attanasio and Weber (1995) thoroughly discuss many econometric issues arising in the Euler equation studies.

[16] Again, Deaton (1992) and Browning and Lusardi (1996) cite many simulation papers in their survey. Particularly useful examples appear in Zeldes (1989) and Hubbard, Skinner and Zeldes (1995).

in their lives.[17] On the other hand, in the latter set, substantial effort is allocated to modeling "realistically" the economic environment facing families, but key preference parameters are imposed from "outside" the specific modeling environment. Furthermore, the simulation papers do not use statistical criteria for evaluating the predictive performance of alternative model specifications, as the econometric papers do.

A motivation for my approach is to bridge these two parallel strands in the existing consumption literature. It involves investigating "fully specified" structural models for retirement expenditure decisions based on intertemporal optimization under uncertainty, but also estimating preference parameters using microdata and evaluating alternative model specifications using statistical criteria. Estimation is particularly important for my application because none of Euler-equation estimation papers consider retired families exclusively, as I care to. Thus, to the extent the current generation of retirees in the U.S. exhibit different preferences, or attitudes toward risk, than Americans of "prime age", appropriate parameters cannot be extracted from the existing literature.

The empirical approach of this paper receives further motivation based on results from two very recent studies. Independently and based on different specific analyses, Carroll (1997) and Ludvigson and Paxson (1997) demonstrate how specification and approximation biases in linearized Euler equations severely damage their ability to produce consistent estimates of structural preference parameters from consumption models using micro data.[18] My approach substitutes an approximate numerical solution to the optimal consumption function from a well-defined dynamic program (based on the exact Euler equation) for apparently biased Euler equations using log-linear consumption growth rates employed in many previous studies. The cost to my approach, relative to estimating linearized Euler equations, is that one must explicitly specify the stochastic environment facing households to derive the relevant intertemporal budget constraints and to numerically solve for optimal consumption rules.

In addition to Hurd (1989), two other recent contributions to the consumption literature implement

---

[17] As discussed by Browning and Lusardi (1996), essentially the same Euler equations arise from many different, and perhaps competing, stochastic specifications of underlying life cycle models.

[18] Virtually all papers employing micro data estimate linear approximations to the exact Euler equations consistent with intertemporal optimization under uncertainty.

estimation strategies similar to mine. First, Lillard and Weiss (1997) estimate parameters from a consumption model in which retired families optimally consider the impact of social security benefit rules for surviving spouses when making their spending decisions. By assuming a quadratic form for household utility, however, Lillard and Weiss (1997) ignore precautionary saving during retirement, which, given the presence of uncertain medical expenses, provides the motivation for this paper. Second, Gourinchas and Parker (1997) use simulation methods to estimate preference parameters for families during the "prime-age" phase of their life cycles. Gourinchas and Parker (1997) model intertemporal allocations between ages 25 and 65 years, but "turn off" the model during retirement to avoid the relevant issues directly faced in this paper. These four papers take the same general empirical approach to studying life cycle consumption allocations, but focus on different specific behavioral motivations for saving during different life-cycle phases and, thus, complement each other well.

To summarize at an intuitive level, the health uncertainty and life cycle models investigated in this paper imply different nonlinear functions relating the state variables (age, wealth and health) to optimal consumption levels during retirement. Furthermore, each model specification implies a nonlinear function relating state and control variables that depends on the preference parameter, γ. Thus, my empirical strategy involves selecting the value of γ for each model specification providing the best fit to the empirical function relating age, health and wealth to levels of expenditure among two samples of retirees in the PSID. Identification of γ comes from sample variation in age, health and wealth and the (implicit) functional form of the optimal consumption policy rule derived from the stochastic life cycle model. Because the empirical approach taken in this paper derives from the simulation branch of the consumption literature, it should be mentioned that I require the same set of ancillary assumptions to estimate the coefficient of relative risk aversion consistently. That is, all other restrictive elements built into the health uncertainty model (deterministic real interest rate, no bequest motive, rational expectations, etc.) play a role in estimation. In presenting the econometric results below, I analyze the sensitivity of estimates for γ with respect to many other modeling elements.

As introduced above, estimation requires micro data on financial wealth, nonasset income, consumption, health status, and the age of the household head. I construct a database for two samples of retired, elderly families from the PSID, observed during the years 1984 through 1986. The first sample contains elderly couples living together (both retired); the second contains elderly, retired individuals living alone. All families are white and none contain individuals other than the respondent and his spouse. The reason for choosing homogeneous samples is that allowing heterogeneous preferences based on additional observed characteristics in the health uncertainty model increases computational complexity (by increasing the dimension of the state vector). A natural alternative to modelling preference-heterogeneity is to select a sample of households most likely to exhibit similar preferences toward the timing of consumption and similar attitudes toward risk. I compare estimation results across the two samples of elderly couples and single individuals.

Total consumption expenditures are not directly available from the PSID, so I modify Skinner's (1987) procedure for imputing this variable to households in the sample. Skinner's idea involves using information available in the PSID (food consumption at and away from home, house value or rent, and utility expenditures) to estimate total consumption for each family based on regressions using the same variables and measured total consumption for families in the 1982-83 Consumer Expenditure Survey. In this paper, I actually undertake the empirical analysis using three different imputation procedures for PSID consumption, all described in Appendix D. First, I follow Skinner's procedure exactly; second, I modify his procedure to allow explicitly for flexible Engel-curves for food and other consumption items; third, I apply estimates from Attanasio and Weber's (1995) structural (AIDS) demand system for food and nonfood expenditures to my PSID sample. For the remainder of this paper, "actual" or "observed" consumption refers to the imputed values from the flexible-Engel-curve modification to Skinner`s procedure. As discussed in part B of Section 5 below, my estimation results are quite robust across these three different imputation procedures.

Predicted consumption from the life cycle model is attained by "turning off" health uncertainty in the model of section II. That is, what I refer to as the life cycle model in this paper simply is the health uncertainty model with future medical expenses equal to their mean values with probability one in all time periods. The only state variables for the life cycle model are wealth and age for each household. Note, all life cycle model specifications examined in this paper include uncertain household longevity, unless specifically mentioned below.

V. Estimates from Health Uncertainty and Life Cycle Consumption Models

This section reports maximum likelihood estimates of the coefficient of relative risk aversion (γ) under several health uncertainty and life cycle model specifications. The relative success of the two consumption models for predicting actual consumption levels among retired elderly families in my PSID sample is determined by comparing their maximized sample log-likelihood values and summary statistics for actual and predicted levels of (log) annual consumption expenditures. First, I discuss estimates from baseline model specifications. Second, I examine the sensitivity of estimation results to several alternative model specifications. Third, I present results from a broad specification search. Across all specifications investigated, the health uncertainty model outperforms the life cycle model for predicting consumption levels during retirement.

A. Estimates from Baseline Model Specifications

First, I report estimation results for baseline specifications of health uncertainty and life cycle models, which maintain parameter values similar to those appearing in the related simulation literature, such as Kotlikoff (1988) and Hubbard, Skinner and Zeldes (1995). Uncertain longevity is parameterized using published life table estimates for age- and sex-specific mortality rates; the real interest rate is three percent; time preference rate is zero percent; and, utility from annual consumption expenditures *per adult equivalent family member* takes the constant relative risk aversion form (health-state-dependent utility is ignored in the baseline case). Thus, the only important departure relative to the simulation literature is that the latter imposes a CRRA utility function for *total family* consumption expenditures, not spending per adult equivalent member. In section B below, I examine the sensitivity of estimation results to alternative specifications for health uncertainty and life cycle models.

Estimates for coefficients of relative risk aversion, sample log-likelihoods values, and predicted consumption levels from baseline health uncertainty and life cycle model specifications appear in Table 1. Panel A contains results from the sample of retired, elderly couples in the PSID; panel B shows results from the sample of retired, elderly individuals living alone. Across both subsamples, baseline health uncertainty specifications predict consumption levels closer to actual values than life cycle models do. Among PSID elderly couples, Table 1 indicates larger log-likelihood values for health uncertainty models (+3.94 vs. -12.70), closer predicted to actual consumption levels ($18,247 vs. $19,225) on average, and smaller average squared deviations of predicted from actual consumption levels (0.134 vs. 0.152). Among PSID sample of elderly who live alone, the best baseline health uncertainty model specifications yield larger log-likelihood values (-186.64 vs. -192.13) and closer consumption predictions (measured as average squared deviations from actual consumption; 0.230 vs. 0.243) than the life cycle model. However, on average, the health uncertainty model with γ=6.25 and the life cycle model with γ=25.00 predict (nearly) identical levels of consumption ($8,597 vs. $8,604).

Table 2 shows how the health uncertainty model yields closer predictions for actual consumption levels than the life cycle model for nearly all observable characteristics among the sample of elderly couples in the PSID.[19] According to Table 2, the baseline health uncertainty model overpredicts actual consumption levels among PSID elderly couples by about thirteen percent, on average. This average discrepancy is small relative to the average nineteen percent overprediction coming from the baseline life cycle model. Average squared deviations between log-actual and log-predicted consumption levels measure the variance of prediction errors or average "distance" between actual and predicted consumption from the two structural models. These columns of results in Table 2 also reveal the dominance of health uncertainty versus life cycle models to predict consumption among retirees in the PSID across nearly all observable characteristics (wealth, income, age and health status). In fact, Table 2 indicates only a single category of families for whom the life cycle model predicts closer consumption levels than the health uncertainty model:

---

[19] To conserve space, I report some results, such as those in Tables 2 and 3, only for the PSID couples' sample when they are qualitatively similar among the sample elderly individuals who live alone. Also, I only discuss results pertaining to 1985 consumption levels in the PSID data because they represent the same patterns as data from the other two years, 1984 and 1986.

families with wealth below $5,000. On the other hand, for all other wealth categories and all age, income and health categories, the baseline health uncertainty model overpredicts consumption levels to a lesser extent than does the baseline life cycle model.

In the baseline specifications, which are chosen for their similarity to models in the existing literature, estimated coefficients of relative risk aversion (7.00 and 24.00) far exceed values used for simulations, which typically fall between 1 and 3. Along with the fact that both dynamic structural models systematically overpredict consumption levels relative to the PSID data, this suggests the possibility of misspecification (or, at least, room for improvement) in the baseline specifications. The following section reports on the sensitivity of CRRA estimates and prediction performance of the health uncertainty relative to the life cycle model for alternative model specifications.

B. Sensitivity of Estimation Results to Alternative Model Specifications

Results from the sensitivity analysis appear in Table 3, which shows how CRRA parameter estimates and log-likelihood values vary across alternative health uncertainty and life cycle model specifications. Each row of Table 3 (after the first) refers to health uncertainty and life cycle model estimates when a single aspect of the models is modified from its baseline value. For example, the second and third rows of Table 3 show how estimation results change if consumption among the sample of PSID elderly couples is imputed using Skinner's (1987) original methodology or the "structural demand system" approach based on Attanasio and Weber's (1995) work. For both the health uncertainty and life cycle models, estimated CRRA parameters vary only slightly relative to the baseline estimates across PSID consumption imputations. Health uncertainty models predict consumption levels better than life cycle models using all three imputation procedures for PSID consumption levels.[20]

The fourth row of Table 3 shows the inability of modelling health-state-dependent utility of consumption to improve consumption predictions relative to the baseline health uncertainty model. This

---

[20] In none of the dozens of health uncertainty or life cycle model specifications estimated, have I discovered important differences in results depending on which of the three procedures is used to impute consumption expenditures to elderly families in the PSID.

occurs despite the fact that poor-health families in the PSID indeed report lower levels of expenditure than those in good health. Though this empirical pattern seems to provide a potentially important role for health-state-dependence in utility and, therefore, in consumption levels, closer examination reveals the problem such a specification encounters. Recall, health deteriorates with age, intuitively and among elderly respondents in the PSID. When dynamically optimizing families forecast a future deterioration in their health, they forecast a future decline in marginal utility when their preferences are health-state-dependent. The rational response, then, is to substitute for more current consumption (when health seems relatively good) at the expense of lower future consumption. Thus, health-state-dependent utility tends to increase expenditures among the young, healthy retired families according this health uncertainty model (relative to the baseline) specification. Of course, Table 2 indicates that such a tendency simply exacerbates problems for the health uncertainty model by resulting in even larger overpredictions for consumption than occur in the baseline specification.

The fifth through tenth rows of Table 3 reflect the same result: alternatives to the baseline specification which tend to reduce predicted expenditures early during the retirement period tend to increase log-likelihood values and reduce estimated CRRA parameters relative to baseline. Removing adult equivalent scales, increasing real interest rates, reducing perceived mortality rates (increasing life expectancy) and introducing negative discount rates represent several ways to shift retirement consumption levels from early to very old ages according to the dynamic models.[21]

The latter three of these effects are standard, but the first one might warrant closer inspection. Modelling utility as a function of expenditures per adult equivalent implies that utility from spending $10,000 is lower (and marginal utility is greater) for a two-person household than a single individual. Optimizing retired couples, therefore, forecast smaller households and lower marginal utility in the future than today. Rationally, they respond by shifting consumption expenditures from the future to the present,

---

[21] Table 3's seventh row, for example, considers an alternative specification in which retired Americans perceive their age- and sex-specific mortality rates to be just a fourth as large each year as the statistics published by the U.S. Department of Health and Human Services (1985).

which increases consumption predictions for health uncertainty and life cycle models during the early periods of retirement. Consistent with the health-state-dependent utility result described above, removing this tendency from the baseline specifications improves the ability of both health uncertainty and life cycle models to predict retirement consumption.

All of the alternative specifications shown in rows five through ten of Table 3 reduce estimated CRRA parameters and increase log-likelihood values relative to baseline. In fact, health uncertainty and life cycle model specifications in which retired Americans have "optimistic" life expectancies (relative to published data based on death certificates) or negatively discount the future, result in CRRA estimates in the range of values from the empirical Euler equation literature. However, none of the alternative specifications in Table 3 (or any of the numerous unreported specifications) changes the primary result that health uncertainty models better predict consumption levels among PSID retirees than life cycle models do.

Table 3 compares estimates from health uncertainty and life cycle model specifications when several alternative specifications are examined one-at-a-time. However, the possibility remains for some combination of alternatives to allow a life cycle model to outperform the best health uncertainty model, in terms of ability to predict PSID retirement consumption levels. To investigate this possibility, I estimate health uncertainty and life cycle models under dozens of combinations of alternatives suggested by Table 3.[22] Table 4 documents significantly greater log-likelihood values, closer average predicted consumption levels and smaller average squared deviations of prediction errors from modified health uncertainty and life cycle models relative to their baseline specifications. However, precautionary saving from explicit uncertainty about future medical expenses implies a reluctance to dissave financial wealth during retirement under the health uncertainty model that cannot be adequately mimicked by any life cycle model permutations -- even those imposing large coefficients of relative risk aversion, negative (subjective) time preference rates and the failure to recognize the possibility of death before a hundred years of age.

---

[22] In fact, I also examined more time preference rates than appear in Table 3 (-10%, -5%, 0%, 5% and 10%) in combination with all other alternative specifications. Also, I tried models in which retirees' perceptions about age- and sex-specific mortality rates were just half the published statistics.

As suggested by Table 3's results, the principle modification of health uncertainty and life cycle model (relative to baseline) to improve their predictive abilities is increasing perceived life expectancies (decreasing perceived annual mortality rates) compared to published statistics. However, the best-fitting health uncertainty model involves a different combination of alternative specifications than does the best-fitting life cycle model. For example, once sufficiently "optimistic" life expectancies are incorporated into the health uncertainty model (dividing published mortality rates by an adjustment factor of ten), no further discounting (via negative subjective time preference rates or positive real interest rates) improves the health uncertainty model's fit to the data. The resulting CRRA estimate, 4.00, is near the range of estimates available from the contemporary Euler equation estimation literature in this case. On the other hand, the best-fitting life cycle specification "turns off" mortality risk altogether (essentially, all individuals think they will live to one hundred years), but still requires a three percent real interest rate and two unusual values for preference parameters: an estimated CRRA equal to 11.00 and a time preference rate of negative ten percent.

Finally, Panel B of Table 4 shows qualitatively similar results hold for the PSID sample of elderly individuals living alone. Under the health uncertainty model, the best specification yields γ=3.00 and LL_N=-181.38. The best specification imposes a real interest rate and a time preference rate equal to zero percent and adjusts individual mortality rates downward by a factor of four relative to their actuarial values. These parameter values are very similar to those reported for PSID elderly couples. For the PSID sample of elderly singles, many different specifications of the life cycle model with uncertain longevity returned log-likelihood values near the baseline case (-192.13) and all required γ to be 25.00. Thus, none of the dozens of alternative specifications investigated outperformed, or differed very significantly from the baseline life cycle model.

To summarize the empirical results, precautionary incentives implied by a dynamic model in which retirees' face uncertain medical expenses in the future helps to reduce predicted consumption levels toward values observed among elderly families, on average, compared to life cycle models. Even when unusual parameters are imposed, such as large coefficients of relative risk aversion and negative time preference rates, life cycle models cannot adequately mimic the reluctance to consume early during retirement to the extent evidenced among families in the data. Note, however, the best fitting health uncertainty model also requires adjusting age- and sex-specific mortality to much less than their actuarial values to overcome a systematic tendency to overpredict consumption levels in the baseline specification.

VI. Conclusions

This paper introduces a dynamic, structural model of household consumption decisions during retirement in which families consider the effects of potential future shocks to their wealth levels when determining how much to spend currently. Explicitly, shocks take the form of exogenous expenses incurred out-of-pocket to finance health care during old age. Additionally, the health uncertainty model incorporates the possibility of a person living past her life expectancy, as well as the possibility of dying prematurely, which also will affect the timing of consumption expenditures. Using panel data from samples of elderly retirees in the PSID, I estimate the coefficient of relative risk aversion under the health uncertainty and life cycle consumption models. In contrast, the existing literature tends to estimate dynamic models without fully specifying the stochastic environment facing forward-looking families or to simulate "complete" life cycle models without subjecting them to standard empirical testing using formal statistical criteria.

The health uncertainty model predicts consumption expenditures much closer, on average, to observed expenditures during retirement than life cycle models with uncertain longevity. This result obtains in the baseline model specifications, which are parameterized to closely match the existing simulation literature, and for every alternative specification investigated in the sensitivity analysis. I conclude, therefore, that uncertain out-of-pocket medical expenses represents an important motive for precautionary saving among elderly Americans. Simulations (unreported) of the health uncertainty model using the estimated coefficient of relative risk aversion imply that precautionary saving arising only from uncertain future out-of-pocket medical expenses amounts to approximately seven percent of annual consumption during the early years of retirement for a typical couple.

On the other hand, improvements in the ability of the health uncertainty model to fit the PSID consumption data can be gained by adjusting downward (substantially) annual mortality rates relative to published actuarial estimates. This result indicates that the substantial precautionary saving implied by the health uncertainty model still is inadequate to explain completely the apparent reluctance of elderly Americans to spend down their financial assets during their early retirement years. Besides being consistent with systematic "optimism" about life expectancy during old-age, this pattern also is consistent with an altruistic motive for inter-generational bequests. In principle, the health uncertainty model can be modified to include such objectives explicitly, but the empirical examination estimation of such a model, which adds complexity relative to the current framework, is left for future research.

Including an explicit parameterization of health uncertainty into the retirement phase of an otherwise standard life cycle consumption model moves us in the correct direction by generating a precautionary motive for saving in old-age, but does not seem to move us far enough toward the data. Additionally, baseline health uncertainty specifications, which most closely follow models previously used in the simulation literature, yield large, and arguably implausible, estimates for the coefficient of relative risk aversion. Alternative health uncertainty model specifications, which arbitrarily adjust perceived mortality rates substantially below their actuarial values, improve the model's fit to the panel consumption data and yield parameter estimates closer to those reported in some previous studies (γ around 3 or 4).[23]

The general estimation methodology developed in this paper, however, is readily applicable to a very wide class of alternative life cycle-type models for family consumption decisions. The paper's contribution, therefore, may be particularly important in light of recent research by others, such as Carroll (1997) and Ludvigson and Paxson (1997), documenting inherent econometric difficulties with the log-linearized Euler equation methodology so frequently applied in the past.

Finally, many simulation studies in public finance (for instance Summers (1981) and Auerbach and Kotlikoff (1987)) examine the implications of fiscal policy reform in life cycle model settings. The economic environment in which the analyses take place are such that families are assumed to make

---

[23] Ludvigson and Paxson (1997) survey parameter estimates from several different approaches to Euler equation estimation in the recent consumption literature.

decisions when all economic variables are known with certainty, all households are assumed to live to their life expectancies and altruistic bequests are ignored. An important lesson to be learned from this paper, however, is that such a simple life cycle model specification is not at all consistent with observed consumption decisions of retired Americans. The standard life cycle model simply cannot explain why elderly families spend so little of their incomes and wealths during their early retirement years. This paper suggests that models incorporating precautionary behavior on the part of rational, forward-looking families, such as Engen and Gale (1993) and Hubbard, Skinner, and Zeldes (1995) and others, are more appropriate for policy analysis in a dynamic context.

References

Altonji, Joseph G. "Intertemporal Substitution in Labor Supply: Evidence from Micro Data," *Journal of Political Economy* 94, Supplement, June, 1986, pp. S176-S215.

Ando, Albert, and Franco Modigliani. "The 'Life Cycle' Hypothesis of Saving: Aggregate Implications and Tests," *American Economic Review* 53, no. 1, March, 1963, p. 55-84.

Attanasio, Orazio P. and Guglielmo Weber. "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy* 103, no. 6, 1995, pp. 1121-1157.

Auerbach, Alan J. and Laurence J. Kotlikoff. *Dynamic Fiscal Policy*. Cambridge: Cambridge University Press, 1987.

Baxter, Marianne, Mario J. Crucini and K. Geert Rouwenhorst. "Solving the Stochastic Growth Model by a Discrete-State-Space, Euler-Equation Approach," *Journal of Business and Economic Statistics* 8, no. 1, January, 1993, pp. 19-22.

Browning, Martin, and Annamaria Lusardi. "Household Saving: Micro Theory and Micro Facts," *Journal of Economic Literature* XXXIV, no. 4, 1996, pp. 1797-1855.

Carroll, Christopher D. "Death to the Log-Linearized Consumption Euler Equation!" mimeo, Johns Hopkins University, September, 1997.

Cohen, Marc A., Eileen J. Tell, and Stanley S. Wallack. "The Lifetime Risks and Costs of Nursing Home Use Among the Elderly," *Medical Care*, vol. 24, no. 12, December, 1986, p. 1161-1172.

Danziger, Sheldon, Jacques Van Der Gaag, Eugene Smolensky, and Michael Taussig. "The Life-Cycle Hypothesis and Consumption Behavior of the Elderly," *Journal of Post-Keynesian Economics*, Winter, 1982-1983, p. 208-227.

Davies, James B. "Uncertain Lifetime, Consumption, and Dissaving in Retirement," *Journal of Political Economy 89*, no. 3, June, 1981, p. 561-577.

Deaton, Angus. *Understanding Consumption*, 1992. Oxford: Clarendon Press.

Deaton, A.S. and J. Muellbauer. "An Almost Ideal Demand System," *American Economic Review* 70, 1980, pp. 312-336.

Diamond, P.A., and J.A. Hausman. "Individual Retirement and Savings Behavior," *Journal of Public Economics* 23, p. 81-114.

Eckstein, Zvi, and Kenneth I. Wolpin. "The Specification and Estimation of Dynamic Stochastic Discrete Choice Models," *The Journal of Human Resources* 24, no. 4, Fall, 1989, p.562-598.

Engen, Eric. "Estimation of Dynamic Models of Precautionary Saving and Life Cycle Consumption Using

Panel Data," mimeo, UCLA, September, 1993.

Engen, Eric M. and William G. Gale. "IRA's and Saving in the Stochastic Life-Cycle Model," mimeo, Brookings Institution, 1993.

Feenberg, Daniel, and Jonathan Skinner. "The Risk and Duration of Catastrophic Health Care Expenditures," *Review of Economics and Statistics* 76, no. 4, November, 1994, pp. 633-47.

Gourinchas, Pierre-Olivier, and Jonathan A. Parker. "Consumption Over the Life Cycle," working paper, November, 1997.

Hamermesh, Daniel D. "Consumption During Retirement: The Missing Link in the Life Cycle," *Review of Economics and Statistics* LXVI, 1984, pp. 1-7.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving," *Carnegie-Rochester Conference Series on Public Policy* 40, June, 1994, pp. 59-125.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. "Precautionary Saving and Social Insurance," *Journal of Political Economy* 103, no. 2, 1995, pp. 360-399.

Hurd, Michael D. "Savings of the Elderly and Desired Bequests," *American Economic Review* 77, no. 3, June, 1987, p. 298-312.

Hurd, Michael D. "Mortality Risk and Bequests," *Econometrica* 57, no. 4, July, 1989, p. 779-814.

Kimball, Miles S. "Precautionary Saving in the Small and in the Large," *Econometrica* 58, no. 1, January, 1990, p. 53-73.

King, M.A., and L. Dicks-Mireaux. "Asset Holdings and the Life-Cycle," *The Economic Journal* 92, June, 1982, p. 247-267.

Kotlikoff, Lawrence J. "Health Expenditures and Precautionary Saving," Chapter 6 in *What Determines Saving?*. Cambridge, MA: MIT Press, 1988.

Leland, Hayne E. "Saving and Uncertainty: The Precautionary Demand for Saving," *Quarterly Journal of Economics* 82, 1968, p. 465-473.

Lillard, Lee, and Yoram Weiss. "Uncertain Health and Survival: Effects on End-of-Life Consumption," *Journal of Business and Economic Statistics*, vol. 15, no. 2, 1996.

Liu, Korbin, and Kenneth G. Manton. "The Characteristics and Utilization Pattern of an Admission Cohort of Nursing Home Patients (II)," *The Gerontologist*, vol. 24, no. 1, 1984, p. 70-76.

Ludvigson, Sydney, and Christina H. Paxson. "Approximation Bias in Linearized Euler Equations," mimeo, Princeton University, July, 1997.

MaCurdy, Thomas E. "Interpreting Empirical Models of Labor Supply in and Intertemporal Framework with Uncertainty," Chapter 3 in J. J. Heckman and B. Singer, eds., *Longitudinal Analysis of Labor Market Data*. Cambridge: Cambridge University Press, 1985.

Mirer, Thad W. "The Wealth-Age Relation Among the Aged," *American Economic Review* 69, no. 3, June, 1979, p. 435-443.

Modigliani, Franco, and Richard Brumberg. "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data," in Kenneth K. Kurihara, ed., *Post-Keynesian Economics*, Rutgers University Press, 1954, p. 388-436.

Palumbo, Michael G. "Health Uncertainty and Optimal Consumption Near the End of the Life Cycle," University of Virginia Ph.D. Dissertation, 1992.

Sandmo, Agnar. "The Effect of Uncertainty on Saving Decisions," *Review of Economic Studies* 37, 1970, p. 353-360.

Skinner, Jonathan. "Variable Lifespan and the Intertemporal Elasticity of Consumption," *Review of Economics and Statistics* 67, 1985, p. 616-623.

Skinner, Jonathan. "A Superior Measure of Consumption from the Panel Study of Income Dynamics," *Economic Letters* 23, 1987, p. 213-216.

Slesnick, Daniel. "Gaining Ground: Poverty in the Postwar United States," *Journal of Political Economy* 101, no. 1, 1993, pp. 1-38.

Slesnick, Daniel. "Is our Data Relevant to the Theory? The Case of Aggregate Consumption," mimeo, September, 1994.

Summers, Lawrence H. "Taxation and Capital Accumulation in a Life Cycle Growth Model," *American Economic Review* 71, September, 1981, pp. 533-554.

Taylor, John B., and Harald Uhlig. "Solving Nonlinear Stochastic Growth Models: A Comparison of Alternative Solution Methods," *Journal of Business and Economic Statistics* 8, no. 1, January, 1993, pp. 1-17.

U.S. Department of Health and Human Services. *U.S. Decennial Life Tables for 1979-81*, DHHS Publication No. (PHS) 85-1150-1, August, 1985.

Viscusi, W. Kip and William N. Evans. "Utility Functions that Depend on Health Status: Estimates and Economic Implications," *American Economic Review* 80, June, 1990, pp. 353-374.

White, Betsy Buttrill. "Empirical Tests of the Life-Cycle Hypothesis," *American Economic Review* 68, no. 4, September, 1978, p. 547-560.

Zeldes, Stephen P. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics*, May, 1989, p. 275-298.

Table 1: CRRA Estimates and Consumption Predictions from Baseline Health Uncertainty and Life Cycle Model Specifications

A. PSID Elderly Couples (N = 142)

	Health Uncertainty Model	Life Cycle Model
γ	7.00	24.00
Standard Error	(1.89)	(3.32)
LLN	3.94	-12.70
Avg. Actual Consumption	15,115	---
Avg. Predicted Consumption	18,247	19,225
Avg. $\eta 2$	0.134	0.152

B. PSID Elderly Singles (N = 144)

	Health Uncertainty Model	Life Cycle Model
γ	6.25	25.00
Standard Error	(1.03)	(2.12)
LLN	-186.64	-192.13
Avg. Actual Consumption	8,999	---
Avg. Predicted Consumption	8,597	8,604
Avg. $\eta 2$	0.230	0.243

Notes: As described in the text, PSID consumption expenditures have been imputed using the CEX-based, flexible-Engel-curve procedure described in Appendix D.

Baseline model specifications include: adult equivalent scales for household consumption; uncertain longevity parameterized using published life table estimates for mortality probabilities; zero percent time preference rate; three percent real interest rate; and, no health-state-dependence in the utility derived from consumption expenditures.

Standard errors are estimated by evaluating the inverse second derivative from a high-order polynomial approximation to the sample log-likelihood function at the estimated parameter value.

$\eta 2$ is the squared difference between the natural log of actual consumption expenditures imputed to PSID families and annual consumption levels predicted according to health uncertainty or life cycle models.

Because log-likelihood functions are extremely flat, grid searches have been truncated at γ=25.00 under life cycle model specifications.

Table 2: Comparing Deviations from Actual Consumption Levels for Predicted Values based on Estimated Baseline Health Uncertainty and Life Cycle Models

1985 Consumption Levels among Sample of PSID Elderly Couples (N = 142)

Household Category	Health Uncertainty Model		Life Cycle Model	
	η	$\eta2$	η	$\eta2$
Total Sample	-0.13	0.13	-0.19	0.15
Asset Bracket				
< 5k	0.07	0.06	0.02	0.08
5 - 15k	0.01	0.07	-0.09	0.09
15 - 30k	-0.18	0.15	-0.27	0.18
30 - 60k	-0.16	0.09	-0.23	0.12
60 - 90k	-0.24	0.14	-0.29	0.15
90 - 120k	-0.17	0.21	-0.21	0.21
120 - 180k	-0.51	0.35	-0.55	0.37
> 180k	-0.59	0.47	-0.60	0.49
Income Bracket				
< 7.5k	0.12	0.09	0.11	0.11
7.5 - 17.5k	-0.06	0.09	-0.16	0.11
> 17.5k	-0.37	0.23	-0.39	0.25
Age Bracket				
65 - 69	-0.15	0.12	-0.21	0.14
70 - 74	-0.12	0.13	-0.16	0.15
75 - 79	-0.09	0.13	-0.18	0.15
80 - 84	-0.22	0.17	-0.33	0.20
85 - 89	-0.20	0.22	-0.33	0.25
Health Status				
Good	-0.21	0.16	-0.26	0.18
Fair	-0.13	0.13	-0.21	0.14
Poor	-0.06	0.13	-0.12	0.16

Notes: η is the average deviation between log-actual and log-predicted consumption based on the estimated baseline health uncertainty or life cycle models. $\eta2$ refers to average squared deviations. Health uncertainty model is best baseline specification, with estimated γ=7.00; life cycle model is best baseline specification, with estimated γ=24.00.

Table 3: CRRA Estimates from Baseline and Several Alternative Health Uncertainty and Life Cycle Model Specifications

Sample of PSID Elderly Couples (N = 142)

	γ, LLn (st. err.)	
Model Specification	Health Uncertainty Model	Life Cycle Model
1. Baseline	7.00, 3.94 (1.89)	24.00, -12.70 (3.39)
2. using Skinner-Imputed Data	6.75, -35.49 (1.80)	25.00, -48.12 (3.81)
3. using Attanasio/ Weber-Imputed Data	9.00, -179.26 (2.09)	25.00, -188.56 (3.42)
4. with health-state-dependent utility	6.75, 0.41 (1.93)	n/a
5. without adult equivalent scales	6.50, 16.90 (2.17)	25.00, 7.78 (2.94)
6. r = 0%	6.75, 9.85 (0.96)	25.00, -6.18 (2.75)
7. mortality rates/4	1.50, 21.57 (0.22)	1.25, 18.47 (0.20)
8. mortality rates/10	2.00, 23.45 (0.23)	1.25, 20.07 (0.19)
9. mortality rates/100	3.25, 27.29 (0.29)	2.00, 21.47 (0.21)
10. b = -10%	4.25, 23.51 (0.31)	3.50, 15.95 (0.25)

Notes: Baseline model specifications include: adult equivalent scales for household consumption; uncertain longevity parameterized using published life table estimates for mortality probabilities; zero percent time preference rate; three percent real interest rate; and, no health-state-dependence in the utility derived from consumption expenditures.

Each row of this table refers to the maximum likelihood estimates based on baseline model specifications with a single parameter adjusted as shown.

Row 4 reports results based on the following function: δ(good health) = 1.00; δ(fair health) = 0.70; δ(poor health) = 0.40. Other functions lead to qualitatively similar results.

Row 5 results refer to CRRA utility as a function of total family expenditures, rather than consumption per adult equivalent member.

Rows 7 through 9 come from model specifications in which perceived mortality rates (age-specific) are smaller than published actuarial rates by factors of four, ten and a hundred, respectively.

Results based on the PSID sample of elderly singles are not reported because of their qualitative similarity to these. Standard error estimates come from evaluating inverse second derivatives for high-order polynomial approximations to the sample log-likelihood function.

Because log-likelihood functions are extremely flat, grid searches have been truncated at γ=25.00 under life cycle model specifications.

Table 4: CRRA Estimates and Consumption Predictions from the Best Health Uncertainty and Life Cycle Model Specifications

A. Sample of PSID Elderly Couples (N = 142)

	Health Uncertainty Model	Life Cycle Model
Estimation Results		
γ	4.00	11.00
Standard Error	(0.62)	(0.93)
LLN	28.43	21.66
Avg. Actual Consumption	15,115	---
Avg. Predicted Consumption	15,698	15,436
Avg. $\eta 2$	0.102	0.114
Best Model Specification		
Adjustment Factor for Annual Mortality Rates	10	100
Adult Equivalence Scales for Utility of Consumption	Off	On
Real Interest Rate	0%	3%
Rate of Time Preference	0%	-10%

B. Sample of PSID Elderly Singles (N = 144)

	Health Uncertainty Model	Life Cycle Model
Estimation Results		
γ	3.00	25.00
Standard Error	(0.51)	(2.12)
LLN	-181.38	-192.13
Avg. Actual Consumption	8,999	---
Avg. Predicted Consumption	8,328	8,604
Avg. $\eta 2$	0.215	0.243
Best Model Specification		
Adjustment Factor for Annual Mortality Rates	4	none
Adult Equivalence Scales for Utility of Consumption	n/a	n/a
Real Interest Rate	0%	3%
Rate of Time Preference	0%	0%

Notes: $\eta 2$ is the squared difference between the natural log of actual consumption expenditures imputed to PSID families and annual consumption levels predicted according to health uncertainty or life cycle models.

These results are based on a broad specification search with respect to alternative parameters imposed under health uncertainty and life cycle consumption models.

Note, many different specifications of the life cycle model yielded nearly the exact same fit for the sample of elderly singles as the baseline case presented above (γ=25.00; LLN=-192.13).