Journal of Economic Dynamics & Control 37 (2013) 535–542

Contents lists available at SciVerse ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Measuring high-frequency income risk from low-frequency data

Paul Klein [a], Irina A. Telyukova [b,*]

[a] Simon Fraser University, Canada
[b] Department of Economics, University of California, San Diego, 9500 Gilman Drive 0508, La Jolla, CA 92093-0508, USA

ARTICLE INFO

Article history:
Received 23 November 2011
Received in revised form
6 July 2012
Accepted 18 September 2012
Available online 5 October 2012

JEL classification:
E21
E24

Keywords:
Idiosyncratic income uncertainty
Frequency
Estimation

ABSTRACT

We estimate a *monthly* income process using *annual* longitudinal household-level income data, in order to understand the nature of income risk faced by households at high frequency, and to provide an input for models that wish to study household decision-making at higher frequency than available data. At both frequencies, idiosyncratic earnings shocks have a highly persistent component. At monthly frequency, transitory shocks account for most of the earnings variance; at annual frequency, the persistent component is dominant. We apply our estimates in the context of a standard incomplete-market model, and show that decision-making frequency per se makes a small difference.

1. Introduction

In the literature on household consumption-saving decisions under exogenously incomplete markets, it is typically assumed that households face some form of idiosyncratic risk. For any applications involving working-age households, a widely studied form of such risk is income uncertainty.

In order to study implications of idiosyncratic income uncertainty, researchers typically assume some process for income that may involve permanent, persistent and/or transitory components. In order to calibrate the models, researchers need to measure these components in the data. There is a large and active literature on estimating income uncertainty in the data; a few recent examples are Guvenen (2007), Guvenen and Smith (2010) and Heathcote et al. (2010). For the estimation of persistent processes, the econometrician needs longitudinal household-level data on income, which leads researchers to use, in most cases, survey data such as the Panel Study on Income Dynamics (PSID) and the Consumer Expenditure Survey (CEX). Alternatively, as in Daly et al. (in press), administrative (register) data are used.

The limitation of all these datasets is that they are annual at best, and sometimes biennial, like the PSID in recent years. This means that the literature typically relies on these once-a-year observations of income to estimate income risk; models then typically use the same period length as in the data.[1] This of course restricts model households to make decisions at an annual frequency. For some decisions, this is an acceptable approximation, but one can think of many other aspects

* Corresponding author. Tel.: +1 858 822 2097.
E-mail address: itelyukova@ucsd.edu (I.A. Telyukova).

[1] An exception is Erosa et al. (2011), who use an indirect inference approach to calibrate the four-monthly wage process on the basis of annual data.

0165-1889/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jedc.2012.10.002

of economic behavior for which we may prefer to model decision-making at quarterly, and even monthly, frequency. Understanding portfolio allocation, especially with respect to liquid assets, studying decisions to revolve or repay secured or unsecured debt, and characterizing demand for money are some issues for which a high-frequency model would be preferred, or even necessary. An example is Telyukova (Forthcoming), who addresses the question of co-existence in household portfolios of expensive credit card debt and low-return checking and savings accounts. In such a model, annual decision-making would be uninformative, as it would obscure the decision to revolve credit card debt each month, or to repay a portion of it using currently available liquid assets.

In this paper, we provide parameter estimates for a yearly and a monthly earnings process designed to match key features of annual PSID data. Based on an extension of Gervais and Klein (2010), we posit a monthly process underlying the observed annual income process; in both cases, we assume that income has a permanent component, a persistent stochastic component, and a transitory component. We estimate the monthly process based on annual data, using a simulated method of moments and moments of the autocovariance function. Our main finding is that a transitory component accounts for about a quarter of overall *annual* variance of earnings, and that this is true whether the model we estimate is monthly or annual. The remainder is mostly accounted for by a component that is highly persistent but far from a random walk. Instead, looking at *monthly* earnings variance, much of that is accounted for by transitory earnings shocks that last less than a year. Our approach to estimating the annual model can be thought of as a contribution to the ongoing debate about how best to estimate wage and earnings processes and the implications of each method for the relative importance of persistent and transitory shocks (see Domeij and Flodén, 2010 and Daly et al., in press for recent contributions), but this is incidental; our main innovation concerns the estimation of the monthly model using annual data.

In addition to providing estimates of monthly income risk, which we believe to be of interest in themselves, we also investigate whether frequency of decision making matters for risk-sharing implications of a standard consumption-saving model. We do this by computing an infinite-horizon version of the Huggett (1993) model at annual and monthly frequency. The main finding here is that the degree of risk sharing is affected by increasing the frequency of decision making, but quantitatively the impact of the frequency is moderate and depends on the size of the borrowing limit. The difference that we do observe – broadly speaking, consumption changes are more responsive to income changes in the monthly model – is due to the higher estimated importance of the permanent component in the annual model, relative to the monthly one. The permanent component is of course impossible to insure against, so there is perhaps a deeper sense in which the monthly model exhibits more risk sharing than the annual model. But a relatively important permanent component does not affect the regression coefficient of household-level consumption changes on income changes in any direct way. It does, however, leave less of the variance to be accounted for by the persistent component. That is, we find that in the annual model, the persistent shocks to income are relatively less important than in the monthly model, and since transitory shocks are easier to self-insure against than persistent shocks, households in the annual model are able to self-insure more completely, at least when the borrowing limit is not too tight.

In any case, we conclude that frequency of decision making per se is not sufficiently important for the degree of risk sharing in the Huggett (1993) framework. Instead, the use of a higher-than-annual frequency model should be driven primarily by specific questions that require the modeling of frequent decision-making, where annual frequency would be insufficient for understanding the issues of interest.

In addition to the literature that estimates income uncertainty in the data, our work is related to the literature on risk sharing in incomplete-market models. Some examples are Krueger and Perri (2004) and Kaplan and Violante (2010), who do this in calibrated models of household decision-making, and Blundell et al. (2008), who use econometric techniques to measure the degree of consumption risk sharing.

The rest of the paper is organized as follows. Section 2 describes the data we use in our estimation, the estimation procedure and results. We then apply these annual and monthly estimates in the context of the Huggett-style model, which we describe, calibrate and compute in Section 3. Section 4 concludes.

2. Estimation of the earnings process

2.1. Data

In order to estimate the earnings process, we rely on the Panel Study of Income Dynamics (PSID). We employ the data from 1968 to 1997, which is the period during which data are available annually; after 1997, PSID becomes biennial. Our sample consists of individuals between the ages of 21 and 62. We consider different subsamples: all men and women, only men, and only male heads of households. For our purposes, all of these samples yield similar results. We also drop those with annual earnings below $2000 in 1968 dollars.

2.2. Procedure

The main challenge in estimating the earnings process is that we have annual data but want to estimate a monthly process. But before we tackle that issue, we have to choose a specification for the time series process that can be made to fit the available facts, regardless of the length of the time period.