



Research Division
Federal Reserve Bank of St. Louis
Working Paper Series



A Leisurely Reading of the Life Cycle Consumption Data

James Bullard
and
James Feigenbaum

Working Paper 2003-017F
<http://research.stlouisfed.org/wp/2003/2003-017.pdf>

July 2003
Revised January 2007

FEDERAL RESERVE BANK OF ST. LOUIS
Research Division
P.O. Box 442
St. Louis, MO 63166

The views expressed are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment. References in publications to Federal Reserve Bank of St. Louis Working Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors.

A Leisurely Reading of the Life Cycle Consumption Data

James Bullard^{a,*}, James Feigenbaum^{b,‡}

^aFederal Reserve Bank of St. Louis; ^bUniversity of Pittsburgh

Received Date; Received in Revised Form Date; Accepted Date

Abstract

A puzzle in consumption theory is the observation of a hump in age-consumption profiles. This paper studies a general equilibrium life-cycle economy with capital in which households include both consumption and leisure in their period utility function. A calibrated version of the model shows that a significant hump in life-cycle consumption is a feature of the equilibrium. Thus inclusion of leisure in household preferences may provide part of the explanation of observed life-cycle consumption humps.

Keywords: Life-cycle consumption and saving, overlapping generations, consumption humps.

JEL classification: D91, E21.

*Corresponding author: bullard@stls.frb.org. Any views expressed are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of St. Louis or the Federal Reserve System.

†The authors thank David Laibson, John Rust, and Filip Vesely for stimulating discussions on this topic, Linnea Polgreen for help analyzing the CPS data, as well as participants at Midwest Macroeconomics and the Society for Economic Dynamics for helpful comments. The authors also thank Pierre-Oliver Gourinchas and Jonathan Parker, as well as Jesus Fernandez-Villaverde and Dirk Krueger for generously providing data from their papers.

‡Email: jfeigen+@pitt.edu.

1. Introduction

An important puzzle in consumption theory is the observation of a hump shape in life-cycle consumption. Gourinchas and Parker (2002), in particular, have documented that household age-consumption profiles, adjusted for both economic growth and family size, have a distinct and statistically significant hump, with actual consumption on the increase during ages 20 – 40 and falling off during ages 50 – 70. One metric that can be used to gauge the nature of this hump is to consider the ratio of peak consumption to age 30 consumption. This ratio is about 1.10 according to the estimates of Gourinchas and Parker (2002), with the peak occurring around age 45. An alternative estimate, due to Fernández-Villaverde and Krueger (2002), suggests a ratio of about 1.18.^{1,2} The hump in lifetime household consumption is a puzzle in the sense that it is inconsistent with a life-cycle model in which households have additively separable utility with a time-invariant structure defined over consumption alone. In that type of model, the optimal consumption plan of households involves smooth age-consumption profiles which either grow (or decay) exponentially over time.³ The recent literature includes investigations of several avenues that might help explain the hump, including precautionary saving, liquidity constraints, variations in family size, and the split between durables and nondurables consumption.

The purpose of this paper is to report on an investigation of another avenue, namely whether the inclusion of leisure in the utility function could help explain the hump in consumption and related phenomena. It has been known since the work of Heckman (1974) that such a theoretical possibility exists in a life-cycle context when household productivity follows a hump-shaped life-cycle pattern, as it does in the data. Accordingly, a general

¹This is their estimate for adult-equivalent nondurables consumption, which is the object which best matches our model concepts, as we detail later in the paper.

²Age 30 turns out to be a convenient normalization for comparison across models, and so it is used as a standard in this paper.

³See, for instance, Yaari (1964).

equilibrium life-cycle economy, with physical capital as the only asset, is studied under preferences consistent with Heckman (1974). Heckman's partial equilibrium findings are verified and extended. The model is calibrated from a macroeconomic perspective, seeking to ensure that the balanced growth path properties match fundamental features of the aggregate postwar U.S. data. The calibration also matches available cross-sectional household evidence concerning profiles of aggregate hours worked by age. To our knowledge, ours is the first quantitative-theoretic general equilibrium analysis focusing directly on this issue.

1.1. Main findings

A hump in life-cycle consumption and related phenomena can be obtained under our general equilibrium calibration of a life-cycle economy in which households bundle consumption and leisure in their period utility function. The consumption hump produced under two such calibrations has a ratio of peak consumption to age 30 consumption of about 1.10, the same as the estimate of Gourinchas and Parker (2002), but smaller than the alternative estimate of Fernández-Villaverde and Krueger (2002). One conclusion is that the leisure-based approach to explaining the hump in life-cycle consumption has merit. The size and shape of the observed life-cycle consumption hump may be at least partly attributable to the effects of leisure in household preferences.

1.2. Related literature

Many theories have posited that some form of market incompleteness might account for observed life-cycle consumption humps, either through borrowing constraints, as suggested by Thurow (1969), or through precautionary saving, as suggested by Nagatani (1972). This type of explanation has received the most attention in the more recent literature, as exemplified by Hubbard, Skinner, and Zeldes (1994), Carroll (1997), and Gourinchas and Parker

(2002).⁴

However, another explanation, first suggested by Heckman (1974) and Becker and Ghez (1975), is that if leisure and consumption are substitutes, then it is not consumption alone but bundles of consumption and leisure that should be smoothed over the life cycle. Given a hump-shaped life-cycle productivity profile, one can then obtain a consumption hump,⁵ along with a U -shaped leisure profile.⁶ Carroll and Summers (1991) as well as Browning and Crossley (2001) hypothesized that a model using this mechanism would require an empirically implausible hump in the life-cycle labor hours profile in order to obtain an empirically plausible hump in consumption. It is well known that empirical life-cycle labor hours profiles are relatively flat between the ages of 30 and 50. They therefore suggested that Heckman's hypothesis was not likely to bear fruit. The present paper aims to provide a quantitative assessment of this idea.

Several authors have used models that combine precautionary saving, substitutability of leisure and consumption, and liquidity constraints in a partial equilibrium environment. Huggett and Ventura (1999), French (2002), and Rust (2002) study questions related to retirement, and life-cycle consumption humps are typically found in these models.⁷ But it is not clear from these papers whether it is buffer-stock saving or leisure-consumption substitution, or some combination, that leads to the life-cycle consumption humps.⁸

⁴There is also a literature suggesting that time-inconsistent or myopic preferences over consumption can play a role. See Laibson (1997) and Caliendo and Aadland (2004).

⁵Bütler (2001) provides a partial equilibrium analysis and exposition of this and related effects in a continuous time life cycle model like the one employed in this paper.

⁶This mechanism was at work in the general equilibrium life-cycle economies of Bullard and Russell (1999) and Rios-Rull (1996). These authors focused on other issues, however.

⁷Heathcote (2002) considers a model which combines home production and buffer-stock saving. In his baseline calibration, leisure and consumption are separable. Although he considers a case where they are substitutable, he does not discuss the life-cycle consumption profiles in that case.

⁸In examining a variant of such a model, Benitez-Silva (2000) did examine (in partial equilibrium) what happens when life-cycle income uncertainty, and hence the precautionary saving feature, is turned off, and found that leisure-consumption substitution alone could produce a hump. However, an institutional maximum on the amount of labor hours that agents could supply was imposed in order to avoid a pronounced life-cycle labor hours hump. Agents are allowed to freely choose labor supply in the present model.

Variation in household size could account for why preferences over consumption by a household might change over the life cycle, as in Attanasio, *et al.* (1999) and Browning and Ejrnaes (2002)—that is, consumption is highest when households are largest. However, several researchers, including Fernández-Villaverde and Krueger (2002) and Gourinchas and Parker (2002), have found that the consumption hump persists even after controlling for household size. Meanwhile, Fernández-Villaverde and Krueger (2001) have shown that the hump persists for both durables and nondurables consumption considered separately, and they have proposed an incomplete markets model in which durables can serve as collateral to explain these stylized facts.

Heathcote, Storesletten, and Violante (2004) consider a buffer-stock saving model with a labor-leisure decision. Their utility function is separable in leisure and consumption but has more degrees of freedom that allow them to obtain a hump in lifecycle consumption along with a relatively flat life-cycle labor hours profile. However, this utility function does not permit a balanced-growth path in general equilibrium.⁹

1.3. Organization

The next section describes a continuous time, general equilibrium life-cycle economy. A version of Heckman's (1974) findings are verified, and the intuition behind the mechanism is described. The analysis then turns to a calibration of the economy to long-run, balanced growth facts for the U.S. since WWII, and in addition to data concerning the life-cycle labor hours profile. The results section discusses our findings and contains some sensitivity analysis, before turning to conclusions and directions for future research.

⁹There are additional puzzles related to household consumption after retirement, but there is little reason to believe that the mechanisms that are most pertinent to resolving those puzzles are the mechanisms most pertinent to accounting for the consumption hump. This paper does not focus on retirement consumption issues.

2. Environment

Time is continuous and is denoted by $t \in \mathbb{R}$. At each instant, or date, t , a new generation, or cohort, of agents is born. The population is constant. Agents are identical within cohorts. The birth date of an agent is often denoted by τ , and an individual agent is called a member of cohort τ . Denote a quantity x relating to an agent alive at time t who was born at τ by $x(t; \tau)$. Aggregate variables have no birth date component and so will simply be denoted in the form $x(t)$. The age of an agent or a cohort is denoted by $s = t - \tau$, and s and $t - \tau$ are used interchangeably. There is no government, and there is a single good which may be either consumed or saved, in which case it is called capital. The only uncertainty in the model, at either the individual or aggregate level, pertains to the life span of each individual. It is important to stress that there are no borrowing constraints of any kind in the model—households may borrow as they wish at any stage of life. Retirement is endogenously determined.

2.1. Household behavior

First consider an individual agent's allocation problem. Let T be the maximum possible life span of an agent. For $s \in [0, T]$, the probability of surviving until age s or beyond is denoted $Q(s)$, a cohort-independent survivor function. A member of cohort τ maximizes expected utility given by

$$\int_{\tau}^{\tau+T} Q(t - \tau) e^{-\rho(t-\tau)} u[c(t; \tau), \ell(t; \tau)] dt, \quad (1)$$

where $c(t; \tau)$ is consumption, $\ell(t; \tau) \in [0, 1]$ is leisure, $\rho > 0$ is the subjective discount rate, and

$$u(c, \ell) = \begin{cases} \frac{1}{1-\sigma} (c^\eta \ell^{1-\eta})^{1-\sigma} & \text{if } \sigma \neq 1 \\ \eta \ln c + (1 - \eta) \ln \ell & \text{if } \sigma = 1 \end{cases} \quad (2)$$

is the specification of the instantaneous period utility function (*a.k.a.* felicity), where $\sigma > 0$ and $\eta \in (0, 1)$. This form of the period utility function is consistent with Heckman (1974) and has proven to be useful in the macroeconomics literature.¹⁰ A household of age $t - \tau$ ($= s$) has an endowment of one time unit, as well as a stream of productivity units over the lifetime denoted $e(s)$ for $s \in [0, T]$. For each unit of time devoted to labor at a date t , $e(s)$ efficiency units are supplied to production. Thus the agent earns income $w(t) [1 - \ell(t; \tau)] e(t - \tau)$ at date t where $w(t)$ is the market-determined real wage per labor efficiency unit, taken as given from the perspective of the individual agent. Based on the empirical evidence discussed later, the stream of productivity units $e(s)$ can be thought of as increasing early in life and decreasing later in life.

There is a consumption loan market in which agents may borrow from or lend to agents in other cohorts. Agents can also lend to firms, in which case we say they are holding claims on capital. Returns in the consumption loan market and in the market for capital are assumed to be equal by arbitrage. Thus an agent born at date τ who is holding $a(t; \tau)$ assets (which may be a negative quantity if the agent is borrowing) from date t to date $t + dt$ will earn (or pay) interest $r(t) a(t; \tau) dt$ during that interval, where $r(t)$ is the market-determined rate of return from date t to date $t + dt$, also taken as given from the perspective of an individual agent. As is typical in the literature, assume there is no market for claims contingent on an agent's lifespan, so an agent cannot insure against mortality risk by purchasing annuities. If an agent dies before the maximal date T , any assets are spread uniformly over the surviving population.¹¹ At time t , a surviving agent receives the bequest $B(t)$. The consumer's problem is therefore to maximize (1) subject to

$$c(t; \tau) + \frac{da(t; \tau)}{dt} = r(t) a(t; \tau) + [1 - \ell(t; \tau)] w(t) e(t - \tau) + B(t), \quad (3)$$

¹⁰These preferences are well-known to be consistent with balanced growth, which is one reason they are popular in the macroeconomics literature.

¹¹Any debts would also be spread uniformly, but the agents will not borrow in the equilibria we study.

and the condition $a(\tau; \tau) = a(T + \tau; \tau) = 0$.

Because this problem is relatively standard, for brevity the solution is not described here.¹²

2.1.1. The dynamics of life-cycle consumption and leisure

Now consider a characterization of the household's life cycle consumption and leisure choices, focusing on the case when the agent chooses to work. Define the hazard rate of dying as $h(s) \equiv -\frac{d \ln Q(s)}{ds}$. Then the dynamics of life-cycle consumption are described by

$$\frac{d}{dt} \ln c(t; \tau) = \frac{r(t) - \rho_{eff}(t - \tau)}{\sigma} + \frac{(1 - \eta)(\sigma - 1)}{\sigma} \frac{d \ln [w(t) e(t; \tau)]}{dt}, \quad (4)$$

where $\rho_{eff}(s) = \rho + h(s)$ is the effective discount rate, taking into account mortality risk. For the moment, let us ignore the effects of mortality, assuming $Q(s) = 1$ for all s .¹³ The standard macroeconomic calibration information has $\sigma > 1$ and $r > \rho$, along with a hump-shaped life-cycle productivity profile. Under these conditions, equation (4) states that if the household's productivity profile $e(s)$ increases as a function of age, consumption will increase as a function of age. On the other hand, if the productivity profile decreases with age at a fast enough rate, consumption will decrease with age. Consequently, a hump-shaped productivity profile can generate a hump-shaped consumption profile.¹⁴ Thus a key question will be whether the generated consumption profiles are empirically reasonable.

Similar statements can be made for the behavior of leisure, and hence hours worked, over the life cycle. The dynamics of life-cycle leisure are described by

$$\frac{d}{dt} \ln \ell(t; \tau) = \frac{r(t) - \rho_{eff}(t - \tau)}{\sigma} - \frac{1 - \eta + \eta\sigma}{\sigma} \frac{d \ln (w(t)e(t; \tau))}{dt}. \quad (5)$$

Under the assumptions $r > \rho$ and $\sigma > 1$, if the household's productivity profile increases at a fast enough rate, then leisure will decrease, while if the productivity profile decreases then

¹²Interested readers may consult Bütler (2001), for example, or a working paper version of this research.

¹³For a more detailed analysis of the effects of mortality risk on the consumption profile, see Feigenbaum (2005) and Hansen and Imrohorglu (2005).

¹⁴This is a version of Heckman's (1974) result.

leisure will increase. Thus one expects to obtain a U -shaped leisure profile over the life cycle under these assumptions, which corresponds to a hump-shaped labor hours profile. Again, a key question is the quantitative importance of this effect.

2.2. Technology and general equilibrium

To close the model in general equilibrium, assume there are a continuum of identical perfectly competitive firms with access to a Cobb-Douglas technology with exogenous growth in labor productivity at the rate g . Since the firms are the same, our analysis proceeds as if there were only a single firm behaving competitively using technology $Y(t) = K(t)^\alpha (e^{gt} N(t))^{1-\alpha}$, where $Y(t)$ is aggregate output at date t , $K(t)$ is the aggregate capital stock at t , and $N(t)$ is the aggregate labor supply at t . Factors are paid their marginal products. Consumption loans will net out in the aggregate, so that the excess demand for assets must equal the total capital stock. This is given by $K(t) = \int_{t-T}^t Q(t-\tau) a(t; \tau) d\tau$. The aggregate labor supply will be the sum over the labor supply for each cohort $N(t) = \int_{t-T}^t Q(t-\tau) [1 - \ell(t; \tau)] e(t; \tau) d\tau$. The focus of the paper is on a balanced growth path equilibrium, where $K(t) = e^{gt} K_0$ and $N(t) = N$. An arbitrage condition equates the rate of return on consumption loans to the rate of return from holding claims to capital, which in turn is the marginal product of capital less the depreciation rate δ . Then the rate of return on capital will be $r = \alpha \left(\frac{K_0}{N}\right)^{\alpha-1} - \delta$. The marginal product of labor dictates the balanced growth path wage as $w(t) = e^{gt} (1-\alpha) \left(\frac{K_0}{N}\right)^\alpha$. To complete the model, the size of bequests is determined by

$$\int_{t-T}^t Q(t-\tau) B(t) d\tau = \int_{t-T}^t Q(t-\tau) h(t-\tau) a(t; \tau) d\tau, \quad (6)$$

which balances the total bequests received by surviving agents at t against the total assets held by agents who die at t .

3. Quantitative analysis

The central question to investigate is whether a balanced growth path equilibrium of this model can be calibrated to be consistent with balanced growth characteristics of the postwar U.S. data along with other data, consistent with hump-shaped life-cycle consumption profiles and other aspects of observed life-cycle behavior documented by Gourinchas and Parker (2002) and other authors. The model's age-profiles of consumption, income, asset-holding, and importantly, hours worked, can be interpreted as representing economy-wide cohort averages. These cohort averages can be compared to available data from the postwar U.S. The model is interpreted as one where agents are "born" at age 25.¹⁵

3.1. Targets in the U.S. data

The macroeconomic balanced growth facts from the U.S. data are standard. The average capital-output ratio is the first target an equilibrium should match. Rios-Rull (1996), calibrating a no-government, general equilibrium life-cycle model like this one, uses a value of 2.94, and so this value is used as a target. Rios-Rull (1996) also suggests targets based on the U.S. data for the ratio of consumption to output 0.748 (which makes the ratio of investment to output equal to 0.252). This provides a second target. The third macroeconomic target is the real interest rate, which is independently determined in the life-cycle framework, as discussed below. Following Gourinchas and Parker (2002), the target real interest rate is set at 3.5 percent.

In addition to these macroeconomic observables, interest also centers on the extent to which there is a calibration of the model that will match empirical data concerning the life-

¹⁵It is well-known that prior to this age, economic behavior is greatly affected by human capital accumulation decisions, among other factors, so that one may not expect the model to match the behavior of these agents well. It would be interesting to include decisions of this type into the model in future research. In addition, the Gourinchas and Parker (2002) estimates apply from age 25, and so this interpretation will facilitate a comparison to their estimated consumption behavior.

cycle profiles of average household consumption of non-durables,¹⁶ average hours worked per week, and average productivity. The data on consumption is available from Gourinchas and Parker (2002). Data on hours worked was obtained from the April 2002 Current Population Survey (CPS). Aggregating across individuals within cohorts in this data yields average hours worked by age cohort. Based on the mean and standard deviation of hours supplied in each cohort, there is a small but statistically significant hump in hours worked by age. The ratio of peak hours, which occurs at age 45, to age 30 hours is about 1.09.

Using cross-sectional data on hourly wages from the 2002 CPS provides a proxy for the household productivity profile. However, this proxy only measures productivity conditional on an agent working—therefore, we call this the “conditional productivity profile.” What enters the model is an unconditional productivity profile. The literature effectively assumes that the difference between the conditional and unconditional profiles is small. This assumption becomes less tenable once a large number of agents begin to retire. Even before age 60, the ratio of conditional to unconditional productivity may vary somewhat with age. One can view the parameters of the unconditional productivity profile $e(s)$ as parameters to be chosen such that, along with other parameters, the unconditional productivity profile of the model is close to the conditional productivity profile in the data prior to age 60.¹⁷

The households devote time to market work, which is most directly affected by the parameter η , and this parameter plays an important role in the life-cycle consumption and labor supply profiles described by equations (4) and (5). According to CPS data, the U.S. average is about 30 hours of market work per person per week. Thus the goal is to induce agents to work, on average across all cohorts, about $30/168 = 17.8$ percent of their available time based on a 24 hour day.¹⁸ Another interpretation of the household’s time endowment

¹⁶Consumer durables are considered part of the capital stock.

¹⁷Hansen (1993) contains estimates of the productivity profile that are also commonly used in the literature. Simply using Hansen’s profile for $e(s)$ and assuming differences between conditional and unconditional profiles are small gives similar results.

¹⁸The utility function states that the marginal disutility from working all hours is minus infinity, a state-

that is popular in the macroeconomics business cycle literature suggests that households spend about one-third of discretionary time in market activity. The main focus is on the first interpretation, but the sensitivity analysis shows that the life-cycle profiles of labor and consumption are similar under both interpretations.

3.2. Calibration

To find out if the model can generate both the macroeconomic, balanced growth targets and the life-cycle hours profile described above, particular calibrations of the model are considered. The parameters that need to be calibrated are α , δ , η , g , ρ , and σ , the unconditional productivity profile $e(s)$, and the survivor function $Q(s)$. Ages and times have been interpreted as being measured in years and agents are born, or enter the economy, at actual age 25. The life-cycle literature includes models that account for mortality risk¹⁹ as well as models that do not.²⁰ To allow comparison with both strands of the literature, two calibrations are considered. In the *fixed lifespan model*, $Q(s)$ is set to 1 for all $s \in [0, T]$ and the terminal date of lifetimes is set at age 80 so that $T = 55$. This choice of T replicates the average age at death in the U.S. postwar data. In the *mortality model*, $Q(s)$ is fit to a polynomial function using data from actuarial life tables (Arias [2004]) which have a maximum lifespan of 100 so that $T = 75$.²¹

The trend growth rate of income per capita for the postwar U.S. economy is observable and is set as $g = 1.56$ percent per annum. Since the other parameters are not directly observable, they are chosen to minimize a loss function. Because the conditional productivity profile is nearly flat between the ages of 30 and 60, the unconditional productivity profile $e(s)$

ment that seems consistent with a 24-hour day. This target value could be pushed lower by accounting for vacation, holidays, sick days, and the 70 percent labor force participation rate.

¹⁹See, for example, Feigenbaum (2005), Fernandez-Villaverde and Krueger (2002), and Hansen and Imrohorglu (2005).

²⁰For example, Attanasio, et al., (1999), Auerbach and Kotlikoff (1987), and Gourinchas and Parker (2002).

²¹See Feigenbaum (2005) for more details.

is approximated by a linear spline, rather than a polynomial, with nodes at (actual) ages $(s_1 + 25, \dots, s_8 + 25) = (25, 30, 40, 45, 52, 60, 70, 80)$. Let $\theta = (\alpha, \delta, \eta, \rho, \sigma, e(s_1), \dots, e(s_8))$, and define the cross-sectional profile of a variable x as $x_{cs} = x(0; -s)$. Then for each model, choose the baseline calibration to minimize the loss function

$$L(\theta) = \sum_{i=0}^{T_c} \frac{1}{\sigma_c^2} \left(\frac{c_{cs}(i)}{c_{cs}(5)} - \frac{\hat{c}_{cs}(i)}{\hat{c}_{cs}(5)} \right)^2 + \sum_{i=0}^{T_e} \frac{1}{\sigma_e^2(i)} (e_{cs}(i) - \hat{e}_{cs}(i))^2 + \sum_{i=0}^{T_n} \frac{1}{\sigma_n^2(i)} (n_{cs}(i) - \hat{n}_{cs}(i))^2 + \frac{1}{\sigma_{K/Y}^6} \left(\frac{K}{Y} - \widehat{\frac{K}{Y}} \right)^6 + \frac{1}{\sigma_{C/Y}^6} \left(\frac{C}{Y} - \widehat{\frac{C}{Y}} \right)^6 + \frac{1}{\sigma_r^6} (r - \hat{r})^6. \quad (7)$$

In this function, \hat{x} denotes the target value for the equilibrium value of the variable x , and σ_x is an estimate of the error in measuring the target value for x . For an age series variable x , T_x is the maximum age in the data set.²² For the labor and endowment profiles, $\sigma_e(i)$ and $\sigma_n(i)$ are just the standard errors for the cross-sectional means for the hourly wages and hours worked of agents of age i . For the consumption profile, σ_c is the root mean squared deviation of Gourinchas and Parker's (2002) raw consumption series from a polynomial fit. If each variable has a normally distributed measurement error with standard deviation equal to its error estimate, then the first three sums of (7) amount to a linear transformation of minus the log likelihood function for θ . In that case, ignoring the last three terms, minimizing L amounts to maximizing the likelihood. The last three terms then serve as penalty functions that are approximately zero for values of θ that give K/Y , C/Y , and r within a neighborhood of their target values and increase rapidly outside this neighborhood, thereby confining our attention to the region of the parameter space where these macroeconomic quantities are in this neighborhood. Since these quantities are not measured perfectly, the literature accepts a range of values for each variable, and the error estimates are chosen so that the acceptable

²²Gourinchas and Parker (2002) report data on consumption from ages 25 to 64, so that $T_c = 39$. For the productivity data, we include data up to age 60 since wage data will not be collected for the large fraction of non-working agents at later ages; so $T_e = 35$. The CPS data on labor hours goes up to age 79, so $T_n = 54$.

neighborhood covers these ranges.²³

The parameters which minimized the loss function were, for the mortality model, $\alpha = 0.40$, $\delta = 0.07$, $\eta = 0.18$, $\rho = -0.0051$, and $\sigma = 10.63$, and for the fixed lifespan model $\alpha = 0.34$, $\delta = 0.06$, $\eta = 0.17$, $\rho = -0.009$, and $\sigma = 10.24$. These values minimized the loss function, but results are presented for alternative combinations in the robustness section below and qualitative results do not change. One remark concerns the parameter σ , which is the inverse of the elasticity of intertemporal substitution in consumption, the value of which is not well-established. It is important to stress that σ^{-1} is the elasticity of intertemporal substitution for an agent who can freely vary both consumption and leisure. Holding leisure constant, as would be the case in most efforts to measure the elasticity, the elasticity of intertemporal substitution will be $(\sigma^*)^{-1}$, where σ^* is defined as $\sigma^* = 1 + \eta(\sigma - 1)$. For the fixed lifespan model, $\sigma^* = 2.55$, and for the mortality model, $\sigma^* = 2.71$. These values are within the range that have been considered in the literature.²⁴

3.3. Results

This section reports on the characteristics of the constructed balanced growth path with respect to the target values adopted based on U.S. data. Table 1 illustrates that the calibrations are successful in matching, or nearly matching, the U.S. data on all of the dimensions that were addressed.²⁵

How does the model fare with respect to the life-cycle consumption evidence presented in Gourinchas and Parker (2002)? The model's predictions regarding the consumption hump are compared to their mean estimates of life-cycle consumption, as summarized in their Figure 2.

²³In particular, $\sigma_{K/Y} = 0.3$, $\sigma_{C/Y} = 0.03$, and $\sigma_r = 0.005$ (50 basis points).

²⁴The calibration of $\rho < 0$ is common in life-cycle models, which do not require positive time-preference. Rios-Rull (1996) used a value of -0.011 .

²⁵Since this is a life-cycle model, nearly all of the parameters play at least some role with respect to each target, and there is no one-to-one mapping between the various parameters and the targets. Thus, there was no guarantee that all targets could be hit, but, in fact, with the suggested calibration the model comes quite close.

According to their estimates, consumption peaks near age 45. The ratio of peak consumption to age 30 consumption is about 1.10, which is used here as a convenient summary statistic. One of the main results is that the constructed balanced growth paths involve life-cycle consumption profiles that are hump-shaped to about the same degree that the estimates of Gourinchas and Parker (2002) suggest. As mentioned earlier, an alternative estimate of the size of the hump in life-cycle consumption is due to Fernández-Villaverde and Krueger (2002). Their comparable ratio of peak to age 30 consumption, which is for their adult-equivalent nondurables consumption, is higher, about 1.18. The baseline steady states then deliver about 55 to 60 percent of the life-cycle consumption hump according to this estimate. The lifetime consumption profiles from the mortality model, the lifespan model, the Gourinchas and Parker data and baseline model, and the Fernandez-Villaverde and Krueger data are all displayed in Figure 1. All consumption profiles are plotted relative to consumption at age 30.

Perhaps the most salient feature of Figure 1 is that both the fixed lifetime and mortality life-cycle model profiles agree quite well with the Gourinchas and Parker data up to age 47. From age 47 to 57, the life-cycle profiles fall midway between the two empirical profiles, before coming close to the Gourinchas and Parker data at age 60.²⁶ One conclusion is that a life-cycle model with leisure is able to produce a consumption profile that falls within the range found by other researchers. Gourinchas and Parker's baseline buffer-stock model produces a profile which is steeper than the life-cycle model, and the available data, early in life. If $r < \rho$, as it is in Gourinchas and Parker's baseline model (there, $r = 3.5\%$ and $\rho = 4.2\%$ at annual rates), then the standard life-cycle permanent-income hypothesis can account for a downward-sloping consumption profile, and this is what the buffer-stock saving model reduces to late in life after the effects of borrowing constraints and income uncertainty

²⁶No attempt was made to directly match the Fernandez-Villaverde and Krueger data; it is included here for comparison purposes only.

have played out. The part of the consumption hump that is often viewed as difficult to explain in the literature is the upward-sloping portion, but this is the portion the life-cycle model with leisure gets about right.²⁷

3.3.1. Labor supply, labor income, and asset holding

The calibration was designed to match the size of the peak in labor hours supplied over the life cycle. Figure 2 shows the life-cycle profile in hours worked in the model economies plotted against the CPS data. The balanced growth path equilibria of the model economies provide a close match to these data.²⁸

The baseline steady states also have implications for lifetime labor income. According to Gourinchas and Parker (2002), the ratio of peak income to age 30 income is about 1.25. The peak occurs near age 50, after the peak in consumption. The balanced growth path of the mortality model economy has the ratio of peak income to age 30 income at about 1.23, about the same as the Gourinchas and Parker (2002) estimate. It occurs between the ages of 45 and 52, coincident with the peak in consumption. In the literature, Attanasio, et al., (1999) and Carroll and Summers (1991) have estimated income peaks that coincide with consumption peaks.

The households in our model also engage in asset-holding over the life cycle. The most striking aspect of both models' equilibrium asset-holding behavior is that households do not borrow early in life. This occurs even though our model allows agents to borrow freely at any time. Thus our model indicates that it may not be necessary to impose borrowing constraints on households in order to match the stylized fact that little borrowing seems to

²⁷It is important to stress that the two models are still not being compared on an entirely equal footing because while the consumption profiles here were obtained in general equilibrium, Gourinchas and Parker's (2002) baseline profile was obtained in a partial equilibrium setting, allowing an additional degree of freedom to exploit in fitting the data.

²⁸Retirement is endogenous in the model. In the data, there is always some fraction of the labor force that is working. Retirement is defined as average cohort labor supply less than 10 hours per week. Then, agents retire at age 67 under both models.

occur for younger households.²⁹

3.3.2. Intuition

How does the model produce a significant consumption hump through leisure consumption substitutability without also obtaining an unrealistically pronounced hump in life-cycle labor hours, as suggested by Carroll and Summers (1991)? The conditional productivity endowment has a fairly constant slope prior to its peak around age 50. Let us denote the slope of $\ln e(t)$ by S . Then equation (5) can be approximated as

$$\frac{d}{dt} \ln \ell(t; \tau) \approx \frac{r - \rho_{eff}}{\sigma} - \frac{1 + \eta(\sigma - 1)}{\sigma} S. \quad (8)$$

This equation shows that the increasing portion of the labor profile (the decreasing portion of the leisure profile) depends on both the interest rate r and the productivity slope S . For a given S , there is a critical interest rate $r_c = \rho + [1 + \eta(\sigma - 1)]S$ such that the labor and leisure profiles will be flat in this region of linear productivity growth.³⁰ The calibrations involve a large enough value of σ to give a pronounced consumption hump, but also yield equilibrium interest rates close to r_c and hence relatively flat labor hours profiles.³¹

3.3.3. Human capital

Gourinchas and Parker (2002) also reported estimates of life-cycle consumption humps by education group, finding that the hump was generally more pronounced as educational

²⁹Deaton (1991, p. 1236) writes, “There has always been somewhat of a puzzle in the consumption literature as to why individuals who anticipate substantial income growth (e.g. students) and who have a preference for smooth consumption ... do not borrow large sums early in life.”

³⁰Under our maintained assumption of $\sigma > 1$, $r_c > \rho$.

³¹The fact that $r > \rho$ in our model economies comes from our need to match balanced growth and related facts in the calibration exercise. Much of the partial equilibrium consumption-savings literature has employed values of $r < \rho$ in contrast to much of the macroeconomics literature. Many of those models emphasize uninsurable idiosyncratic labor income risk and precautionary savings. It is an open question whether general equilibrium calibrations exist for those models which allow for a match of the balanced growth facts and life cycle facts, particularly the labor hours data, with $r < \rho$. For the model presented in this paper, we could not find a calibration that matched our targets with $r < \rho$. One advantage of a general equilibrium approach is that it places some discipline on the choice of an interest rate, helping to justify the fact that r is close to the value r_c described above.

attainment increases. A natural question is whether the same finding would hold here for the life-cycle model with leisure. To address this question, consider the behavior of a single agent in our economy who has a productivity profile which is different from the average productivity profile. A measure of this profile was obtained using wage data by educational group from the 2002 CPS, ignoring the distinction between conditional and unconditional productivity. Then, consider how this agent would behave within the constructed equilibrium. The results are presented in Figure 3. These results are similar to the empirical estimates reported by Gourinchas and Parker (2002), as the life-cycle consumption hump is more pronounced for agents with higher levels of educational attainment. The data concerning life-cycle consumption by educational attainment is shown in Figure 4.

3.3.4. Robustness

Sensitivity analysis is reported in Table 2, in which parameters are allowed to differ one at a time from the baseline values of the mortality model, keeping all other parameters constant. This approach is intended to allow one to find out which aspects of the calibration have a large influence on the results. The primary message from Table 2 is that the life-cycle peak in consumption is robust to a variety of perturbations on the various calibration choices. Other aspects of the calibrated economies tend to fall away from target (since other parameters are not being adjusted simultaneously), but the consumption hump itself is not very sensitive to these parameter variations.

The last four lines of Table 2 present some alternative economies in which all parameters can vary relative to the mortality model baseline. One is the fixed lifespan model which was discussed earlier, presented here for comparison purposes. Another is a “separable preferences” case in which the period utility function is given by $\eta \ln c + (1 - \eta) (1 - \nu)^{-1} \ell^{1-\nu}$. For this case, the parameter values are those which minimized the loss function (7).³² In this

³²The parameters then are $\alpha = .41$, $\delta = 0.07$, $\eta = 0.27$, $\rho = 0.027$, $\nu = 3.82$.

case the life-cycle consumption hump disappears because there is no consumption-leisure substitutability.³³ The third case labelled “log-linear hazard” uses period utility (2) and baseline parameter values, but replaces the polynomial fit for $Q(s)$ with a survivor function estimated using a log-linear hazard rate model; results are similar to the baseline economy. The fourth case has period utility (2) and simply uses parameter values from Rios-Rull (1996).

4. Conclusions

A general equilibrium life-cycle model with capital has been analyzed in which both consumption and leisure enter the household’s period utility function in a manner consistent with Heckman (1974). The main contribution has been to provide a quantitative assessment of Heckman’s (1974) hypothesis in a general equilibrium setting. A macroeconomic view has been adopted and so the model has been calibrated to match key balanced growth features of the U.S. data, and also to match the relatively flat life-cycle profile in hours worked in the data. Such a model produces an equilibrium with a life-cycle consumption hump similar to that documented by Gourinchas and Parker (2002) and Fernández-Villaverde and Krueger (2002). Thus a key conclusion is that taking account of the effects of consumption-leisure substitutability in household preferences may help explain the life-cycle consumption data that has been viewed as puzzling.

References

- Arias, E., 2004. United States life tables, 2001. National Vital Statistics Reports 52, 14.
- Attanasio, O., Banks, J., Meghir, C., Weber, G., 1999. Humps and bumps in lifetime consumption. *Journal of Business and Economic Statistics* 17, 22-35.
- Auerbach, A., Kotlikoff, L., 1987. *Dynamic fiscal policy*. Cambridge University Press, Cambridge.

³³The slight hump that remains is due to mortality risk alone.

- Benitez-Silva, H., 2000. A dynamic model of labor supply, consumption/saving, and annuity decisions under uncertainty. Working Paper, State University of New York at Stony Brook.
- Browning, M., and Crossley, T., 2001. The life-cycle model of consumption and saving. *Journal of Economic Perspectives* 15, 3-22.
- Browning, M., and Ejrnaes, M., 2002. Consumption and children. Center for Applied Microeconomics Working Paper 2002-06.
- Bullard, J., and Russell, S., 1999. An empirically plausible model of low real interest rates and unbacked government debt. *Journal of Monetary Economics* 44, 477-508.
- Bütler, M., 2001. Neoclassical life-cycle consumption: a textbook example. *Economic Theory* 17, 209-221.
- Caliendo, F., and Aadland, D., 2004. Short-term planning and the life-cycle consumption puzzle. Manuscript, Northern Arizona University.
- Carroll, C., 1997. Buffer-stock saving and the life-cycle/permanent income hypothesis. *Quarterly Journal of Economics* 112, 1-55.
- Carroll, C., and Summers, L., 1991. Consumption growth parallels income growth: some new evidence. In: Douglas, B., Shoven, J. (Eds.), *National Saving and Economic Performance*. University of Chicago Press, Chicago.
- Deaton, A., 1991. Saving and liquidity constraints. *Econometrica* 59, 1221-1248.
- Fernández-Villaverde, J., and Krueger, D., 2001. Consumption and saving over the life cycle: how important are consumer durables? Working Paper, University of Pennsylvania.
- Fernández-Villaverde, J., and Krueger, D., 2002. Consumption over the life cycle: some facts from consumer expenditure survey data. Working Paper, University of Pennsylvania.
- Feigenbaum, J., 2005. Can mortality risk account for the consumption hump? Working Paper, University of Pittsburgh.
- French, E., 2002. The effects of health, wealth, and wages on labor supply and retirement behavior. Working Paper, Federal Reserve Bank of Chicago.
- Gourinchas, P., and Parker, J., 2002. Consumption over the life cycle. *Econometrica* 70, 47-89.
- Hansen, G., 1993. The cyclical and secular behaviour of the labour input: comparing efficiency units and hours worked. *Journal of Applied Econometrics* 8, 71-80.
- Hansen, G., and Imrohoroglu, S., 2005. Consumption and hours over the life cycle: the role of annuities. Manuscript, University of Southern California.
- Heathcote, J., 2002. Home production and retirement. Working Paper, Georgetown University.
- Heathcote, J., Storesletten, K., and Violante, G., 2004. The macroeconomic implications of rising wage inequality in the U.S. Working Paper, Georgetown University.

- Heckman, J., 1974. Life cycle consumption and labor supply: an explanation of the relationship between income and consumption over the life cycle. *American Economic Review* 64, 188-94.
- Hubbard, R.G., Skinner, J., and Zeldes, S., 1994. The importance of precautionary motives in explaining individual and aggregate saving. *Carnegie-Rochester Conference Series on Public Policy* 40, 59-125.
- Huggett, M., and Ventura, G., 1999. On the distributional effects of social security reform. *Review of Economic Dynamics* 2, 498-531.
- Laibson, D., 1997. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112, 443-477.
- Nagatani, K., 1972. Life cycle saving: theory and fact. *American Economic Review* 62, 344-353.
- Rios-Rull, J-V., 1996. Life-cycle economies and aggregate fluctuations. *Review of Economic Studies* 63, 465-489.
- Rust, J., 2002. Using a life cycle model to predict induced entry effects of a government experiment. Working Paper, University of Maryland.
- Yaari, M., 1964. On the consumer's lifetime allocation process. *International Economic Review* 5, 304-317.

TABLE 1. BASELINE BALANCED GROWTH PATH CHARACTERISTICS.

Target	U.S. Data	Fixed Lifespan Model	Mortality Model
Capital-output ratio	2.94	3.38	3.53
Consumption-output ratio	0.748	0.743	0.700
Real interest rate, percent	3.50	4.10	4.42
Fraction of time working	0.190	0.191	0.190
Life-cycle hours peak age	47	45	45
Life-cycle hours peak size	1.09	1.06	1.09
Consumption peak age	44	45 – 52	45 – 52
Consumption peak size	1.10 – 1.18	1.11	1.10

Table 1: The baseline economies compared to targets from the U.S. data. The fraction of time working is calculated for agents aged 25 to 55.

TABLE 2. ROBUSTNESS TO ALTERNATIVE CALIBRATIONS.

Deviation from baseline	Balanced growth			Life-cycle peaks			Work time
	r (%)	K/Y	C/Y	c	y	n	$\bar{n}_{25:55}$
Baseline mortality	4.42	3.52	0.700	1.10	1.23	1.09	0.190
$\sigma = 8$ ($\sigma^* = 2.24$)	3.55	3.81	0.675	1.09	1.24	1.10	0.196
$\sigma = 12$ ($\sigma^* = 2.95$)	4.88	3.38	0.711	1.10	1.22	1.09	0.187
$\rho = 0.005$	5.42	3.24	0.724	1.10	1.23	1.10	0.185
$\alpha = 0.36$	3.87	3.32	0.717	1.09	1.27	1.14	0.191
$\delta = 0.04$	5.42	4.26	0.763	1.12	1.16	1.03	0.185
$\eta = 0.185, \rho = 0.03, g = 0.0$	4.19	3.59	0.750	1.11	1.19	1.07	0.192
$\eta = 0.33, \rho = -0.035$	4.19	3.59	0.694	1.09	1.18	1.06	0.347
Fixed lifespan	4.10	3.38	0.743	1.11	1.18	1.05	0.190
Separable preferences	4.51	3.57	0.695	1.01	1.32	1.16	0.189
Log-linear hazard	4.49	3.50	0.702	1.10	1.22	1.09	0.189
Rios-Rull parameters	3.48	4.05	0.707	1.17	1.46	1.15	0.357

Table 2: Sensitivity of key aspects of the model economy to alternative parameter values. The first column shows the aspect of the alternative calibration which is different from the mortality model baseline, keeping all else constant. Life-cycle peaks refers to the ratio of peak to age 30 values of a variable. Working time is measured as the average across cohorts aged 25 to 55.

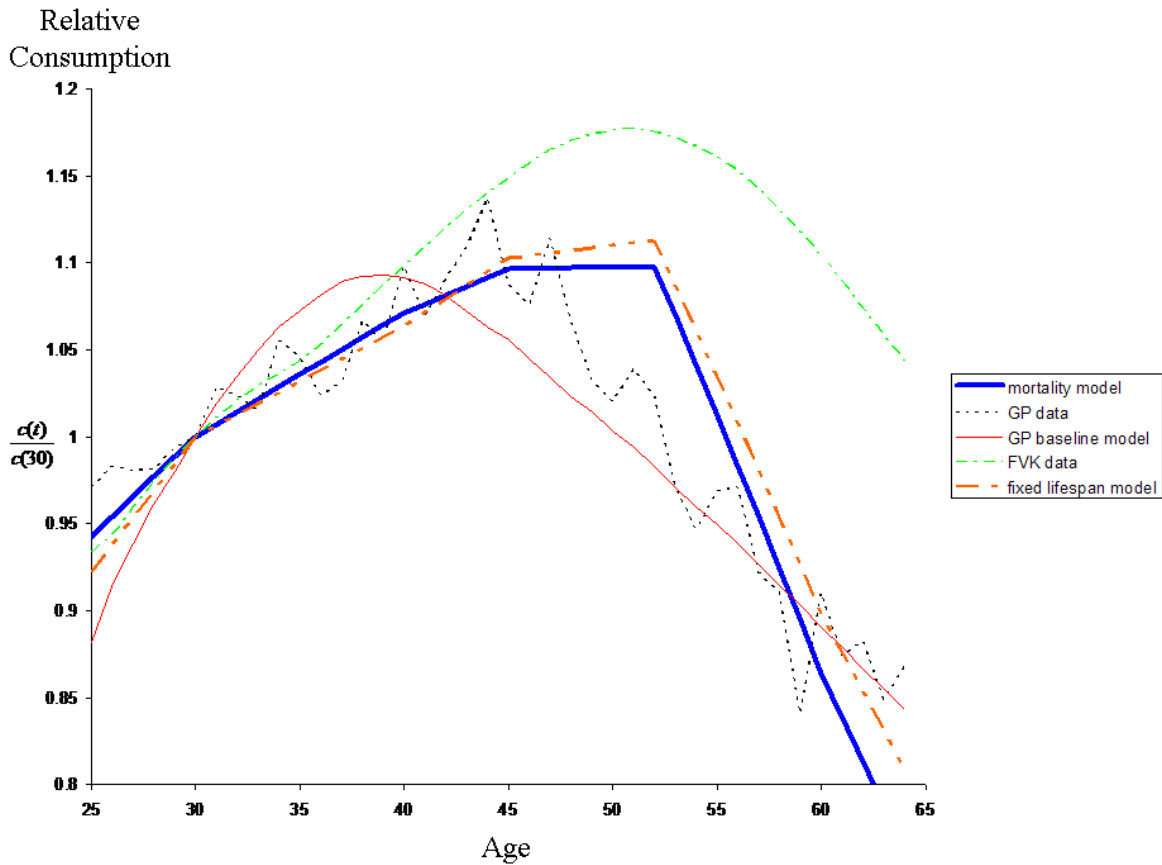


Figure 1: Comparison of the life-cycle consumption profile from both the fixed lifespan model and the mortality model with the empirical profiles of Gourinchas and Parker (2002) and Fernandez-Villaverde and Krueger (2002), as well as with the baseline buffer-stock model of Gourinchas and Parker (2002). The consumption level for each profile is normalized to be 1 at age 30.

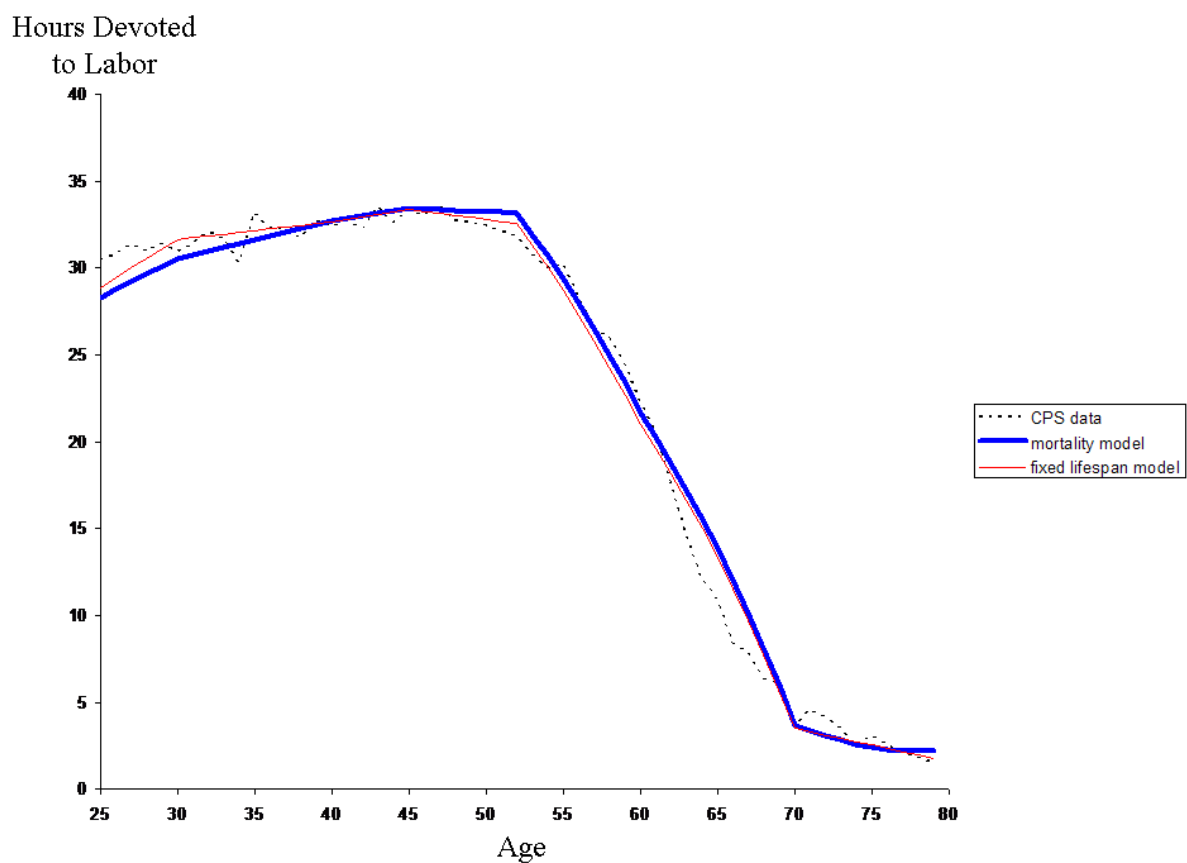


Figure 2: Hours worked by age, in the CPS data and according to both the mortality model and the fixed lifespan model.

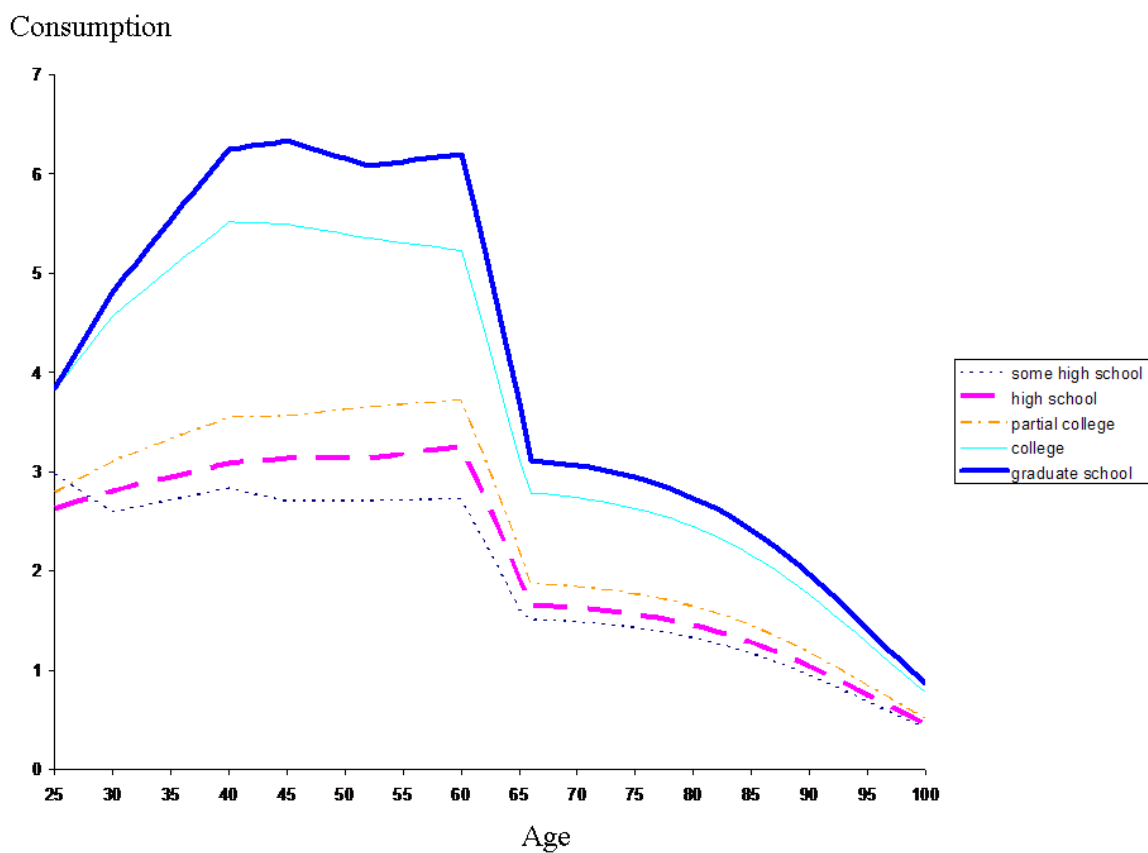


Figure 3: Consumption humps for the mortality model. Consumption humps are larger when educational attainment is higher.




Figure 4: The consumption data by education group between ages 25 and 65. Humps are more pronounced as educational attainment increases.