Measuring High-Frequency Income Risk from Low-Frequency Data*

Paul Klein
Simon Fraser University

Irina A. Telyukova
University of California, San Diego

First draft: August 18, 2009
This version: September 27, 2012

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 ce makes a small difference.

JEL classification: E21, E24
Keywords: idiosyncratic income uncertainty, frequency, estimation

*We thank the editor, Chris Otrok, as well as two referees and participants of the 2011 CEF San Francisco conference for their helpful comments and suggestions. Corresponding author: Irina A. Telyukova, UCSD Department of Economics, 9500 Gilman Drive 0508, La Jolla, CA 92093-0508. (858) 822-2097. itelyukova@ucsd.edu

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. (2011), 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

[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.

of many other aspects 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 (2011), 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. (2011) 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 ce 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.

Our choice of specification is designed to capture the key statistical properties in the micro data on earnings. Following many other authors, we extract the idiosyncratic component of log earnings by regressing log-earnings on a cubic in age, dummies for education, gender, race, marital status and birth cohort, and retaining the residuals. The question is how to model this residual.

Our specification is as follows. The residual $y_{i,t}$ of monthly or annual log earnings is assumed to have three distinct components according to

$$y_{i,t} = \alpha_i + z_{i,t} + x_{i,t} \tag{1}$$

where we call α_i the *permanent* component (since it doesn't change as a household ages) with variance σ_α^2, where the *persistent* component $z_{i,t}$ satisfies

$$z_{i,t-1} = \rho z_{i,t-1} + \varepsilon_{i,t} \tag{2}$$

and where $\varepsilon_{i,t}$ and $x_{i,t}$ are i.i.d., with variances σ_ε^2 and σ_x^2 respectively. In contrast to Guvenen (2007), we do not allow for heterogeneous predictable age-earnings profiles.

Since in our model agents have infinite lives, it makes sense to ignore the age dimension of observations, and we do not allow the variances of the shocks $\varepsilon_{i,t}$ and $x_{i,t}$ to depend on age. For reasons of parsimony, we also do not allow them to depend on cohort or on calendar time. Thus the only parameters that need to be estimated are σ_ε^2, ρ, σ_x^2 and σ_α^2. This is done by GMM where the moments are the autocovariances $\Gamma_k = \mathsf{E}[y_{i,t}y_{i,t+k}]$, where the kth covariance is computed as the average over all possible products $y_{i,t}y_{i,t+k}$ for which data are available and regardless of age.

The choice of specification described in Equation (1) is based on some striking features of the autocovariance function of the residuals, displayed in Figure 1. What we see there is that Γ_k falls steeply as k goes from 0 to 1 and then very gradually, with a near-constant rate of decay, as k increases further. This is evidence in favor of the view that purely transitory shocks (or possibly measurement error) accounts for a large fraction of the total variance of earnings. It also points in the direction of idiosyncratic earnings having a component that is persistent but not quite a random walk.

These conclusions—that the persistent component does not have a unit root and the transitory component is important—contrast somewhat with the influential work of Storesletten et al. (2004), who argue that the persistent component is a random walk and that it accounts

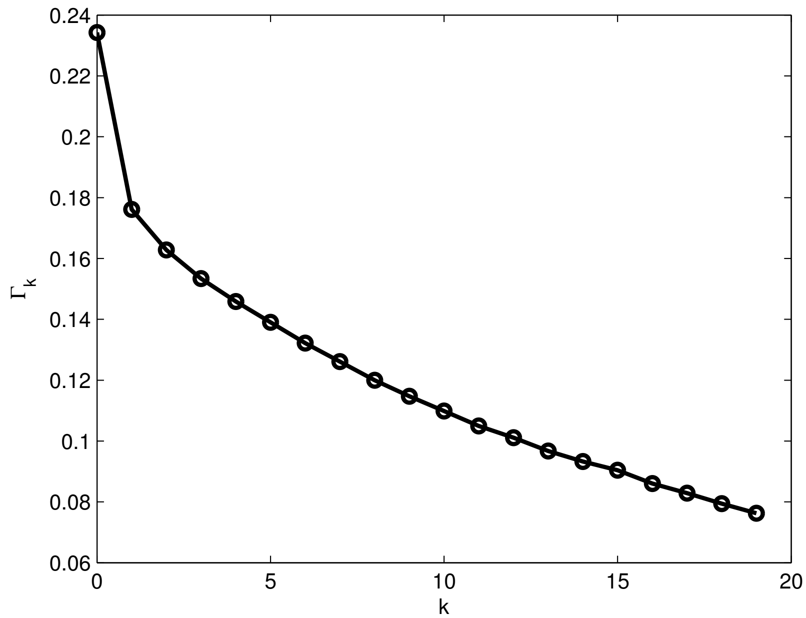

Figure 1: The autocovariance function of log earnings residuals.

for most of the variance. This approach is followed in Blundell and Preston (1998), who model earnings as the sum of a random walk and a white noise process.[2] As Figure 1 shows, the autocovariance function of earnings is not consistent with such a representation. When $k \geq 1$, Γ_k tends to decline more or less geometrically and at a non-negligible rate. This is why our specification allows for a persistent component that is not a random walk.

Moreover, it is not obvious (though it is possible) that Γ_k tends to zero as it would if the earnings residual consisted of just a transitory and a persistent (but unit root) component. In fact Γ_{k+1}/Γ_k starts out at about 0.92 at $k = 1$ and then increases somewhat towards unity as k increases so that Γ_k appears to tend a strictly positive limit. More data would be required

[2] As discussed in Domeij and Flodén (2010) and Daly et al. (2011), given the specification in Blundell and Preston (1998), the estimation results are dramatically different depending on whether the estimation is done in levels or differences. Here we bypass that issue by using a different specification.

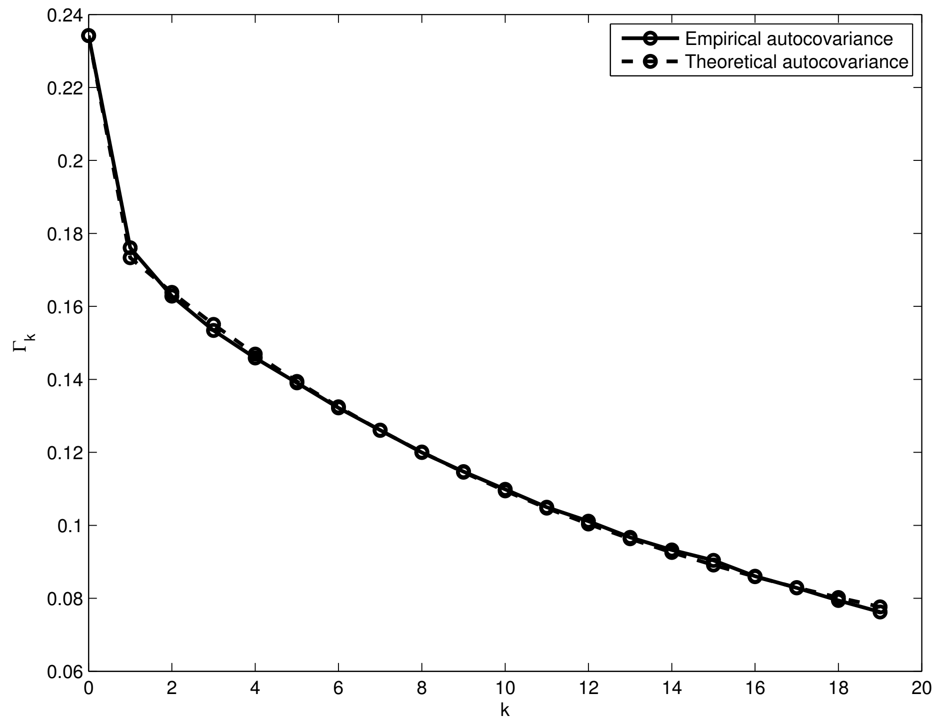

Figure 2: Yearly model. Empirical and estimated theoretical autocovariances.

for any firm conclusion about this, but in the absence of any strong evidence that Γ_k tends to zero as $k \to \infty$, it makes sense to allow for a truly permanent individual-specific component as well.

Given the specification (1) it is straightforward to derive the theoretical autocovariance function. It is

$$\Gamma_k = \sigma_\alpha^2 + I_{\{k=0\}} \cdot \sigma_x^2 + \frac{\rho^k}{1-\rho^2}\sigma_\varepsilon^2$$

where $I_{\{k=0\}}$ is an indicator function that equals one if $k = 0$ and zero otherwise. The GMM estimation chooses ρ, σ_α^2, σ_ε^2 and σ_x^2 so as to minimize the (unweighted) distance between the first 20 theoretical autocovariances and their empirical counterparts; the fit of the model is shown in Figure 2. The estimation results are summarized in the first row of Table 1 and discussed in Section 2.4 below. In particular, notice that $\rho \approx 0.93$, a number quite far from unity.

2.3 A Monthly Earnings Process

We now discuss how to estimate a monthly model. The statistical specification is still given by Equation (1). If we had monthly data, we could proceed exactly as above. But we do not. Nevertheless, a monthly model has implications for the annual autocovariance function and we can use these implications to estimate the parameters of the monthly model. This approach is similar to but extends that of Gervais and Klein (2010).

Computing the theoretical moments presents something of a challenge. They can be computed by simulation, following Lee and Ingram (1991). However, this is very time-consuming in terms of CPU time. On the other hand, exact analytical formulas are not available because of Jensen's inequality. To see this, denote (residual) log earnings in month s by $\upsilon_{i,s}$. We then define annual earnings via

$$y_{i,a} = \ln\left(\frac{1}{12}\sum_{s=0}^{11} \exp\left\{\upsilon_{i,12a+s}\right\}\right) \tag{3}$$

Thus our statistical model implies that

$$y_{i,a} = \ln\left(\frac{1}{12}\sum_{s=0}^{11} \exp\left\{\alpha_i + x_{i,s} + \sum_{k=0}^{s} \rho^{s-k}\varepsilon_{i,k}\right\}\right). \tag{4}$$

If it weren't for the ln followed by exp, analytical formulas would be available. If we remove these functions, we obtain approximate results, and the errors are fairly small for the autocovariances, with the exception of the 0th autocovariance (the variance). This exception, however, is sufficiently important, so that we have chosen to compute moments by simulation.

Specifically, we draw $N = 1000$ individual histories of length $T = 1000$ years (12000 months). We draw the line there in order not to slow down the estimation too much but at the same time to maintain acceptable precision. If $N = 6000$ and $T = 6000$, the largest change in an autocovariance is about 0.6 percent compared to the $N = 1000$, $T = 1000$ case. If N is then maintained at 6000 but T is increased to 10000 years, the autocovariances change by a further 0.2 percent for a total maximum change of 0.8 percent compared with the $N = 1000$, $T = 1000$ year case. Essentially this means that the autocovariances in our simulation are accurately computed to two decimal places. The fit of the monthly model, in terms of the annual autocovariance function, is shown in Figure 3.

2.4 Estimation results

In Table 1 we report the estimation results for the stationary versions of the annual and the monthly models, respectively. The key thing to notice is that the relative importance of the three components of earnings is quite different in the two cases. In the monthly model, about 75 percent of the variance is accounted for by transitory shocks and about 21 percent by persistent shocks, leaving about 4 percent for the permanent shock. In the annual model, about 22 of the total variance is accounted for by the transitory shocks and 59 percent by the

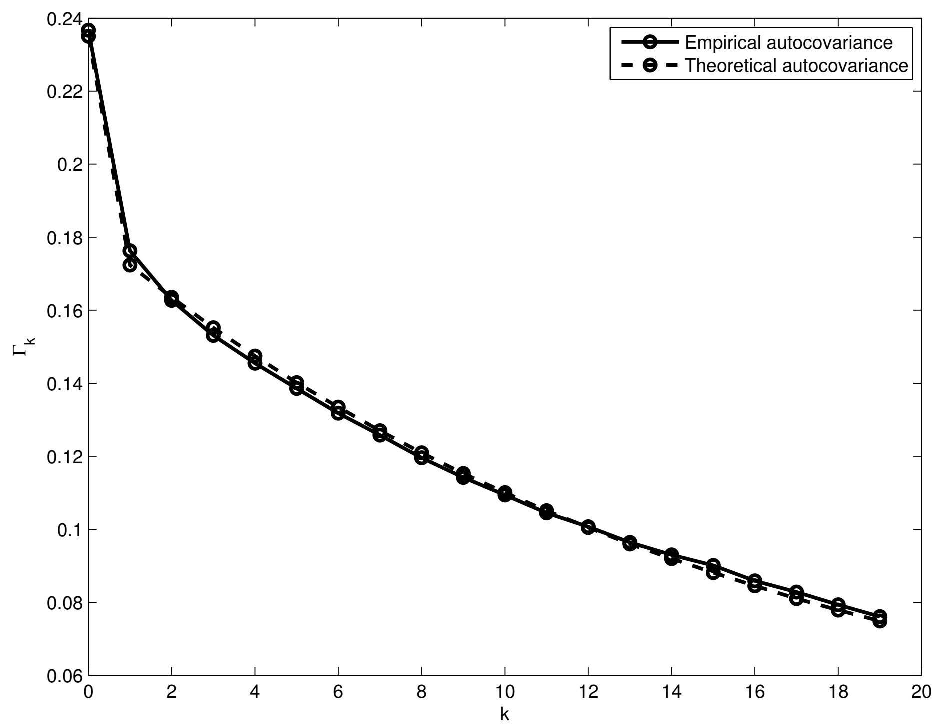

Figure 3: Monthly model. Empirical and estimated theoretical yearly autocovariances.

persistent component, leaving about 20 percent for the permanent component. What explains this difference? Our findings are consistent with significant monthly noise that is to some extent washed out at the annual frequency. In any case, our first result is that it makes a difference whether we use an annual or a monthly model to measure the relative importance of transitory and persistent shocks to earnings.

We have just said that, in the monthly model, 75 percent of the variance of earnings is accounted for by transitory shock. That is the *monthly* variance. What about the yearly variance? Clearly much more of the variance of the transitory component washes out than that of the persistent component when we add up monthly earnings to annual frequency. As it turns out, in the monthly model, about 24 percent of the annual variance of earnings is accounted for by the transitory component, and about 64 percent by the persistent component, leaving almost 13 percent for the permanent component. Thus in the monthly model, the transitory component has essentially the same relative importance in the monthly as in the annual model. The relative importance of the permanent component is less robust; this suggests that the relative importance of the permanent and the persistent component are not so easily identifiable by the data.

2.5 Robustness and identification

What aspects of our specification are forced upon us by the data and what aspects have to be assumed more or less arbitrarily?

To shed light on this issue, consider a more general specification that will nest our baseline.

As before, log monthly income is the sum of three components: a permanent component, a persistent component and a transitory component. But now, assume that the persistent component remains constant with some probability $1-p$ and that the transitory component equals zero with some probability $1-q$, these events being independent. Thus, the probability of income remaining constant from one month to the next is $(1-p)(1-q)$, while persistent and transitory shocks arrive with probability p and q respectively.

It turns out that with this specification, p and q cannot be identified. Indeed, the model can be made to fit the empirical autocovariance function extremely well regardless of the precise choice of p and q, provided that neither is zero. This is perhaps not surprising; with annual data, there is a limit to how much information can be extracted about events at a monthly frequency. However, certain very important features of the data generating process are strikingly robust to different assumptions about p and q. To make this point, consider four possibilities, (1) $p = 1/2$ and $q = 1/2$, (2) $p = 1/2$ and $q = 1/4$, (3) $p = 1/4$ and $q = 1/2$ and, finally, (4) $p = 1$ and $q = 1$. Notice that case (4) corresponds to our baseline specification of the monthly earnings process.

In the baseline case (4), we concluded that transitory shocks account for 24 percent of the annual variance of log earnings. What is that number in cases (1)-(3)? In case (1) the fraction is 22 percent, in case (2) it is 23 percent, in case (3) it is also 23 percent. Thus this fraction is rather robust to the exact specification, and we conclude that any model capable of replicating the empirical annual autocovariance function must assign about a quarter of the total variance to shocks that last no longer than a year.

3 Application: Implications for Risk-Sharing

3.1 General-Equilibrium Incomplete-Market Model

In our computational experiments we use the canonical infinite-horizon model of Huggett (1993). Each period, households choose the level of consumption c_t and saving a_{t+1} in a real risk-free bond, given their current earnings and asset states. The earnings state is stochastic and idiosyncratic, and consists of three components as described in the estimation section: the permanent component α, the persistent component z_t, and the transitory component x_t. The permanent component is known to the household and does not change period to period. The persistent income state z_t is discrete and evolves according to a Markov process with associated transition function $\Gamma_z(z'|z)$. The transitory shock x is likewise discrete and i.i.d. Denote by $\mathbb{P}_x$ the probability of realization of a given shock x.

The household problem in recursive formulation is

$$
\begin{aligned}
V(\alpha, z_t, x_t; a_t) &= \max_{c_{t+1}, a_{t+1}} u(c_t) + \beta \sum_{z'} \sum_{x'} \Gamma_z(z'|z_t) \mathbb{P}_{x'} V(\alpha, z_{t+1}, x_{t+1}; a_{t+1}) && (5) \\
\text{s.t. } c_t + a_{t+1} &= \exp(\alpha + z_t + x_t) + a_t(1 + r_t) && (6) \\
c_t &\geq 0 && (7) \\
a_{t+1} &\geq \underline{a} && (8)
\end{aligned}
$$

Here, the household expectation of future shock realizations is written in terms of discretized shocks z_t and x_t. The risk-free bond pays real return r_t. The transition matrices give probabilities of future shock realizations conditional on current realizations. In the budget constraint, the current earnings realization is given by the product of the three components of (log) earnings. Households can borrow subject to the borrowing constraint (8) where $\underline{a} < 0$.

We will consider the stationary equilibrium of this economy, where the distribution of agents along the earnings and assets dimensions is constant over time. Define the transition function $\Pi(s, B)$ from state $s_t = (\alpha, z_t, x_t; a_t) \in S$ to the subset of state space B in the standard way.[3] Denote by $\Psi(s_t)$ the distribution of agents across the state space. Denote by $g_a(s_t)$ and $g_c(s_t)$ the decision rules with respect to saving and consumption in some period t.

The stationary equilibrium for this economy is the set of functions $(g_c(s), g_a(s), \Psi(s), V(s), r(s))$ such that: (a) $g_c(s)$ and $g_a(s)$ are optimal decision rules for the household given the price $r(s)$; (b) consumption and asset markets clear, so that $\int g_c(s) d\Psi = \int \alpha z x d\Psi$ and $\int g_a(s) d\Psi = 0$; (3) the distribution Ψ is a stationary probability measure, so that $\Psi(B) = \int \Pi(s, B) d\Psi \; \forall B$.

See, e.g., Huggett (1993) for a detailed discussion of this equilibrium concept, which is standard in the literature.

3.2 Calibration and Computation

We calibrate those parameters of the model that pertain to our earnings process from the estimation that we described above. The estimation yields parameters of the AR(1) processes for the persistent and transitory components of earnings which we then discretize using the method of Rouwenhorst (1995).

Following Huggett (1993), we test a variety of borrowing limits, from 0.5 to 4 times average

[3]Formally, B is a subset of the Borel σ-algebra on the state space S.

annual earnings, computing the equilibium interest rate in each case. We choose the standard CRRA functional form for the utility function: $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$. The remaining parameters are calibrated within the range in Huggett (1993). For the annual calibration, the discount factor β^a is set at 0.97, which at monthly frequency yields $\beta^m = 0.9975$. The coefficient of risk aversion is $\sigma = 2$.

The algorithm to compute the model is standard: given a guess of the interest rate r, we solve the household problem, then compute the stationary distribution of agents across the state space, and then check the asset market clearing condition. We iterate on the price until the market clears given household optimization. To solve for the decision rules of the household, we use the endogenous grid method of Carroll (2006). The search for the market-clearing price is done using a bisection method.

3.3 Measures of Risk-Sharing

We devise several measures of risk sharing, all of which are inspired by considering two polar opposite cases: perfect risk-sharing and autarky. The measures are:

1. The regression coefficient β of log-consumption changes on log-earnings changes:

   $\Delta \log(c_{it}) = \beta \Delta \log(e_{it}) + \epsilon_{it}$, where $e_{it} = \alpha z_{it} x_{it}$. (See Krueger and Perri (2004).)

2. Variance of log-consumption relative to variance of log-earnings.

3. Autocorrelation of log-consumption.

The first two measures share the property that they are one in autarky and zero under perfect risk sharing. All three form a convenient metric through which model predictions are summarized and compared across the models of differing frequencies. For both models, the measures are computed for *annual* consumption and earnings, which in the monthly case requires adding up monthly data to annual frequency.

3.4 Results

In Table 2, we report the results from the computational experiment. First, notice that as the borrowing limit increases, the interest rate increases as well, since the bond becomes less valuable at the margin as a hedge against earnings risk when the borrowing constraint is relaxed. This results is consistent with the findings of Huggett (1993). Second, as the borrowing limit increases, the regression coefficient falls, suggesting increasing ability to self-insure against risk, as we would expect (see, e.g., Krueger and Perri (2004)); moreover, our regression coefficients are in line with previous literature. The other measures of risk-sharing confirm this tendency as well.

If we now compare the models of differing frequency, the main result is that frequency of decision-making on its own does not make a quantitatively dramatic difference. All of our measures of risk-sharing are close in the two models, as are the implied interest rates. Whether or not self-insurance is better at annual or monthly frequency depends on the borrowing limit; being able to borrow more improves the extent of self-insurance regardless of frequency. However, we also see that for higher borrowing limits, agents are better able to self-insure in the annual model than in the monthly one.

The reason is that as we discussed above, in the annual model, the persistent component accounts for 59% of the annual income variance, and the permanent component explains 20%; in the monthly model, the persistent component contributes 64% of the annual income variance, while the permanent component contributes 13%. Because of this, households in the monthly model are overall better able to share risk than their counterparts in the annual model, simply because the permanent and therefore uninsurable component is less important. But the *measured* degree of risk sharing is another matter. The responsiveness of consumption changes to income changes depends on how important the persistent component is relative to the transitory component, not on the importance of the permanent component, because the measures of risk-sharing are conditional on the realization of the permanent component of earnings. Since the persistent component is relatively more important in the monthly model than in the annual model, there is less measured risk sharing at monthly frequency of decision-making.

4 Conclusion

In this paper, we used standard annual longitudinal household-level income data from PSID to study the nature of income risk that households face at monthly frequency. In particular, we estimated an annual income process typical of the literature, but also posed and estimated an underlying monthly income process. We view the results of our estimation as interesting in their own right, as they shed light on the properties of risk that households face in the data between the times that we observe them in our surveys. We find that monthly earnings, though quite persistent, have a very significant transitory component.

We used our estimates to test whether frequency of decision-making is important in deriving implications of incomplete-market models for risk-sharing between households. In the context of the Huggett (1993) model, we find that frequency alone does not lead to significant differences in risk-sharing, although interesting variation results from differing implied contributions of permanent versus persistent shocks to earnings. We view this as encouraging for the literature that has studied risk-sharing predominantly in annual models.

We believe our results to be valuable to anyone who is interested in studying household decisions that are not usefully modeled at low frequency, so that an annual model would be too restrictive, and even uninformative, for the question of interest. For example, aspects of portfolio allocation may be best studied at the frequency at which households are paid their labor income, especially if the interest might be in money demand or liquid assets more generally. The same goes for applications pertaining to household decisions to borrow or revolve debt. Our results can be used directly as an input to such calibrated models, or our methodology can be applied more broadly to study high-frequency risk of other types or from other data sets.

References

Blundell, R., Pistaferri, L., Preston, I., 2008. Consumption inequality and partial insurance. American Economic Review 98 (5), 1887–1921.

Blundell, R., Preston, I., May 1998. Consumption inequality and income uncertainty. Quarterly Journal of Economics 113 (2), 603–640.

Carroll, C. D., 2006. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. Economics Letters 91 (3), 312–20.

Daly, M., Hryshko, D., Manovskii, I., 2011. Reconciling estimates of income processes in growth rates and levels, in Progress.

Domeij, D., Flodén, M., 2010. Inequality trends in sweden 1978-200. Review of Economic Dynamics 13 (1), 179–208.

Erosa, A., Fuster, L., Kambourov, G., 2011. Towards a micro-founded theory of aggregate labor supply, working Papers 2011-13, Instituto Madrileño de Estudios Avanzados (IMDEA) Ciencias Sociales.

Gervais, M., Klein, P., 2010. Measuring consumption smoothing in CEX data. Journal of Monetary Economics 57 (8).

Guvenen, F., 2007. Learning your earning: Are labor income shocks really very persistent? American Economic Review 97 (3), 687–712.

Guvenen, F., Smith, A., 2010. Inferring labor income risk from economic choices: An indirect inference approach, mimeo, University of Minnesota.

Heathcote, J., Storesletten, K., Violante, G., 2010. The macroeconomic implications of rising wage inequality in the united states. Journal of Political Economy 118 (4), 681–722.

Huggett, M., 1993. The risk-free rate in heterogenenous-agent incomplete-insurance economies. Journal of Economic Dynamics and Control 17, 953–969.

Kaplan, G., Violante, G. L., 2010. How much insurance beyond self-insurance? American Economic Journal: Macroeconomics 2, 53–87.

Krueger, D., Perri, F., 2004. Understanding consumption smoothing: Evidence from US consumer expenditure data. Journal of the European Economic Association 3 (2-3), 340–350.

Lee, B. S., Ingram, B. F., 1991. Simulation estimation of time-series models. Journal of Econometrics 47, 197–205.

Rouwenhorst, K. G., 1995. Asset pricing implications of equilibrium business cycle models. In: Cooley, T. F. (Ed.), Frontiers of Business Cycle Research. Princeton University Press, pp. 294–330.

Storesletten, K., Telmer, C., Yaron, A., 2004. Consumption and risk-sharing over the life cycle. Journal of Monetary Economics 51, 609–633.

Telyukova, I. A., 2011. Household need for liquidity and the credit card debt puzzle. Review of Economic Studies, forthcoming.

Table 1: Parameters for the Earnings Process

	σ^2_α	ρ	σ^2_ε	σ^2_x
Annual	0.0461	0.9254	0.0197	0.0506
Monthly	0.0268	0.9947	0.0016	0.5395

Table 2: Risk-Sharing Implications of the Annual and Monthly Models

Frequency	$\underline{a}$	r	β	$\frac{var(\ln(c))}{var(\ln(e))}$	$\rho(\ln c)$
Annual	0.5	1.6	0.73	0.94	0.75
Monthly		1.4	0.72	0.96	0.78
Annual	1.5	1.8	0.60	0.90	0.79
Monthly		1.6	0.62	0.93	0.80
Annual	2.5	2.0	0.49	0.88	0.82
Monthly		1.7	0.57	0.92	0.81
Annual	4	2.1	0.40	0.89	0.85
Monthly		1.8	0.52	0.92	0.82

Note: $\underline{a}$ is the borrowing limit given as a factor multiplying average annual earnings. β is the regression coefficient of log consumption changes on log earnings changes. $\rho(\ln c)$ is the autocorrelation of log consumption. The interest rate r is the annual rate.