The Insurance Value of Public Insurance Against Idiosyncratic Income Risk*

Christopher Busch† Rocio Madera‡

February 15, 2024

preliminary—please do not circulate

Abstract

The tax and transfer system partially insures households against individual income risk. We discuss under which assumptions differences between income processes estimated for household gross income and disposable income are informative about the (welfare) value of this partial insurance. Our approach works directly with income processes estimated separately on the two income measures, and does not require the specification (nor estimation) of a tax function. Instead we use an incomplete markets framework that links an estimated income process to consumption. Its key feature is that the degree of partial insurance is directly parameterized: Technically, this allows to solve for the degree of insurance provided by the tax and transfer system as a fixed point. The approach works with standard restrictions on income processes and preferences, and it further enables us to explore the role of higher-order risk for the value assigned to public insurance.

Keywords: Idiosyncratic income risk, tax and transfer system, social insurance policy, incomplete markets.

1 Introduction

Most individuals face individual income risk, against which they are partially insured. This insurance comes from various sources, both public and private in nature. Public insurance typically comes through a combination of various policy instruments, and

*For detailed discussions in early stages of the project we thank David Domeij and Fatih Guvenen. For helpful comments we thank Jeanne Commault, Bulent Guler, Burhan Kuruscu, Jonas Loebbing, Alex Ludwig, Sergio Ocampo, Gonzalo Paz-Pardo, Kjetil Storesletten, Mirko Wiederholt, and Peter Zorn.

†LMU Munich, CESifo and CEPR; `chris.busch.econ@gmail.com`; `https://chrisbusch.eu`

‡Southern Methodist University, CESifo; `romadera@smu.edu`; `http://s2.smu.edu/romadera`

the overall mix of the tax and transfer system translates household gross income into household disposable income. For one, progressivity of the income tax system implies insurance from an ex ante perspective as it compresses the distribution of possible income changes. Other policies are explicitly designed to dampen downside risk: most notably, unemployment insurance dampens temporary income losses due to job loss. Further, as is well-studied in the literature, individual risk moves with aggregate risk and exhibits strong cyclical patterns: In times of economic downturns individuals face larger risk of income losses (i.e., increased downside risk) along with smaller risk of income gains (i.e., reduced upside risk). This translates into asymmetric shifts of the distribution of individual income changes over the business cycle—the resulting relationship between aggregate and idiosyncratic income changes is well-captured by procyclical skewness of the distribution of idiosyncratic changes. Such procyclical skewness is well-documented for a large set of diverse countries (see, e.g., Busch *et al.*, 2022; Guvenen *et al.*, 2014). At the same time, there is evidence that existing tax and transfer schemes are successful in providing partial insurance against individual risk in general (see, e.g., Blundell *et al.*, 2014; De Nardi *et al.*, 2021), and against increased downside risk in contractions in particular: Busch *et al.* (2022) document that taxes and transfers dampen both the extent of income changes, downside risk, and its cyclicality—in the United States, Germany, or Sweden: household disposable income exhibits less dispersed and more symmetrically distributed changes than household gross income, and the distribution is less responsive to the aggregate business cycle.

In this paper, we discuss which assumptions permit to assign an *insurance value* to the smoothing power of the tax and transfer system, and to interpret it in welfare terms. Given that insurance capabilities against transitory income variations and those against permanent ones starkly differ, we first follow the route of separating the two components using the structure of an income process. Next, many rich (administrative) data sources that allow for detailed exploration of household income trajectories do not cover consumption data. While income is observed, consumption eventually translates into welfare. Thus, some structure needs to be formulated in order to evaluate the insurance value of the tax and transfer system. In particular, one needs to model how disposable income maps into household consumption. In a data context where consumption data is available, this can be used to directly inform this relationship empirically. Given consumption, one can explore the link between gross income and consumption, and relate it to the link between disposable income and consumption. The comparison is in principle informative about the size of insurance coming from the tax and transfer system (cf. Blundell *et al.*, 2008, henceforth BPP). To be more precise, we formulate a consumption function that takes one input—the income shocks $\boldsymbol{\phi}$ (which collect a permanent shock η and a transitory shock ε)—and which is parameterized by the degree of insurance $\boldsymbol{\lambda}$ against the two types of shocks. Finally, assigning welfare to the obtained insurance measure requires the specification of preferences. In sum, there are three steps: (i.) assumptions on the income process give a distinction between transitory and permanent components, (ii.) a consumption function that links transitory and permanent shocks into consumption, (iii.) a preference specification that allows for welfare interpretation. In our approach, which we sketch next, we derive the consumption function of step (ii.) from the same model that we then use to assign a welfare value in step (iii.).

We use an incomplete markets model, with two key ingredients: exogenous idiosyncratic productivity risk, and access to only partial insurance against it (as in Aiyagari, 1995; Huggett, 1993; Imrohoroglu, 1989a). Standard incomplete markets models explicitly feature one source of insurance: self-insurance through savings in a riskfree asset. Our approach is agnostic with respect to the exact source of insurance (as discussed, e.g., in Blundell *et al.*, 2008). Our objective is to obtain a consumption function that relates idiosyncratic shocks to consumption for a given degree of partial insurance, while being able to directly control this degree of insurance. To this end, we adopt a model structure that features the abstraction of two distinct types of risks—one perfectly insurable, and one uninsurable. This is implemented with an island structure (a la Heathcote *et al.*, 2014, henceforth HSV), where a fraction of the overall idiosyncratic risk is shared among a group of individuals, and a fraction is purely idiosyncratic.[1] There is an equilibrium with no asset trade across islands, such that island-level shocks

[1] Reflecting the idiosyncratic nature of the considered shock, there is a continuum of such groups (islands), and a continuum of agents on each island.

pass-through one-for-one to consumption. Within an island, risk is perfectly shared, and individual shocks do not pass-through to individual consumption. This implies that the fraction of total shocks attributed to within-island idiosyncratic shocks can be interpreted as the degree of partial insurance against (total) idiosyncratic risk—and also translates into the coefficient of a BPP-type regression. This fraction is the key parameter of the model. For a given degree of partial insurance, the (exogenous) income process maps into an (endogenous) consumption process.

We use this structure of the model in order to trace out the degree of partial insurance provided by the tax and transfer system in a flexible way that does not require the specification of a tax function. To this end, we consider two scenarios in the model. In the first scenario, we treat a stochastic income process that captures regularities of post-government earnings as fundamental. Given an assumed amount of insurance against risk remaining after taxes and transfers, the model delivers a consumption process. In the second scenario, we treat an income process that captures regularities of pre-government earnings as fundamental. We then search for the degree of partial insurance $\boldsymbol{\lambda}$ (i.e., the fraction of shocks that is purely idiosyncratic and thus insured) as a fixed point for which the model yields a consumption process that makes households indifferent to the consumption process obtained under the post-government income process. This way, we obtain a measure of the overall amount of partial insurance against pre-government income fluctuations, which we translate into the degree of partial insurance provided by the tax and transfer system. Thus, the model serves as a device for the measurement of the insurance value of taxes and transfers, which takes as inputs (estimated) income processes for pre-government income and post-government income, and makes a minimal set of structural assumptions (on preferences) and practical assumptions (on insurance beyond taxes and transfers).

In the benchmark calculation, we assume that permanent shocks to post-government income cannot be insured against. Varying the degree of partial insurance against permanent risk in post-government income does not have relevant effects on the measured degree of partial insurance provided by the tax and transfer system. Given the catchall nature of the partial-insurance parameter $\boldsymbol{\lambda}$ in our employed model framework, we consider full insurance against transitory shocks as a plausible benchmark. This is motivated by quantitative insights from calibrated incomplete markets models, in which households savings in a riskfree asset already generate a very low pass-through of transitory shocks to consumption (e.g., Busch and Ludwig, 2023; De Nardi *et al.*, 2020; Kaplan and Violante, 2010a).

We apply our approach to estimated income processes for income moments from Swedish tax data. To generate non-zero skewness and excess kurtosis consistent with the data, we specify the econometric model for (log) income as the sum of a permanent process and a transitory component, where all innovations (shocks) are drawn from mixtures of normals. We estimate two sets of parameters of this process separately for pre- and post-government household labor income by matching moments that capture the salient features of household income change distributions and their cyclical properties as documented in Busch *et al.* (2022).[2] We find that the degree of partial insurance provided by the tax and transfer system amounts to about 43%, which translates into a welfare gain, expressed as a consumption equivalent variation (CEV), of about 14.3% under log utility. We then focus on the part of that gain that is attributable to smoothing business cycle variation of the distribution. Taxes and transfers insure about 6% of the cyclical changes in the distribution which translates into a CEV of about 1.3%. However, the remaining risk (in post-government household-level income) is still substantial: households are willing to pay 4.6% of their consumption to completely eliminate procyclical fluctuations in skewness.

We then explicitly explore the role of taking into account higher-order risk in comparison to a Gaussian distribution that shares the same variance with the distributions considered in the benchmark analysis. We find that with log utility it does not matter much for the overall insurance gain of the tax and transfer system, which is not surprising as risk attitudes against skewness and kurtosis are relatively weak in this specification. Still, the insurance against cyclical variations in the distribution is valued twice as much under a Gaussian distribution. Thus, when not taking into account

[2]Note that the specific parametric form of the distribution is not essential, as long as relevant moments of the distribution are matched; see, e.g., Busch and Ludwig (2023), who illustrate how central moments of the distribution map into choices of agents in a life-cycle model.

higher-order risk one would overestimate the insurance value of the tax and transfer system against cyclical variations in idiosyncratic risk. Similarly, one would overestimate the potential gain of further smoothing.

Related Literature. There is an extensive literature on the welfare benefits of tax and transfer systems across the globe. For the case of Sweden, Floden and Linde (2001) found large welfare gains from both redistribution and insurance against total uninsurable income risk. In addition, certain public insurance instruments act as automatic stabilizers against aggregate fluctuations (McKay and Reis, 2016). Drawing on recent empirical findings by Busch *et al.* (2022), we aim to bridge the gap between these two lines of research and gain insights into the welfare implications of tax and transfer systems for mitigating the pass-through of aggregate fluctuations to individual income.

In doing so, we contribute to the literature on the welfare costs of business cycles, which has a long history, tracing its origins to the pioneering work of Lucas (1987) but widely generalized to the context of heterogeneous agents facing idiosyncratic income risk and incomplete markets (Imrohoroglu, 1989b; Storesletten *et al.*, 2001; Krusell *et al.*, 2009). This literature emphasizes the role of distributions and cyclical variation in idiosyncratic income risk as a source of amplification of the welfare costs of cyclical fluctuations. The distributional changes considered in these works are symmetric and following a Normal distribution. In contrast, we pose a flexible distribution that allows for asymmetric fluctuations of idiosyncratic risk that also capture the fact that changes are more likely to be very small or very large compared to a Normal distribution (see evidence in, e.g., Guvenen *et al.*, 2014; Busch *et al.*, 2022). Importantly, our main goal lies on quantifying the success of the existing tax and transfer system in smoothing the extent and business cycle variation of idiosyncratic risk; different to Busch and Ludwig (2023), who explore the role of remaining higher-order risk in a quantitative model.

To highlight the importance of our channels, we also adopt a less quantitatively rich but very transparent modeling framework linking cyclical idiosyncratic risk to consumption dynamics and welfare. We build on Heathcote *et al.* (2014)'s partial insurance framework. In this sense, we bridge results from the life-cycle literature that focuses on a bundle of self- and family-insurance channels beyond the traditional savings instruments considered in the business-cycle literature (Blundell *et al.*, 2008; Krueger and Perri, 2006; Kaplan and Violante, 2010b). In contrast to HSV, who pose a process for wages and explicitly model two insurance channels—endogenous labor supply and a progressive tax function—we treat household income as the fundamental source of risk, and incorporate a rich income process with time-varying risk in the spirit of McKay (2017) into the model framework, while retaining analytical tractability.

The paper is organized as follows. Section 2 outlines the incomplete markets model. Section 3 introduces the income process. Section 4 discusses the quantitative results, and Section 5 concludes.

2 Quantitative Model: A Measurement Device

2.1 Model Economy

Endowment structure and preferences. We consider a stochastic endowment economy, which is populated by a continuum of islands, each of which is in turn populated by a continuum of agents. There are two types of shocks: one common to all members of an island and the other purely idiosyncratic. The within-island shocks wash out on the island, the island-level shocks wash out across islands, such that there is no aggregate risk to total endowment. An island refers to a group of agents that are described by the same history of island-level shocks (common to all members of the group).

Islands can be thought of as a network of family members, who perfectly share the risks faced by each individual. If, for example, all family members work in the same industry and live in the same region, there will be shocks that hit every member equally and hence cannot be insured within the family network. Importantly for the quantitative analysis, there is no need to define empirical counterparts to the model islands.

Specifically, individual income (endowments) is assumed to follow

$$\begin{aligned} y_t &= y_t^{island} + y_t^{idio} \\ y_t^i &= z_t^i + \varepsilon_t^i, \quad \varepsilon_t^i \sim F_{\varepsilon,t}^i, \quad \text{for } i \in \{island, idio\} \\ z_t^i &= z_{t-1}^i + \eta_t^i, \quad \eta_t^i \sim F_{\eta,t}^i, \quad \text{for } i \in \{island, idio\} \end{aligned} \tag{1}$$

where z_t^i and ε_t^i for $i \in \{island, idio\}$ denote the island-level and idiosyncratic permanent and transitory components of income. All stochastic components of income are independent and normalized such that $\int \exp(x_t^i)\, dF_{x,t}^i = 1$ for $i \in \{island, idio\}$ and $x \in \{\varepsilon, \eta\}$.

Agents live finite lives. Each period a mass $(1-\delta)$ of newborns enters the economy with age 0. The probability of survival from age a to age $a+1$ is constant at δ. Agents maximize discounted lifetime utility, whereby we assume time- and state-separable preferences. For the per period utility function, we use log utility as the benchmark: $U(c_t) = \log(c_t)$. We also study the importance of this assumption and inspect the role of stronger risk attitudes by using an alternative specification with a CRRA per period utility function with parameter of relative risk aversion larger than 1. The discount factor β is constant across the population.

Age 0 agents entering in year τ hold zero financial wealth and are allocated to an island of agents which then share the same sequence of island-level shocks $\left\{\eta_t^{island}, \varepsilon_t^{island}\right\}_{t=\tau}^{\infty}$.

Asset markets and equilibrium. Every period agents engage in asset trade. There is a full set of state-contingent claims available to agents within islands. Claims are in zero net supply. Across islands, agents cannot trade claims contingent on the island-level shocks. This restriction on available assets implies that in equilibrium island-level shocks remain uninsured, while within-island shocks are fully insured, and risk is shared by all individuals on an island. In other words, a no-trade equilibrium in the spirit of Constantinides and Duffie (1996) exists. While in their model, idiosyncratic endowment shocks remain fully uninsured in this no-trade equilibrium, the equilibrium in our model entails partial insurance: there is no asset trade across islands, while agents within an island insure themselves perfectly against the individual-specific shocks. This mimics the result in Heathcote *et al.* (2014).

In equilibrium, log consumption and consumption change are given by[3]

$$\log c_t\left(\mathbf{x}_t, y_t^{idio}\right) = y_t^{island} + \log \int \exp\left(y_t^{idio}\right) dF^a_{y^{idio},t} \quad (2)$$

$$\begin{aligned}\Delta \log c_t &= \Delta y_t^{island} + \log \frac{\int \exp\left(\eta_t^{idio}\right) dF^{idio}_{\eta,t} \int \exp\left(\varepsilon_t^{idio}\right) dF^{idio}_{\varepsilon,t}}{\int \exp\left(\varepsilon_{t-1}^{idio}\right) dF^{idio}_{\varepsilon,t-1}} \quad (3)\\ &= \eta_t^{island} + \Delta\varepsilon_t^{island} + \log \frac{\int \exp\left(\eta_t^{idio}\right) dF^{idio}_{\eta,t} \int \exp\left(\varepsilon_t^{idio}\right) dF^{idio}_{\varepsilon,t}}{\int \exp\left(\varepsilon_{t-1}^{idio}\right) dF^{idio}_{\varepsilon,t-1}}\\ &= \eta_t^{island} + \Delta\varepsilon_t^{island},\end{aligned}$$

where $\Delta \log c_t = \log c_t - \log c_{t-1}$ and $\Delta\varepsilon_t^{island} = \varepsilon_t^{island} - \varepsilon_{t-1}^{island}$. The above equation summarizes the major advantage—relative to standard incomplete market models—of introducing the partial insurance framework by the abstraction of islands: it allows for an analytical solution in which consumption changes are expressible as an explicit function of idiosyncratic shocks. Note that the uninsurable, island-level shocks translates one-for-one to consumption. The individual realizations of the two insurable shocks, however, do not affect consumption: given perfect risk-sharing, all members of an island consume the mean realization of these shocks.

Degree of partial insurance. Blundell *et al.* (2008) introduce a *pass-through coefficient* for a given shock to denote the fraction of the shock that translates into consumption changes. We use the model equivalent of the pass-through to define the insurance coefficients against transitory and permanent shocks, where in the same spirit as, e.g., Kaplan and Violante (2010a) we assume that shock components are observed, i.e., transitory and permanent shocks can be told apart by model agents. As is clear from (3), in our model framework, *island*-shocks translate one-for-one to consumption—the pass-through of shock to consumption is one—and *idio*-shocks do not translate into consumption—the pass-through of shock to consumption is zero. We thus consider the

[3]The derivation of consumption outlined in Heathcote *et al.* (2014) carries over one-for-one to our model version, simplified by the fact that we do not have a tax function nor endogenous labor supply.

pass-through not to one of those shocks, but to the combined *island* and *idio*-shocks. First, consider the pass-through of overall income changes, i.e., Δy_t:

$$1-\lambda = \frac{\mathrm{cov}(\Delta \log c_t, \Delta y_t)}{\mathrm{var}(\Delta y_t)} = \frac{\mathrm{cov}(\Delta y_t^{island}, \Delta y_t)}{\mathrm{var}(\Delta y_t)} = \frac{\mathrm{var}(\Delta y_t^{island})}{\mathrm{var}(\Delta y_t)}. \tag{4}$$

Thus, the ratio of the variance of *island* shocks to total shocks determines the pass-through of income changes to consumption—and conversely, the ratio of the variance of *idio* shocks to total shocks gives the degree of partial insurance, λ.

We then split the above up, explicitly considering different degrees of insurance against transitory and permanent shocks, which gives

$$\begin{aligned} 1-\lambda_{trans} &= \frac{\mathrm{cov}(\Delta \log c_t, \varepsilon_t)}{\mathrm{var}(\varepsilon_t)} \\ &= \frac{\mathrm{cov}(\Delta \varepsilon_t^{island}, \varepsilon_t)}{\mathrm{var}(\varepsilon_t)} = \frac{\mathrm{cov}(\varepsilon_t^{island} - \varepsilon_{t-1}^{island}, \varepsilon_t^{island} + \varepsilon_t^{idio})}{\mathrm{var}(\varepsilon_t)} \\ &= \frac{\mathrm{var}(\varepsilon_t^{island})}{\mathrm{var}(\varepsilon_t^{island} + \varepsilon_t^{idio})} = \frac{\mathrm{var}(\varepsilon_t^{island})}{\mathrm{var}(\varepsilon_t^{island}) + \mathrm{var}(\varepsilon_t^{idio})} \end{aligned} \tag{5}$$

and

$$\begin{aligned} 1-\lambda_{perm} &= \frac{\mathrm{cov}(\Delta \log c_t, \eta_t)}{\mathrm{var}(\eta)} \\ &= \frac{\mathrm{cov}(\eta_t^{island}, \eta_t)}{\mathrm{var}(\eta_t)} = \frac{\mathrm{cov}(\eta_t^{island}, \eta_t^{island} + \eta_t^{idio})}{\mathrm{var}(\eta_t)} \\ &= \frac{\mathrm{var}(\eta_t^{island})}{\mathrm{var}(\eta_t^{island} + \eta_t^{idio})} = \frac{\mathrm{var}(\eta_t^{island})}{\mathrm{var}(\eta_t^{island}) + \mathrm{var}(\eta_t^{idio})}, \end{aligned} \tag{6}$$

such that the degree of partial insurance against permanent shocks, λ_{perm} is given by the fraction of the variance of permanent shocks attributable to the *idio*-component, and the degree of partial insurance against transitory shocks is given by the fraction of the variance of transitory shocks attributable to the *idio*-component.

Tax and transfer system. We then introduce a tax and transfer system that alters the endowment stream faced by agents. We do not explicitly model the tax system, but retain full flexibility about its nature—i.e., we do not make any functional form assumption. Instead, we consider a second scenario in which agents face income stream (1), but with different distributions of shocks. Importantly, we maintain the normalization that $\int \exp\left(x_t^i\right) dF_{x,t}^i = 1$ for $i \in \{island, idio\}$ and $x \in \{\varepsilon, \eta\}$, which means that we consider a tax and transfer system that cross-sectionally redistributes endowments, which means that, e.g., we do do not allow for wasteful government consumption nor for transfer payments financed by debt.

2.2 Insurance Value of Tax and Transfer System

We now use the model structure outlined above in order to back out the degree of partial insurance provided by the tax and transfer system. To this end, we consider the following experiment. Agents live in one of two possible scenarios that differ in the endowment streams that agents face. In the first, the endowment stream describes pre-government incomes. In the second, the endowment stream describes post-government incomes. We then assume a degree of partial insurance against (total) individual shocks in the post-government scenario—i.e., we assume values for λ_{trans}^{post} and λ_{perm}^{post}. Given this assumed amount of partial insurance, we obtain stochastic consumption streams per equation (3).

We then find the degree of partial insurance in the pre-government scenario that makes agents ex ante indifferent to living in the post-government scenario (for the given degree of insurance in the latter). Given that there are two types of shocks, in principle multiple combinations of $\{\lambda_{trans}^{pre}, \lambda_{perm}^{pre}\}$ can exist that make agents indifferent. We assume that $\lambda_{trans}^{pre} = \lambda_{trans}^{post} = 1$. Thus, we assume that transitory shocks are well insured and do not pass-through to consumption. In the abstraction of the island model this shows by having no island-level shocks, and instead all transitory shocks happen purely within islands, and thus can be insured away fully by agents trading state-contingent claims. As mentioned in the introduction, note that incomplete market models typically find very high insurance against transitory shocks through private savings alone (Busch and Ludwig, 2023; De Nardi *et al.*, 2020; Kaplan and Violante, 2010a), thus full insurance appears to be a plausible assumption.

This leaves partial insurance against permanent risk as the relevant margin of the model. Given an assumption regarding λ_{perm}^{post}, we find the λ_{perm}^{pre} that makes agents indifferent. In our benchmark calculation, we assume that $\lambda_{perm}^{post} = 0$. This assumption is motivated by empirical results in Blundell *et al.* (2016), who find that the degree of partial insurance on top of government and family transfers is very close to zero. The obtained λ^{pre} can then be interpreted as the degree of partial insurance provided by the government under the assumption that there is no additional partial insurance. We show below that the assumption on λ_{perm}^{post} is not strong.

3 Pre- and Post-Government Income in Sweden

Given the measurement device provided by the model outlined above, we are set for evaluating the degree of partial insurance provided by the tax and transfer system. The two empirical ingredients necessary are two stochastic income streams: one that captures the regularities of pre-government income, and one that captures the regularities of post-government income. We estimate these using Swedish data moments.

3.1 Estimated Income Processes

Let y_t^{pre} and y_t^{post} denote log of pre- and post-government household income, respectively. We assume that it follows the following permanent-transitory process (where we drop the explicit reference to pre or post-government income):

$$\begin{aligned} y_t &= z_t + \varepsilon_t \qquad (7) \\ z_t &= z_{t-1} + \eta_t \end{aligned}$$

where ε_t is an *iid* transitory shock, and η_t denotes a permanent shock with time-varying and business-cycle-dependent distribution, modeled as in McKay (2017). We specify the distribution functions such that the process can match excess kurtosis and skewness found in the data.

In particular, the transitory component ε_t is drawn from a mixture of two normals:

$$\varepsilon_t \sim \begin{cases} \mathcal{N}(\bar{\mu}_\varepsilon, \sigma^2_{\varepsilon,1}) & \text{with prob. } p_{\varepsilon,1} \\ \mathcal{N}(\bar{\mu}_\varepsilon, \sigma^2_{\varepsilon,2}) & \text{with prob. } 1 - p_{\varepsilon,1} \end{cases} \tag{8}$$

where $p_{\varepsilon,1}$ denotes the probability of drawing from component 1; $\bar{\mu}_\varepsilon$ is chosen such that $\mathbb{E}\left[\exp(\varepsilon)\right] = 1$. The permanent component η_t follows a mixture of three normals:

$$\eta_t \sim \begin{cases} \mathcal{N}(\bar{\mu}_{\eta,t} + \mu_{\eta,1} + \phi_1 x_t, \sigma^2_{\eta,1}) & \text{with prob. } p_{\eta,1} \\ \mathcal{N}(\bar{\mu}_{\eta,t} + \mu_{\eta,2} + \phi_2 x_t, \sigma^2_{\eta,2}) & \text{with prob. } p_{\eta,2} \\ \mathcal{N}(\bar{\mu}_{\eta,t} + \mu_{\eta,3} + \phi_3 x_t, \sigma^2_{\eta,3}) & \text{with prob. } p_{\eta,3} \end{cases} \tag{9}$$

where $p_{\eta,i}$, $i = 1, 2, 3$, denotes the probability of drawing from component i, where $\sum_{i=1}^{3} p_{\eta,i} = 1$. The parameters ϕ_i determine how strongly aggregate risk as captured by x_t translates into changes of the distribution of idiosyncratic earnings risk. x_t is standardized log GDP growth. As part of our goal is to capture the business-cycle fluctuations of idiosyncratic income risk, we choose $\bar{\mu}_{\eta,t}$ such that $\mathbb{E}\left[\exp(\eta_t)\right] = 1$. In the estimation, we then shift the distribution so as to impose the mean of medium-run (3-year) income changes to be as in the data. We use GDP growth as the empirical measure of aggregate fluctuations in order to make the quantitative results easily interpretable. Over the period of estimation, the average GDP growth rate is 2.15% with a standard deviation of about 2.35%.

Estimation of process. We estimate the set of parameters $\chi = \{\chi_{trans}, \chi_{perm}\}$ where

$$\chi_{trans} = \{\sigma_{\varepsilon,1}, \sigma_{\varepsilon,2}, p_{\varepsilon,1}\} \tag{10}$$

$$\chi_{perm} = \{\mu_{\eta,2}, \mu_{\eta,3}, \sigma_{\eta,1}, \sigma_{\eta,2}, p_{\eta,1}, p_{\eta,2}, \phi_2, \phi_3\} \tag{11}$$

by the simulated method of moments (SMM).[4] We target the time series of L9050 and L5010[5] of the 1, 3, and 5-year earnings changes distribution, the average of the Crow-Siddiqui measure of kurtosis of 1-, 3-, and 5-year changes, as well as the age profile of the cross-sectional variance from ages 25 to 60. The Crow-Siddiqui measure of kurtosis (Crow and Siddiqui, 1967) is defined as $\mathcal{CS} = \frac{(P97.5-P2.5)}{(P75-P25)}$. This gives 213 moments for our estimation of the income process for Sweden.

To construct the simulated income profiles over time, we write earnings growth as a function of the shocks, using equation (7):

$$y_t - y_{t-s} = \varepsilon_t - \varepsilon_{t-s} + \sum_{j=0}^{s-1} \eta_{t-j}, \tag{12}$$

for different horizons $s = 1, 3, 5$. The simulated series of the life-cycle variance profile of log earnings is computed as follows. We assume a time-invariant distribution of shocks by imposing $x_t = 0 \; \forall t$. Notice that this assumes that the variance accumulates linearly over the life cycle. We then normalize the series so that the variance at age 25 in the simulation is 0. Finally, we rescale the resulting simulated profile to exhibit the same mean as its empirical counterpart.

We simulate these profiles $R = 10$ times for $I = 100,000$ individuals and compute the moments corresponding to the aforementioned targets. To find $\hat{\chi}$, we minimize the average scaled distance between the simulated and empirical moments. A weighting matrix is used to scale the life-cycle profile. In particular, we weight the variance profile with 20% and the remaining moments with 80%. For the optimization part, we use a global version of the Nelder-Mead algorithm with several quasi-random restarts, as described in Guvenen (2011).

Let c_n^m denote the empirical moment n ($n = 1, \cdots, N$) that corresponds to cross-sectional target $m \in \big\{L5010(\Delta^1 y_t), L5010(\Delta^3 y_t), L5010(\Delta^5 y_t), \ldots, var(y_{age=25}), \ldots,$ $var(y_{age=60})\big\}$. In each simulation, we draw a matrix of random variables $X_r = \big\{\varepsilon_1^i, \varepsilon_2^i, \ldots,$

[4] For identification purposes, we impose $\mu_{\eta,2} \geq 0$, $\mu_{\eta,3} \leq 0$, and $\phi_1 = 0$. With this assumption, the time-varying means of the three mixtures will control the center, right tail, and left tail of the distribution of η, respectively. For practical purposes, we further assume $p_{\eta,2} = p_{\eta,3}, \sigma_{\eta,2} = \sigma_{\eta,3}$.

[5] $L9050 = P90 - P50$ denotes the difference between the 90^{th} and 50^{th} percentiles, and likewise $L5010 = P50 - P10$.

$\varepsilon_T^i, \eta_1^i, \ldots, \eta_T^i\}_{i=1}^I$ where T denotes the last year available in the data. For each simulation, we calculate the respective simulated moments $d_n^m(\chi, X_r)$ given the parameter vector χ.

We minimize the scaled deviation $F(\chi)$ between each data and simulated moment

$$min_\chi F(\chi)' W F(\chi)$$

where F is defined as

$$F_n(\chi) = \frac{d_n^m(\chi) - c_n^m}{|c_n^m|}$$

$$d_n^m(\chi) = \frac{1}{R}\sum_{r=1}^{R} d_n^m(\chi, X_r)$$

Parameter estimates. Table 1 shows the parameter estimates. To illustrate the magnitude of the estimated swings in the distribution of idiosyncratic risk, consider the time period around the Great Recession. During those years, the GDP growth rate plummets to a negative GDP growth of -5.04% in 2008 (about three standard deviations below the average), recovers to a strong 6.59% in 2009 (about 2 standard deviations above the average), followed by an about average growth year in 2010 with 2.49%. Over the course of these three years, the distribution of individual earnings changes is estimated to vary markedly as shown in Figure 1, which plots the distribution of the permanent component of income changes, η_t for both pre-government and post-government income. Each panel shows the simulated distribution for the estimated mixture of Normals in red (histogram), as well as the density of the corresponding Normal distribution with the same mean and variance (but with a skewness of 0 and a kurtosis of 3). In the plots, we use the normalization such that $E[\exp(\eta_t)] = 1$. For completeness, Figures 3 and 4 in Appendix A show the simulated moments at these parameters together with the empirical moments.

Table 1: Estimated Parameter Values

Parameter	Description	Pre-Gov.	Post-Gov.
$p_{\varepsilon,1}$	Mixture prob. of ε distribution	0.892	0.877
$\sigma_{\varepsilon,1}$	Std. dev. of ε distribution mix. comp. 1	0.055	0.047
$\sigma_{\varepsilon,2}$	Std. dev. of ε distribution mix. comp. 2	0.628	0.401
$p_{\eta,1}$	Mixture prob. of η distribution mix. comp. 1	0.981	0.981
$p_{\eta,2}$	Mixture prob. of η distribution mix. comp. 2	0.010	0.009
$p_{\eta,3}$	Mixture prob. of η distribution mix. comp. 3	0.010	0.009
$\sigma_{\eta,1}$	Std. dev. of η distribution mix. comp. 1	0.086	0.057
$\sigma_{\eta,2}$	Std. dev. of η distribution mix. comp. 2	0.020	0.009
$\sigma_{\eta,3}$	Std. dev. of η distribution mix. comp. 3	0.020	0.009
$\mu_{\eta,2}$	Mean of mixt. comp. 2 of η distribution	0.002	0.008
$\mu_{\eta,3}$	Mean of mixt. comp. 3 of η distribution	-0.158	-0.065
ϕ_2	Aggregate risk transmission mixt. comp. 2	1.186	1.240
ϕ_3	Aggregate risk transmission mixt. comp. 3	0.467	0.229
M	# moments targeted in estimation	213	213

Note: Estimated parameters for gross household labor income (Pre-Gov.) and household income after taxes and transfers (Post-Gov.) in Sweden.

As captured in the three rows of the figure, the distribution of permanent income changes varies over the cycle in an asymmetric way for both measures of income (pre- and post-government). Strong negative GDP growth (as in 2008) goes hand-in-hand with a left-skewed distribution as captured by a negative coefficient of skewness, while strong positive GDP growth (as in 2009) comes with a right-skewed distribution. In about average growth times, the idiosyncratic distribution turns out to be well captured by a Gaussian distribution—and while it is already very narrow for pre-government income, the tax and transfer system compresses the distribution even more: the variance is halved.

In an expansionary year (like 2009) the distribution turns into a right-skewed distribution, as captured by a positive coefficient of skewness (the third standardized moment), and also a positive measure of Kelley's skewness, which is a percentile-based measure of skewness calculated as $\mathcal{KS} = \frac{(P90-P50)-(P50-P10)}{P90-P10}$. $\mathcal{KS}$ takes on values $\in (-1, 1)$, and captures the relative size of the left and right tails in overall dispersion. Kelley's skewness is a useful statistic to interpret the magnitude of the change in the distribution over the cycle. For pre-government income, the value of $\mathcal{KS} = 0.44$ for

2009 indicates that $(P90 - P50)$ accounts for 72% of the $(P90 - P10)$ dispersion.[6] On the other hand, in 2008 the value of $\mathcal{KS} = -0.67$ indicates that $(P90 - P50)$ accounts for only 16.5% of the $(P90 - P10)$ dispersion.

[6] Note that $P90 - P50 = 0.5 + \frac{\mathcal{KS}}{2}$.

Figure 1: Cross-Sectional Distribution of Permanent Income Changes

2008 (*GDP growth:* -5.04%):

(a) Pre-Government Income

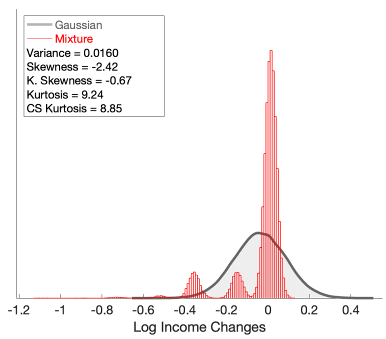

(b) Post-Government Income

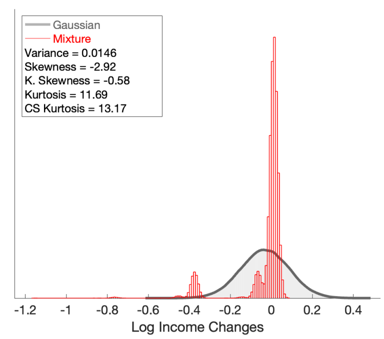

2009 (*GDP growth:* 6.59%):

(c) Pre-Government Income

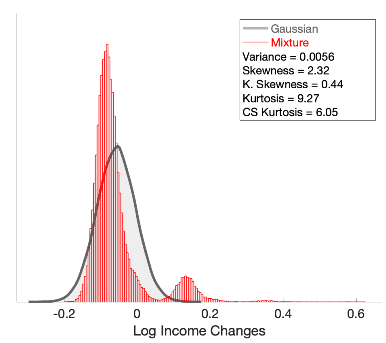

(d) Post-Government Income

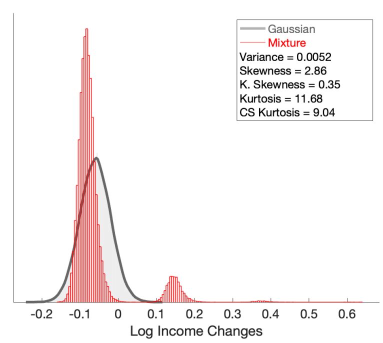

2010 (*GDP growth:* 2.49%):

(e) Pre-Government Income

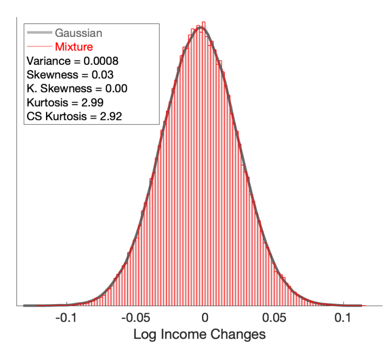

(f) Post-Government Income

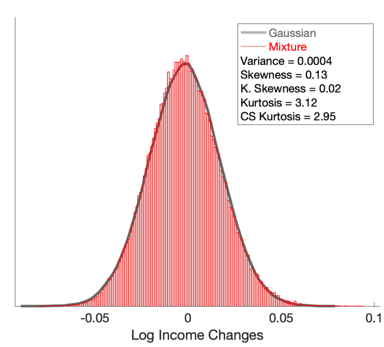

Note: Each figure shows the distribution of simulated (pre-government) permanent income changes η in red (histogram), as well as the density of the corresponding normal distribution with same mean and variance but with 0 Skewness and a Kurtosis of 3. The different figures correspond to three years that are representative of different states of the business cycle.

The tax and transfer system dampens the swings of skewness over the business cycle, which is captured in the parameter estimates for ϕ_2 and ϕ_3 in Table 1. This is reflected in the distributions plotted for years 2008 and 2009. Also for post-government income, $\mathcal{KS}$ changes from negative in 2008 to positive in 2009. However, the difference is less pronounced than for pre-government income. In 2008, $\mathcal{KS} = -0.58$ indicates that $(P90 - P50)$ accounts for 21% of the $(P90 - P10)$ dispersion. In 2009, $\mathcal{KS} = 0.35$ indicates that $(P90 - P50)$ accounts for 67.5%. Furthermore, the distribution is leptokurtic for both income measures in 2008 and 2009, with a somewhat higher kurtosis for post-government income, which implies that the tax and transfer system overall increases the concentration of the distribution.

To sum up, taxes and transfers, (i.), reduce overall dispersion of income changes, (ii.), reduce the cyclicality of dispersion and skewness, (iii.), increase concentration of income changes in both contractionary and expansionary years. The question we turn to now is: how do households value this?

3.2 Cross-Sectional Distribution Over the Life Cycle

From an ex-ante perspective, the distribution of possible income streams that can realize over the life cycle are relevant when it comes to the assessment of different risk scenarios. Given our assumption on full insurance against transitory shocks, the permanent shocks are relevant for the consumption distribution, and thus for welfare. Those shocks accumulate and generate a distribution that widens as a cohort ages.

We now consider a cohort of agents that lives through the Swedish macroeconomic history captured by the process of x_t that is used in the estimation of the income process. This way, we obtain a distribution of possible paths of the permanent income component for individuals entering the Swedish economy in year 1978, who subsequently receive stochastic income according to the estimated income process. For a given degree of partial insurance, we can then generate the corresponding consumption paths by simulating forward using equation (3).

Consider the blue line in panel (a) in Figure 2: it shows how the variance of (the model-constructed) cross-sectional permanent income component of pre-government income evolves for the cohort living through the Swedish macroeconomic history. During the contractions of the early 1990s and the late 2000s, the distribution of shocks becomes more dispersed, and thus the increase of the cross-sectional variance gets steeper. Panels (c) and (d) show that this increase in contractions happens stronger in the lower tail, which reflects an asymmetric swing of the distribution, that also manifests itself in the evolution of cross-sectional skewness, which is shown in panel (b): it tends to get more negative in contractions, and more positive in expansions.

Figure 2: Cross-Sectional Distribution of Permanent Income

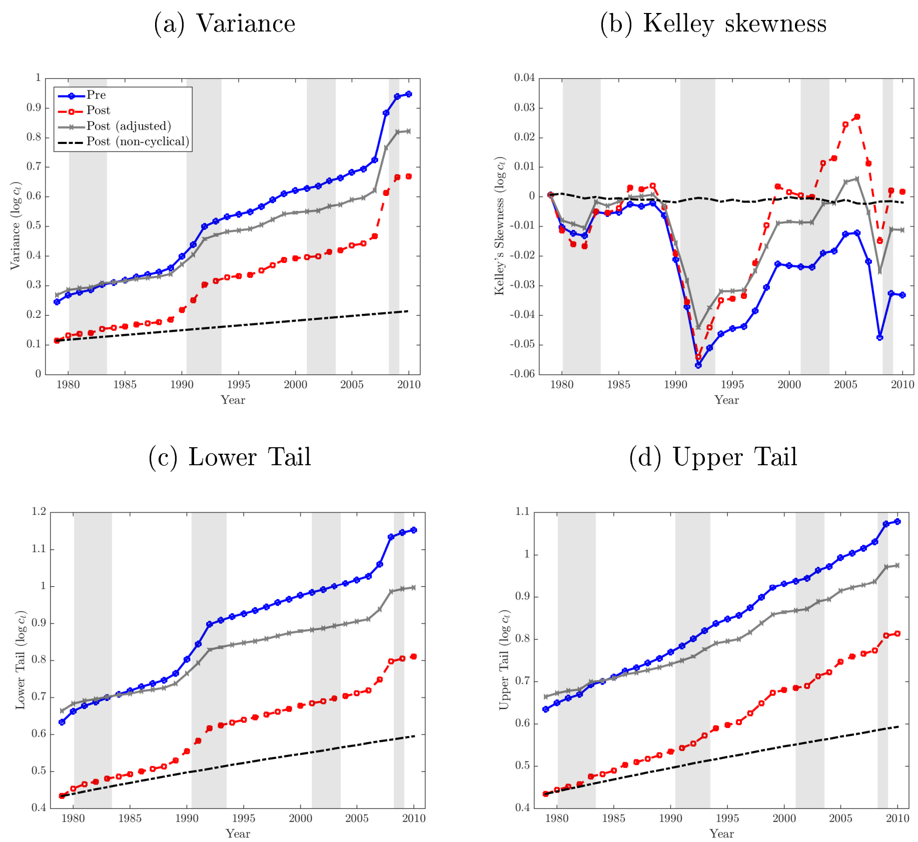

Note: Each figure shows a moment of the simulated cross-sectional distribution of permanent income for a cohort that lives through the Swedish macroeconomic history and faces, (i), the estimated pre-government income process; (ii), the estimated post-government income process; (iii), the post-government income process adjusted for initial variance; or, (iv), the post-government income process that eliminates all cyclicality of the distribution of shocks.

In each of the four panels of Figure 2, the red line reports the cross-sectional moments of the permanent component of post-government income for the same cohort. In line with the discussion of the estimated permanent income change component in the previous section, the first key difference is that the overall dispersion at every age is smaller (see panel a). Second, in the years leading up to the recession of the early 1990s, the asymmetry as measured by Kelley's skewness behaves very similarly; in the subsequent recovery Kelley's skewness of post-government income gets less and less negative and turns positive around the mid-2000s.

4 Measured Insurance Value of Taxes and Transfers

4.1 Overall Level of Insurance

We now employ the model measure derived in Section 2.2 to derive the degree of partial insurance against permanent income risk, λ_{perm}^{pre}, implied by the tax and transfer system. Thus, in line with the description in Section 2.2, we the goal is to find the λ_{perm}^{pre}, which yields a consumption stream that makes households indifferent to facing the post-government income stream—with a given amount of partial insurance when facing the latter.

For a given λ_{perm}^{pre}, we scale the estimated parameters of the permanent shocks such that the variance of the resulting distribution for η_t^{idio} is equal to fraction λ_{perm}^{post} of the overall variance of the permanent shock η. The scaling is such that the shape of the distribution as captured by the coefficient of skewness remains the same. We normalize such that $E\left[exp\left(\eta^{island}\right)\right] = E\left[exp\left(\eta^{idio}\right)\right] = 1$. Under log utility we find $\lambda_{perm}^{pre} = 0.43$, which means that the existing tax and transfer schedule in Sweden corresponds to insuring households against 43% of permanent shocks to household labor income, as shown in Table 2.

In order to assess the magnitude of this degree of partial insurance in terms of welfare, we use the model to calculate the consumption equivalent variation (CEV) that makes agents in the scenario with the *pre-government income stream and no partial* *insurance* indifferent to the world with the *pre-government income stream and partial insurance of the size given by* λ^{pre}_{perm}. The 43% partial insurance translates into a CEV of 14.3% when assuming log utility. Hence, the existing tax and transfer system provides sizable insurance. Note that this calculation abstracts from any first-order effects: both a potential level effect of the tax and transfer system on the aggregate income of a given cohort and the cyclical variation in average income changes are taken out of the equation.

4.2 Decomposition of Insurance Channels

Initial dispersion. When interpreting these results, it is important to notice that government policy reduces the overall level of cross-sectional dispersion, and the cyclicality of shocks. In order to differentiate those two smoothing effects, we impose in a second run of the same experiment that the cross-sectional variance at age 25 (when agents are born in the model) is the same as for the pre-government process. The moments of the resulting permanent income process are shown as the gray lines in Figure 2. We now obtain $\lambda^{pre}_{perm} = 0.06$, i.e., moving from the pre- to the post-government income stream adjusted to the same initial variance amounts to partial insurance of 6%, which translates into a CEV of about 1.3%.

Gain of eliminating cyclicality. Given the already sizable insurance, what is the scope of additional government policy as a means of insurance against cyclical risk? In order to approach this question, we consider the same experiment for a counterfactual income process. Assume that on top of what the government already does, cyclicality is completely shut down for the post-government income stream. For this experiment, we set $\phi_2 = \phi_3 = 0$, thus imposing the distribution of idiosyncratic income changes that corresponds to periods of average GDP growth. This yields the profiles of cross-sectional moments shown by the dashed lines in Figure 2. This implies an even stronger degree of insurance of about 64% (or 27% when adjusting for initial variance at age 25).

Considering the CEV connected to those insurance parameters, the scope of additional insurance is sizable: through the lens of the model, when adjusting for initial variance effects, an additional welfare gain of about 4.6 percentage points is possible.

Table 2: Partial Insurance and Welfare Gains of the Tax and Transfer System

Scenario	λ^{pre}_{perm}	CEV	λ^{pre}_{perm} (cycl.)	CEV (cycl.)
			log utility	
Pre to Post	**43%**	**14.26%**	**6%**	**1.28%**
Gaussian	*43%*	*15.52%*	*7%*	*2.97%*
Pre to Post*	**64%**	**17.53%**	**27%**	**5.91%**
	65%	*20.60%*	*29%*	*11.15%*
			$CRRA\ w/\ Risk\ Aversion = 2$	
Pre to Post	**36%**	**32.65%**	**5%**	**3.03%**
	42%	*34.80%*	*8%*	*7.23%*
Pre to Post*	**66%**	**46.34%**	**34%**	**19.13%**
	65%	*47.57%*	*30%*	*25.80%*

Note: The term λ^{pre}_{perm} denotes the degree of partial insurance against permanent shocks. * indicates that the cyclicality of the permanent shocks is shut down. See text for details on the scenarios. The CEV columns denote the corresponding consumption equivalent variation associated with the change from the world with the pre-government income stream and no partial insurance to a world with the pre-government income stream and partial insurance of the size given by λ^{pre}_{perm}.

Role of higher-order moments. In the estimation of the income process we were careful to match not only the dispersion of income changes, but also measures of skewness and kurtosis, i.e., higher-order moments of the distributions of individual income changes over the business cycle. As discussed in Section 3, those moments capture salient features of how the distribution varies over the business cycle, as it becomes more left-skewed in contractions. Thus, the next question we ask is whether for our model measure of partial insurance it is relevant to take those higher-order moments (and their cyclical changes) into account or not.

Thus, we now reconsider the exercise, but now assume that agents are exposed to Gaussian earnings processes that share the first and second-moment properties with the estimated pre- and post-government income processes, respectively, but have zero skewness and a kurtosis of 3. Notably, the variance still co-moves with the aggregate state of the economy, as it does in the benchmark case. Note that this implies that the dispersion evolves as displayed in panel (a) of Figure 2, but Kelley's skewness is zero throughout.

The gray rows in Table 2 show the results that correspond to the exact same exercises as in the benchmark analysis, but for the Gaussian shock distributions. There are two take-aways. First, the measured insurance values (and their reflections in CEVs) are of roughly the same magnitude for the overall insurance value of taxes and transfers. Second, the insurance gain against cyclical risk translates into about twice under a Gaussian distribution (2.97% vs. 1.28%). Thus, not taking into account skewness and kurtosis of the distribution of idiosyncratic risk, one would overestimate the insurance value of the existing tax and transfer system. Likewise, the potential additional gain of completely eliminating cyclical variation of idiosyncratic risk is about twice as high: a total gain of 11.15% vs. a total gain of 5.91%.

Role of risk attitudes. So far, we made the assumption that agents have log utility (relative risk aversion of 1). Preferences that feature a constant relative risk aversion larger than 1 are widely used in macroeconomics, and in incomplete market models in particular. The bottom half of Table 2 reports the results for the case of a parameter of relative risk aversion of 2, a standard value. In the context of the analysis it is important to note that this parameter pins down relative risk attitidues also against higher-order risk, which are relevant in order for skewness and kurtosis of the distribution to matter for utility (see detailed discussions in, e.g., Eeckhoudt, 2012; Busch and Ludwig, 2023). Three patterns emerge. First, the insurance value of the tax and transfer system against total earnings is in general smaller than under risk aversion of 1. Second, however, the CEV of insuring income risk is larger. Third, when focusing on the cyclical component of earnings shocks, both the insurance and welfare gains from taxes and transfers are larger than in the benchnmark counterpart. The importance of taking into account higher-order moments (vs. a Gaussian distribution) holds for the stronger risk attitudes.

Role of full pass-through of post-government income. In our benchmark analysis, we derive the consumption profile for households facing the post-government income stream under the assumption of no further partial insurance, i.e., $\lambda^{post}_{perm} = 0$. Given this assumption, we then derive the degree of partial insurance that delivers a consumption stream that makes households indifferent when they face the pre-government income stream. We now explore robustness of the approach with respect to this assumption. For this, we assume that instead, 10% of permanent shocks to post-government income are insured. This delivers a slightly somewhat less dispersed consumption profile. We then evaluate the degree of partial insurance against pre-government income that makes households indifferent; and also repeat the same additional calculations we did for the benchmark case. Results are reported in Table 3.

The obtained partial insurance parameters λ^{pre}_{perm} now combine both, the partial insurance provided by the tax and transfer system, and the additional partial insurance that comes from other insurance channels. Therefore, of course, the obtained λ^{pre}_{perm} reported in Table 3 are larger than the ones reported in the benchmark exercise of Table 2. In order to back out the degree of partial insurance that is provided by taxes and transfers, note that $(1-\lambda^{post}_{perm})$ is exactly equal to the ratio between the variance of (log) consumption growth and the variance of (log) permanent shocks—see equation (6). Thus, we scale up the consumption variance obtained under pre-government income with partial insurance λ^{pre}_{perm} accordingly and obtain the government-provided insurance as

$$\lambda^{gov} = 1 - \frac{1-\lambda^{pre}_{perm}}{1-\lambda^{post}_{perm}} = \frac{\lambda^{pre}_{perm} - \lambda^{post}_{perm}}{1-\lambda^{post}_{perm}} \tag{13}$$

Table 3: Partial Insurance and Welfare Gains of the Tax and Transfer System ($\lambda^{pos}_{perm} = 0.1$)

Scenario	λ^{gov}	λ^{pre}_{perm}	CEV	λ^{gov}(cycl.)	λ^{pre}_{perm} (cycl.)	CEV (cycl.)
				log utility		
Pre to Post	**43%**	**49%**	**15.13%**	‘tbf7%	**16%**	**3.29%**
Gaussian	*43%*	*49%*	*16.33%*	*7%*	*16%*	*4.88%*
Pre to Post*	**64%**	**68%**	**18.09%**	**28%**	**35%**	**7.53%**
	64%	*68%*	*20.92%*	*29%*	*36%*	*12.36%*
				CRRA w/ Risk Aversion = 2		
Pre to Post	**37%**	**43%**	**35.87%**	**6%**	**15%**	**8.35%**
	42%	*48%*	*36.80%*	*8%*	*17%*	*11.32%*
Pre to Post*	**66%**	**69%**	**47.81%**	**33%**	**40%**	**22.81%**
	66%	*69%*	*48.34%*	*30%*	*37%*	*28.46%*

Note: The term λ^{pre}_{perm} denotes the degree of partial insurance against permanent shocks. * indicates that the cyclicality of the permanent shocks is shut down. See text for details on the scenarios. The CEV columns denote the corresponding consumption equivalent variation associated with the change from the world with the pre-government income stream and no partial insurance to a world with the pre-government income stream and partial insurance of the size given by λ^{pre}_{perm}.

Note that for $\lambda^{post}_{perm} = 0$ (our benchmark case), equation (13) implies that $\lambda^{gov} = \lambda^{pre}_{perm}$. For $\lambda^{post}_{perm} = 0.1$, we show the resulting values for λ^{gov} alongside their λ^{pre}_{perm} counterparts in Table 3. Up to rounding error the obtained measures for partial insurance provided by the tax and transfer system are effectively idential to the ones obtained in the benchmark case.

5 Conclusion

The tax and transfer system partially insures households against individual income risk. We discuss under which assumptions differences between income processes estimated for household gross income and disposable income are informative about the (welfare) value of this partial insurance. Our approach works directly with income processes estimated separately on the two income measures, and does not require the specification (nor estimation) of a tax function. Instead we use an incomplete markets framework that links an estimated income process to consumption. Its key feature is that the degree of partial insurance is directly parameterized: Technically, this allows to solve for the degree of insurance provided by the tax and transfer system as a fixed point. The approach works with standard restrictions on income processes and preferences, and it further enables us to explore the role of higher-order risk for the value assigned to public insurance.

Through the lens of our structural model, the degree of overall insurance amounts to 43%, corresponding to 14% in consumption-equivalent terms under log-utility in Sweden. After isolating the gains from a lower initial variance at age 25, the degree of partial insurance amounts to 6% (CEV of about 1.3%). However, the remaining risk in post-government household-level income is still substantial. If cyclical variation of risk was completely eliminated, the partial insurance value would amount to 64%, or a CEV of 16.5%—and thus individuals would be better off by about 3 percentage points of consumption equivalent variation. While the partial insurance value of public insurance is very similar against *skewed* and *symmetric* income risk, the corresponding CEV gain would be overstated (3%—more than twice as large—against the cyclical component) if the pro-cyclicality in skewness of idiosyncratic risk is ignored.

References

Aiyagari, S. R. (1995). Optimal Capital Income Taxation with Incomplete Markets, Borrowing Constraints, and Constant Discounting. *Journal of Political Economy*, **103**, 1158–1175.

Bénabou, R. (2002). Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency? *Econometrica*, **70** (2), 481–517.

Blundell, R., Graber, M. and Mogstad, M. (2014). Labor income dynamics and the insurance from taxes, transfers, and the family. *Journal of Public Economics*.

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

—, — and Saporta-Eksten, I. (2016). Consumption inequality and family labor supply. *American Economic Review*, **106** (2), 387–435.

Busch, C., Domeij, D., Guvenen, F. and Madera, R. (2022). Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance. *American Economic Journal: Macroeconomics*, **14** (2), 207–42.

— and Ludwig, A. (2023). Higher-Order Income Risk over the Business Cycle. *International Economic Review (forthcoming)*.

Constantinides, G. M. and Duffie, D. (1996). Asset Pricing with Heterogeneous Consumers. *Journal of Political Economy*, **104** (2), 219–240.

Crow, E. L. and Siddiqui, M. M. (1967). Robust estimation of location. *Journal of the American Statistical Association*, **62** (318), 353–389.

De Nardi, M., Fella, G., Knoef, M., Paz-Pardo, G. and Van Ooijen, R. (2021). Family and government insurance: Wage, earnings, and income risks in the Netherlands and the U.S. *Journal of Public Economics*, **193**, 104327.

De Nardi, M., Fella, G. and Paz-Pardo, G. (2020). Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare. *Journal of the European Economic Association*, **18** (2), 890–926.

Eeckhoudt, L. (2012). Beyond risk aversion: Why, how and what's next. *The Geneva Risk and Insurance Review*, **37** (2), 141–155.

Floden, M. and Linde, J. (2001). Idiosyncratic risk in the united states and sweden: Is there a role for government insurance?1. *Review of Economic Dynamics*, **4**, 406–437.

Guvenen, F. (2011). Macroeconomics with heterogeneity: A practical guide. *Federal Reserve Bank of Richmond Economic Quarterly*, **97** (3), 255–326.

—, Ozkan, S. and Song, J. (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy*, **122** (3), 621–660.

Heathcote, J., Storesletten, K. and Violante, G. L. (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review*.

Huggett, M. (1993). The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies. *Journal of Economic Dynamics and Control*, **17**, 953–969.

Imrohoroglu, A. (1989a). Cost of Business Cycles with Indivisibilities and Liquidity Constraints. *Journal of Political Economy*, **97(6)**, 1364–1383.

— (1989b). Cost of business cycles with indivisibilities and liquidity constraints. *Journal of Political Economy*, **97** (6), 1364–1383.

Kaplan, G. and Violante, G. L. (2010a). How Much Consumption Insurance beyond Self-Insurance? *American Economic Journal: Macroeconomics*, **2** (4), 53–87.

— and — (2010b). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics*, **2** (4), 53–87.

Krueger, D. and Perri, F. (2006). Does income inequality lead to consumption inequality? evidence and theory. *Review of Economic Studies*, **73** (1), 163–193.

Krusell, P., Mukuyama, T., Sahin, A. and Smith, A. A. (2009). Revisiting the welfare effects of eliminating business cycles. *Review of Economic Dynamics*, **12**, 393–404.

Lucas, R. E. (1987). *Models of Business Cycles*. New York: Basil Blackwell.

McKay, A. (2017). Time-varying idiosyncratic risk and aggregate consumption dynamics. *Journal of Monetary Economics*, **88**, 1–14.

— and Reis, R. (2016). The role of automatic stabilizers in the u.s. business cycle. *Econometrica*, **84** (1), 141–194.

Storesletten, K., Telmer, C. I. and Yaron, A. (2001). The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk. *European Economic Review*, **45** (7), 1311–1339.

A Estimated Income Processes

The figures show the estimated income processes for pre- and post-government household income along with the data counterparts of the targeted set of moments.

Figure 3: Pre-Government Income Fit

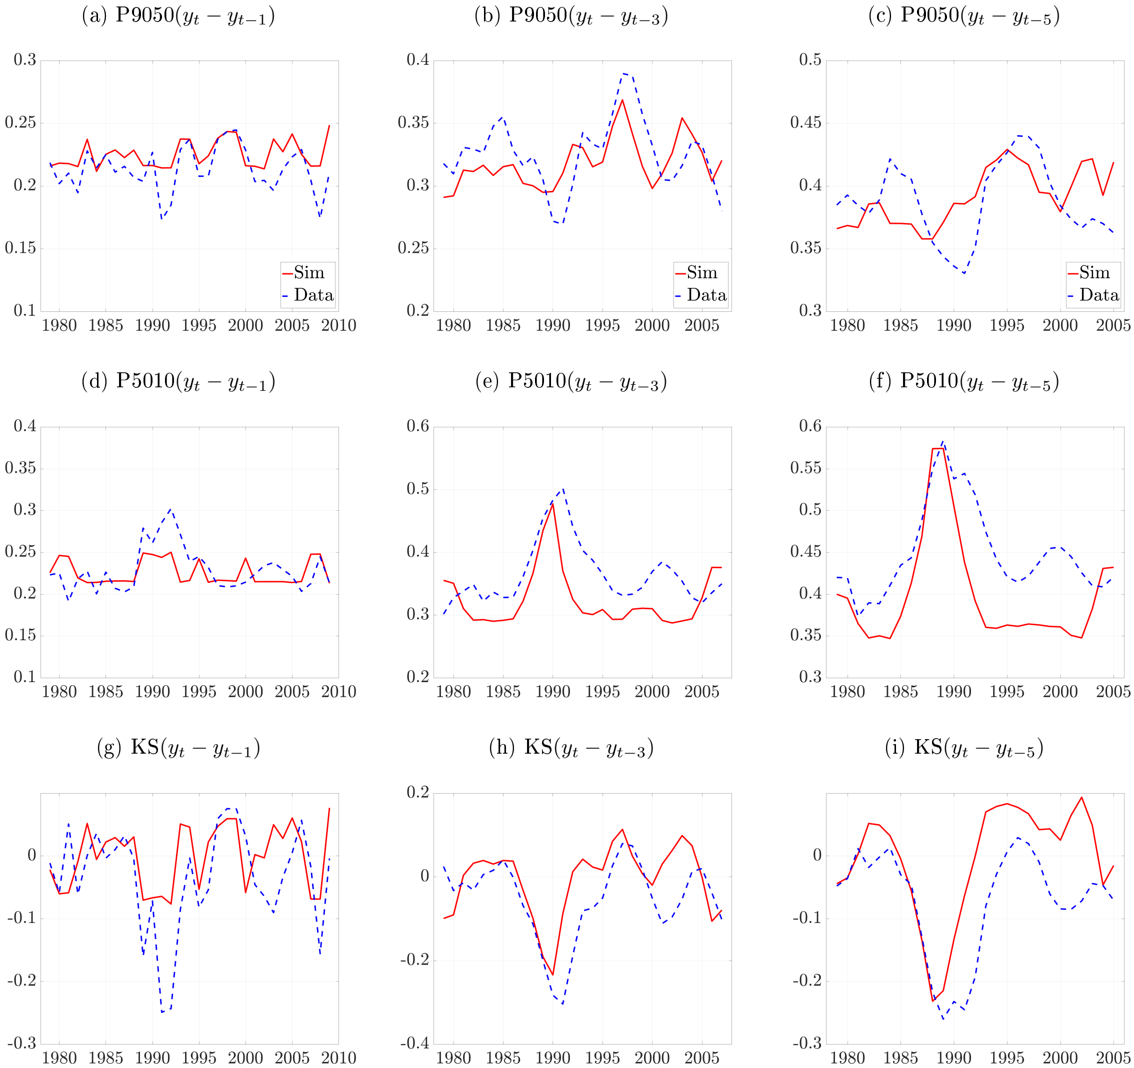

Note: Each panel shows the time series of a moment of short-run, medium-run, or long-run income changes together with the corresponding moment implied by the estimated income process.

Figure 4: Post-Government Income Fit

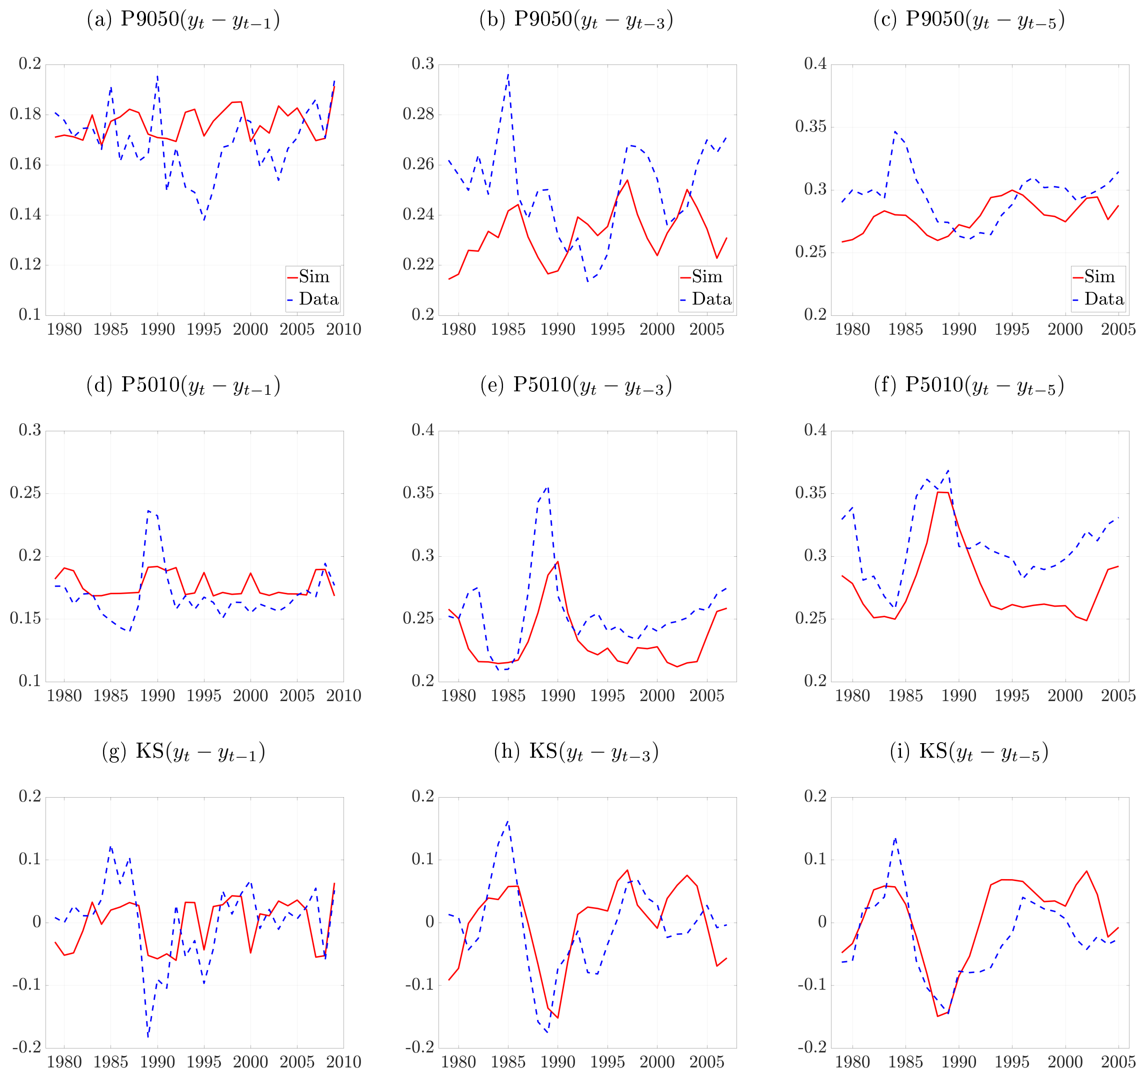

Note: See notes to Figure 3.

B Scaling Income Processes

Given estimates of the income process, we scale the parameters of the permanent shocks η to feed them into the model; fraction λ is insurable and the rest is uninsurable. This scaling implies that the first three standardized moments of the distribution of insurable shocks are given as below: for the first three moments of the uninsurable shocks, simply replace λ with $1 - \lambda$.

$$
\begin{aligned}
E\left[\eta_t^{idio}\right] &= \sum_{i=1}^{3} p_{\eta,i}\mu_{\eta^{idio},i,t} = \sum_{i=1}^{3} p_{\eta,i}\lambda^{1/2}\mu_{\eta,i,t} = \lambda^{1/2}\sum_{i=1}^{3} p_{\eta,i}\mu_{\eta,i,t} = \lambda^{1/2}E\left[\eta_t\right] \equiv \lambda^{1/2}\mu_{\eta,t} \\
var\left[\eta_t^{idio}\right] &= \sum_{i=1}^{3} p_{\eta,i}\left(\sigma^2_{\eta^{idio},i} + \mu^2_{\eta^{idio},i,t}\right) - \left(E\left[\eta_t^{idio}\right]\right)^2 = \sum_{i=1}^{3} p_{\eta,i}\left(\lambda\sigma^2_{\eta,i} + \lambda\mu^2_{\eta,i,t}\right) - \left(\lambda^{1/2}E\left[\eta_t\right]\right)^2 \\
&= \lambda\left(\sum_{i=1}^{3} p_{\eta,i}\left(\sigma^2_{\eta,i} + \mu^2_{\eta,i,t}\right) - E\left[\eta_t\right]^2\right) = \lambda var\left[\eta_t\right] \\
skew\left[\eta_t^{idio}\right] &= \frac{1}{var\left[\eta_t^{idio}\right]^{3/2}} \sum_{i=1}^{3} p_{\eta,i}\left(\mu_{\eta^{idio},i,t} - E\left[\eta_t^{idio}\right]\right)\left[3\sigma^2_{\eta^{idio},i} + \left(\mu_{\eta^{idio},i,t} - E\left[\eta_t^{idio}\right]\right)^2\right] \\
&= \frac{1}{\lambda^{3/2}var\left[\eta_t\right]^{3/2}} \sum_{i=1}^{3} p_{\eta,i}\left(\lambda^{1/2}\mu_{\eta,i,t} - \lambda^{1/2}E\left[\eta_t\right]\right)\left[3\lambda\sigma^2_{\eta,i} + \left(\lambda^{1/2}\mu_{\eta,i,t} - \lambda^{1/2}E\left[\eta_t\right]\right)^2\right] \\
&= \frac{1}{\lambda^{3/2}var\left[\eta_t\right]^{3/2}} \sum_{i=1}^{3} p_{\eta,i}\lambda^{1/2}\left(\mu_{\eta,i,t} - E\left[\eta_t\right]\right)\left[\lambda\left(3\sigma^2_{\eta,i} + \left(\mu_{\eta,i,t} - E\left[\eta_t\right]\right)^2\right)\right] \\
&= \frac{1}{var\left[\eta_t\right]^{3/2}} \sum_{i=1}^{3} p_{\eta,i}\left(\mu_{\eta,i,t} - E\left[\eta_t\right]\right)\left[\left(3\sigma^2_{\eta,i} + \left(\mu_{\eta,i,t} - E\left[\eta_t\right]\right)^2\right)\right] \\
&= skew\left[\eta_t\right]
\end{aligned}
$$