

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

**A Semiparametric Characterization of Income Uncertainty Over  
the Life Cycle**

**James A. Feigenbaum and Geng Li**

**2010-42**

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.

# A Semiparametric Characterization of Income Uncertainty over the Life Cycle

James Feigenbaum\*

and

Geng Li †‡

July 14, 2010

---

\*Department of Economics and Finance; Utah State University; Logan, UT 84322-3565. E-mail: J.Feigen@aggiemail.usu.edu.

†Federal Reserve Board, Mail Stop 93, Washington, DC 20551. E-mail: geng.li@frb.gov. The views presented in this paper are those of the authors and are not necessarily those of the Federal Reserve Board or its staff.

‡We thank Frank Caliendo, Chris Carroll, Juan Carlos Cordoba, David DeJong, Karen Dynan, Jonathan Heathcote, Michael Palumbo, Stephen Shore, and conference and seminar participants at the NBER Summer Institute, Conference on Computation in Economics and Finance, the Econometric Society North American Summer Meetings, Society of Economic Dynamics Annual Meetings, Federal Reserve Board, SUNY Stony Brook, University of Pittsburgh, and University of Texas-Dallas for helpful discussions and comments. All remaining errors are our own.

# A Semiparametric Characterization of Income Uncertainty over the Life Cycle

## **Abstract**

We propose a novel approach to estimate household income uncertainty at various future horizons, and characterize how the estimated uncertainty evolves over the life cycle. We measure income uncertainty as the variance of linear forecast errors conditional on information available to households prior to observing the realized income. This approach is semiparametric because we impose essentially no restrictions on the statistical properties of the forecast errors. Relative to previous studies, we find lower and less persistent income uncertainties that call for a life cycle consumption profile with a less pronounced hump.

JEL Classification: D12, D91, E24

Keywords: income uncertainty, forecast errors, semiparametric estimates, life cycle dynamics

# 1 Introduction

Income uncertainty plays an important role in household decisions regarding consumption, saving, and investment. First, in the presence of incomplete financial markets, households facing uninsurable income risk have to save for precautionary reasons (Aiyagari (1994)). How much they need to save crucially depends on their level of future income uncertainty. Second, people choose their investment portfolios to achieve optimal exposure to risks. How much people should invest in risky assets, such as stocks, also depends on their idiosyncratic exposure to income risks (Viceira (2001)). Finally, choices of financial contracts and durable goods purchasing decisions are critically affected by household income uncertainty. For instance, Campbell and Cocco (2003) show that homeowners with more risky income should choose fixed-rate mortgages over adjustable-rate mortgages.

Many of these household financial and economic decisions exhibit strong life cycle patterns. It is well known that nondurable goods consumption has a hump-shaped life cycle profile that closely tracks income (Gourinchas and Parker (2002)). Meanwhile, stock market participation rates are particularly high among middle-aged households. Homeownership and durable goods consumption also tend to vary along the life cycle (Fernández-Villaverde and Krueger (2010)). Despite its pivotal importance, the existing literature has not yielded a concrete measure of income uncertainty, and to the best of our knowledge, even less work has been done to characterize how income uncertainty evolves over the life cycle. This paper helps bridge that gap.

The most frequently used approach to characterizing income uncertainty adopts a dichotomy that assumes income shocks can be decomposed into a permanent (random walk) and a transitory (typically i.i.d.) component.<sup>1</sup> The identification of the variances of income shocks is obtained by examining either the variances of income growth over periods of differ-

---

<sup>1</sup>The exceptions include Storesletten, Telmer, and Yaron (2004), which modeled the nonpermanent shocks as an AR process.

ent lengths in longitudinal data (e.g., Carroll and Samwick (1997)) or the changes of income distributions of the same cohort over time in repeated cross-sectional data (e.g., Deaton and Paxson (1994)).

Deviating from this paradigm, we first define income uncertainty as the variance of forecast errors of future income conditional on the information available to the households at the time of the forecast. Altogether, this approach is semiparametric because on the one hand we adopt a parsimonious linear projection equation to construct the forecast errors. On the other hand, we impose essentially no restrictions on the statistical properties of deviations from this linear projection (i.e., the forecast errors). Instead, we construct the forecast errors and estimate their variances for each forecast horizon up to 25 years ahead. Then, we study how these forecast errors are correlated across different forecast horizons and how their variances evolve over the life cycle.

Broadly speaking, our approach shares the similar spirit of using cross-sectional variations to characterize income riskiness. However, relative to existing literature, this semiparametric approach has three important advantages. First, it allows us to estimate the persistence of income shocks without resorting to any *a priori* parameterizations of these shocks, especially the unit root process that is typically assumed in previous studies. Indeed, examining the correlations of income shocks over forecast horizons, we find that the typical income process is more accurately characterized as a persistent autoregressive (AR) process rather than a unit-root process. Consequently, our approach generally implies substantially lower cumulative income risks over the life cycle compared to the existing results in the literature.

Second, our approach allows us to study how income uncertainty evolves over the life cycle. For a given age in the future, we consistently find, as one would expect, that uncertainty about income diminishes as the consumer approaches this age. More strikingly, for a fixed forecast horizon, income uncertainty demonstrates a U-shaped or J-shaped profile. When consumers are young, income uncertainty at a fixed future horizon gradually declines with age, presumably as decisions on career, human capital development, and fertility are

resolved. Income uncertainty reaches and stays at its low levels during middle age. Afterwards, uncertainty rises to levels higher than at the early stage of the life cycle for some types of income—potentially due to uncertainty about working hours and health risks. This dynamic of income uncertainty is consistent with Jaimovich and Siu (2009) who find that the volatility of the business cycle component of hours worked exhibits a U-shaped pattern with respect to age.

Third, the existing methodology typically filters the predictable component of income conditional on demographic and work-related information observed concurrent to income, even when longitudinal data are available. Essentially, the filtered out part is what is ex post accountable to econometricians, and not necessarily what is ex ante predictable to households when they assess the riskiness of their future income. In contrast, our semiparametric approach adopts more realistic and flexible information specifications. Our baseline specification assumes that households in year  $t$  forecast their income at year  $t + s$  conditional on only the demographic and work-related information about the household in year  $t$ . This assumption is rather restrictive and in some sense allows for only the information available to econometricians, not to the households, in year  $t$ . Conceivably, households may know more about their future selves in year  $t + s$ . For example, Cunha, Heckman, and Navarro (2005) point out individuals may have private information that is relevant to predicting their future income. In this spirit, in augmented specifications, we assume that households in year  $t$  possess some demographic and work-related information about their future selves in year  $t + s$ . Not surprisingly, we find that expanding the information set further reduces the estimated income uncertainties, especially at farther horizons.

To further illustrate the economic significance of the refined measure of income uncertainty proposed in this paper, we examine its implications for precautionary saving and the mean consumption profile over the life cycle. Specifically, we calibrate a parsimonious income process to (partially) match our semiparametric estimates and solve for the optimal life cycle consumption profile. Because our semiparametric approach yields significantly lower

income risks over the life cycle, households accumulate less precautionary saving than previous estimates would indicate. Although the magnitude of the consumption hump implied by the model is somewhat sensitive to the calibration of technology and preference parameters, using our semiparametric estimates of uncertainty, we typically obtain an appreciable consumption hump but one that is much smaller than existing income uncertainty estimates imply and is much smaller than the one empirically observed. Given that the consumption hump is unlikely accounted for only by precautionary saving, we view our simulation results as more plausible than those derived using previous uncertainty estimates.

We proceed as follows. Section 2 reviews the existing methods of measuring income variability and discusses dimensions in which these methods can be improved. Section 3 introduces a semiparametric approach to measuring income uncertainty. Section 4 describes the data and sample selection procedures. Section 5 presents the main results and sensitivity analysis. Section 6 contrasts our empirical results with those reported in earlier contributions. Section 7 studies the precautionary saving and consumption over the life cycle under various estimates of income uncertainty. Section 8 concludes and sets an agenda for future work.

## **2 A Critical Review of the Existing Approaches to Studying Income Uncertainty**

Traditionally, when rich microdata were not readily available, researchers had to infer the volatility of household income from aggregate time-series data. The estimated measures of income volatility were often interpreted as income uncertainty. Recent studies of aggregate output (see, among others, Kim and Nelson (1999); and McConnell and Perez-Quiros (2000)) have led to findings of the so-called Great Moderation, a sharp decline in the growth volatility of GDP and its components since the mid-1980s. Fogli and Perri (2006) attribute a substantial fraction of the decrease in personal savings and the rise of external balances to this reduction in aggregate volatility. However, aggregate data can mask important variations and correlations in income at the household level, so a decrease in the volatility of

aggregate income does not necessarily imply that income at the household level has become less risky. 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. Because of this gambling, though the aggregate income of the economy remains constant and bears no uncertainty, individual households may be exposed to substantial income risk.

Studying household income directly avoids this shortcoming and became feasible, as large surveys of household income have been conducted. However, income uncertainty is typically studied via other aspects of variability, such as income inequality and volatility. This is not surprising, on the one hand, uncertainty per se is difficult to observe. On the other hand, greater uncertainty can often lead to greater inequality or volatility. For example, in a seminal early contribution, Deaton and Paxson (1994) examine the repeated cross-sectional data and document that the within-cohort inequality of consumption and income increases with age. They argue that the increased inequality is the “cumulative differences in the effects of luck,” but they have not provided a quantitative analysis of income uncertainty.

More quantitative results are obtained using longitudinal surveys, such as the Panel Study of Income Dynamics. Similarly, much of the work has focused on the volatility and inequality of household income. For example, Dynan, Elmendorf, and Sichel (2007) examine the standard deviation of the percent change in income across households and conclude that household income growth has become more volatile over the past several decades.<sup>2</sup> Alternatively, some researchers take the mean income of a household over a given period as a proxy for permanent income and treat the difference between the observed income and the mean as the transitory component of income. Using this method, Gottschalk and Moffitt (1994) find that an increase in the variance of transitory earnings could explain a large portion of the widening of income inequality during the 1970s and 80s. In general, it is hard

---

<sup>2</sup>Dynan, Elmendorf, and Sichel (2007) also provide an elegant and comprehensive survey of the literature studying household income volatilities using household data.

to make an inference about income uncertainty (and its changes) from these results because the variabilities they studied may include income changes that are predictable to households.

More elaborate models often postulate some parsimoniously parameterized income process that includes a predictable part and a stochastic part, which in turn is a combination of permanent and transitory shocks. A battery of econometric techniques have been developed to study the properties of the stochastic part of income shocks. For example, Carroll and Samwick (1997) and Gourinchas and Parker (2002) study the following specification.<sup>3</sup> The logarithm of income,  $y_t$ , is decomposed into

$$y_t = p_t + \varepsilon_t, \tag{1}$$

where  $\varepsilon_t$  is a transitory shock and  $p_t$  is a permanent income shock that follows

$$p_t = g_t + p_{t-1} + \eta_t, \tag{2}$$

in which  $g_t$  is a predictable component and  $\eta_t$  is the shock to permanent income. Although this specification is closer to the goal of measuring income uncertainty, we argue this methodology can be refined in the following respects.

First, the predictable component of income is typically constructed by regressing the log of observed income of year  $t$  on household demographics and work-related variables observed in the same year. The fitted value of the dependent variable is then interpreted as the predictable component. To the extent that households do not have perfect foresight on this information, the fitted value constructed in such a regression is essentially the component *ex post* explainable to econometricians rather than the component *ex ante* predictable to households. Put differently, constructing the predictable income by conditioning on information concurrent with the realization of income potentially assumes too much information, and consequently may lead to underestimation of the income uncertainty perceived by the household.

---

<sup>3</sup>The models used to estimate income uncertainty are almost the same in the two papers. We use the same notations as in Carroll and Samwick (1997).

Second, if we fix the calendar time of a given future income, the predicted value should converge to the realized value of income at this time as the household progresses in time and uncertainties are resolved. However, because the predictable component is not constructed specifically for any projection horizon in the standard method, it does not directly capture this evolution of income uncertainty.

Third, the above parametrization assumes the income process is subject to permanent shocks. As we show later in this paper, household income shocks do have a persistent component. However, this component is best characterized as a fairly persistent, but not a permanent (or unit-root), process. This distinction has a substantial effect on the level of the cumulative income uncertainty over the life cycle and subsequently affects the implied precautionary saving motives (in particular for younger consumers).

Fourth, fitting individual income with a trend that is only driven by demographic and work-related characteristics requires assuming that all households similar in these dimensions share a common life cycle trend. This assumption is consistent with the model introduced by MaCurdy (1982). A competing view is that consumers in fact face individual-specific income profiles, as proposed by Lillard and Weiss (1979). More recently, Guvenen (2007) presents evidence that consumption data are more consistent with the view of Lillard and Weiss (1979). If the underlying income process is better characterized in this way, fitting it with a trend common to seemingly similar households will increase the residual variance and hence exaggerate the underlying income uncertainty.

Finally, apart from a few exceptions (e.g., Baker and Solon (2003)), previous studies have not fully characterized how household income risks evolve over the life cycle. It is both theoretically and empirically appealing to study whether household income uncertainty does stay constant, and, if not, how it evolves over the life cycle. Heuristically, a single 22-year-old college graduate entering the labor market should have more uncertainties about his income five or ten years down the road than a 40-year-old man with a family and a settled career path over the same time horizon.

### 3 A Forecast-Based Semiparametric Approach

We propose a semiparametric approach to estimating household income uncertainty that attempts to address the concerns outlined in the previous section. This approach is flexible enough to incorporate various assumptions about the household’s information set and can be used to characterize the evolution of income uncertainty over the life cycle. Our key insight is that greater income uncertainty should make future income more difficult to forecast, conditional on the information available to the household at the time of the forecast. Accordingly, we use forecast accuracy as a metric of the underlying income uncertainty, or riskiness. The larger the variances of forecast errors are, the greater uncertainty, or risk, a household faces regarding its future income.

For a household whose head’s age is  $t$ , let the logarithm of its income  $s$ -year-ahead be  $y_{t,s}$ , which can be decomposed as

$$y_{t,s} = E [y_{t,s} | \mathcal{I}_t^H] + \varepsilon_{t,s}, \quad (3)$$

where  $E [y_{t,s} | \mathcal{I}_t^H]$  is the mathematical expectation of age  $t+s$  income,  $y_{t,s}$ , conditional on age- $t$  household information,  $\mathcal{I}_t^H$ , and  $\varepsilon_{t,s}$  is an error term orthogonal to  $\mathcal{I}_t^H$ . We characterize life cycle income uncertainties using a variance matrix,  $\Omega$ , and a sequence of correlation matrices,  $\Theta^q$ . The element  $\omega_{t,s}$  of the  $\Omega$  matrix is the variance of the  $s$ -year-ahead forecast errors of age- $t$  households,<sup>4</sup> i.e.,

$$\omega_{t,s} = Var[\varepsilon_{t,s}] = Var[y_{t,s} - E(y_{t,s} | \mathcal{I}_t^H)]. \quad (4)$$

Elements of the  $\Theta^q$  matrix,  $\theta_{t,s}^q$ , are the correlation coefficients between the  $s$ -year-ahead forecast errors,  $\varepsilon_{t,s}$ , and the  $q$ -year-ahead forecast errors,  $\varepsilon_{t,q}$ , of the age- $t$  households, i.e.,

$$\theta_{t,s}^q = Corr(\varepsilon_{t,s}, \varepsilon_{t,q}). \quad (5)$$

For example, elements of  $\Theta^1$  are the correlation coefficients between the  $s$ -year-ahead forecast errors and the 1-year-ahead forecast errors. Although the forecast errors are constructed

---

<sup>4</sup>For reference convenience, we refer to a household whose head’s age is  $t$  as an age- $t$  household.

using standard linear projections, our treatment of the forecast errors is nonparametric because we do not presume that the income shocks,  $\varepsilon_{t,s}$ , follow any specific process. In examining the unrestricted  $\Omega$  and  $\Theta^q$  matrices, we can directly study the persistence of income uncertainty as well as its dynamics over the life cycle. In the remainder of the paper, we will refer to the year when the forecast is carried out as the base year.

One hurdle to implementing this strategy is that we do not know the joint distribution of  $y_{t,s}$  and  $\mathcal{I}_t^H$ . As a result, 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 establish a benchmark and to assess the bias introduced by ignoring any additional (superior) information potentially possessed by households, we experiment with two specifications. First, in what we label as the restricted information specification (RIS), we project  $y_{t,s}$  conditional on  $\mathcal{I}_t^R$ , the information set that includes only what an econometrician can observe regarding a household as of age  $t$ . This information is what the households *certainly* possess at age  $t$ . Second, in what we label as the augmented information specification (AIS), we projected  $y_{t,s}$  conditional on the augmented information set  $\mathcal{I}_t^A$ , where

$$\mathcal{I}_t^A = \mathcal{I}_t^R \cup \mathcal{I}_t^F. \quad (6)$$

The augmenting information set,  $\mathcal{I}_t^F$ , contains the household's future demographic and work related characteristics as of year  $t + s$ , the year to be forecasted. This information is what the households *likely or possibly* know at age  $t$ . To fix the idea, we estimated the following RIS equation,

$$y_{i,t,s} = \alpha + \beta_0 y_{i,t} + \beta_1 y_{i,t-1} + \beta_2 y_{i,t-2} + \gamma Z_{i,t} + \xi Trend_{i,t,s} + \varepsilon_{i,t,s}, \quad (7)$$

and AIS equation,

$$y_{i,t,s} = \alpha + \beta_0 y_{i,t} + \beta_1 y_{i,t-1} + \beta_2 y_{i,t-2} + \gamma Z_{i,t} + \delta Q_{i,t,s} + \xi Trend_{i,t,s} + \varepsilon_{i,t,s}. \quad (8)$$

In the above specifications,  $Z_{i,t}$  is a vector of variables that belongs to  $\mathcal{I}_t^R$  for household  $i$ .

We only include information about the head of household.<sup>5</sup> This vector contains race, educational attainment, marital status, family size, a currently laid-off or unemployed dummy, a currently self-employed dummy, and a vector of occupation and industry dummies, all evaluated at age  $t$ . In addition,  $Z_{i,t}$  includes a fourth-order age polynomial, evaluated at age  $t + s$ .  $Q_{i,t,s}$  is a vector of variables that belongs to the augmenting information set,  $\mathcal{I}_t^F$ . We assume  $Q_{i,t,s}$  includes family size, marital status, a retirement dummy, a part-time dummy, a self-employed dummy, and a vector of occupation and industry dummies, all evaluated at age  $t + s$ .<sup>6</sup>

In addition to  $\mathcal{I}_t^R$  and  $\mathcal{I}_t^A$ , our specification deviates from most previous specifications in that we include both current and lagged income in the projection equations. In principle, if we have a very long income history for a given household, a univariate time series model could potentially have some decent forecasting power. Such a long time series of household income is also useful for identifying the income process heterogeneity that Lillard and Weiss (1979) and Guvenen (2007) have stressed. Including some recent income history in the projection equation can help tease out information about recent income shocks and capture part of the individual-specific information of income growth that is not revealed by current income and other observable characteristics. In practice, our model includes two lags to preserve degrees of freedom. Finally, we added a simple calendar year trend to control for aggregate economic growth.<sup>7</sup>

Several important caveats apply to our specifications. First, as most households do make plans about their family and career ahead of time, 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, or whether they will change occupation and

---

<sup>5</sup>Our data source does not have information for spouses that is as complete as heads. Including the available spousal information does not qualitatively change the results.

<sup>6</sup>We do not include  $t+s$  educational attainment because the data we use do not regularly update household education attainment information.

<sup>7</sup>We also estimate the model with a vector of year dummies, assuming households have perfect foresight on future aggregate growth and business cycle fluctuations. The results are very much similar to the model with a linear trend.

industry. However, it is less likely that households know all this information when  $s$  is large.<sup>8</sup> That said, the information we have is limited to what is collected in surveys. There are other information elements  $\iota$  such that  $\iota \in \mathcal{I}_t^H$ , but  $\iota \notin \mathcal{I}_t^R \cup \mathcal{I}_t^F$ . Therefore, it is generally not true that the AIS estimates yield the lower bound of household income uncertainty. However, as we will present later, the comfortably large value of  $R^2$  of Eq. (8) reassures us that the unobservable information only has a limited effect. Second, it is worthwhile to point out that Eqs. (7) and (8) are merely forecasting equations, and they are estimated solely on the basis of maximizing  $R^2$ . We are not estimating a structural model, so the coefficients estimated for Eqs. (7) and (8) should not be interpreted as structural parameters. Third, because we let the households use the estimated coefficients to project their future income, we implicitly assume that either these coefficients are stable over time or that households have foresight on the future values of these coefficients (e.g., the household would know in advance the return to schooling and experience). We follow Carroll (1994) and resort to the first assumption in the baseline analysis. As a robustness check, we alternatively allow for time-varying coefficients and estimate the coefficients using a lagged sample. By this exercise, we found that the so-estimated income uncertainty over all projection horizons is somewhat higher than in our baseline specifications. However, all key results are qualitatively unaltered. Finally, although we allow income uncertainty to vary with age, we compute these forecast errors by applying the same coefficients of the forecasting model to all households. As another robustness check, we also estimated 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 any age dependence in the accuracy of the forecasting model.

---

<sup>8</sup>Our study forecasts future income up to twenty-five years ahead.

## 4 Data Description and Sample Construction

We used data from the Panel Study of Income Dynamics (PSID), which is a nationwide longitudinal survey of households conducted by the Institute of Social Research at the University of Michigan. Before 1997, the PSID was an annual survey; after 1997, it became a biennial survey. We use the PSID data from 1968 to 2005, totaling 34 waves, covering nearly four decades. As of the 2005 wave, the PSID surveyed 8,002 households.<sup>9</sup>

In addition to extensive information about work status, employment history, and demographic characteristics, the PSID collects detailed information on income, which allowed us to study various components of total family income. Moreover, the longitudinal structure of the PSID permits us to link the information collected for the same household in two waves that are  $s$  years apart. For instance, consider a household that had been surveyed for ten years from 1971 through 1980. When we projected the five-year-ahead income, this household rendered five current and future income pairs  $(t, t + s) = (1971, 1976), \dots, (1975, 1980)$ . This structure is essential to estimate the projection equation and to experiment with alternative information specifications.

We estimated Eq. (7) and (8) for each forecast horizon. Accordingly, the samples are constructed separately for each forecast horizon. Several rules apply when we construct the sample. First, we include only the nationwide representative sample and exclude the households in the low-income supplemental sample.<sup>10</sup> Second, to focus on the working-age households, we restrict household heads to be older than 23 in the base year and younger than

---

<sup>9</sup>The PSID not only surveys the households in the original sample stratified in 1968, but also the households headed by the grown-up children or other moved-out members of the households in the original sample. The PSID survey has a very high retention rate. The vast majority of households surveyed in one year will continue to participate in the next wave—more than 1,200 households have 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 4,802 households, whereas the 1994 wave surveyed more than 10,000 households. The PSID subsequently stopped surveying households in its noncore sample.

<sup>10</sup>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.

65 in the year to be forecasted. For example, in the sample we use to forecast five-year-ahead income, we restrict the heads of sample households to be younger than 60. Consequently, the sample we use to study forecast errors at farther horizons is typically smaller than the sample used for a closer horizon. Third, we remove households that reported zero or negative income in the base year, the two lagged years, or the year to be forecasted. Fourth, we remove households whose heads were disabled or retired, were primarily keeping house, or were students in the base year. Fifth, we remove households whose heads reported zero working hours in the base year and the two lagged years.<sup>11</sup> Finally, in order to minimize the bias caused by outliers and measurement errors, we trim off households with very high- or very low-income levels and growth rates between age  $t$  and  $t + s$ .<sup>12</sup>

To explore the precautionary saving implication of the estimated income uncertainty, we primarily focus on the most relevant income definition, family non-asset income. Because Carroll and Samwick (1997) also study family non-asset income, examining the same income definition also facilitates comparison with their results. We discuss our estimates of uncertainty profile for other income measures in the robustness analyses. In our study, we estimate the variance of forecast errors of future family non-asset income up to 25 years ahead. For illustration purposes, Table 1 lists the number of observations used in the estimation for each forecast horizon. We notice that the sample constructed for the two-year-ahead forecast has the largest number of observations (nearly 46,000) because the PSID became a biennial survey after 1997. The sample sizes largely decline with forecast horizons. However, even at the 25-year-ahead horizon, we maintain a decent sample size that is above 2,200. The income variables are deflated using 1982-1984 dollars. In addition, all waves but 2003 and 2005 of the PSID data have a 1970 census industry and occupation code. The 2003 and 2005 PSID

---

<sup>11</sup>We do not require positive hours in the year to be forecasted because we allow that the head may be retired at age  $t + s$ .

<sup>12</sup>We 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  $t + s$ . This strategy may appear to be aggressive and lead to an underestimate of income risks. However, when comparing the results, we apply the previous methods on a sample after the same trimming and find that the semiparametric estimate of income uncertainty is substantially lower than the estimates of other methods.

Table 1: Number of Observations for Each Sample

<i>s</i>	1 Year	2 Years	3 Years	4 Years	5 Years
<i>N</i>	43,313	47,109	37,583	39,782	32,926
<i>s</i>	6 Years	7 Years	8 Years	9 Years	10 Years
<i>N</i>	33,421	28,761	27,740	24,947	22,736
<i>s</i>	11 Years	12 Years	13 Years	14 Years	15 Years
<i>N</i>	20,553	18,372	16,500	14,640	12,970
<i>s</i>	16 Years	17 Years	18 Years	19 Years	20 Years
<i>N</i>	11,331	9,957	8,572	7,452	6,293
<i>s</i>	21 Years	22 Years	23 Years	24 Years	25 Years
<i>N</i>	5,342	4,376	3,603	2,834	2,244

data used the industry and occupation code derived from the 2000 census. To be consistent, we regroup the 2000 census industry and occupation code to match the 1970 categories as close as we can.

## 5 Main Empirical Results

We construct five-year centered moving averages of the estimated variances of forecast errors to smooth out noises in the series. Recall that we have from one-year-ahead to twenty-five-year-ahead projections, and our sample household heads were between the age of 23 to 64 in the base year. The centered moving average uncertainty matrix  $\Omega$  therefore has a dimension of  $38 \times 25$ , with each row corresponding to the same age and each column corresponding to the same projection horizon. To demonstrate how income uncertainty evolves over the life cycle, for each model specification (RIS and AIS), we plot uncertainty age profiles at four forecast horizons (four columns of  $\Omega$ ): one-year and two-year ahead, representing the near future, as well as five-year and ten-year ahead, representing the medium and more remote future. To illustrate the statistical significance of our results, we also plot the 95% confidence interval band for each profile .

## 5.1 Income Uncertainty Changes over the Life Cycle

We begin with examining the uncertainty of total family non-asset income estimated under the more restrictive information specification, the RIS equation. The results are shown in the top four panels of Figure 1. Consistent with our intuition that uncertainty regarding income at a fixed future date resolves as the date approaches, the level of uncertainty regarding income two years ahead is considerably higher than the income uncertainty one year ahead. The confidence intervals associated with both uncertainty profiles indicate this gap is statistically significant. Notably, both profiles share a similar U-shaped pattern over the life cycle. Income uncertainty is high when the head is young; it declines during his late twenties and early thirties. Afterward, uncertainty fluctuates within a narrow range at relatively low levels before rising in his mid-fifties, which continues as households approach retirement. The ratios between the maximum and minimum levels of uncertainty over the life cycle (the max-min ratio) are both about 1.5 for the one-year- and two-year- ahead future income. The plotted 95% confidence intervals suggest that both profiles are rather tightly estimated, which reinforces the statistical significance of the U-shaped pattern. A similar convex profile of income uncertainty is documented in Baker and Solon (2003). They decompose the stochastic part of income into a transitory and a permanent component and find that the age profile of the variance of the transitory innovations exhibits a pronounced U-shaped pattern.<sup>13</sup> Our results, in contrast, do not rely on the dichotomy of transitory and permanent shocks.

Profiles of similar shape repeat in the left panel on the second row, which shows income uncertainties at the medium (five-year-ahead) future. The standard errors of the estimated profiles are somewhat larger than in the top two panels because we have a smaller sample over these longer projection horizons. The uncertainty associated with income five years ahead hits bottom in the mid-thirties and rises gradually through the early fifties. After that,

---

<sup>13</sup>A similar pattern was also noticed by Gordon (1984).

income uncertainty rises more rapidly as households approach retirement ages. The profile of income uncertainty ten years ahead does not exhibit a pronounced decline during young ages. The level of uncertainty stays low through the mid-thirties before rising in the mid-fifties. Comparing with the near future, the variations of the five-year-ahead income uncertainty over the life cycle are about as pronounced, with the max-min ratio equal to 1.5. However, the variations of income uncertainties at the more remote future are less pronounced, with the max-min ratio equal to merely 1.3 for the ten-year horizon. Moreover, although the level of uncertainty in the medium future is much higher than in the near future, the increase is not proportional to the forecast horizon. The two-year-ahead uncertainty is on average 60% higher than the one-year-ahead uncertainty. If these uncertainties rise proportionally with forecast horizons, the five-year-ahead uncertainty, for example, should be 240% higher than the one-year-ahead uncertainty. However, it is only 140% higher. As we will further elaborate later in the paper, this nonlinearity suggests the persistent component of income shocks is not exactly permanent.

Income-uncertainty profiles at the very remote future (up to 25 years) are also estimated (not shown). Because of smaller sample sizes at these horizons, these profiles are estimated with wider confidence intervals. Like the ten-year-ahead uncertainty age profile, the very remote income uncertainty age profiles all show a pronounced upward trend throughout the life cycle but do not exhibit clear declines during young ages. Finally, for overlapping ages, the twenty-five-year-ahead profile is on average 450% higher than the one-year-ahead profile, to a much less extent than what proportional increases imply.

Now, we turn our attention to the relaxed information specification, AIS, in which we allow households to have foresight on some of their future demographic and work-related characteristics. The results are shown in the bottom four panels of Figure 1. Two features are noteworthy. First, because these future income projections are conditioned on a larger information set, the AIS forecast error variances are significantly smaller than the variances estimated using the RIS equation. This pattern holds for all forecast horizons and for

households of all ages. Moreover, the discrepancy between the RIS and AIS results widens with the forecast horizon. The estimated one-year-ahead uncertainty under the RIS is only 6% higher than that estimated under the AIS, whereas the gap is about 25% at the ten-year horizon. It widens to 36% at the 25-year horizon. This widening trend is consistent with our intuition because the difference between  $\mathcal{I}^R$  and  $\mathcal{I}^A$  is what households are assumed to know about themselves at age  $t + s$  as of age  $t$ . Given that these characteristics typically change only gradually, if  $s$  is small, the correlation between the elements in  $\mathcal{I}^R$  and  $\mathcal{I}^F$  will be high, discounting the net value of  $\mathcal{I}^F$ . Conversely, if  $s$  is large,  $\mathcal{I}^R$  should have less predictive power on  $\mathcal{I}^F$ . Consequently, augmenting with  $\mathcal{I}^F$  adds more new information and improves the forecasting performance more significantly. Second, the U-shaped contours at the near and the intermediate future horizons are similar to those shown in the top panels. However, the max-min ratios are somewhat lower for the AIS estimates, especially for the profiles of more remote future horizons. This result is not surprising because a good portion of the rise in income uncertainty after middle age comes from retirement and related risks, which are included in the augmented information set.

## 5.2 Correlations of Forecast Errors

Besides the levels of income uncertainties over the life cycle, to completely characterize the income uncertainty, we must know how the stochastic components of income at different horizons are correlated. Figure 2 presents the correlations between the one-year-ahead forecast errors and the forecast errors at other horizons. To keep the graph readable, we only plot the correlations for the households whose heads are 30, 40, and 50 years old that are estimated using the RIS equation. The AIS results are very similar. We note that, on the one hand, because the correlations are positive at all forecast horizons, income shocks must have a persistent component. On the other hand, the correlations decline substantially with forecast horizon. The correlation between the one-year-ahead and the two-year-ahead forecast errors is about 0.5, whereas the correlation between the one-year-ahead and ten-year-ahead

forecast errors is only about 0.2. This rapid decline suggests that the persistent component of income shocks is more likely to be an AR process than a unit-root process. Although we cannot rule out the existence of a unit-root component, if one exists, its innovations should have a relatively small variance. The chart also reveals that the correlations of households of various ages are very similar, suggesting little changes in the persistence of income shocks over the life cycle, in contrast to the evolution of shock variances.

### 5.3 The (In)significance of Lagged Income

We include lagged income in the projection equation. The motivation is to capture some heterogeneity that cannot be accounted for by a common trend determined by other observable characteristics. Naturally, this method will not capture all the heterogeneity set forth by Lillard and Weiss (1979), but it does acknowledge that past income history can be useful in predicting future income, even after controlling for current income levels and the standard household demographic and work-related characteristics. Indeed, for almost all family income components and forecast horizons, the estimated coefficients of the lagged income are highly significant (with  $p$ -value typically smaller than 0.001). In addition, the sum of the coefficients of current and lagged income is larger than the coefficient of the current income estimated in a model with no lagged income included, implying that with lagged income included, the projection picks up more signal from the income history. However, what is somewhat puzzling to us is that although the coefficients of the two lagged incomes are statistically significant, adding them does not beef up the overall fitness of the projection equation to a great extent. The  $R^2$  of the models increase only slightly, on average 2%, when lagged income is included.

### 5.4 Robustness Tests

To examine whether the dynamics of income uncertainty over the life cycle that we have shown are a spurious consequence of our choice of income definition, model specifications,

or the shrinking sample sizes over projection horizons, we conduct a series of robustness analyses. First, we redo the analysis above using household head labor income. Figure 3 plots the variances of forecast errors over the life cycle. Like in Figure 1, household head labor income uncertainty also changes substantially over the life cycle. However, different from family non-asset income, labor income uncertainty follows a J-shaped pattern, with a much more pronounced rise of uncertainty as the head approaches retirement age. The late life cycle rise is considerably smaller under the AIS specification, in which we allow the head to have some foresight about his retirement (part-time) status in the year to be forecasted. Figure 4 plots the correlations of forecast errors over horizons, which are largely similar to what Figure 2 has shown. Similar analyses done for total family income and head and wife combined non-asset income yield results similar to Figures 1 and 2, whereas analyses done for head and wife combined labor income and head wages and salaries yield results similar to Figures 3 and 4.

Second, we examine whether changes in uncertainty over the life cycle are due to model misspecifications. Recall that we project the future income of households at different ages using the same set of coefficients. If the projection-equation coefficients should be age specific and the coefficients we use are closer to the true parameters for the middle-aged households than for the younger and older households, the one-size-fits-all approach will reduce fitness for younger and older households and artificially increase the estimated income uncertainties for these age groups. We divide our sample into five subgroups by household head age and estimate Eqs. (7) and (8) separately for each subgroup. Then, we calculate variances of forecast errors constructed using the revised projection coefficients. We find that both the uncertainty age profiles and correlation curves are similar to what Figures 1 and 2 show.

Third, in the same spirit, we examine the effects of allowing for time-varying projection coefficients that households are assumed to have no foresight on. We split the sample into two parts and estimate the coefficients using the earlier sample. We then use the estimate coefficients to project future income in the later sample. This method assumes that house-

holds are backward looking in figuring out the coefficients applicable to the future and can be exceedingly restrictive. For example, one does not have to infer returns to college education in the 1990s using exactly the returns in the 1970s. The purpose of the analysis is to illustrate how much bias can potentially be introduced if we ignore the time dependence of the projection coefficients. We only do this analysis up to ten years ahead because we need long enough subsamples to estimate the coefficients. We find that the contours of the life cycle profiles of income uncertainty and the correlations of the forecast errors estimated using this method are very similar to those estimated assuming time-invariant coefficients. Not surprisingly, the levels of income uncertainty are somewhat higher. However, this gap narrows consistently with projection horizons. At the near term, the backward-looking estimates of variances of forecast errors are about 17% higher than in the baseline specification, and at the ten-year ahead horizon, the gap becomes 7%.

Finally, we examine whether changes in the sample size as we vary forecast horizons (as given in Table 1) might drive the shape of the uncertainty profile. We estimate the forecast equations using a smaller common sample that spans at least ten years and reconstruct the income uncertainty measures up to ten years ahead. Again, all results are qualitatively unchanged.

## 6 Comparison with Earlier Results

How substantive are the innovations we have introduced with the semiparametric measures of income uncertainty? How different are our results compared to previous studies? We answer these questions by contrasting our results to the income uncertainty estimates in the influential work of Carroll and Samwick (1997). Three reasons lead us to choose these results to compare with. First, their specific interest was precautionary saving, so their estimates focus explicitly on income uncertainty rather than inequality, volatility, or other types of income variability. Second, they also use the PSID data, making a comparison and

interpreting the differences easier. Third, apart from the specified innovations, our models share many similarities with theirs.

As we summarized in Eqs. (1)-(2), Carroll and Samwick (1997) decompose  $y_t$  into a permanent component,  $p_t$ , and a transitory shock,  $\varepsilon_t$ . The permanent component,  $p_t$ , is further assumed to follow a random walk with predictable income growth  $g_t$  such that

$$p_t = g_t + p_{t-1} + \eta_t, \tag{9}$$

where  $\eta_t$  is the shock to permanent income. Let  $\sigma_\eta^2$  and  $\sigma_\varepsilon^2$  be the variance of the permanent and transitory shocks. Define  $VAR_d$  as the variance of the  $d$ -year income difference. Filtering out  $g_t$ , it is easy to show that

$$VAR_d = Var[y_{t+d} - y_t] = d\sigma_\eta^2 + 2\sigma_\varepsilon^2, \tag{10}$$

noting that the econometrician does not know how either  $y_t$  or  $y_{t+d}$  decomposes into their permanent and transitory parts. The innovation variances  $\sigma_\eta^2$  and  $\sigma_\varepsilon^2$  can be estimated by evaluating  $Var[y_{t+d} - y_t]$  at various difference lengths,  $d$ . Using a PSID sample from 1981 to 1987, Carroll and Samwick (1997) report  $\sigma_\eta^2 = 0.022$ , and  $\sigma_\varepsilon^2 = 0.044$ . Gourinchas and Parker (2002) report almost identical results. These parameters are changed noticeably when we estimate their model using an extended PSID sample that covers a longer period, from 1968 to 2005. The updated estimates call for a significantly larger variance of the transitory income shocks,  $\sigma_\varepsilon^2 = 0.054$ , and a significantly smaller variance of permanent shocks,  $\sigma_\eta^2 = 0.012$ .

Figure 5 contrasts the total family non-asset income uncertainty at various forecast horizons implied by the original and updated Carroll and Samwick estimates with our semiparametric profiles (pooled across all ages). The  $\sigma_\eta^2$  and  $\sigma_\varepsilon^2$  reported in Carroll and Samwick (1997) imply a higher and steeper linear profile, whereas the  $\sigma_\eta^2$  and  $\sigma_\varepsilon^2$  that we estimated using the same methodology with the extended PSID sample imply a lower and flatter profile. The linearity between income uncertainties and forecast horizons arises because of the

random walk assumption imposed. The slope of the linear profile is equal to the variance of permanent shocks, and the intercept is equal to two times the variance of transitory shocks. The higher concave curve is the semiparametric profile estimated using the RIS equation, while the lower concave profile is estimated using the AIS equation.

Several features appear in this chart are noteworthy. First, we notice that the Carroll-Samwick profile is substantially higher than the semiparametric profiles estimated even under the restrictive information specifications, the RIS equation. At the near horizons, the gap is about 20%, whereas at remote horizons, because the RIS semiparametric profile is concave, the wedge widens substantially to about 60%. Second, the updated Carroll-Samwick profile has a higher intercept but a flatter slope due to the larger  $\sigma_\epsilon^2$  but smaller  $\sigma_\eta^2$ . Third, the updated Carroll and Samwick profile is higher than the RIS profile in the near term but coincides with the RIS profile beyond five years. The updated Carroll and Samwick profile remains uniformly and substantially higher than the AIS profile. This gap is more striking, taking into account that the predictable component of income is constructed using age  $t + s$  information in the Carroll and Samwick (1997) approach, whereas the semiparametric approach conditions on a mix of age  $t$  and age  $t + s$  information. Moreover, not shown in the chart, the updated Carroll-Samwick profile is also higher than the profile allowing for time-varying projection coefficients and assuming backward-looking households (the third robustness check).

Why are the uncertainty profiles implied by the original and updated Carroll and Samwick estimates so different? Two potential factors may contribute to this change. First, we notice that in Carroll and Samwick (1997), the maximum difference length corresponds to  $d = 6$ , which is relatively small. Using a longer panel not only adds data collected in more years, but also allows us to study the difference in income over longer intervals. To assure the comparability of our PSID sample with theirs, we first set  $d = 6$  and found the estimated transitory and permanent shocks variances to be 0.036 and 0.023, respectively. These estimates are very close to Carroll and Samwick (1997). Subsequently, we update

these estimates for various choices of  $\max[d]$  to examine whether the estimates are sensitive to variation in  $d$ .

Figure 6 contrasts how the estimated variances of transitory and permanent shocks, identified by the specification of Eqs. (9)-(10), vary with  $\max[d]$ . Because we project future income up to 25 years ahead in our nonparametric estimates, we choose the largest value of  $\max[d]$  to be 25. The permanent shock variance decreases, whereas the transitory shock variance increases monotonically with  $\max[d]$ . This phenomenon implies either that the variance of the permanent shocks—if the persistent shock process is indeed a random walk—is estimated with a significant upward bias using a short panel, or that the persistent shocks do not follow a random walk process. To see this, suppose the true model of the shocks to “permanent” income is

$$p_t = \rho p_{t-1} + \eta_t, \tag{11}$$

instead of a unit-root process  $p_t = p_{t-1} + \eta_t$ . After some algebra, we can show that  $VAR_d$  is not equal to  $d\sigma_\eta^2 + 2\sigma_\epsilon^2$  as in Eq. (10). Rather, we have

$$VAR_d = Var[y_{t+d} - y_t | y_t] = \frac{1 - \rho^{2d}}{1 - \rho^2} \sigma_\eta^2 + [1 + \rho^{2d}] \sigma_\epsilon^2. \tag{12}$$

Using L'Hôpital's Rule, it is easy to verify that (10) is the limiting case when  $\rho \rightarrow 1$ . If Eq. (10) holds,  $\sigma_\eta^2$  can be calculated by taking the difference  $VAR_{d+1} - VAR_d$ . The Carroll and Samwick estimates can be viewed as a weighted average of such variances across  $d \leq \max[d]$ . However, if Eq. (12) is the true model, taking the difference between  $VAR_{d+1}$  and  $VAR_d$ , we get the presumed estimate of the variance of permanent shocks,

$$VAR_{d+1} - VAR_d = \rho^{2d} [\sigma_\eta^2 + (\rho^2 - 1) \sigma_\epsilon^2]. \tag{13}$$

Assuming  $\rho$  is close to 1 (e.g.,  $\rho = 0.9$ ) and  $\sigma_\eta^2$  and  $\sigma_\epsilon^2$  are of the same order of magnitude, then  $VAR_{d+1} - VAR_d$  will be a decreasing function of  $d$ . Therefore, when we increase  $\max[d]$ , the average of  $VAR_{d+1} - VAR_d$  over  $d \leq \max[d]$  also decreases. Meanwhile, a downward biased estimate of the permanent shock variance leads to an upward biased estimate of the transitory shock variance.

Finally, Carroll and Samwick (1997) also estimate the variance of the permanent and transitory income shocks by age. However, their estimates do not imply a U-shaped life cycle pattern. Indeed, their transitory shock variance exhibits a hump-shaped pattern over the life cycle, peaking in the early forties, whereas their permanent shock variance demonstrates more irregular life cycle dynamics.

To summarize, in contrast to Carroll and Samwick (1997) and later work using similar methods (e.g., Gourinchas and Parker (2002)), the semiparametric estimates of income uncertainty introduced in this paper reveal a uniformly lower level of income uncertainty over all future horizons, and suggest that the persistent component of income shocks is likely not a permanent shock. Our approach also implies life cycle dynamics of income uncertainty that are more consistent with layman's intuition and the results of Baker and Solon (2003) and Gordon (1984).

## **7 Consumption and Precautionary Saving over the Life Cycle: An Application of the Semiparametric Estimates**

As we discussed in the beginning of the paper, a battery of household decisions critically depend on households' uncertainty about their future income. We revisit one such question, the optimal quantity of consumption and precautionary saving over the life cycle, as an example that illustrates how household decisions under the semiparametric estimate of income uncertainty introduced in this paper differ from those derived using the preexisting income uncertainty estimates.

### **7.1 Background**

The canonical Rational-Expectations Life Cycle/Permanent-Income Hypothesis (RE-LCPIH) predicts that rational consumers should allocate consumption over the life cycle in such a way as to maximize lifetime utility, which in turn implies a monotonic consumption profile

over the life cycle. However, since Thurow (1969), many empirical studies have documented a hump-shaped pattern of life cycle consumption. In light of this apparent inconsistency, a large volume of literature has added various features to the standard RE-LCPIH model to account for a hump-shaped consumption profile.<sup>14</sup>

We focus on one of the main factors raised to explain the hump-shaped consumption profile—precautionary saving—for it directly speaks to the effect of income uncertainty. Nagatani (1972) first suggested that precautionary saving would reduce consumption early in the life cycle, and Skinner (1988) and Feigenbaum (2008a) 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 life cycle, this will lead to a concave consumption profile. It has been shown extensively that, in partial-equilibrium models calibrated against the measures of uncertainty described in Section 2, precautionary motives can induce a sizable hump (see for example, Carroll and Summers (1991); Hubbard, Skinner, and Zeldes (1994); Carroll (1997); Gourinchas and Parker (2002); and Feigenbaum (2008b)).

However, it is not likely that precautionary saving single-handedly causes the hump, for other factors, such as leisure and consumption substitution (Heckman (1974); and Becker and Ghez (1975); and Bullard and Feigenbaum (2007)) or time-varying mortality risk (Feigenbaum (2008c) and Hansen and Imrohoroglu (2008)), can also account for the hump. This insight imposes a challenge to the existing parametric estimates of income uncertainty as the models calibrated with these estimates often yield a consumption hump that is larger than observed in data. We introduce

$$S = \frac{E[c_{t^*}]}{E[c_0]}$$

as a measure of the magnitude of the consumption hump, where  $t^*$  is the age that maximizes  $E[c_t]$ . The data constructed in Gourinchas and Parker (2002) suggest  $S = 1.17$ . However, under fairly plausible parameterizations, the income risks reported in Gourinchas and Parker

---

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

(2002) call for  $S = 1.28$ . Conceivably, adding other factors to these models would further exaggerate the size of the hump.

In contrast, here we show that with the semiparametric estimates of income uncertainty introduced in this paper, precautionary saving typically yields a significant but substantially smaller hump than is observed in the data. This is because our estimates call for smaller variances and lower persistence, reducing the lifetime income uncertainty households face, in particular at a young age.

## 7.2 A Life Cycle Consumption-Saving Model

Ideally, we would like to model the income process with age-varying uncertainties that replicate all the variances and correlations that we estimate in Section 5. However, the parsimonious income processes typically employed in this literature assume constant shock variances. Generalizing such income processes to allow for age-dependent shock variances is a difficult task that is beyond the scope of this paper. (See Feigenbaum and Li (2010a) for more on this issue.) Thus, we consider a general-equilibrium overlapping-generations model with incomplete markets à la Huggett (1996) and an income process with age-independent shock variances.

In each period, a continuum of unit measure of consumers is born. Each household lives for  $T_w$  working periods and  $T_r$  retirement periods. Income is the only source of uncertainty in the model. A consumer maximizes

$$E \left[ \sum_{t=0}^{T_w-1} \beta^t u(c_t; \gamma) + \beta^{T_w} V_{T_w}(x_{T_w}) \right], \quad (14)$$

where  $c_t$  is consumption at age  $t$ ,

$$u(c; \gamma) = \begin{cases} \frac{c^{1-\gamma}}{1-\gamma} & \gamma \neq 1 \\ \ln c & \gamma = 1 \end{cases} \quad (15)$$

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

We follow the literature on precautionary saving by assuming that at age  $t$  the logarithm of income decomposes into

$$y_t = w + a_t + p_t + \varepsilon_t, \quad (16)$$

$$p_t = \rho p_{t-1} + \eta_t, \quad (17)$$

where  $a_t$  is a deterministic age-dependent factor,  $\eta_t$  is a persistent shock,  $\varepsilon_t$  is a temporary shock, and  $w$  is the logarithm of the economy-wide wage per efficiency unit,  $W$ . Specifically, we assume that  $\eta_t$  and  $\varepsilon_t$  are independent processes with variances  $\sigma_\eta^2$  and  $\sigma_\varepsilon^2$ , respectively, and such that, unconditionally,

$$E[e^{z_t}] = E[e^{p_t}] = 1. \quad (18)$$

There is one intertemporal asset, a risk-free bond, that pays the fixed gross interest rate  $R$ . Let  $b_{t+1}$  denote the quantity of bonds an agent purchases at age  $t$  that would pay  $Rb_{t+1}$  at  $t + 1$ . The budget constraint for a household at age  $t$  is given by

$$c_t + b_{t+1} = y_t + Rb_t. \quad (19)$$

Since our intent is to focus on the extent to which uncertainty and precautionary saving help explain the consumption hump as opposed to borrowing constraints, we allow borrowing but with full commitment to debt contracts. Thus, the Aiyagari (1994) borrowing limit is in force.<sup>15</sup>

To complete the model, we assume there is a production sector with the Cobb-Douglas production function

$$Y = K^\alpha N^{1-\alpha}, \quad (20)$$

---

<sup>15</sup>This borrowing limit has negligible effect on the results in this finite-lifetime model, as most agents will not have time to accumulate so much debt. Although much of the literature includes a tiny probability of a very large negative income shock, this assumption would obscure the effects of the uncertainty we study here, so no such shock is included in our parametrization.

where  $K$  is the aggregate capital stock, equal to the stock of savings in the economy

$$K = E \left[ \sum_{t=0}^{T-1} b_t \right], \quad (21)$$

and  $N$  is the labor supply, equal to the stock of labor efficiency units supplied to the market:

$$N = \sum_{t=0}^{T-1} e^{at}. \quad (22)$$

In a steady-state general equilibrium, the factor prices  $R$  and  $W$  will be constant and must satisfy the profit-maximization conditions for firms:

$$R = \alpha \left( \frac{K}{N} \right)^{\alpha-1} + 1 - \delta \quad (23)$$

and

$$W = (1 - \alpha) \left( \frac{K}{N} \right)^{\alpha} \quad (24)$$

### 7.3 Model Calibration and the Results

Note that Eq. (17) implies that persistent shock follows an AR(1) process as in Huggett (1996) and Storesletten, Telmer, and Yaron (2004) as opposed to the unit-root process considered by Carroll and Samwick (1997) and Gourinchas and Parker (2002). The model produces a conditional variance matrix with elements

$$Var[\ln y_{t+h}|y_t] = \frac{1 - \rho^{2h}}{1 - \rho^2} \sigma_{\eta}^2 + (1 + \rho^{2h}) \sigma_{\varepsilon}^2. \quad (25)$$

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

$$corr(\ln y_{t+h}, \ln y_{t+1}|y_t) = \frac{\rho^{h-1} \sigma_{\eta}^2 + \rho^{h+1} \sigma_{\varepsilon}^2}{\sqrt{V[\ln y_{t+h}|y_t] V[\ln y_{t+1}|y_t]}}. \quad (26)$$

We calibrate the parameters  $\rho$ ,  $\sigma_{\eta}$ , and  $\sigma_{\varepsilon}$  in order to best replicate the variance and correlation matrices presented in Section 5, minimizing a root mean squared deviation of the values generated by the model according to (25)-(26) from their estimated values. We

implement the calibration separately for the RIS and AIS specifications. Finally, we choose  $a_t$  to match the income profile mean in Gourinchas and Parker (2002). We will compare the life cycle consumption profile implied by the income processes calibrated, using four income uncertainty estimates—the RIS, the AIS, the Gourinchas and Parker (2002) estimates, and the updated Carroll and Samwick estimates presented in Section 6.

The remaining parameters are calibrated as follows. We set the share of capital to a typical value  $\alpha = 0.34$  from the literature. We set  $\gamma = 3$  to match properties of the life cycle consumption profile as in Feigenbaum (2008b), leaving  $\beta$  and  $\delta$  to be calibrated. We experiment with two parametrization strategies. First, we exploit the common presumption that the ratio  $C/Y$  and  $K/Y$  can be measured. In equilibrium, we must have

$$C + \delta K = Y, \tag{27}$$

where

$$C = E \left[ \sum_{t=0}^{T-1} c_t \right]$$

is aggregate consumption and  $\delta K$  is steady-state investment. Therefore,  $\delta$  is determined by (27) and  $\beta$  can be accordingly calibrated to match the target ratio of  $K/Y$ . We follow Rios-Rull (1996) and set  $C/Y = 0.75$ . The value of  $K/Y$  is a subject of more debate in the literature. Therefore, we experiment with two values for  $K/Y$ —3 and 3.5.<sup>16</sup> The calibration results and the computed values for consumption peak ratio  $S$  are reported in Table 2. Figures 7 and 8 present the consumption profile over the life cycle under the two parameterizations of  $K/Y = 3$  and  $K/Y = 3.5$ , respectively.

---

<sup>16</sup> $K/Y = 3.5$  is close to the upper end of the range of the capital-to-output ratio reported previously.

Table 2: Calibration Results and the Magnitude of the Consumption Hump

Income process	$\rho$	$\sigma_\eta$	$\sigma_\varepsilon$	Target $K/Y$	$\beta$	$S$
RIS	0.918	0.207	0.168	3.0	0.923	1.021
AIS	0.906	0.191	0.177	3.0	0.931	1.000
GP	1.000	0.146	0.210	3.0	0.905	1.275
CS updated	1.000	0.110	0.232	3.0	0.926	1.125
RIS	0.918	0.207	0.168	3.5	0.937	1.081
AIS	0.906	0.191	0.177	3.5	0.944	1.043
GP	1.000	0.146	0.210	3.5	0.920	1.342
CS updated	1.000	0.110	0.232	3.5	0.939	1.180
Memo	$\alpha = 0.34$	$\gamma = 3$	$\delta = 8.3\%$ if $K/Y = 3$	$\delta = 7.1\%$ if $K/Y = 3.5$		

Note: The  $\sigma_p$  and  $\sigma_z$  parameters for GP are reported in Gourinchas and Parker (2002); those for the CS updated are estimated in Section 6.

Several observations are noteworthy in the table and the charts. First, the calibrated income process under the AIS specification involves more volatile temporary shocks but persistent shocks that are both less persistent and less volatile than under the RIS specification. Second, it is somewhat surprising that the computed magnitude of consumption hump,  $S$ , is sensitive to the choice of  $K/Y$ , with higher  $K/Y$  inducing great consumption humps. Specifically, it is perhaps counterintuitive that the size of the consumption hump should increase with  $K/Y$ . This is purely a general-equilibrium effect. In partial equilibrium, with fixed interest rates,  $S$  is an increasing function of the ratio of the amount of uncertainty, measured by the variance of income, to total wealth, including financial wealth and current income, squared. Assuming the standard deviation of income scales linearly with income,  $S$  will be a decreasing function of the wealth-to-income ratio, the microeconomic analog of  $K/Y$ . This is the reason why the standard incomplete-markets theory predicts that people will build up a buffer stock of wealth—to insure themselves against income risk. However, in general equilibrium, precautionary saving increases  $K$  in the aggregate. For a fixed risk aversion, low  $K/Y$  implies low precautionary saving and high  $K/Y$  implies high precautionary saving. Third, and most importantly, the consumption humps called for by the semiparametric estimates of income uncertainties are much smaller than those implied by both the Gourinchas and Parker (2002) and the updated Carroll-Samwick estimates of income un-

certainties. The  $S$  values derived under the latter two income processes often exceed the size of the hump in the empirical data, whereas the  $S$  values derived under the AIS and the RIS specifications are always smaller than the empirical observation, even under the rather extreme parametrization of  $K/Y = 3.5$ .<sup>17</sup>

Because the target ratio of  $K/Y$  is at the best imprecisely observed in the data, we also implement an alternative parametrization strategy. We take the calibrated value for  $\beta = 0.923$  under the RIS specification with  $K/Y = 3$  and apply this specific value of  $\beta$  to models with other specifications of income process. We then evaluate, when the discount factor is fixed, whether the semiparametric estimates of income uncertainty imply a more plausible size of the consumption hump and the ratio of  $K/Y$ . As shown in Table 3, when  $\beta$  is fixed to equal to 0.923, the RIS and AIS income processes continue to yield smaller values of the  $S$  ratio, whereas the ratios derived from the updated Carroll-Samwick and the Gourinchas-Parker estimates of income uncertainties are close to or even above the ratio observed in the data. In addition, the Gourinchas-Parker (2002) estimates also imply an equilibrium ratio of  $K/Y$  that is much larger than the target ratio when  $\beta$  is calibrated under the RIS model.

Table 3: Comparison of the Magnitude of the Consumption Hump—Fixed Discount Factor

Income process	$K/Y$	$S$
RIS	3.0	1.021
AIS	2.8	1.000
GP	3.5	1.280
CS updated	2.9	1.128

<sup>17</sup>Because  $S$  is defined as the ratio between  $C_{max}$  and  $C_0$ , not  $C_{min}$ ,  $S_{AIS}$  is only slightly higher than one despite the appreciable local maximum shown in the graph.

## 8 Conclusion

We propose a novel semiparametric approach of measuring household income uncertainty and study the dynamics of income uncertainty over the life cycle. Our estimates of income uncertainty are typically smaller than previous studies have documented and imply less persistence in income shocks. Moreover, we find that income uncertainties evolve noticeably over the life cycle. Young and old consumers on average have more risky future income relative to middle-age consumers. Furthermore, our estimates of income uncertainty generally call for fewer precautionary savings than previous research has found, in particular at younger ages. In a companion paper, Feigenbaum and Li (2010b) study how income uncertainty has changed over time. In contrast to what researchers have found in the aggregate data, the semiparametric estimates of household income uncertainty rose noticeably in the last several decades, and the increase was widespread across demographic and income groups.

A wide variety of theoretical questions related to household financial decisions over the life cycle can be revisited using the metric of income uncertainty we introduce. For example, in this paper we show that our estimate of income uncertainty implies that precautionary saving can create a significantly hump-shaped life cycle profile of mean consumption, but one that is smaller than observed in the data. In contrast, the precautionary saving implied by previous estimates of income uncertainties often calls for a consumption hump that is significantly greater than observed in the data.

## References

- [1] Aiyagari, S. Rao (1994), “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics*, vol. 109, pp. 659-84.
- [2] Baker, Michael, and Gary Solon (2003), “Earnings Dynamics and Inequality among Canadian Men, 1976-1992: Evidence from Longitudinal Income Tax Records,” *Journal of Labor Economics*, vol. 21, pp. 289-321.
- [3] Becker, Gary S., and Gilbert R. Ghez (1975), *The Allocation of Time and Goods over the Life Cycle*, New York: NBER and Columbia Press.
- [4] Browning, Martin, and Thomas F. Crossley (2001), “The Life-Cycle Model of Consumption and Saving,” *Journal of Economic Perspectives*, vol. 15, pp. 3-22.
- [5] Bullard, James, and James Feigenbaum (2007), “A Leisurely Reading of the Life Cycle Consumption Data,” *Journal of Monetary Economics*, vol. 54, pp. 2305-20.
- [6] Campbell, John Y., and Joao F. Cocco (2003). “Household Risk Management and Optimal Mortgage Choice,” *Quarterly Journal of Economics*, vol. 118, pp. 1449-94.
- [7] Carroll, Christopher D. (1994), “How Does Future Income Affect Current Consumption?” *Quarterly Journal of Economics*, vol. 104, pp. 111-47.
- [8] Carroll, Christopher D. (1997), “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics*, vol. 112, pp. 1-55.
- [9] Carroll, Christopher D., and Andrew A. Samwick (1997), “The Nature of Precautionary Wealth,” *Journal of Monetary Economics*, vol. 40, pp. 41-71.
- [10] Carroll, Christopher D., and Lawrence H. Summers (1991), “Consumption Growth Parallels Income Growth: Some New Evidence,” in *National Saving and Economic Perfor-*

*mance*, eds. B. Douglas Bernheim and J. Shoven, Chicago: University of Chicago Press, pp. 305-48.

- [11] Cunha, Flavio, James Heckman, and Salvador Navarro (2005), “Separating Uncertainty from Heterogeneity in Life Cycle Earnings,” *Oxford Economic Papers*, vol. 57, pp. 191-261.
- [12] Deaton, Angus, and Christina Paxson (1994), “Intertemporal Choice and Inequality,” *Journal of Political Economy*, vol. 102, pp. 437-67.
- [13] 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.
- [14] Feigenbaum, James (2008a) “Information Shocks and Precautionary Saving,” *Journal of Economic Dynamics and Control*, vol. 32, pp. 3917-38.
- [15] Feigenbaum, James, (2008b) “Precautionary Saving Unfettered,” working paper.
- [16] Feigenbaum, James, (2008c), “Can Mortality Risk Explain the Consumption Hump?” *Journal of Macroeconomics*, vol. 30, pp. 844-72.
- [17] Feigenbaum, James, and Geng Li (2010a), “Lifecycle Dynamics of Income Uncertainty and Consumption,” working paper.
- [18] Feigenbaum, James, and Geng Li (2010b), “Household Income Uncertainty over Thirty Years: We Do Not See a Great Moderation Here,” working paper.
- [19] Fernández-Villaverde, Jesús, and Dirk Krueger (2010), “Consumption and Saving over the Life Cycle: How Important are Consumer Durables?” *Macroeconomic Dynamics*, forthcoming.

- [20] Fogli, Alessandra, and Fabrizio Perri (2006), “The Great Moderation and the U.S. External Imbalance,” *Monetary and Economic Studies*, vol. 26, pp. 209-25.
- [21] Gordon, Roger H. (1984), *Differences in Earnings and Ability*, New York: Garland.
- [22] Gottschalk, Peter, and Robert Moffitt (1994), “The Growth of Earnings Instability in the U.S. Labor Market,” *Brookings Papers on Economic Activity*, vol. 1994, pp. 217-72.
- [23] Gourinchas, Pierre-Olivier, and Jonathan A. Parker (2002), “Consumption over the Life Cycle,” *Econometrica*, vol. 70, pp. 47-89.
- [24] Guvenen, Fatih (2007), “Learning Your Earning: Are Labor Income Shocks Really Very Persistent?” *American Economic Review*, vol. 97, pp. 687-712.
- [25] Hansen, Gary D., and Selahattin İmrohoroğlu (2008), “Consumption over the Life Cycle: The Role of Annuities,” *Review of Economic Dynamics*, vol. 11, pp. 566-83.
- [26] 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*, vol. 64, pp. 188-94.
- [27] 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*, vol. 40, pp. 59-125.
- [28] Huggett, Mark (1996), “Wealth Distribution in Life-Cycle Economies,” *Journal of Monetary Economics*, vol. 38, pp. 469-94.
- [29] Jaimovich, Nir, and Henry E. Siu (2009), “The Young, the Old, and the Restless: Demographics and Business Cycle Volatility,” *American Economic Review*, vol. 99, pp. 804-26.

- [30] Kim, Chang-Jin, and Charles Nelson (1999), “Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle,” *Review of Economics and Statistics*, vol. 81, pp. 608-16.
- [31] Lillard, Lee, and Yoram Weiss (1979), “Components of Variation in Panel Earnings Data: American Scientists 1960-70,” *Econometrica*, vol. 47, pp. 437-54.
- [32] MaCurdy, Thomas (1982), “The Use of Time-Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis,” *Journal of Econometrics*, vol. 18, pp. 83-114.
- [33] McConnell, Margaret, and Gabriel, Perez-Quiros (2000), “Output Fluctuations in the U.S.: What Has Changed Since the Early 1980s?” *American Economic Review*, vol. 90, pp. 1464-76.
- [34] Nagatani, Keizo (1972), “Life Cycle Saving: Theory and Fact,” *American Economic Review*, vol. 62, pp. 344-53.
- [35] Rios-Rull, Jose-Victor (1996), “Life-Cycle Economies and Aggregate Fluctuations,” *Review of Economic Studies*, vol. 63, pp. 465-89.
- [36] Skinner, Jonathan (1988), “Risky Income, Life Cycle Consumption, and Precautionary Savings,” *Journal of Monetary Economics*, vol. 22, pp. 237-55.
- [37] Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron (2004), “Consumption and Risk Sharing Over the Life Cycle,” *Journal of Monetary Economics*, vol. 51, pp. 609-33.
- [38] Thurow, Lester C. (1969), “The Optimum Lifetime Distribution of Consumption Expenditures,” *American Economic Review*, vol. 59, pp. 324-30.
- [39] Viceira, Luis M. (2001), “Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income,” *Journal of Finance*, vol. 56, pp. 433-70.

Figure 1: Income uncertainty over the life cycle—family noncapital income

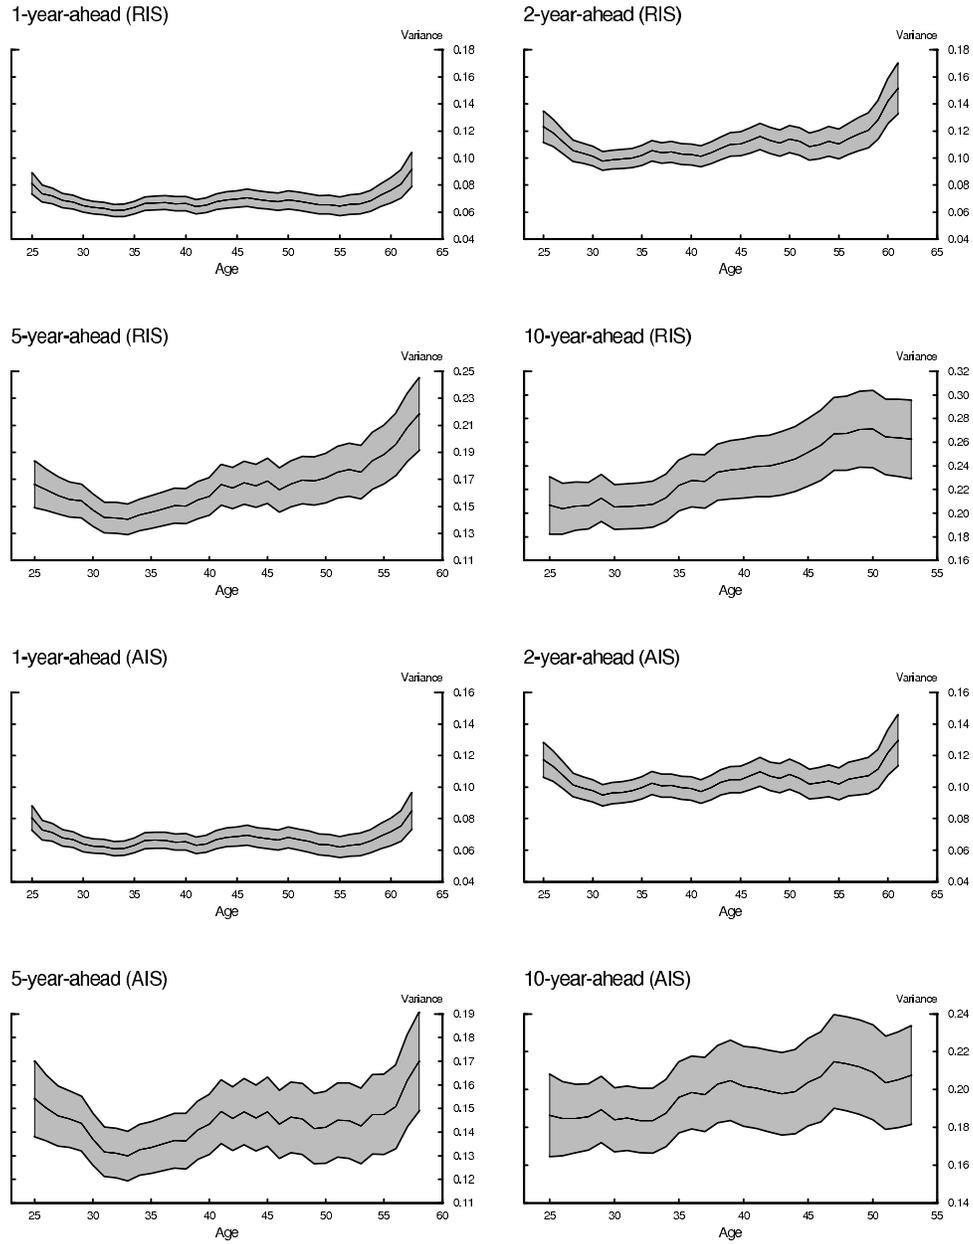


Figure 2: Correlation structure of forecast errors—family noncapital income

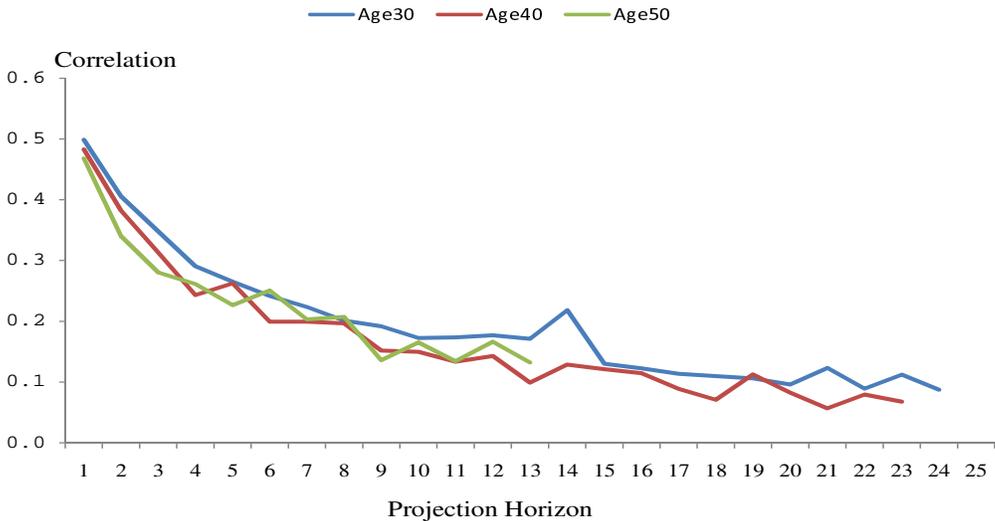


Figure 3: Income uncertainty over the life cycle—head labor income

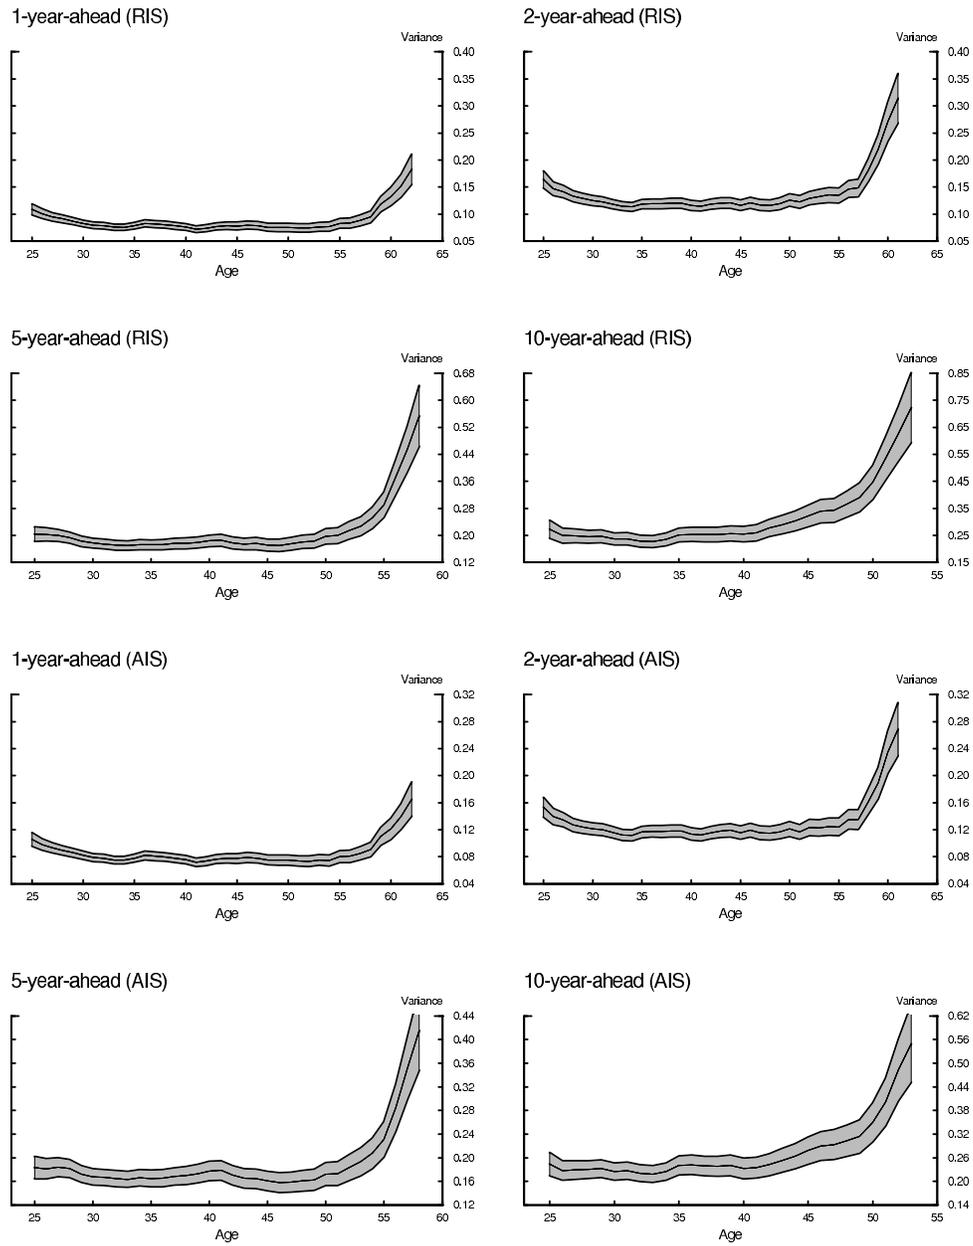


Figure 4: Correlation structure of forecast errors—head labor income

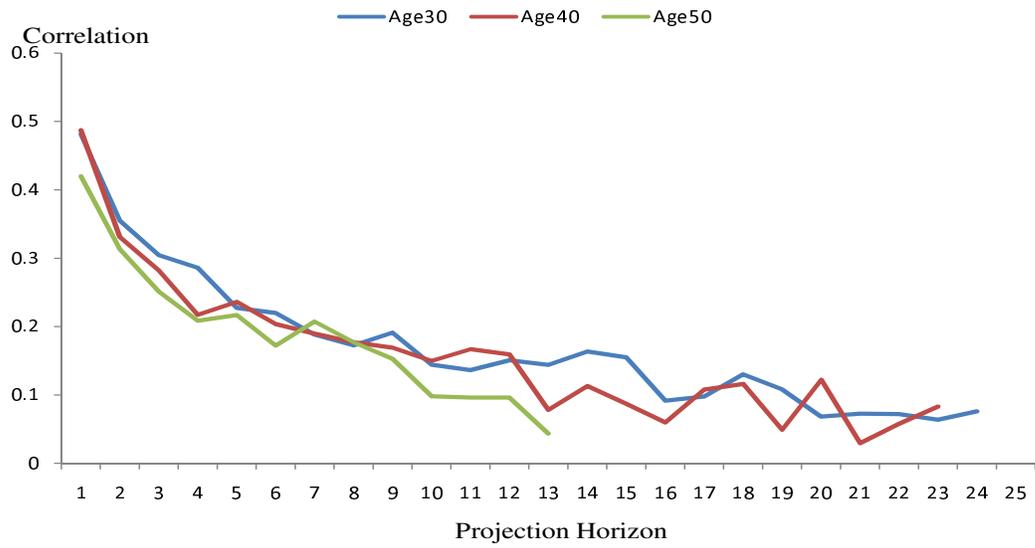


Figure 5: Comparison of semiparametric estimates of income uncertainty with preexisting estimates (pooled across all ages)

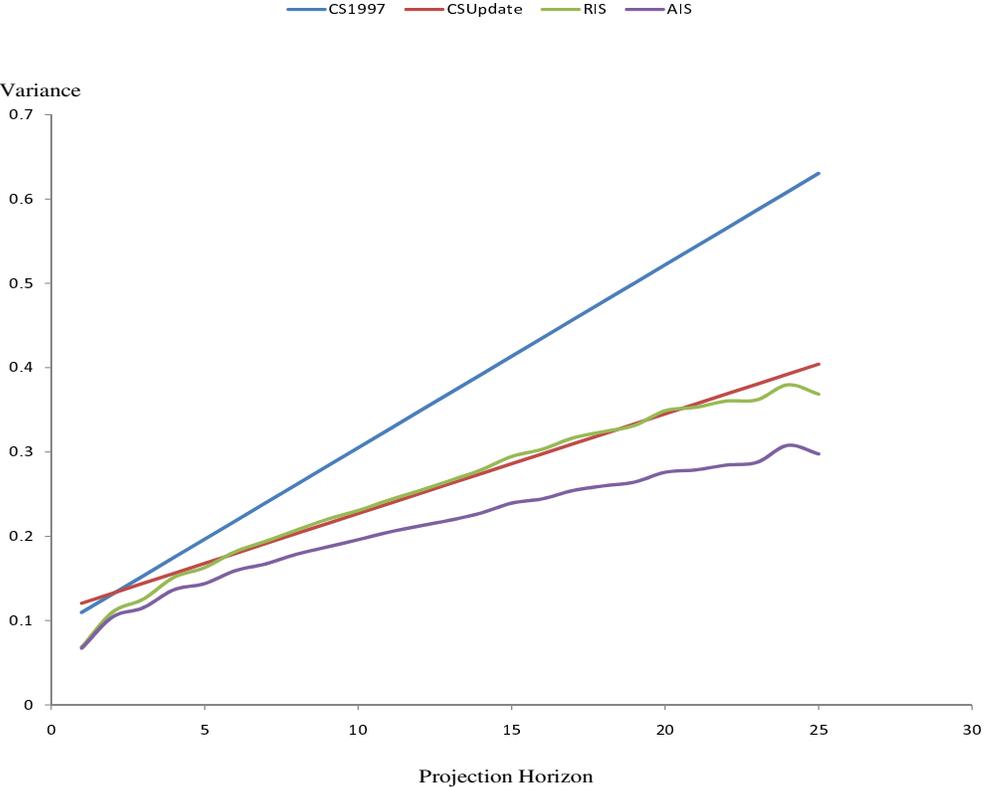


Figure 6: Panel length and the estimated variances of permanent and transitory shocks

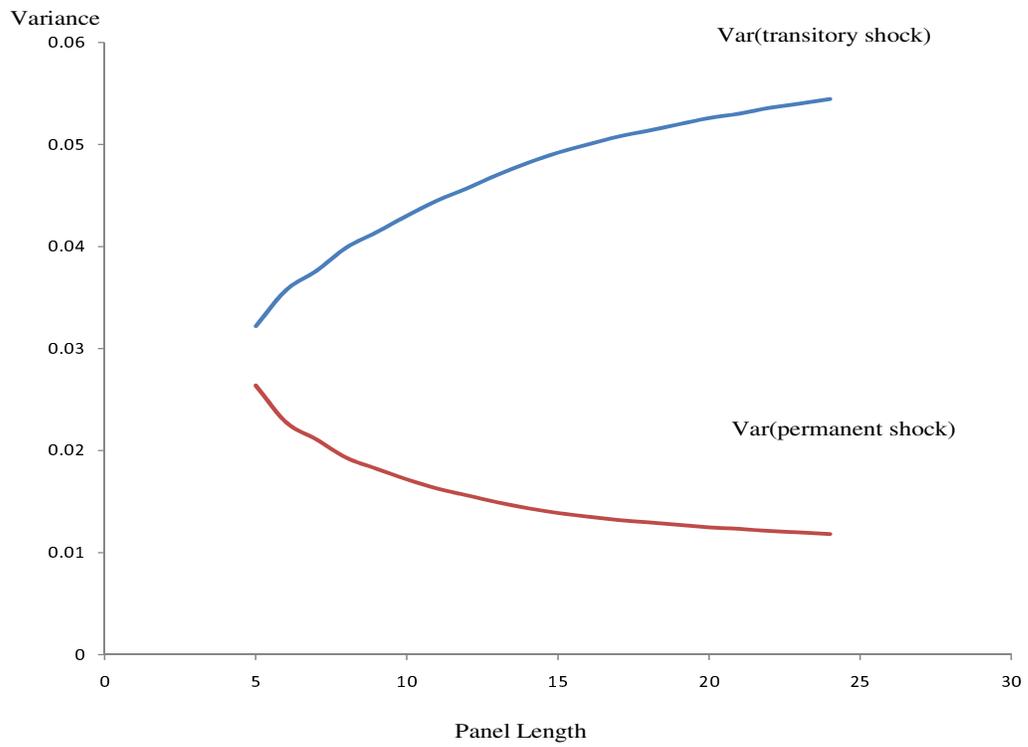


Figure 7: Compare the life cycle consumption profiles under different income processes— $K/Y = 3$

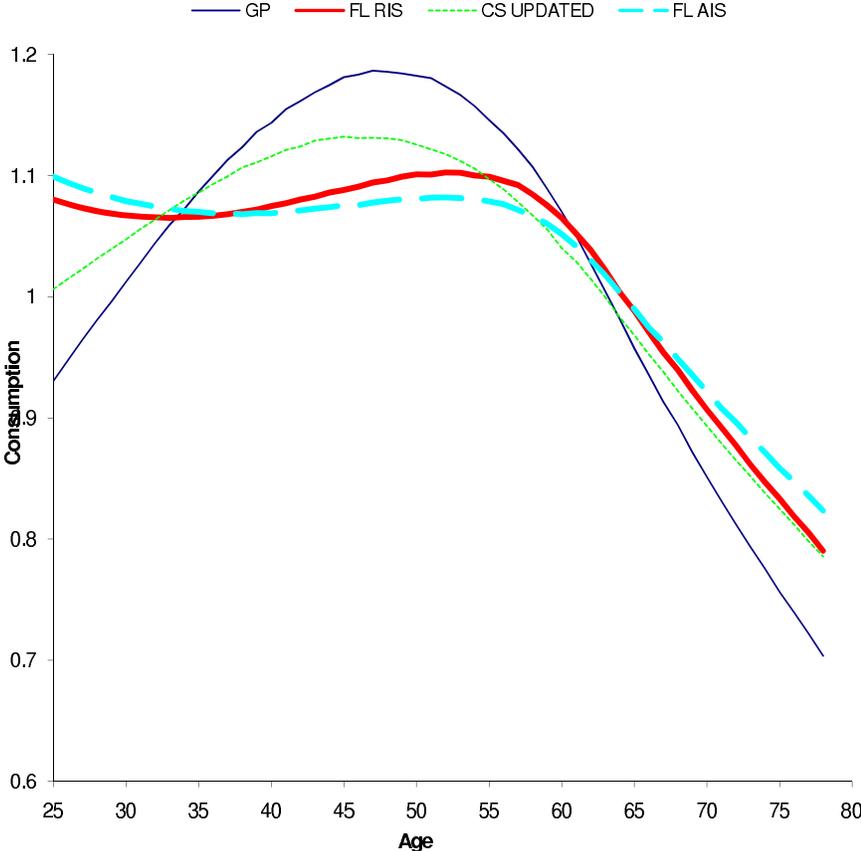


Figure 8: Compare the life cycle consumption profiles under different income processes— $K/Y = 3.5$

