Journal of Macroeconomics 32 (2010) 169–184

Contents lists available at ScienceDirect

Journal of Macroeconomics

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

Idiosyncratic labour income risk and aggregate consumption: An unobserved component approach

Lorenzo Pozzi *

Tinbergen Institute and Department of Economics, Erasmus University Rotterdam, Burgemeester Oudlaan 50, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands

ARTICLE INFO

Article history:
Received 4 March 2008
Accepted 14 December 2009
Available online 24 December 2009

JEL classification:
E21

Keywords:
Labour income uncertainty
Consumption
Precaution
State space models
GARCH errors
Unobserved component
Bayesian

ABSTRACT

We investigate the importance of aggregate and consumer-specific or idiosyncratic labour income risk for aggregate consumption changes in the US over the period 1952–2001. Theoretically, the effect of labour income risk on consumption changes is decomposed into an aggregate and into an idiosyncratic part. Empirically, aggregate risk is modelled through a GARCH process on aggregate labour income shocks and individual risk is modelled as an unobserved component and obtained through Kalman filtering. Our results suggest that aggregate labour income risk explains a negligible fraction of the variance of aggregate consumption changes. A more important part of aggregate consumption changes is explained by the unobserved component. The interpretation of this component as reflecting idiosyncratic labour income risk is supported by the finding that it is negatively affected by received consumer transfers. Idiosyncratic labour income risk thus matters for the aggregate economy.

1. Introduction

In this paper we investigate the effects of labour income risk on aggregate consumption changes using quarterly data for the US over the period 1952–2001. The contribution of this paper to the literature is threefold.

First, using the theoretical results of Caballero (1990) as a starting point, we present a theoretical framework in which the effect of labour income risk on the change in aggregate consumption is decomposed into two parts: the impact of aggregate labour income risk and the impact of consumer-specific or idiosyncratic labour income risk.[1] This decomposition is useful because limiting income risk to aggregate income risk is too restrictive. The reason is that the variance of aggregate labour income is low. As a result, in permanent income models with time-separable preferences, the magnitude of the average growth in consumption can only be explained by values of risk aversion that are much higher than what is widely believed (see Gourinchas and Parker, 2001). Another reason is that some authors argue that there is no theoretical a priori justification (e.g. Deaton, 1992, p. 37) or empirical evidence (e.g. Banks et al., 2001) to suggest that risk pooling mechanisms that effectively eliminate individual-specific labour income risk actually do exist.

* Tel.: +31 10 408 12 56; fax: +31 10 408 91 61.
E-mail addresses: pozzi@ese.eur.nl, pozzi@few.eur.nl.
URL: http://people.few.eur.nl/pozzi.

[1] See Banks et al. (2001) for a comparable decomposition at the individual/cohort level but without the possibility to aggregate analytically since preferences are of the constant relative risk aversion type.

0164-0704/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.jmacro.2009.12.003

Second, to complement studies that use a micro-based approach,[2] we follow a pure aggregate time series approach. Aggregate labour income risk is estimated through a GARCH process on aggregate labour income shocks. Idiosyncratic income risk is modelled as an unobserved component and estimated through Kalman filtering techniques. The identification of the unobserved component is based on the use of two determinants of income risk suggested in the literature: (government) transfers received by consumers and the business cycle. The main contribution of the paper is that we investigate whether changes in the amount of transfers received by consumers affect the unobserved component and the change in consumption. From the papers by Hubbard et al. (1995) and Engen and Gruber (2001), we know that transfers provided by the social security system (e.g. health and unemployment insurance) may reduce idiosyncratic labour income risk by providing insurance against bad draws of labour income. As far as the business cycle is concerned, Carroll (1992) notes that "the most drastic fluctuations in household income are those associated with spells of unemployment". We investigate whether changes in the unemployment rate affect the unobserved component and therefore the change in aggregate consumption. There is a literature that investigates whether idiosyncratic labour income risk shows cyclical behaviour (see Storesletten et al., 2004; Parker and Preston, 2005; Primiceri and van Rens, 2004). We compare the results of our macro-based time series approach to the results of these papers that use micro data.

Third, we use a Bayesian approach to parameter estimation. A Bayesian approach allows us to incorporate prior knowledge into our estimations. Priors are also useful to estimate GARCH effects in state space systems in the presence of outliers in the data. Existing procedures to detect and correct for outliers in GARCH models (see e.g. Franses and van Dijk, 1999) cannot be applied in a state space context.

Our results suggest that aggregate labour income risk explains only a negligible fraction of the variance of aggregate consumption changes. The unobserved component explains a more important part of consumption changes. This component can be interpreted as idiosyncratic labour income risk since it is negatively affected by the trend in transfers received by consumers (i.e. transfers can be considered a proxy for idiosyncratic labour income risk). This suggests that idiosyncratic labour income risk matters for the aggregate economy. We argue that the trend change in transfers received by consumers can explain low frequency movements in consumption changes. We find no hard evidence that the unobserved component is driven by the change in the unemployment rate however. We then extend the model by adding rule-of-thumb consumers who base their consumption decisions on current income. The existence of rule-of-thumb consumers may be due to liquidity constraints and capital market imperfections (see e.g. Campbell and Mankiw, 1990) or myopia (see e.g. Flavin, 1985). Estimation of the extended model suggests that our results are robust when excess sensitivity of consumption to anticipated disposable income is taken into account.

The paper is structured as follows. In Section 2 we present a consumption model with time-varying aggregate and time-varying idiosyncratic labour income risk. In Section 3 we present our basic empirical specification which is put into state space form to estimate the unobserved component. We discuss the Bayesian estimation of the unknown parameters in the model. In Section 4 we present the estimation results for our basic model. In Section 5 we investigate whether our conclusions are affected by the introduction of rule-of-thumb consumers in the model. Section 6 concludes.

2. A consumption model with time-varying aggregate and time-varying idiosyncratic labour income risk

In this section we derive an expression for the change in aggregate private consumption that takes into account uncertainty with respect to aggregate labour income and uncertainty with respect to the idiosyncratic component of labour income. The latter type of risk is present because insurance markets are assumed to be incomplete (i.e. there is no risk pooling across consumers). The model uses the results of Caballero (1990) in a setting where consumers are heterogeneous in the sense that they experience different income draws. As a result, given the absence of insurance mechanisms, consumption trajectories and wealth levels may diverge considerably over consumers.

The economy consists of n consumers, each having an infinite planning horizon. Each consumer i (where $i = 1, \ldots, n$) has a utility function of the constant absolute risk aversion (*CARA*) type, namely $u(c_{it}) = (-1/\gamma)e^{-\gamma c_{it}}$ where c_{it} is real consumption of consumer i in period t and where γ is the coefficient of absolute risk aversion ($\gamma > 0$) which also equals the coefficient of absolute prudence. We use this type of utility function instead of the more usual utility function of the constant relative risk aversion (*CRRA*) type because of its analytical convenience, i.e. it facilitates aggregation. We further assume that all consumers can freely lend and borrow, i.e. capital markets are perfect. In Section 5 we relax this assumption. We also assume that all consumers face the same constant real interest rate r which equals their rate of time preference.[3] Unlike capital markets, insurance markets are incomplete. That is, consumers cannot insure themselves through the use of so-called Arrow securities (see Deaton, 1992, pp. 35–36) that could be traded among them to smooth consumption across different states of the world.

Given the stated assumptions, the first-order condition in period $t+1$ for consumer i is,

$$E_{it}(e^{-\gamma \Delta c_{it+1}}) = 1, \tag{1}$$

[2] The use of micro-based uncertainty measures has certain drawbacks that can be avoided by following an aggregate approach. First, there is the small length of the available time series. Second, decomposing these measures into an aggregate and an idiosyncratic part is not straightforward. Third, the use of these measures can be problematic in the presence of measurement errors or "self-selection" problems (see Attanasio, 1999, for a discussion).

[3] This implies that intertemporal substitution effects caused by the (anticipated) interest rate are ruled out. There is evidence that the ex ante real interest rate has no impact on consumption growth (see Hall, 1988; Campbell and Mankiw, 1990 or Ludvigson, 1999) even when time-varying income risk is taken into account (see Parker and Preston, 2005).