

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

Lifecycle Dynamics of Income Uncertainty and Consumption

James Feigenbaum and Geng Li

2008-27

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Lifecycle Dynamics of Income Uncertainty and Consumption

James Feigenbaum* and Geng Li†

May 20, 2008

Abstract

Uninsurable income risk is often cited as an explanation for empirical deviations from the Lifecycle/Permanent-Income Hypothesis such as the observation that the life-cycle profile of mean consumption is hump-shaped. Most methods used for estimating income uncertainty essentially measure the cross-sectional variance of a subpopulation rather than the true uncertainty or riskiness perceived by consumers. In this paper, we employ a nonparametric approach to estimate idiosyncratic income uncertainty. We measure income uncertainties as the variance of income forecasting errors at different ages and over different time horizons. The estimated life-cycle income uncertainty profile is U-shaped and generally implies a lower degree of income uncertainty relative to the previous literature. We subsequently use these nonparametric estimates to calibrate a (time-inconsistent) lifecycle model to assess whether a consumption hump can be generated by precautionary saving given more robust measures of income uncertainty. We show that, with plausible risk aversion coefficient and discounting factors and an endogenous, rarely active borrowing limit, our refined measure of income uncertainty is large enough to generate a significant consumption hump that peaks around age 55 and closely matches with the observed magnitude of the consumption hump. We also notice that the variation in the volatility of income shocks with respect to both age and forecast horizon has a significant impact on the size and peak age of the consumption hump.

JEL Classification: E21

Keywords: consumption hump, income risk, time-inconsistent expectations, forecasting errors

*Department of Economics; University of Pittsburgh; 4906 W. W. Posvar Hall; 230 South Bouquet St.; Pittsburgh, PA 15260. E-mail: jfeigen@pitt.edu. URL: www.pitt.edu/~jfeigen.

†Federal Reserve Board. Email address: geng.li@frb.gov. The views presented in this paper are those of the author's and are not necessarily those of the Federal Reserve Board and its staff. We thank Dave DeJong, Karen Dynan, Jonathan Heathcote, Michael Palumbo, and seminar participants at the Federal Reserve Board, SUNY Stony Brook, and the University of Pittsburgh for helpful discussions and comments. All remaining errors are our own.

1 Introduction

It is well documented that household nondurable-good consumption exhibits a hump-shaped profile over the lifecycle. On average, consumption increases when the consumer is young, peaks in middle age, and gradually declines until retirement (Fernandez-Villaverde and Krueger (2007), Gourinchas and Parker (2002), Thurow (1969)).¹ However, this pattern of consumption dynamics is not consistent with the standard Rational-Expectations Lifecycle/Permanent-Income Hypothesis (RE-LCPIH) under the assumption of complete financial markets, geometric discounting, and additively separable preferences for consumption. In this framework, the lifecycle consumption profile should be monotonic (Yaari (1964)). In the case when the consumer's discount rate equals the market interest rate, the consumption profile should be flat as the consumer can perfectly smooth his consumption over the lifecycle.

One popular explanation for a hump-shaped consumption profile is that households may face an income flow with a component of uninsurable risk (Carroll (1997), Feigenbaum (2007), Gourinchas and Parker (2002)). As shown by Leland (1968) and Sandmo (1970), a consumer with risky future income will save for precautionary reasons. This reduces current consumption and, on average, increases future consumption. The rate of consumption growth increases with the variance of income perceived by the consumer (Skinner (1988), Feigenbaum (2008b)). After all uncertainty about income is resolved, the profile will coincide with the path predicted by the RE-LCPIH, which will be decreasing if the interest rate is less than the discount rate. If uncertainty decreases over the lifecycle and is sufficiently large at the beginning, the consumption profile will be concave. Thus the dynamics of uncertainty over the lifecycle have an effect on the shape of the consumption profile. The paper has two objectives. First, we construct a refined measure of uncertainty over the lifecycle. Second, we assess whether this refined measure implies enough variation in lifecycle income uncertainty to account for the observed deviation of lifecycle consumption from a monotonic profile.

We propose a novel way of measuring income uncertainty and apply our estimates to a model of lifecycle consumption. Previous researchers have typically taken the cross-sectional variance of income over a subpopulation as the measure of income uncertainty. However, this approach assumes that households and econometricians have identical information sets. In contrast, Cunha, Heckman, and Navarro (2005) have pointed out individuals may have additional, private information that is relevant to predicting their future income. In other words, one has to be mindful of the distinction between income heterogeneity and uncertainty. The possibility that econometricians may have overestimated the amount of uncertainty faced by individuals creates doubt about the ability of precautionary motives to account for the consumption hump.² Our approach

¹After retirement, a sharp decline in nondurables consumption typically occurs (Bernheim, Skinner, and Weinberg (2001)).

²It also puts into question estimates of how much precautionary motives contribute to

is designed to address this critique.³

Using panel data that provide information about a large number of households over many years, we estimate an income forecasting equation that includes additional information available in the data than is usually used to estimate the predictable component of income. Specifically, we assume that households take into account their lagged and current income as well as impending life events when forecasting their future income. We interpret the sample variance of forecasting errors as the income uncertainty faced by households. Our approach therefore allows us to estimate a volatility matrix that summarizes the uncertainty a consumer should expect as a function of age and forecasting horizon. We also estimate the correlation between forecasting errors at various horizons. Our methodology is nonparametric in that we do not assume any functional form for the income process. We focus on the moments of the income distribution without trying to infer any underlying parameters.

We then incorporate these estimates of the volatility and correlation matrices into a lifecycle model of consumption and saving to assess whether it can account for the hump in lifecycle consumption.⁴ Since we have not estimated the unconditional probability of reaching any possible history of income draws, the consumer cannot update his expectations about future income conditional on his current state using Bayes' rule, as is normally done in the context of rational expectations. Instead, for each age we specify an income process that is consistent with our estimates of the volatility and correlation of forecast errors at each future time horizon. The consumer solves for his optimal consumption function in the present via a backwards recursion while assuming this posited income process governs his income dynamics in all future periods. However, in each ensuing period the consumer will posit a different income process governed by the volatility and correlation matrix estimates for that age. Thus the consumer has time-inconsistent expectations.

The estimated income volatility matrix suggests that income uncertainty perceived by households evolves substantially over the lifecycle. For a given age in the future, we consistently find, as one would expect, that uncertainty about income at that age diminishes as the consumer approaches that age. For a fixed time horizon, when consumers are young, income uncertainty gradually declines with age, presumably as decisions on career, human capital development and fertility are resolved. Income uncertainty reaches its lowest level during middle age. Afterwards, income uncertainty rises again, potentially due to uncertainty about the working hours and health risks. Our U-shaped lifecycle profile of income uncertainty is consistent with Jaimovich and Siu (2006).

aggregate saving (Carroll and Samwick (1998), Feigenbaum (2007b), Gourinchas and Parker (2001)).

³There is also a literature that seeks to address the separate question of how much income is dispersed, irrespective of whether households anticipate that their income may deviate from the mean, so as to address issues of inequality and the sources thereof. See, for example, Heathcote, Storesletten, and Violante (2004), Huggett, Ventura, and Yaron (2007), Krueger and Perri (2001), and Storesletten, Telmer, and Yaron (2000).

⁴Heckman and Navarro (2005) use a similar approach to investigate the effect of income uncertainty on the decision of how much human capital to accumulate.

Studying the CPS data, they construct the business cycle component of hours worked for workers of various age groups and find a U-shaped pattern in the volatility of hours worked by age

We also estimate for each age and horizon the correlation between the forecasting error in the next period (to the current age) and the error in the forecast at that horizon. As one would expect, for a given age these correlations decrease with the forecasting horizon. Except for very young ages, our results are also consistent with the hypothesis that the correlation between the one-year ahead forecast error and the forecast error at some fixed horizon is independent of age.

In comparison to the existing literature, we do find less uncertainty than has been assumed in most previous work on precautionary saving. Nevertheless, for a plausible calibration of the preference parameters, we still find that precautionary saving will lead to a hump-shaped lifecycle consumption profile. Indeed, one criticism of Gourinchas and Parker (2002) is that their estimates of income uncertainty imply so much precautionary saving that they require a risk aversion of one half or a discount factor of 0.84 (Feigenbaum (2007)) to obtain a lifecycle consumption profile with a hump that is not significantly larger than what is seen in the data. Our more modest estimates of income uncertainty produce a consumption hump consistent with the data under more standard preference parameters. We also find that the dependence of income uncertainty on both age and forecast horizon has a significant impact on the size and shape of the consumption hump.

We proceed as follows. Section 2 describes the empirical methodology that we employ. Section 3 discusses the time-inconsistent consumption-saving model and our theoretical results. We also review the pertinent literature in each section. We conclude with some discussion and remarks in Section 4.

2 Measuring Income Uncertainty Over the Lifecycle

2.1 Common Practice and Some Critique

Income uncertainty plays an important role in household decisions regarding consumption, saving, and investment. Traditionally, when rich micro data were not readily available, researchers had to infer the volatility of personal income from aggregate time-series data on GDP, and this volatility was taken as a proxy for income uncertainty. More recent studies of aggregate output and income such as McConnell and Perez-Quiros (2000) have documented the so-called “Great Moderation”, a sharp decline in the volatility of GDP growth since the mid 1980s. However, aggregate data can mask important variations and correlations in income at the household level, so a decrease in the volatility of aggregate income volatility does not necessarily imply that income at the

household level has become less uncertain. To see this, consider a hypothetical economy populated by two consumers. Every period, each consumer receives one unit of endowment and then they engage in a zero-sum game of chance with their endowments. After they finish gambling, the aggregate income of the economy remains fixed at two with no uncertainty at all, but one cannot say the income of each individual consumer is risk-free, though their incomes are perfectly (negatively) correlated.

Studying household income directly avoids this difficulty. To this end, large panel surveys of household income, such as the Panel Study of Income Dynamics (PSID), have become increasingly available. Previous work has routinely used either cross-sectional or time-series variances of income as a proxy for income uncertainty. For example, Dynan, Elmendorf and Sichel (2007) simply focus on the standard deviations across households of the percent change in household income.⁵ More elaborate models often postulate some parsimoniously parameterized income process that includes a predictable part and a stochastic part, which in turn is some combination of permanent and transitory shocks. In practice, econometricians estimate the predictable part of income as a function of an age polynomial and other demographic variables as well as year dummies to filter out economy-wide time variation (Carroll 1994, Carroll and Samwick 1997, Gourinchas and Parker 2002). The residual of observed income relative to estimated income is then interpreted as the stochastic component. Various techniques can be applied to investigate the variance and persistence of the stochastic part to quantify income riskiness.

As an alternative approach, instead of estimating an income trend, some researchers take the mean of the income of a household over a given period as a proxy for permanent income and treat the gap between observed income and the mean as the transitory component. In this manner, Gottschalk and Moffitt (1994) argued an increase in the variance of transitory earnings could explain a large portion of the widening of income inequality during the 1970s and 80s.

Essentially, what these methods attempt to measure is either the cross-sectional variance or the time-series volatility of a specific component of income that is orthogonal to specified information sets. Conceptually, it is true that greater income uncertainty will create larger cross-sectional variance and/or higher time-series volatility. However, it is not generically true that a larger variance or higher volatility always implies greater income uncertainty. Several important caveats must be kept in mind when interpreting such variations as a measure of the underlying income uncertainty.

First, fitting individual income with a trend driven by age and other demographic characteristics requires the assumption that all individuals share a common life-cycle income trend. This is the view introduced by, among others, MacCurdy (1982). A competing view is that individuals face individual-specific income profiles, as proposed by Lillard and Weiss (1979). Guvenen (2007) labels the MacCurdy-type income process as a restricted income profile, and

⁵Dynan, Elmendorf and Sichel (2007) also provide an elegant and comprehensive survey of the literature studying household income volatilities using household data.

the Lillard-and-Weiss-type income process as a heterogeneous income profile. He presents evidence that consumption data is more consistent with the view of the heterogeneous income process. If the underlying income process is better characterized in this way, fitting it with a trend common to all households will increase the residual variance and hence exaggerate the apparent income uncertainty.

Second, the presumed information set on which the stochastic component is estimated could be only a potentially small subset of the full information set possessed by households. Consequently, what an econometrician treats as unpredictable variations may be predictable to households. If so, measured variations will also exaggerate apparent income uncertainty.

As a more concrete example, consider the following simple hypothetical economy where household i 's income at time t is given by

$$y_{i,t} = e^{\psi_i t},$$

where ψ_i , the rate of income growth, is a random number assigned by nature to household i . Suppose ψ_i is known to the household but is not observed by the econometrician. The common practice in the literature is to take the mean of ψ_i over the population as the rate of predictable income growth for all households. If an econometrician measures income uncertainty as the dispersion of income relative to this assumed trend, he will find that income uncertainty grows quadratically with respect to time, but this is a spurious conclusion since all households have a deterministic income process. Likewise, if the econometrician cannot identify the income trend for each household, the estimated time series volatility will also exaggerate the apparent income uncertainty.

Finally, previous studies have not fully characterized how household income uncertainty evolves over the life cycle. The variance and persistence of income shocks are rarely postulated to depend on age. However, it is both theoretically and empirically appealing to study whether the income uncertainty perceived by households does stay constant over the life cycle. Intuitively, a single 22-year old college graduate first entering the labor market should have much greater uncertainty about his income five or ten years down the road than a 40-year old with a spouse, both of whom have settled on career paths.

2.2 A Forecast-Based Non-Parametric Approach

What method can more consistently and accurately measure income uncertainty, and also shed light on its evolution over the life cycle? Uncertainty stems from risk. Heuristically, greater uncertainty should make future income more difficult to forecast. Intuitively then, for a given income process the income level at a remote future should bear more uncertainty than the income at a near future. Conversely, we may use the forecast accuracy to approximate and evaluate the underlying uncertainty. The larger the variances of forecast errors are the greater uncertainty a household has about its future income.

Let income at age t be y_t . Then at age t , the s -period-ahead income

can be decomposed as

$$y_{t+s} = E[y_{t+s}|\mathcal{I}_t^H] + \varepsilon_{t+s}^t, \quad (1)$$

where $E[y_{t+s}|\mathcal{I}_t^H]$ is the mathematical expectation of y_{t+s} conditioned on age- t household information, \mathcal{I}_t^H , and ε_{t+s}^t is an error term orthogonal to \mathcal{I}_t^H . We define the realized uncertainty associated with the s -period-ahead income for age- t consumers, $\omega_{t,s}$, as the variance of this error term across all the consumers with same age, t , or

$$\Omega_s^t = \text{Var}[\varepsilon_{t+s}^s|\mathcal{I}_t^H] = \text{Var}[y_{t+s} - E[y_{t+s}|\mathcal{I}_t^H]|\mathcal{I}_t^H]. \quad (2)$$

This approach is nonparametric because it does not presume that income shocks follow any specified process. We characterize lifecycle income uncertainties using two $L \times P$ matrices, Ω and Θ , where we have data for L ages over P time horizons.. Element $\omega_{t,s}$ of the Ω matrix is the variance of the s -year-ahead forecast errors of all age- t households whereas element $\theta_{t,s}$ of the Θ matrix is the correlation between the s -year-ahead forecast error and the 1-year-ahead forecast error of age- t households. In examining the matrices Ω and Θ , we can explicitly study the lifecycle dynamics of income uncertainty and its implications on consumption and asset holding patterns over the lifecycle.

One obstacle to implementing this strategy is we do not know the joint distribution of y_{t+s} and \mathcal{I}_t^H , and, therefore, we cannot compute $E[y_{t+s}|\mathcal{I}_t^H]$ directly. Indeed, we do not even know exactly what \mathcal{I}_t^H encompasses. To understand how severe this superior information problem could be and to provide some alleviation, we experiment with two specifications. First, in what we call the restricted information specification (RIS), we project y_{t+s} conditional on \mathcal{I}_t^R , the information set that households *certainly* have at age t since econometricians collect this data from the households at that time. Second, in what we call the augmented information specification (AIS), we project y_{t+s} condition on the augmented information set \mathcal{I}_t^A , where

$$\mathcal{I}_t^A = \mathcal{I}_t^R \cup \mathcal{I}_t^F. \quad (3)$$

The augmenting information set, \mathcal{I}_t^F , contains elements that econometricians observe as of age $t + s$, the target age of the forecast, but that households *likely or possibly* know at age t . Put differently, \mathcal{I}_t^F approximates the superior information that households possess but econometricians do not observe concurrently. To fix the idea, we estimate the following RIS equation

$$y_{i,t+s} = \alpha + \beta_0 y_{i,t} + \beta_1 y_{i,t-1} + \beta_2 y_{i,t-2} + \gamma Z_{i,t} + \xi Trend_{t+s} + \varepsilon_{i,t+s}^t, \quad (4)$$

and AIS equation

$$y_{i,t+s} = \alpha + \beta_0 y_{i,t} + \beta_1 y_{i,t-1} + \beta_2 y_{i,t-2} + \gamma Z_{i,t} + \delta Q_{i,t+s} + \xi Trend_{t+s} + \varepsilon_{i,t+s}^t. \quad (5)$$

In the above equations, $Z_{i,t}$ is a vector of variables that belong to \mathcal{I}_t^R . This

includes race, education level, marital status, family size, a dummy of whether household members are currently laid off or unemployed, a dummy of whether household members are self-employed, and a vector of occupation and industry dummies, all evaluated at age t , when expectations about income at $t + s$ are formed. In addition $Z_{i,t}$ includes a fourth-order polynomial of imputed years of working experience, evaluated at the forecast horizon, $t + s$. Q_{t+s} is a vector of variables that belong to the augmenting information set, \mathcal{I}_t^F . We assume Q_{t+s} includes family size, marital status, a retirement dummy, a self-employment dummy, a vector of occupation and industry dummies, all evaluated at $t + s$.

Besides \mathcal{I}_t^R and \mathcal{I}_t^A , our specification departs from most previous specifications for estimating the predictable income component in that we include not only standard demographic and employment variables, but also the current and lagged income in our projection equation. In principle, if we have a very long income history for a given household, even a univariate ARIMA model could potentially have decent forecasting power. Including some recent income history can help to tease out information about recent income shocks, and capture part of the individual-specific component of income growth that is emphasized by Lillard and Weiss (1979) and Guvenen (2007). In our specification, we include two lags to preserve degrees of freedom. Finally, we add a simple calendar year trend to control for aggregate economic growth.

There are several important caveats we need to point out. First, because most households do plan ahead regarding their family and career, it is not unreasonable to assume households know several years ahead of time what their family size and marital status will be; whether they will be working, retired, or self-employed; and whether they will change occupation and industry. However, it is more difficult to justify that households know all this information as the time horizon gets long.⁶ Therefore if s is large, we might not always have $\mathcal{I}_t^F \subset \mathcal{I}_t^H$. Second, we try our best to project y_{t+s} on an information set as close to \mathcal{I}_t^H as possible, but it is still possible that there exist some information elements ι such that $\iota \in \mathcal{I}_t^H$, but $\iota \notin \mathcal{I}_t^R \cup \mathcal{I}_t^F$. Third, equations (4) and (5) should be interpreted only as a forecasting equation, instead of a structural income equation. Finally, although we allow for age-varying income uncertainty, we compute these forecast errors by applying the same parameters of the forecasting model to all households of different ages. As a robustness check, we estimate the forecasting model separately for each age group. Our results are qualitatively preserved, which reassures us that the age profile of income uncertainty is not driven by the across-age inconsistency of the fitting accuracy of the forecasting model.

2.3 Data Description and Sample Construction

We use data from the PSID. The PSID is a nationwide household longitudinal survey conducted by the Institute of Social Research at the University of Michigan. Before 1997, the PSID was an annual survey, while after 1997 it

⁶Our study forecasts future income up to twenty-five years ahead.

became a biannual survey. As of year 2008, there are 34 waves of data that cover 37 years. The first wave of data was collected in 1968, and the latest wave was collected in 2005.

The PSID not only surveys the households in the original sample constructed in 1968, but also the households headed by the grown-up children of the original sample of households. The sample of households in the PSID has a very high retention rate. The vast majority of the households that the PSID surveys in one year will continue to participate in the next wave. There are more than 1,200 households that stayed in the survey for more than 30 years. Consequently, the sample size of the survey has grown considerably since 1968. The first wave of the PSID had only 4802 households, whereas the 1994 wave had more than 10,000 households. The PSID subsequently stopped surveying households in its non-core sample. As of the most recent wave of 2005, the survey has 8002 households.

Besides extensive information about work status, employment history and demographic characteristics, the PSID has detailed information on household income. In the current paper we want to study the relationship between the household's income uncertainty and consumption dynamics over the life cycle. For our definition of income, we will focus on the household total income, which is the most relevant measure of income vis-a-vis household consumption.

The longitudinal structure of PSID data allows us to link a household in time t to the same household in time $t + s$. We will treat the same household in different waves as independent households and not exploit the econometric properties of the longitudinal structure of the data. Assume a household was surveyed over ten years from 1971 through 1980. When we project the five-year-ahead income, this household renders five current and future income pairs $(t, t + s) = (1971, 1976), \dots, (1975, 1980)$. We will simply pool these pairs together without controlling for the household fixed effect.

Several selection rules apply when we construct the sample. First, we restrict the heads of our sample households to be those younger than 66 years old at the year to be forecasted. Thus, if we project income of year $t + s$ as of year t , we keep only the households whose heads are younger than $66 - s$ as of year t . For example, in the sample we use to forecast five-year-ahead income, we restrict the heads of sample households to be younger than 61 years old. Consequently, the sample we use to study forecast errors at farther horizons is smaller than the sample used for a closer horizon. On the other side of the age restriction, we remove all households whose heads are younger than 23 years old in the base year. In addition, we remove households whose heads are disabled or retired, are primarily keeping house, or are students in the base year. We further remove households whose heads report zero working hours in the base year and the previous two years or did not report valid industry or occupation information. Finally, in order to minimize the effects of outliers, we trim off the households with very high or very low income levels and growth rates ⁷.

⁷We trim off the top and bottom 1% of lagged, current, and future income level distributions and the distribution of income growth between year t and $t + s$.

Forecast Horizon	1 Year	2 Years	3 Years	4 Years	5Years
Number of Observations	72905	71722	62027	59547	53479
Forecast Horizon	6 Years	7 Years	8 Years	9 Years	10 Years
Number of Observations	49365	44424	40132	36049	32137
Forecast Horizon	11 Years	12 Years	13 Years	14 Years	15 Years
Number of Observations	28614	25320	22348	19505	16934
Forecast Horizon	16 Years	17 Years	18 Years	19 Years	20 Years
Number of Observations	14504	12491	10514	8921	7331
Forecast Horizon	21 Years	22 Years	23 Years	24 Years	25 Years
Number of Observations	6034	4801	3758	2831	2003

Table 1: Number of Observations for Each Forecast Horizon

Variable	Mean	StdDev	Variable	Mean	StdDev
Log(Family Income)	10.21	0.62			
Headage	39.81	10.68	Family Size	3.31	1.75
< High School	29.1		High School	23.3	
Some College	29.4		College	18.2	
White	66.6		Married	71.1	

Table 2: Summary Statistics of the One-Year-Ahead Forecast Sample

In our study, we estimate the variance of forecast errors for time horizons up to 25 years. Because we include two lags of income in our forecasting equation and the PSID started in 1968, the first base year is 1970. We do not use the PSID 2003 and 2005 data because these waves used different occupation and industry codes that cannot be mapped to those used in the previous waves.⁸ Table 1 lists the number of observations used in estimation for each forecast horizon. Table 2 provides summary statistics of key variables at the base year for the largest sample, the one-year-ahead forecast sample. The family income variable is deflated using 1982-1984 dollars.

⁸All waves but 2003 and 2005 of the PSID data have 1970 census industry and occupation code. The 2003 and 2005 PSID data used the industry and occupation code derived from the 2000 census.

2.4 Estimation Results

We estimate forecast-error variances for households with heads between 23 and 64 years old. To further remove possible noise in the series, we report five-year-centered moving averages for households between 25 and 62 years old. Because we have from one- to twenty-five-year-ahead income projections, our income uncertainty matrix Ω has dimensions 38×25 ⁹. Figures 1 and 2 present several lifecycle income uncertainty profiles at various forecast horizons. Figure 1 plots the AIS results and Figure 2 plots the RIS results.

We first focus on the AIS results. The top panel of Figure 1 shows near-term income uncertainties. Not surprisingly, the level of uncertainty regarding the two-year-ahead income is higher than the uncertainty associated with the one-year-ahead income. Apart from the magnitude, both uncertainty profiles share a similar U-shaped pattern over the lifecycle. Income uncertainty is high when consumers are in their mid to late twenties. Income uncertainty continues to decline through the mid thirties. After that uncertainties stay at a relatively stable level before rising again in the mid forties. This rise continues into ages close to retirement. The ratio between the maximum and minimum uncertainties over the lifecycle (the max-min ratio) is 1.46 for the one-year-ahead income and 1.37 for the two-year ahead income.

Similar, but less pronounced, U-shaped profiles repeat in the middle panel, which shows medium-term income uncertainties. The uncertainty associated with the five-year-ahead income hits bottom in the mid thirties and stays quite flat until the mid forties. The uncertainty profile of the ten-year-ahead income also exhibits a U-shape, but it bottoms at a somewhat earlier age. In addition, the max-min ratio is 1.30 for the five-year horizon and 1.22 for the ten-year horizon. Both are appreciably lower than the max-min ratios of uncertainties at nearer horizons.

Finally, income-uncertainty profiles at remote horizons, e.g. fifteen and twenty-five years ahead, are plotted in the bottom panel. These curves also show appreciable U-shaped patterns. Because of the age restrictions we impose, these uncertainty profiles cover a much shorter age-span. The max-min ratios are 1.25 for both series, lower than the ratio in the near term.

Now turn to Figure 2, the RIS results. Three features of this figure are noteworthy. First, the contours of these life cycle income uncertainty profiles are very similar to those in the AIS results, shown in Figure 1. Second, because these future income projections are conditioned on a smaller and more restrictive information set, the RIS forecast error variances are greater than the AIS variances across all forecast horizons and for consumers of all ages. Third, the discrepancy between the RIS and the AIS results widen with the forecasting horizon, s . On average, the RIS one-year-ahead uncertainty is only 4% higher than the AIS uncertainty, whereas the margin is above 30% at the twenty-five-year horizon. This is not surprising because the difference between \mathcal{I}^R and \mathcal{I}^A is what households might know in the base year but which econometricians only observe after a s -year lag. If s is small, the correlation between the elements

⁹The lower triangle of the matrix is not populated.

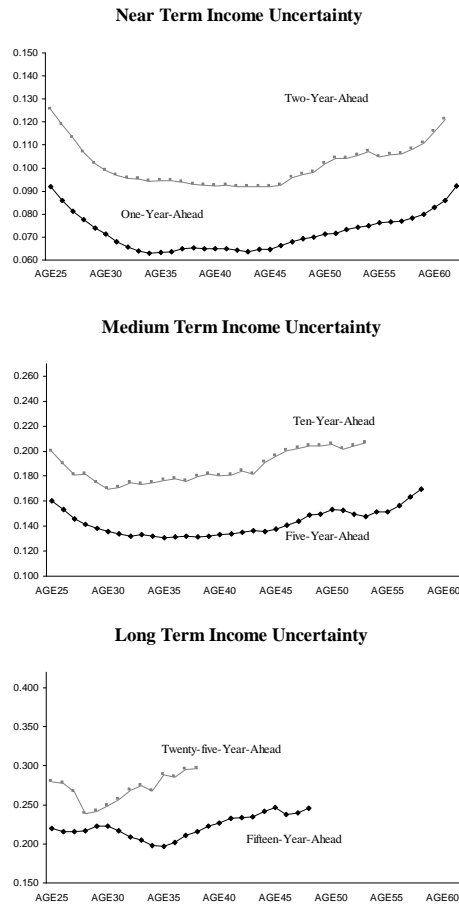


Figure 1: Variance of forecast errors for different forecast horizons as a function of age for the augmented information set (AIS) estimates.




Figure 2: Variance of forecast errors for different forecast horizons as a function of age for the restricted information set (RIS) estimates.

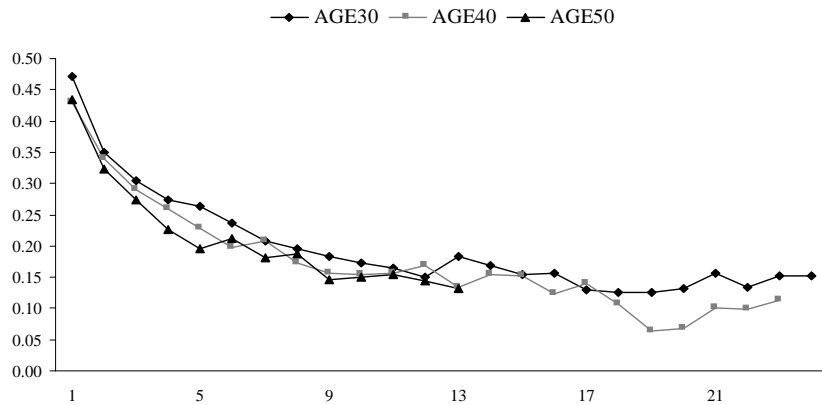


Figure 3: AIS estimates of the correlations between one-year-ahead forecast errors and forecast errors h years ahead for ages 25, 35, and 45.

in \mathcal{I}^R and \mathcal{I}^F is very high. That is to say the net value of adding \mathcal{I}^F is quite limited. If s is large, \mathcal{I}^R has little predictive power on \mathcal{I}^F , so introducing \mathcal{I}^F adds much more new information and consequently beefs up the forecasting performance.¹⁰

Beside the magnitudes of income uncertainties over the life cycle, we are also interested in the correlations among the stochastic components of income at various horizons. We compute for each age group the correlations between the one-year ahead forecast errors and the forecast errors at other horizons. Figure 3 presents the correlations of consumers that are thirty, forty, and fifty years old in the AIS model. The chart shows that, apart from the longest forecast horizons, the correlations of consumers of various ages are quite similar. We note the following patterns in this graph. First, the correlations decline with the forecast horizon. Second, for all age groups, the correlations between forecast errors at one- and two-year horizons are about the same and slightly below 0.5.

¹⁰We will further illustrate this effect in Figure 4

2.5 Robustness Test

To verify whether the U-shaped uncertainty profile over the lifecycle is a spurious consequence of our model specification, sample size, or sample selection, we conduct a series of robustness tests. First, we examine whether the changes in forecast-error variances are due to model misspecification. Remember we project the future income of households of different head ages using the same set of coefficients. If the projection equations should be age-specific and the projection equation we use is closer to the true equations for middle-aged households than to the true equations for younger and older households, the one-size-fit-all approach will reduce fitness for younger and older households and artificially increase income uncertainties for these age groups. We divide our sample into five subgroups by head of household age and reestimate Eqs. (4) and (5) separately for each subgroup. Then we calculate forecast-error variances as we did before and we find the U-shaped uncertainty profiles are qualitatively preserved.

Second, we examine whether changes in the sample size as we vary the forecast horizon (as given in Table 1) might drive the shape of the uncertainty profile. We reestimate the forecast equations using a smaller common sample and reassess income uncertainties over the lifecycle. Apart from the fact that more matrix elements cannot be estimated accurately because of the smaller sample size, the magnitude and dynamics of income uncertainty are very similar to what we presented above.

Finally, we test if our results are driven by low-income households, which are oversampled by the PSID. The core PSID sample consisted of two independent samples – a nationwide representative sample and a sample of low-income families. In the first wave of the survey, the nationwide representative sample has about 3,000 households and the low-income sample has about 2,000 households. We redo our analysis using only the households in the nationwide representative sample and their offspring. The results are very similar to those obtained using the entire PSID core sample.¹¹

2.6 Discussion and Comparison with Earlier Results

How substantive are the innovations we have introduced into these nonparametric measures of income uncertainty? How different are our results compared to previous results in the literature? We answer this question by contrasting our results to the income uncertainty estimates in the influential work of Carroll and Samwick (1997) and Gourinchas and Parker (2002).

We first briefly review the methodology adopted in Carroll and Samwick (1997)¹². The logarithm of income, y_t , is decomposed into a permanent component, p_t , and a transitory shock, ϵ_t , where p_t is further assumed to follow a

¹¹For more information about the PSID sample design, see the online documentation at <http://psidonline.isr.umich.edu/Guide/Overview.html>.

¹²We use the same notations as in their paper.

random walk with predictable income growth g_t such that

$$p_t = g_t + p_{t-1} + \eta_t. \tag{6}$$

and η_t is the permanent income shock. Let σ_η^2 and σ_ϵ^2 be the variance of the permanent and transitory shocks. It is easy to show that

$$\text{Var}[y_{t+d} - y_t] = d\sigma_\eta^2 + 2\sigma_\epsilon^2, \tag{7}$$

noting that the econometrician does not know how either y_t or y_{t+d} decomposes into their permanent and transitory parts. To estimate g_t , Carroll and Samwick (1997) fit actual income data with age, occupation, education, industry, household demographic variables, and age-interaction terms, adjusting for economy-wide growth in income. Subsequently, σ_η^2 and σ_ϵ^2 can be estimated by evaluating $\text{Var}[y_{t+d} - y_t]$ at various time horizons d . They report that for the full sample $\sigma_\eta^2 = 0.022$, and $\sigma_\epsilon^2 = 0.044$. Gourinchas and Parker (2002) employ the same methodology and redo the estimation using an updated sample and report almost identical results.¹³ We denote their results CS-GP. Figure 4 contrasts the income uncertainty implied by Eq. (7) at various future horizons using CS-GP parameters and our nonparametric estimates from both the AIS and RIS specifications. To be consistent with CS-GP measures, in Fig. 4 our variances are calculated at each forecast horizon using the whole sample of households of all relevant ages. For example, the ten-year-ahead variance is computed using the forecast errors of all households whose heads are between 23 and 55 years old.

The graph shows that at near horizons, the CS-GP uncertainty estimates are about 30% to 40% greater than the estimated income uncertainty under the AIS assumptions, and 15% to 20% greater than estimates under the RIS assumptions. The gap widens substantially at farther horizons. Beyond a twenty-year horizon, the CS-GP variance estimates more than double the AIS estimates and are over 50% higher than the RIS estimates.

2.7 Summary of Empirical Findings

We construct a nonparametric measure of income uncertainty and study its dynamics over the lifecycle. Our estimates of income uncertainty are typically smaller than previous studies have documented. Our estimates also imply less persistence in income shocks. Over the lifecycle, we find robust U-shaped patterns in the evolution of income uncertainty. Young and old consumers on average have more risky future income relative to middle-age consumers. This U-shaped pattern is robust to a number of sample and model specifications and prevails at almost all horizons.

¹³When we re-estimate the CS-GP model using our sample, which is much longer and allows us to include the first-order difference of fitting residuals many years apart, we find somewhat larger estimate of variance of transitory shock but smaller estimate of variance of permanent shock. We interpret this as an evidence indicating that the persistent component of the income shocks are not exactly a random walk process.

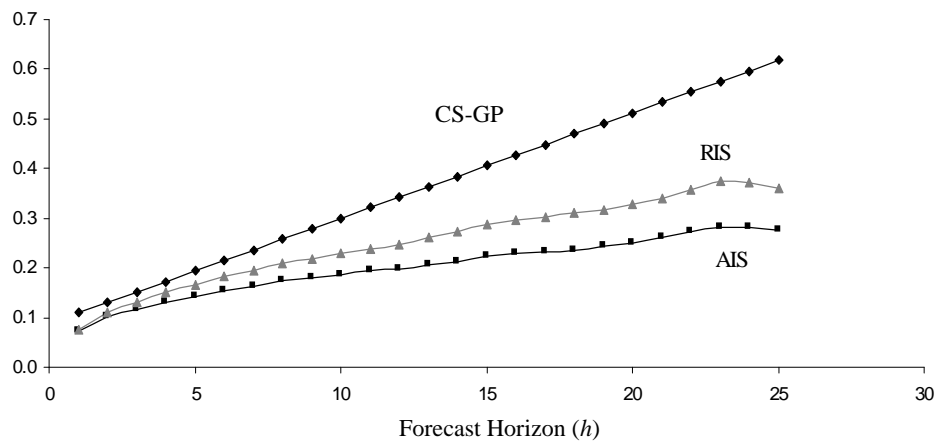


Figure 4: Comparison of estimates of the variance of forecast errors at different forecast horizons for the entire relevant sample using the approach of Carroll and Samwick (1997) and Gourinchas and Parker (2002) versus our approach with both the RIS and AIS specifications.

3 A Lifecycle Model with Time-Varying Income Uncertainty

3.1 Existing Theory of the Consumption Hump

Modern consumption theory begins with the standard Rational-Expectations Lifecycle/Permanent-Income Hypothesis: if consumers are rational then they should allocate consumption over the lifecycle in such a way as to maximize lifetime utility rather than according to simple rules of thumb, such as to consume a constant fraction of disposable income. Thurow (1969) first noted that empirical patterns of lifecycle consumption are hump-shaped and quite similar to the profile of income over the lifecycle. On the face of it, this would appear to refute the RE-LCPIH. However, there are several modifications to the standard model that can account for a hump-shaped consumption profile.¹⁴ Each of these modifications introduces testable predictions outside the realm of nondurable consumption. It is an open question how well competing models can simultaneously account for the mean dynamics of lifecycle consumption and the relevant data from other areas of economics, especially since different modifications do not necessarily complement each other.

Often studied in conjunction with borrowing constraints, precautionary saving is the most popular explanation for the consumption hump and the one we focus on in this paper. Precautionary saving will arise if we dispense with the assumption of complete markets so agents are unable to perfectly insure themselves against idiosyncratic risk (Leland (1968), Sandmo (1970)). Nagatani (1972) first suggested that precautionary saving would reduce consumption early in the lifecycle, and Skinner (1988) and Feigenbaum (2008b) have fleshed out how the growth rate of mean consumption from one period to the next increases with income uncertainty. Consequently, if income uncertainty decreases over the lifecycle, this will lead to a concave consumption profile. Using previous measures of uncertainty as described in Section 2.1, several researchers, including Carroll (1997), Carroll and Summers (1991), Gourinchas and Parker (2002), and Hubbard, Skinner, and Zeldes (1994), have documented with calibrated partial-equilibrium models that a combination of borrowing constraints and precautionary saving can produce a hump-shaped consumption profile similar to the data. Feigenbaum (2007) showed that, in general equilibrium, precautionary saving could better account for the hump absent no-borrowing constraints with Gourinchas and Parker's (2002) specification of the income process, which in particular assumes that permanent income shocks are highly persistent and follow a unit-root process.

Nevertheless, precautionary saving is far from the only explanation for a hump-shaped consumption profile, but it is the only one that crucially depends on income uncertainty, which is why it is essential to document precisely how much income uncertainty consumers actually face. Continuing the explo-

¹⁴For a more detailed review of this literature, see Browning and Crossley (2001).

ration of incomplete markets, Fernandez-Villaverde and Krueger (2005) have established there is also a hump in durable-goods consumption, and they have found that if durable goods have to be used as collateral for loans then the need to purchase durable goods before borrowing can also produce a hump-shaped profile for consumption of both types of goods. Heckman (1974) and Becker and Ghez (1975) suggested that if leisure and consumption are substitutes then a hump-shaped profile of wages over the lifecycle would induce a hump-shaped profile of consumption. This mechanism also has the side effect of a hump-shaped profile of labor hours, but Bullard and Feigenbaum (2007) recently found there are reasonable calibrations that can match both the consumption hump and the labor-hours profile. Time-varying mortality risk can also explain the hump (Feigenbaum (2008a), Hansen and Imrohoroglu (2006)), though only for a specific set of parameters.

The difficulty of separating out the individual effects of age and other confounding variables means that some portion of the consumption hump may be artificial and due to measurement error. Attanasio et al (1999) and Browning and Ejrnæs (2002) have argued that variations in household size over the lifecycle could play such a confounding role while Aguiar and Hurst (2003, 2007) have shown that the substitution of home production for market consumption may also vary with age.

Of course, a consumption hump can also arise if we abandon the assumption of perfectly rational decision-makers with time-consistent models. Laibson (1997) proposed that hypergeometric discounting could explain the hump while Caliendo and Aadland (2004) have shown that if households can only make plans over a ten- to twenty-year horizon then this can also account for the hump. While the main purpose of this paper is to study the lifecycle dynamics of income uncertainty and its impact on precautionary saving, the theoretical model we introduce in the next section also, for practical reasons, has an element of time-inconsistency.

3.2 The Model

Our characterization of uncertainty, detailed in Section 2, is fundamentally nonparametric since we only measure moments of the income process and do not estimate a parametric specification. Defining uncertainty in nonparametric terms has the advantage that our results are model-independent, but it has the disadvantage that we have no ready-made model of the income process that can be incorporated into lifecycle behavioral models. If we ignore the correlation data we have collected, we could suppose that the income process is simply a sequence of independent shocks with variances given by our volatility matrix. However, it is well established that, for plausible calibrations of utility, income uncertainty will only have a significant impact on consumption and saving if income shocks are persistent (Skinner (1988)). It is not enough for the realization of income to be an uncertain event. Shock in earlier periods have to reveal prior information about this income, so each piece of information the

household anticipates it will receive about this income carries some uncertainty also (Feigenbaum (2008b)). It is the combined effect of all these information shocks that accounts for the large effects of precautionary saving in models such as Feigenbaum (2007) or Gourinchas and Parker (2002). In order for income in one year to convey information about income in later years, these income shocks must be dependent.

In a rational expectations model, a household will, by definition, use Bayes' Rule to update its beliefs about the future as its income path is revealed. This requires a complete specification of the probability for every possible income history. Calibrating a time-consistent model that matches our volatility and correlation matrices is a complicated, multidimensional problem that we eschew in this paper. Instead, we calibrate a separate Markov process for each age group that describes how the household believes its income will evolve from that period on. The income process for each age group is calibrated to match the moments of forecast errors for that age group, but we do not impose any restrictions on the collection of these processes that would be necessary to insure time-consistent expectations. Thus we allow households to have time-inconsistent expectations.

We consider a partial-equilibrium model with a consumer who lives for T_w working periods and T_r retirement periods. For each age $t = 0, \dots, T_w - 2$, the consumer believes future income is determined by a stochastic process that matches the moments obtained from age- t income forecasts, and we do not require these beliefs to be consistent with Bayes' Law. Let $E^{(t)}$ be the expectation operator with respect to the consumer's beliefs at age t . An age- t consumer then maximizes

$$E^{(t)} \left[\sum_{s=t}^{T_w-1} \beta^{s-t} u(\tilde{c}_s(t); \gamma) + \beta^{T_w-t} V_{T_w}(\tilde{x}_{T_w}(t)) \right],$$

where $c_s(t)$ is consumption planned for period s as of age t ,

$$u(c; \gamma) = \begin{cases} \frac{c^{1-\gamma}}{1-\gamma} & \gamma \neq 1 \\ \ln c & \gamma = 1 \end{cases} \quad (8)$$

for $\gamma > 0$, $x_{T_w}(t)$ is financial wealth at retirement, and $V_{T_w}(x_{T_w})$ is the retirement value function.¹⁵ Tildes denote random variables.

Let us suppose an age- t consumer believes income at age $s \geq t$, $y_s(t)$, can take on one of n values $Y_s^1(t) < Y_s^2(t) < Y_s^3(t) < \dots < Y_s^n(t)$. We assume the probability distribution is a first-order Markov process with transition probabilities

$$\Pr^{(t)}[y_{s+1}(t) = Y_{s+1}^j(t) | y_s(t) = Y_s^i(t)] = \Pi_s^{ij}(t) > 0 \quad (9)$$

for $i, j = 1, \dots, n$. The expectation operator $E^{(t)}$ computes expectations with

¹⁵See the appendix for details.

respect to this probability distribution.

To compute aggregates for the population, we must also specify the actual probability distribution for income. We assume, for $0 \leq t < T_w$, that actual income at t , $y_t = y_t(t) \in \{Y_t^1, \dots, Y_t^n\}$, where

$$Y_t^i = Y_t^i(t)$$

for $i = 1, \dots, n$ and $t = 0, \dots, T_w - 1$. Since the $\{Y_{t+1}^1(t), \dots, Y_{t+1}^n(t)\}$ may differ from $\{Y_{t+1}^1(t+1), \dots, Y_{t+1}^n(t+1)\}$, it is not obvious what probability distribution for y_{t+1} conditional on y_t we should impose on the model. However, the specification of the actual probability distribution should not have any great effect on the shape of the consumption hump.¹⁶ Since we are not concerned with replicating the wealth distribution and since we also have not estimated the actual income process, for the purpose of studying the consumption hump it is sufficient to simply assume the actual income distribution is governed by some first-order Markov process with an age-independent transition matrix that governs mobility between the n rankings of the income states. For example, if there are two states in each period, the probability of going from the low to the high state is age-independent, although the income values corresponding to the low and high states do depend on age. For $0 \leq t < T_w$, we denote the actual transition matrix by

$$\Pr[y_{t+1} = Y_{t+1}^j | y_t = Y_t^i] = \Pi_{ij} > 0 \quad (10)$$

for $i, j = 1, \dots, n$. Denoting the invariant probability distribution of Π by π , we assume the actual initial probability distribution is

$$\Pr[y_0 = Y_0^i] = \pi_i.$$

Households can reallocate income across the lifecycle using the one intertemporal asset in the economy, a risk-free bond that pays the fixed gross interest rate R . Let $b_{s+1}(t)$ denote the quantity of bonds an agent at age t plans to purchase at age s that would then pay $Rb_{s+1}(t)$ at age $s + 1$. Thus the budget constraint an agent at age t expects to face at age s is

$$c_s(t) + b_{s+1}(t) = y_s(t) + Rb_s(t) \equiv x_s(t),$$

where $x_s(t)$ is cash on hand as defined by Deaton (1991).

¹⁶The shape of the mean consumption profile is ultimately determined by the mean rate of consumption growth. Feigenbaum (2007a) and Skinner (1988) have shown that the expected rate of consumption growth depends nonlinearly on wealth, so the mean rate of consumption growth will depend on the distribution of wealth. However, this wealth effect is proportional to the variance of income and so is second-order in terms of deviations of income from the mean. Since the first-order effect of deviations from the mean will, by definition, vanish after taking expectations, the lowest order effect of any dispersion in the wealth distribution is also second order. Thus the aggregate effect of deviations of wealth from mean wealth is a second-order correction to a second-order effect, making it a fourth-order effect. The effect of the wealth distribution on the shape of the lifecycle consumption profile should therefore be negligible.

Since our intent is to focus on the ability of uncertainty and precautionary saving to explain the consumption hump, we allow borrowing but with full commitment to debt contracts. Since consumption is required to be non-negative, the consumer faces the endogenous borrowing limit (Aiyagari (1994)) that he would never borrow more than the minimum possible present value of income he might earn in the future, i.e. the minimum that he could possibly pay back in the future. However, there is potentially a disconnect between what the consumer believes this borrowing limit is and what the borrowing limit actually is. At age t an agent will believe at age s that the borrowing limit he faces is

$$B_{s+1}(t) = \sum_{i=s+1}^{T_w-1} \frac{Y_i^1(t)}{R^{i-s}}.$$

At age t , the actual borrowing limit should be

$$B_{t+1} = \sum_{i=t+1}^{T_w-1} \frac{Y_i^1}{R^{i-t}}.$$

To prevent the consumer from borrowing more than he actually can pay back, we impose the exogenous borrowing constraint¹⁷

$$b_{t+1}(t) \geq -B_{t+1}.$$

Thus, given y_t and b_t , the consumer's problem at age t is

$$\max_{\{c_s(t)\}_{s=t}^{T_w-1}, \{b_{s+1}(t)\}_{s=1}^{T_w-1}} E_t \left[\sum_{s=t}^{T_w-1} \beta^{s-t} u(\tilde{c}_s(t); \gamma) + \beta^{T_w-t} v_{T_w}(R\tilde{b}_{T_w}(t)) \right] \quad (11)$$

subject to

$$\begin{aligned} c_s(t) + b_{s+1}(t) &= y_s(t) + Rb_s(t) \quad s = t, \dots, T_w - 1 \\ b_{t+1}(t) &\geq -B_{t+1}. \end{aligned}$$

To simulate the model, at age t we assume that $c_t = c_t(t)$ and $b_{t+1} = b_{t+1}(t)$. See the Appendix for details on the computational procedure.

Virtually every theoretical explanation for the hump described in Section 3.1 can quantitatively account for the hump if both the interest rate R and discount factor β are free parameters to be calibrated. Bullard and Feigenbaum (2007) and Feigenbaum (2008a) have emphasized the importance of studying the consumption hump with general-equilibrium models that put some discipline on the choice of β and R . Since our actual income process is chosen for convenience rather than empirical veracity and the market-clearing interest rate will be very

¹⁷While imposing this constraint is necessary to ensure the model has a well-defined solution, the distinction between $B_{t+1}(t)$ and B_{t+1} is usually small, so the overwhelming majority of households will not be affected by the constraint.

sensitive to the actual income process, we do not endogenize R . The purpose of the theoretical exercise that follows is to see whether, for any plausible values of β , γ , and R , there is sufficient uncertainty to generate enough precautionary saving to replicate the consumption hump.

3.3 Specification for the Income Process

Suppose that at age t , the consumer believes for $s \geq t$ that

$$y_s(t) = a_s p_s(t) z_s(t), \quad (12)$$

where a_s is an age-dependent factor, $p_s(t)$ is a (semi)permanent shock, and $z_s(t)$ is a temporary shock. Specifically, we assume that $p_s(t)$ and $z_s(t)$ are independent, unconditionally unit-mean processes such that

$$\text{corr}(\ln(p_{s+1}(t)), \ln(p_s(t))) = \rho < 1,$$

$$V[\ln p_{s+1}(t) - \rho \ln p_s(t) | p_s(t)] = \sigma_p^2(t),$$

and

$$V[\ln z_s(t)] = \sigma_{z,s-t}^2(t).$$

Note that the income specification of (12) differs from the standard specification of Carroll and Samwick (1997) and Gourinchas and Parker (2002) in two important respects. First, the permanent shocks do not follow a unit-root process, although this modification is less remarkable since many papers in the literature have considered an AR(1) income process, including Feigenbaum (2008b) and Huggett (1996). The second and more important difference is that the variance of the permanent shocks depends on the age t when forecasting occurs while the variance of the temporary shocks depends both on t and the forecasting horizon $s - t$.

We assume that a_t is consistent across different ages since we are focusing on changes in the perception of the variance of income rather than on changes in the perception of the mean. We also assume that the autocorrelation ρ is consistent across ages. This is consistent with Fig. 3, which shows that the correlation between one-year ahead forecasts and h -year ahead forecasts is essentially independent of age.¹⁸

Thus

$$\ln y_s(t) = \ln a_s + \ln p_s(t) + \ln z_s(t).$$

¹⁸We also experimented with income specifications where there is no correlation between income shocks and where the forecasted income at $t + s$ is only correlated with the forecasted income at $t + 1$. However, there was not enough persistence in shocks with these specifications to generate a significant hump in consumption profiles (Feigenbaum (2007a)).

For $h \geq 1$,

$$\begin{aligned}\ln y_{t+h}(t) &= \ln a_{t+h} + \ln p_{t+h}(t) + \ln z_{t+h}(t) \\ &= \rho^h \ln y_t + \ln a_{t+h} - \rho^h \ln a_t + \sum_{i=1}^h \rho^{h-i} (\ln p_{t+i}(t) - \rho \ln p_{t+i-1}(t)) \\ &\quad + \ln z_{t+h}(t) - \rho^h \ln z_t(t).\end{aligned}$$

Thus

$$V[\ln y_{t+h}(t)|y_t] = \frac{1 - \rho^{2h}}{1 - \rho^2} \sigma_p^2(t) + \sigma_{z_h}^2(t) + \rho^{2h} \sigma_{z_0}^2(t) \quad (13)$$

For $h > 1$, the correlation between $\ln y_{t+h}(t)$ and $\ln y_{t+1}(t)$ conditional on y_t is

$$\begin{aligned}\text{corr}(\ln y_{t+h}(t), \ln y_{t+1}(t)|y_t) & \quad (14) \\ &= \frac{\rho^{h-1} \sigma_p^2(t) + \rho^{h+1} \sigma_{z_0}^2(t)}{\sqrt{\left(\frac{1-\rho^{2h}}{1-\rho^2} \sigma_p^2(t) + \sigma_{z_h}^2(t) + \rho^{2h} \sigma_{z_0}^2(t)\right) (\sigma_p^2(t) + \sigma_{z_1}^2(t) + \rho^2 \sigma_{z_0}^2(t))}}.\end{aligned}$$

To discretize this process, we restrict $\ln p_s(t)$ to take on values such that

$$\frac{\ln p_s(t)}{\sigma_p(t)} \in \{P^1, \dots, P^{n_p}\}$$

and $\ln z_s(t)$ to take on values such that

$$\frac{\ln z_s(t)}{\sigma_{z, s-t}(t)} \in \{Z^1, \dots, Z^{n_z}\}.$$

Defining

$$Q_s(t) = \frac{\ln p_s(t)}{\sigma_p(t)}$$

and

$$Z_s(t) = \frac{\ln z_s(t)}{\sigma_{z, s-t}(t)},$$

we specify an i.i.d. probability distribution for $Z_s(t)$,

$$\Pr[Z_s(t) = Z^i] = \pi_i^z \quad (15)$$

for $i = 1, \dots, n_z$, and a Markov distribution for $Q_s(t)$,

$$\Pr[Q_{s+1}(t) = P^j | Q_s(t) = P^i] = \Pi_{ij}^p \quad (16)$$

for $i, j = 1, \dots, n_p$ and $s \geq t$. Thus the total number of income states at age t and horizon s is $n = n_p n_z$. For

$$k = n_z(i-1) + j, \quad (17)$$

where $i = 1, \dots, n_p$ and $j = 1, \dots, n_z$, the k th income state is

$$Y_s^k(t) = a_s \exp(\sigma_p(t)P^i(t)) \exp(\sigma_{z,s-t}(t)Z^j).$$

Note that this specifies the time-inconsistent probability distribution as perceived by the household. We must also specify the Markov process of the actual probability distribution for realized income y_t . As described above, we assume that the household's perceptions of its current temporary and permanent shocks are correct, so

$$y_t = y_t(t) = a_t \exp(\sigma_p(t)Q_t) \exp(\sigma_{z0}(t)Z_t),$$

where

$$\begin{aligned} Q_t &= Q_t(t) \\ Z_t &= Z_t(t). \end{aligned}$$

This implies the set of actual income states for y_t coincides with the perceived states at t , so

$$Y_t^k = Y_t^k(t)$$

for $k = 1, \dots, n$.

Let π_i^p denote the invariant distribution of Π^p . We then assume the actual probability distribution of Q_0 is

$$\Pr[Q_0 = P^i] = \pi_i^p$$

for $i = 1, \dots, n_p$ and the actual probability distribution of Z_0 is

$$\Pr[Z_0 = Z^i] = \pi_i^z$$

for $i = 1, \dots, n_z$. For $0 \leq t < T_w - 1$, we assume that (15) and (16) correctly specify the probability distribution for Q_{t+1} and Z_{t+1} , so

$$\Pr[Z_{t+1} = Z^i] = \pi_i^z$$

for $i = 1, \dots, n_z$ and

$$\Pr[Q_{t+1} = P^j | Q_t = P^i] = \Pi_{ij}^p.$$

Note that this implies the unconditional distribution of Q_t is π^p for all t .

3.4 Calibration

For the income process, we assume both the permanent and temporary shocks are governed by two-state processes.¹⁹ The temporary shocks are parameterized by $Z^2 = -Z^1 = 1$ and $\pi_1^z = \pi_2^z = 1/2$. Likewise, the permanent

¹⁹We consider the robustness of this assumption in Section 3.6.

shocks are parameterized by

$$Q^2 = -Q^1 = \frac{1}{\sqrt{1-\rho^2}}$$

and $\pi_1^p = \pi_2^p = 1/2$. The transition matrix for the permanent shocks is

$$\Pi^p = \frac{1}{2} \begin{bmatrix} 1 + \rho & 1 - \rho \\ 1 - \rho & 1 + \rho \end{bmatrix}.$$

We calibrate ρ , $\sigma_p(t)$, and $\sigma_{zh}(t)$ so as to minimize the distance of the correlation and volatility matrices relative to their predicted values from (13) and (14). Specifically, we parameterize the permanent-shock standard deviation as a d_p -degree polynomial

$$\sigma_p(t) = \sum_{i=0}^{d_p} D_i^p t^i \quad (18)$$

and the temporary-shock standard deviation as a tensor product of d_z^t and d_z^h -degree polynomials

$$\sigma_{zh}(t) = \sum_{i=0}^{d_z^t} \sum_{j=0}^{d_z^h} D_{ij}^z t^i h^j. \quad (19)$$

Then we set ρ , $D_0^p, \dots, D_{d_p}^p$, $D_{00}^z, \dots, D_{0d_z^h}^z$, $D_{10}^z, \dots, D_{1d_z^h}^z, \dots, D_{d_z^t 0}^z, \dots, D_{d_z^t d_z^h}^z$ to minimize the sum of the squares of the deviations between the predicted values of the matrix elements and their measured values. Note that, as is standard in this literature (see for example Feigenbaum (2007) or Gourinchas and Parker (2002)), the Bellman Eq. (20) implies that the household precisely knows its income state. Thus, in solving the model, we assume the consumer has additional information that the econometrician does not have when computing the volatility and correlation matrices (13) and (14), for the consumer knows how his current income breaks down into permanent versus temporary income shocks.

In Section 2, we measure the volatility and correlation matrices for horizons up to $H_{\max} = 25$. For $t < T_w - H_{\max} - 2$ we will also need to specify $\sigma_{zh}(t)$ for $h \in \{H_{\max} + 1, \dots, T_w - t - 1\}$, but these standard deviations are not identified by the available data. We will consider what happens both if we linearly extrapolate (19) for $h > H_{\max}$ and if we assume a flat extrapolation where $\sigma_{zh}(t) = \sigma_{zH_{\max}}(t)$ for $h > H_{\max}$. We also do not have information about $\sigma_{z0}(T_w - 1)$ since we have no data at forecasts in the last working period. However, we do have the volatility matrix element $V[\ln y_{T_w-1}(T_w - 2) | y_{T_w-2}]$, so it is reasonable to assume (19) will still be valid at $t = T_w - 1$. Likewise, we simply extrapolate (18) to obtain $\sigma_p(T_w - 1)$.

We consider the nonparametric estimates obtained with both the AIS and RIS specifications of Section 2 with a cubic approximation that sets $d_p = d_z^t = d_z^h = 3$. To assess the importance of the age and forecast-horizon de-

Model	ρ
AIS TIME CON	0.910
AIS TIME INC	0.958
RIS TIME CON	0.925
RIS TIME INC	0.964

Table 3: Correlation ρ for both the time-inconsistent (TIME INC) and time-consistent (TIME CON) income processes and both the AIS and RIS estimates of the volatility and correlation matrices.

pendence of uncertainty we also consider a time-consistent calibration of the income process where σ_p and σ_z are constants independent of t and h . Thus the time-consistent income process falls into the class of Markov income processes that have previously been studied in the literature (for example in Feigenbaum (2008b) and Huggett (1996)). The time-consistent calibration is obtained, both for the AIS and RIS specifications, as above but with $d_p = d_z^t = d_z^h = 0$.

The time-inconsistent and time-consistent calibrations of $\sigma_p(t)$ are plotted as a function of age t in Fig. 5 for both the AIS and RIS specifications. Likewise, the four calibrations of $\sigma_{zh}(t)$ are plotted as a function of age for representative horizons in Fig. 6. The correlation for each calibration is given in Table 3. The variance of permanent income shocks is uniformly larger for the time-consistent calibration than the time-inconsistent calibration, and for most horizons the time-consistent variance is twice as large as the time-inconsistent variance. The correlations for the time-consistent calibrations are modestly smaller than the corresponding time-inconsistent calibrations. At short time horizons the variance of temporary income shocks is comparable between the time-consistent and time-inconsistent calibrations, but the variance of temporary income shocks increases with the forecast horizon in the time-inconsistent model while necessarily remaining constant in the time-consistent model. Thus permanent income shocks will have greater emphasis in the time-consistent models whereas temporary income shocks will have more emphasis in the time-inconsistent models.

For the AIS estimates, the root-mean-squared deviation between the time-inconsistent model's predictions for the volatility and correlation matrices and the corresponding empirical estimates is 0.022. For the time-consistent model, the root-mean-squared deviation is 0.051. Fig. 7 shows how the variances of forecast errors for these two models compare to the nonparametric estimates as a function of age at different forecast horizons. Since the time-consistent model assumes that the standard deviation of temporary and permanent income shocks is independent of age, its variance graphs are flat whereas the time-inconsistent model is able to capture the U-shaped dependence of the variances with respect to age. For short horizons of one to two years, the time-consistent model overpredicts the uncertainty at all ages. For long horizons, the time-consistent model significantly underpredicts the uncertainty. Consistent with the root-mean-squared deviations, the time-inconsistent model almost




Figure 5: Permanent shock standard deviation $\sigma_{zh}(t)$ as a function of age t for the time-inconsistent (TIME INC) and time-consistent (TIME CON) income processes with both the RIS and AIS estimates of the volatility and correlation matrices.




Figure 6: Temporary shock standard deviation $\sigma_{z_h}(t)$ as a function of age t for horizons h of (a) one year, (b) five years, (c) fifteen years, and (d) twenty-five years for both the time-inconsistent (TIME INC) and time-consistent (TIME CON) income processes and both the AIS and RIS estimates of the volatility and correlation matrices.




Figure 7: Variance of forecast errors for different forecast horizons as a function of age for the augmented information set (AIS) estimates and both the time-consistent and time-inconsistent income processes.

uniformly does better at matching the forecast-error variances, although it does underpredict the variance at the two-year horizon. The comparison is similar for the RIS specification.

Fig. 8 shows how the correlations between one-year ahead forecast errors and h -year ahead forecast errors compare at ages 30, 40, and 50 between the nonparametric estimates and the time-inconsistent income process under the AIS specification. Fig. 9 shows the same comparison for the time-consistent income process. The time-inconsistent model matches the correlations slightly better as the time-consistent model overpredicts the correlation at short horizons and underpredicts it at long horizons.

In addition to the income process, we also have to calibrate the preference parameters β and γ , and since this is a partial-equilibrium model the gross interest rate R . Following Feigenbaum (2008b), we set the discount factor and interest rate to common values from the literature: $\beta = 0.96$ and $R = 1.035$. Feigenbaum (2007) found that a risk aversion of $\gamma = 3$ could best account for the lifecycle consumption profile under Gourinchas and Parker's (2002) income

Time-Inconsistent Correlations




Figure 8: Correlation of one-year ahead forecast errors and h -year ahead forecast errors as a function of forecast horizon h at ages 30, 40, and 50 for the augmented information set (AIS) estimates and the time-inconsistent income process.

Time-Consistent Correlations




Figure 9: Correlation of one-year ahead forecast errors and h -year ahead forecast errors as a function of forecast horizon h at ages 30, 40, and 50 for the augmented information set (AIS) estimates and the time-consistent income process.

Model	$\frac{E[c_{t_{\max}}]}{E[c_0]}$	$t_{\max}(+25)$
GP CONS. DATA	1.15	45
AIS TIME INC	1.15	56
RIS TIME INC	1.25	60
AIS TIME CON	1.28	61
RIS TIME CON	1.43	60
GP TIME CON	2.04	57

Table 4: Ratio of peak to initial consumption and peak age for the lifecycle consumption profile as measured by Gourinchas and Parker (2002) and as predicted by the model with time-inconsistent (TIME INC) and time-consistent (TIME CON) income processes calibrated for both the RIS and AIS specifications. For comparison, we also include an income process similar to Gourinchas and Parker’s (2002) baseline income process.

process, so we maintain this value.

3.5 Theoretical Predictions for Consumption Hump

Fig. 10 shows lifecycle profiles for mean consumption (normalized by mean initial income) for the time-inconsistent and time-consistent calibrations under both the AIS and RIS specifications along with Gourinchas and Parker’s empirical measurements of the mean consumption profile^{20,21} For comparison with the previous literature, we also include the results for a time-consistent model with Gourinchas and Parker’s estimates of the shock variances: $\sigma_p^2 = 0.0212$ and $\sigma_z^2 = 0.0440$. In their model, the permanent income shocks follow a unit-root process. Since a unit-root process cannot be nested in our model, we set $\rho = 0.99$. In Table 4, we also report the peak to initial consumption of the lifecycle profile of mean consumption and the age of peak consumption.

All five models produce a hump-shaped lifecycle consumption profile for our chosen calibration of the preference parameters and the interest rate. The consumption-saving model with the time-inconsistent income process based on the AIS specification most closely matches Gourinchas and Parker’s estimates of mean consumption as a function of age, although this ranking might change with other calibrations.

Our primary interest here lies in the degree to which different measures of uncertainty impact the consumption profile for a given calibration. Since $\beta R < 1$, in the absence of uncertainty the lifecycle consumption profile should be monotonically decreasing with a peak at the initial age of 25, so the peak to initial consumption ratio should be 1. Thus the peak to initial consumption

²⁰The graphs in Fig. 10 and ensuing figures were obtained by simulating one million lifecycle paths per age group.

²¹The results for the time-inconsistent models were obtained with a linear extrapolation of the temporary shock variances as a function of the forecast horizon. Results obtained with a flat extrapolation are given in Section 3.6.




Figure 10: Mean consumption profile (normalized by mean initial income) as measured by Gourinchas and Parker (2002) and as predicted by the model with time-inconsistent (TIME INC) and time-consistent (TIME CON) income processes calibrated for both the RIS and AIS specifications. For comparison, we also include an income process similar to Gourinchas and Parker's (2002) baseline income process.

ratio can be viewed here as a measure of how much precautionary saving causes these models to deviate from the LCPIH. Since we are also holding the degree of risk aversion fixed, the variation in precautionary saving over different income processes should reflect the amount of uncertainty faced over the lifecycle under each process.

Not surprisingly, since the AIS specification assumes households have more information than the RIS specification, the two AIS models have smaller peak to initial consumption ratios than their corresponding RIS models. Moreover, the Gourinchas and Parker-based model with its near unit-root process has a substantially larger peak to initial consumption ratio than the four models we estimate since it has the most uncertainty. Indeed, for $\beta = 0.96$ and $R = 1.035$, Gourinchas and Parker (2002) have to dial the risk aversion all the way down to one half to get a lifecycle consumption profile that resembles the data, but this would not be necessary for our more robust uncertainty measure.

For both the RIS and AIS specifications, we also find that the time-inconsistent model has a significantly smaller peak to initial consumption ratio than the corresponding time-consistent model. This can be explained in terms of Figs. 5 and 6. Because the time-consistent model assumes a constant variance for the permanent and temporary shocks, independent of age and forecast horizon, the time-consistent model needs a large permanent shock variance to best match the volatility and correlation matrices. In contrast, the time-inconsistent model has a smaller permanent shock variance and exploits its ability to increase the temporary shock variance at longer forecast horizons to better match the forecast-error moments. As Constantinides and Duffie (1996) argued, permanent shocks will have a substantially larger impact on the behavior of consumers, and Feigenbaum (2007,2008b) confirm that more persistent income shocks lead to greater precautionary saving. This intuition is further corroborated here. Thus, failing to account for the time and forecast-horizon dependence of uncertainty can bias upward estimates of the importance of precautionary saving.

3.6 Robustness Checks

One cause for concern about the calibration of the income process in Section 3.4 is that we have no data on the variance and correlation of forecast errors beyond a twenty-five-year horizon, but the model requires us to specify the household's beliefs about income at all future horizons. In Section 3.5, we reported results for the baseline case where we assume a linear extrapolation of Eq. (19) for $h \geq H_{\max}$. In Fig. 11, we consider what happens with the alternative assumption of a flat extrapolation such that $\sigma_{zh}(t) = \sigma_{z,H_{\max}}(t)$ for $h \geq H_{\max}$. For both the AIS and RIS specifications, we find a negligible difference between the two extrapolations. The effect of a temporary income shock twenty-five years in the future is going to be heavily discounted, so a change in the temporary shock variance at such large horizons has little to no effect on consumption behavior.

Another potential cause for concern about our results is that we use only




Figure 11: Mean consumption profile (normalized by mean initial income) as measured by Gourinchas and Parker (2002) and as predicted by the model with time-inconsistent income processes calibrated for both the RIS and AIS specifications with linear and flat extrapolations of the temporary shock variances as a function of time horizon.

two-state processes to model both the permanent and temporary income shocks. Would a finer distribution produce different results? Feigenbaum (2008b) suggests that only the second-order moments should have a significant effect on the shape of the mean consumption profile. Given the common assumption that shocks are normally distributed, the next moment of interest is the kurtosis. For a normal distribution, the kurtosis should be 3. However, an unskewed two-state process must have the minimum kurtosis of 1. For the temporary shock process, we can easily replicate both the variance and kurtosis of a normal distribution with a three-state process characterized by $Z^3 = -Z^1 = 1$ and $Z^2 = 0$ with $\pi_1^z = \pi_3^z = 1/6$ and $\pi_2^z = 2/3$. We can also obtain an unconditional distribution of permanent shocks that matches the variance and kurtosis of a normal distribution in this way, but with $\rho \approx 1$ we cannot do the same for the conditional distribution of permanent shocks, which is what should matter most for precautionary saving. With three states, the best we can do is choose a transition matrix that minimizes the kurtosis of the conditional distribution. If the unconditional kurtosis is K_p , the minimum conditional kurtosis goes as $1.5K_p/(1 - \rho)$ in the limit as $\rho \rightarrow 1$.

In Fig. 12, we plot the mean consumption profile for the AIS specification of the time-inconsistent model for various choices of K_p and the kurtosis of the temporary shock distributions. To get a sense of scale, we also plot the empirical profile of Gourinchas and Parker (2002). Three of the model profiles, for which the unconditional kurtoses of both shocks are between 1 and 3, are virtually indistinguishable. The ratio of peak consumption to initial consumption for all three curves is between 1.125 and 1.146. Only when we increase K_p to 10 do we get any significant departure, and even then the peak to initial ratio only increases to 1.17. Comparing Fig. 12 to Fig. 10, we see that the impact of using a finer distribution is much smaller than the impact of ignoring the time-dependence of uncertainty or adding more information to the model.

As a final robustness check, let us consider how sensitive the model is to the risk aversion γ . In Fig. 13 and Table 5 we replicate Fig. 10 and Table 4 but for $\gamma = 2$. We still get hump-shaped consumption profiles, but not surprisingly with less risk aversion there is less precautionary saving so the humps are more modest in size. Interestingly, while the consumption profile for the AIS specification of the time-inconsistent model now has a peak to initial consumption ratio smaller than Gourinchas and Parker (2002) find, this calibration does match the age of the peak.

4 Conclusion

We introduce a new method of measuring income uncertainty and apply the estimates obtained via this approach to investigate the extent to which variations in income uncertainty over the lifecycle could be responsible for the hump-shaped consumption profile. Our measurement technique articulates the distinction between income heterogeneity and uncertainty, and acknowledges

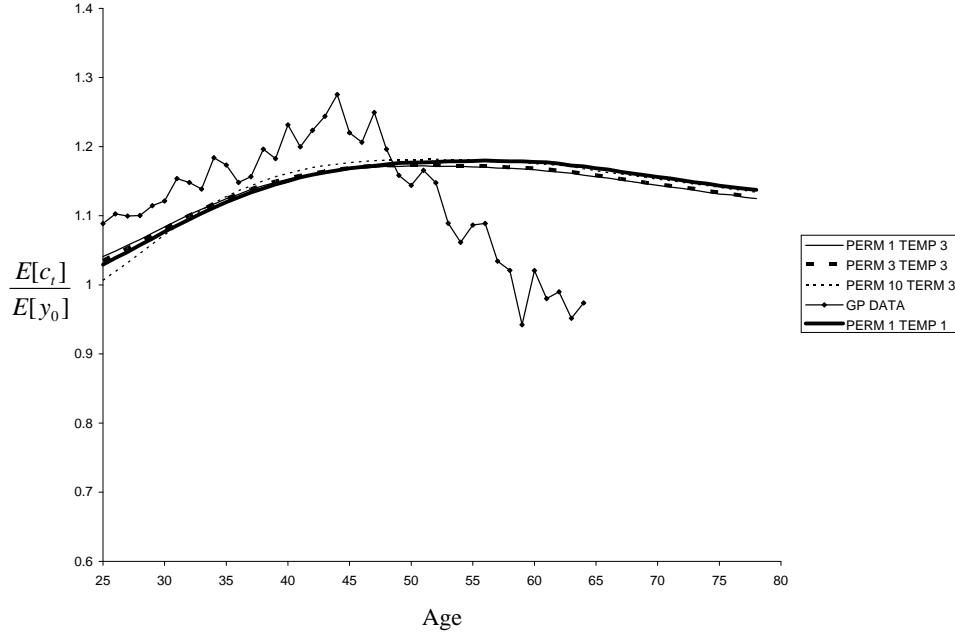


Figure 12: Mean consumption profile (normalized by mean initial income) as measured by Gourinchas and Parker (2002) and as predicted by the model with time-inconsistent income processes calibrated for the AIS specifications with three-state processes for the permanent (PERM) and temporary (TEMP) shocks with various kurtoses.

Model	$\frac{E[c_{t_{\max}}]}{E[c_0]}$	$t_{\max}(+25)$
GP DATA	1.15	45
AIS TIME INC	1.07	46
RIS TIME INC	1.13	55
AIS TIME CON	1.16	59
RIS TIME CON	1.28	58
GP TIME CON	1.58	53

Table 5: Ratio of peak to initial consumption and peak age for the lifecycle consumption profile as measured by Gourinchas and Parker (2002) and as predicted by the model with time-inconsistent (TIME INC) and time-consistent (TIME CON) income processes calibrated for both the RIS and AIS specifications when $\gamma = 3$. For comparison, we also include an income process similar to Gourinchas and Parker’s (2002) baseline income process.




Figure 13: Mean consumption profile (normalized by mean initial income) as measured by Gourinchas and Parker (2002) and as predicted by the model with time-inconsistent (TIME INC) and time-consistent (TIME CON) income processes calibrated for both the RIS and AIS specifications with $\gamma = 2$. For comparison, we also include an income process similar to Gourinchas and Parker's (2002) baseline income process.

that households may have superior information to econometricians. Our estimation reveals an income uncertainty level that is lower than what has been presented in the existing literature and shows that income uncertainty does evolve over the lifecycle in a fashion consistent with conventional wisdom. In addition, we show that for a plausible calibration of the preference parameters, our estimate of income uncertainty does imply a hump-shaped lifecycle consumption profile that matches consumption data very well.

Two lines of research are worth pursuing in the future. First, it is important to study whether we can find a time-consistent, parametric income process consistent with our nonparametric estimates. This is necessary in order to study what happens in general equilibrium for this model. Second, we will study more carefully the source of the variation of income uncertainty over the lifecycle.

A Computational Procedure

We can write the problem (11) as a recursive sequence of Bellman equations as follows. Let us denote the state variable $I_s(t)$ such that $y_s(t) = Y_s^{I_s(t)}(t)$. For each base age $t \in \{0, \dots, T\}$, the consumer solves for the age- t perceived value functions $V_s(x_s, I_s(t); t)$ at current and future ages $s = t, \dots, T_w - 1$. For $s = t, \dots, T_w - 1$, the Bellman equation is

$$V_s(x_s, I_s(t); t) = \max_{c_s(t), b_{s+1}(t)} u(c_s(t); \gamma) + \beta E^{(t)} [V_{s+1}(Y_{s+1}^{\tilde{I}_{s+1}^{(t)}}(t) + Rb_{s+1}(t), \tilde{I}_{s+1}(t); t) | I_s(t)] \quad (20)$$

subject to

$$c_s(t) + b_{s+1}(t) = x_s$$

For the case when $s = t$, we impose the additional constraint

$$b_{t+1}(t) \geq -B_{t+1}.$$

The retirement value function that terminates the sequence of value functions is simply the perfect-foresight utility of an agent who will live for T_r periods with no further income. With CRRA utility (8), this is

$$V_{T_w}(x_{T_w}) = \left(\frac{\phi^{-T_r} - 1}{\phi^{-1} - 1} \right)^\gamma u(x_{T_w}; \gamma), \quad (21)$$

where

$$\phi = (\beta R^{1-\gamma})^{-1/\gamma}$$

is the inverse marginal propensity to consume in the limit of large lifetimes. The sequence of Bellman equations (20) can easily be translated into the sequence of Bellman equations solved in Feigenbaum (2008b), and the computational procedure we use here is described in the technical appendix for that paper.²² Note that only the policy functions $c_t(x_t, I_t; t)$ for $t = 0, \dots, T_w - 1$ are actually relevant to the behavior of the model.

²²This is available online at <http://www.pitt.edu/jfeigen/infoshockappendix.pdf>.

References

- [1] Aguiar, Mark and Erik Hurst, (2003), "Consumption vs Expenditure," Working Paper.
- [2] Aguiar, Mark and Erik Hurst, (2007), "Life-Cycle Prices and Production," *American Economic Review* 97: 1533-1559.
- [3] Aiyagari, S. Rao, (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics* 109: 659-684.
- [4] Attanasio, Orazio P., James Banks, Costas Meghir, and Guglielmo Weber, (1999), "Humps and Bumps in Lifetime Consumption," *Journal of Business & Economic Statistics* 17: 22-35.
- [5] Becker, Gary S. and Gilbert R. Ghez, (1975), *The Allocation of Time and Goods Over the Life Cycle* (New York: NBER and Columbia Press).
- [6] Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg, (2001), "What Accounts for the Variation in Retirement Wealth Among U.S. Households?" *American Economic Review* 91: 832-857.
- [7] Browning, Martin and Thomas F. Crossley, (2001), "The Life-Cycle Model of Consumption and Saving," *Journal of Economic Perspectives* 15(3): 3-22.
- [8] Browning, Martin and Mette Ejrnæs, (2002), "Consumption and Children," Centre for Applied Microeconometrics Working Paper 2002-06.
- [9] Bullard, James and James Feigenbaum, (2007), "A Leisurely Reading of the Lifecycle Consumption Data," forthcoming in *Journal of Monetary Economics*.
- [10] Caliendo, Frank and David Aadland, (2004), "Short-Term Planning and the Life-Cycle Consumption Puzzle," *Journal of Economic Dynamics and Control* 31: 1392-1415.
- [11] Carroll, Christopher D., (1997), "Buffer-Stock Saving and the Life-Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics* 112: 1-55.
- [12] Carroll, Christopher D., (2004), "Theoretical Foundations of Buffer Stock Saving," Working Paper.
- [13] Carroll, Christopher D. and Andrew A. Samwick, (1997), "The Nature of Precautionary Wealth," *Journal of Monetary Economics* 40: 41-71.
- [14] Carroll, Christopher D. and Andrew A. Samwick, (1998), "How Important is Precautionary Saving?" *Review of Economics and Statistics* 80(3): 410-419.

- [15] Carroll, Christopher D. and Lawrence H. Summers, (1991), "Consumption Growth Parallels Income Growth: Some New Evidence," in *National Saving and Economic Performance* eds. B. Douglas and J. Shoven (Chicago: University of Chicago Press).
- [16] Constantinides, George M. and Darrell Duffie, (1996), "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy* 104: 219-240.
- [17] Cunha, Flavio, James Heckman, and Salvador Navarro, (2005), "Separating Uncertainty from Heterogeneity in Life Cycle Earnings," *Oxford Economic Papers* 57: 191-261.
- [18] Deaton, Angus, (1991), "Saving and Liquidity Constraints," *Econometrica* 59: 1221-1248.
- [19] Dynan, Karen, Douglas W. Elmendorf and Daniel E. Sichel, (2007), "The Evolution of Household Income Volatility," Federal Reserve Board Finance and Economics Discussion Series 2007-61.
- [20] Feigenbaum, James, (2007), "Precautionary Saving Unfettered," Working Paper.
- [21] Feigenbaum, James, (2008a), "Can Mortality Risk Explain the Consumption Hump?" Forthcoming in *Journal of Macroeconomics*.
- [22] Feigenbaum, James, (2008b), "Information Shocks and Precautionary Saving," Forthcoming in *Journal of Economic Dynamics and Control*.
- [23] Fernández-Villaverde, Jesús and Dirk Krueger, (2005), "Consumption and Saving over the Life Cycle: How Important are Consumer Durables?" Working Paper.
- [24] Fernández-Villaverde, Jesús and Dirk Krueger, (2007), "Consumption over the Life Cycle: Some Facts from Consumer Expenditure Survey Data," *Review of Economics and Statistics* 89: 552-565.
- [25] Gottschalk, Peter and Robert Moffitt, (1994), "The Growth of Earnings Instability in the U.S. Labor Market," *Brookings Papers on Economic Activity* 217-272.
- [26] Gourinchas, Pierre-Olivier and Jonathan A. Parker, (2001), "The Empirical Importance of Precautionary Saving," *American Economic Review* 91: 406-412.
- [27] Gourinchas, Pierre-Olivier and Jonathan A. Parker, (2002), "Consumption over the Life Cycle," *Econometrica* 70: 47-89.
- [28] Guvenen, Fatih, (2007), "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" *American Economic Review* 97: 687-712.

- [29] Hansen, Gary D. and Selahattin İmrohoroğlu, (2006), “Consumption over the Life Cycle: The Role of Annuities,” NBER Working Paper 12341.
- [30] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante, (2004), “The Cross-Sectional Implications of Rising Wage Inequality in the United States,” Working Paper.
- [31] Heckman, James, (1974), “Life Cycle Consumption and Labor Supply: An Explanation of the Relationship Between Income and Consumption Over the Life Cycle,” *American Economic Review* 64: 188-194.
- [32] Heckman, James and Salvador Navarro, (2005), “Using Dynamic Models to Measure Uncertainty, Learning and Psychic Costs to Understand Schooling Choices and Schooling Outcomes,” Institute of Computational Economics Presentation.
- [33] Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, (1994), “The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving,” *Carnegie-Rochester Conference Series on Public Policy* 40: 59-125.
- [34] Huggett, Mark, (1996), “Wealth Distribution in Life-Cycle Economies,” *Journal of Monetary Economics* 38: 469-494.
- [35] Huggett, Mark, Gustavo Ventura, and Amir Yaron, (2007), “Sources of Lifetime Inequality,” Working Paper.
- [36] Jaimovich, Nir and Henry E. Siu, (2006), “The Young, the Old, and the Restless: Demographics and Business Cycle Volatility,” Working Paper.
- [37] Krueger, Dirk and Fabrizio Perri, (2001), “Does Income Inequality Lead to Consumption Inequality? Empirical Findings and a Theoretical Explanation,” Working Paper.
- [38] Laibson, David, (1997), “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics* 112: 443-477.
- [39] Leland, Hayne E., (1968), “Saving and Uncertainty: The Precautionary Demand for Saving,” *Quarterly Journal of Economics* 82: 465-473.
- [40] Lillard, Lee, and Yoram Weiss, (1979), “Components of Variation in Panel Earnings Data: American Scientists 1960-70,” *Econometrica* 47: 437-454.
- [41] MaCurdy, Thomas (1982), “The Use of Time-Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis,” *Journal of Econometrics* 18: 83-114.
- [42] McConnell, Margaret and Perez-Quiros, Gabriel, (2000), “Output Fluctuations in the U.S.: What Has Changed Since the Early 1980s?” *American Economic Review* 90 : 1464-1476.

- [43] Nagatani, Keizo, (1972), "Life Cycle Saving: Theory and Fact," *American Economic Review* 62: 344-353.
- [44] Sandmo, A., (1970), "The Effect of Uncertainty on Saving Decisions," *Review of Economic Studies* 37: 353-360.
- [45] Skinner, Jonathan, (1988), "Risky Income, Life Cycle Consumption, and Precautionary Savings," *Journal of Monetary Economics* 22: 237-255.
- [46] Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, (2000), "Consumption and Risk Sharing Over the Life Cycle," Working Paper.
- [47] Thurow, Lester C., (1969), "The Optimum Lifetime Distribution of Consumption Expenditures," *American Economic Review* 59: 324-330.
- [48] Yaari, Menahem E., (1964), "On the Consumer's Lifetime Allocation Process," *International Economic Review* 5: 304-317.