The Insurance Value of Public Insurance Against Idiosyncratic Income Risk*

Christopher Busch[†] Rocio Madera[‡]

April 14, 2025

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: This allows us 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, partial insurance, social insurance policy, incomplete markets.

*We thank David Domeij and Fatih Guvenen for detailed discussions, as well as 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

1 Introduction

Most individuals face individual income risk: firm closures, unemployment spells, or health shocks are just some of the sources of downside risk, which together with typical career ups and downs lead to volatile income trajectories. The institutional framework in many economies offers a mix of various policy instruments that jointly cushion disposable income against this volatility. Some of these policies are explicitly designed as a buffer for specific sources of downside risk: most notably, the unemployment insurance system dampens temporary income losses due to job loss. Similarly, disability insurance insures against severe health risk. Progressivity of the income tax system implies insurance from an ex ante perspective as it compresses the distribution of possible income changes, without targeting explicit reasons for income losses (or gains). Existing evidence for various countries documents that, on average, the overall tax and transfer schemes, in tandem with private insurance mechanisms, are successful in providing partial insurance against individual risk (see, e.g., Blundell *et al.*, 2014; De Nardi *et al.*, 2021).

In this paper, we assess the *insurance value* of the existing tax and transfer system in Sweden. We do so in an extensive panel data set from tax register data, and develop an approach that is generally applicable in a setting where data is available on household-level incomes before and after taxes and transfers. This is a typical setting as many rich (administrative) data sources allow for the detailed exploration of household income trajectories, but do not cover equally reliable data on consumption or consumption expenditures.

Approach. Our approach rests on an analytical model framework, which allows us to parameterize the degree of partial insurance. We account for three main features of household-level income, and for the difference between gross and net incomes along these features. First, income risk is in part transitory and in part permanent—with compressed distributions for net incomes. Second, income risk is distributed asymmetrically, where positive and negative income changes of the same magnitude are not equally likely—with this asymmetry less pronounced for net incomes. Third, income risk changes systematically over the business cycle—with less pronounced swings for net incomes. Those features (or a subset of the list) are well-documented for a large set of diverse countries (see, e.g., Blundell *et al.*, 2014; Busch *et al.*, 2022). We then use our model framework to interpret the differences between gross and net income. While we exploit systematic differences between the two income measures, we do so in a flexible manner that does not require the specification of a parametric function to link gross and net income.

Overview. The first step of our analysis is thus the formulation of a statistical income process that captures these data aspects. The second step is the setup of a model as a device for the measurement of the insurance value of taxes and transfers. In this context, we take the common perspective that consumption (not income directly) eventually translates into welfare. Given household-level income processes that capture the distributional regularities of gross income and net income, and given a measure of consumption, one can use some measurement device to trace out the degree of insurance coming from the tax and transfer system (as prominently discussed in Blundell *et al.*, 2008). In the absence of consumption data, we thus need to impose a model that maps disposable income into household consumption. To this end, we employ a consumption function which takes permanent and transitory income shocks as inputs, and which is parameterized by the degrees of insurance against the two types of shocks. The third step is to assign welfare (and thus an insurance value) to the obtained insurance measure, which requires the specification of preferences.

In sum, there are three main elements of the analysis: (i.) an income process which gives a distinction between transitory and permanent components, (ii.) a consumption function which translates transitory and permanent shocks into consumption, and (iii.) a preference specification which allows for a welfare interpretation. In our approach, which we sketch next, we explicitly derive the consumption function of step (ii.) from the same model that we then use to assign a welfare value in step (iii.).

Our Measurement Device. We adopt an island model structure that features the abstraction of two distinct types of risks—one perfectly insurable (within island), and one uninsurable (island-level) (a la Heathcote *et al.*, 2014). The model features an equilibrium with no asset trade across islands. We then use the fact that the ratio of island-level shocks to total shocks is exogenous, which makes the degree of partial insurance a key parameter of our measurement device: For a given degree of partial insurance, an (exogenous) income process maps into an (endogenous) consumption process.

We use this feature of the model to trace out the degree of partial insurance provided by the tax and transfer system in a flexible way that does not require us to parametrically restrict this insurance, e.g., through specification of a tax function. In particular, within the model, we consider households that face an income process that captures regularities of pre-government earnings. We then find the degree of partial insurance that they need to receive, in order to be indifferent to instead facing the post-government earnings process—with some degree of partial insurance against post-government risk given. 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 measurement device for the degree of insurance coming from taxes and transfers, which takes as inputs (estimated) income processes for gross income and net income, and makes a minimal set of structural assumptions—on preferences and on the degree of partial insurance beyond taxes and transfers.

Application to Sweden. We apply our approach to estimated income processes for income moments from Swedish tax data that covers the period 1976–2011. 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).[1] 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 (CEV: 1.3%). However, 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. We find that under 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

[1]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 (2024), who illustrate how central moments of the distribution map into choices of agents in a life-cycle model.

against skewness and kurtosis are relatively weak with log utility. Still, insurance against cyclical variation of risk is valued twice as much under Gaussian shocks. Thus, not taking into account 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.

After a short literature discussion that follows, the rest of the paper is organized as follows. Section 2 builds up the meaurement of partial insurance from established definitions. Section 3 outlines the quantitative model used as our measurement device. Section 4 first introduces the income process, which disentangles transitory and permanent components of income, and serves as the input for the measurement. It then goes on to discuss the measured insurance value of taxes and transfers in Sweden, and Section 5 concludes.

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 redistribution and insurance against 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 gain insights into the welfare implications of tax and transfer systems for mitigating the pass-through of aggregate fluctuations to individual income.

Our study of cyclical risk links our analysis 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, 1989a; 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 (2024), who explore the role of remaining higher-order risk in a quantitative model.

We build on standard incomplete market models (SIM), which feature two key ingredients: exogenous idiosyncratic productivity risk, and access to only partial insurance against it through savings in a riskfree asset (cf. Aiyagari, 1995; Huggett, 1993; Imrohoroglu, 1989b). Relative to this, our model is agnostic with respect to the exact source of insurance. The abstraction to capture partial insurance through an island structure a la Heathcote *et al.* (2014) enables us to derive the consumption function analytically given a degree of partial insurance. In this sense, we relate to a strand of the life-cycle literature that focuses on a bundle of self- and family-insurance channels beyond the traditional savings instruments (Blundell *et al.*, 2008; Krueger and Perri, 2006; Kaplan and Violante, 2010). In contrast to Heathcote *et al.* (2014), 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.

2 Partial Insurance by the Public Insurance System

Before introducing our full measurement framework, it is useful to briefly discuss established definitions of partial insurance. Theoretical measures and their empirical counterparts involve linking consumption changes to individual income changes. In the absence of complete markets that allow for full insurance against idiosyncratic risk, individual risk is partially insurance against. If shocks were directly observable in the data, one could pin down *pass-through* coefficients β_{perm} and β_{trans} using OLS regressions of consumption changes on permanent and transitory shocks, respectively:

$$\beta_{perm} = \frac{\operatorname{cov}(\Delta \ln c_{it}, \eta_{it})}{\operatorname{var}(\eta_{it})}, \qquad \beta_{trans} = \frac{\operatorname{cov}(\Delta \ln c_{it}, \varepsilon_{it})}{\operatorname{var}(\varepsilon_{it})}. \tag{1}$$

Given these coefficients, the parameters of partial insurance are given by

$$\lambda_{perm} = 1 - \beta_{perm} \tag{2}$$

$$\lambda_{trans} = 1 - \beta_{trans}. \tag{3}$$

Given that transitory and persistent shocks are not observable in the data, empirical evidence builds around specifying consumption as a function of income shocks, and estimating this consumption function together with a stochastic income process (cf., Blundell *et al.*, 2008). The pass-through coefficients β_1 and β_2 are then identified together with the variances of the shocks from a set of population moments. Blundell *et al.* (2008) estimate pass-through coefficients for different measures of income in the Panel Study of Income Dynamics. Of particular interest to us are the estimations using household-level gross income and disposable income, respectively.

These measures of pass-through for the two income measures are directly linked in the following sense: consumption is a function of disposable income, and thus consumption reacts to changes in gross income through adjustments of disposable income. In other words, the total pass-through from gross income to consumption combines the pass-through from gross income to disposable income with the pass-through from disposable income to consumption. To structure this, think of a tax function a la Bénabou (2000) and Bénabou (2002), as more recently used by, e.g., Heathcote *et al.* (2017), where net tax revenues at income level y are given by $T(y) = y - \phi y^{1-\tau}$, so that disposable income is given by:

$$y^{disp} = \phi y^{1-\tau}. \tag{4}$$

The progressivity parameter, τ,[2] directly translates into the elasticity of disposable income with respect to gross income, $1 - \tau$: $\Delta \ln y^{disp} = (1 - \tau)\Delta \ln y$, and thus

$$\frac{\text{cov}(\Delta \ln c, \Delta \ln y^{disp})}{\text{var}(\Delta \ln y^{disp})} = \frac{(1-\tau)\text{cov}(\Delta \ln c, \Delta \ln y)}{(1-\tau)^2 \text{var}(\Delta \ln y)}. \tag{5}$$

[2]If $\tau > 0$, marginal tax rates exceed average rates and hence the tax and transfer system is considered progressive. Conversely, when $\tau = 0$ the tax and transfer scheme is flat.

This implies that $\lambda = 1 - (1-\tau)(1-\lambda^{disp})$. It is useful to consider some bounds for common reference values:

$$\lambda = \begin{cases} \tau & \text{if } \lambda^{disp} = 0 \text{ (no self-insurance)} \\ 1 & \text{if } \lambda^{disp} = 1 \text{ (full self-insurance)} \end{cases} \tag{6}$$

where λ^{disp} is the degree of partial consumption insurance against shocks to disposable income.

In other words, if agents are able to fully self-insure, then the degree of public insurance is irrelevant. If agents have no ability to self-insure (they are hand-to-mouth), then total insurance is equal to public insurance, which is exactly equal to the degree of progressivity.[3] Previous studies show that most agents are somewhere in between (e.g., Blundell *et al.*, 2008, estimate a $\lambda^{disp} = 0.64$ using panel data for the United States from the PSID).

Estimation of the pass-through coefficients uses panel data on all three measures: gross income, disposable income, and consumption. While administrative sources of income, before and after taxes and transfers, have become widely available in recent years, data on consumption is still scarce and subject to measurement issues.[4] The sketched tax function on the other hand links pre- and post-tax incomes. This implies that, without resorting to consumption data, it captures pass-through from gross to net income, and as such can serve as the basis of evaluating the degree of insurance from taxes. However, this advantage in terms of a lower data requirement comes with strong parametric assumptions regarding exactly this link: The variance of changes is scaled by $(1-\tau)^2$, see (5), while, e.g., the skewness of changes is the same. In contrast, in our analysis, we flexibly capture distributional features of gross income changes and disposable income changes without posing any such restriction on the relationship between these income measures and without resorting to consumption data.

[3] Of course, generally the degree of self-insurance is endogenous with respect to existing policies that determine the degree of public insurance.

[4] In survey data, measurement error and low frequency pose challenges. In administrative data, imputations are required. In bank records, samples are rarely representative. And, in all cases, a pervasive measurement issue regardless of the source is the disconnect between expenditures and consumption, particularly serious for durable consumption.

3 A Quantitative Model as Measurement Device

3.1 The 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.

We use the perpetual youth framework of Yaari (1965) and assume that each period a mass $(1-\delta)$ of newborns enters the economy with age 0. At any age, the probability of survival to the next period is constant at $\delta \in (0,1)$. Individual income (endowment) of agent i in period t is given by

$$\begin{aligned} y_{i,t} &= y_{i,t}^{island} + y_{i,t}^{idio} \\ y_{i,t}^{x} &= z_{i,t}^{x} + \varepsilon_{i,t}^{x}, \quad \varepsilon_{i,t}^{x} \sim F_{\varepsilon,t}^{x} \\ z_{i,t}^{x} &= z_{i,t-1}^{x} + \eta_{i,t}^{x}, \quad \eta_{i,t}^{x} \sim F_{\eta,t}^{x}, \quad \text{for } x \in \{island, idio\} \end{aligned} \tag{7}$$

where $z_{i,t}^{x}$ and $\varepsilon_{i,t}^{x}$ for $x \in \{island, idio\}$ denote the island-level and idiosyncratic permanent and transitory components of income. All stochastic components of income are independent and normalized such that $\int \exp\left(\eta_{i,t}^{x}\right) dF_{\eta,t}^{x} = 1$ for $x \in \{island, idio\}$ and likewise for ε. 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 $\left\{\eta_{i,t}^{island}, \varepsilon_{i,t}^{island}\right\}_{t=\tau}^{\infty}$.

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: $U(c_{i,t}) = \ln(c_{i,t})$. 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.

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. 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, in our model, there is partial insurance: island-level shocks remain uninsured, while within-island shocks are fully insured, and risk is shared by all individuals on an island. This mimics the result in Heathcote *et al.* (2014).

In this equilibrium with no trade across islands, the period t log consumption of an agent i of age $a_{i,t}$, with income components $(y_{i,t}^{island}, y_{i,t}^{idio})$ are given by[5]

$$\ln c_{i,t}\left(a_{i,t}, y_{i,t}^{island}, y_{i,t}^{idio}\right) = y_{i,t}^{island} + \ln \int \exp\left(y^{idio}\right) dF_{y^{idio},t}^{a_{i,t}}. \tag{8}$$

The main information carried by this consumption equation is that the individual realization of the island-level income component is consumed, while, instead, all agents consume the mean realization of the idiosyncratic income component. Here, the distribution of this idiosyncratic component depends on both time and age, as is captured by $F_{y^{idio},t}^{a_{i,t}}$. It depends on time t, because the cross-sectional distributions of $\varepsilon_{i,t}^{idio}$ and $\eta_{i,t}^{idio}$ depend on t; it further depends on age a, because the permanent shocks $\eta_{i,t}^{idio}$ accumulate over age, resulting in a widening distribution of the permanent component $z_{i,t}^{idio}$. The consumption equation also 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 analyt-

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

ical solution in which consumption can be expressed explicitly as a function of idiosyncratic shocks. That is, given an endowment process, we can directly calculate the consumption level (and changes) implied by the model.

Degree of partial insurance. We now consider the model equivalent of the pass-through coefficient a la Blundell *et al.* (2008) in equation (1) to capture insurance against transitory and permanent shocks, respectively. Within the model, we make the common assumption that agents can observe transitory and permanent shocks directly.[6]

The consumption function translates into consumption change

$$\begin{aligned}\Delta \ln c_{i,t}\left(a_{i,t}, y_{i,t}^{island}, y_{i,t}^{idio}\right) &= \Delta y_{i,t}^{island} + \ln \frac{\int \exp\left(\eta_{i,t}^{idio}\right) dF_{\eta,t}^{idio} \int \exp\left(\varepsilon_{i,t}^{idio}\right) dF_{\varepsilon,t}^{idio}}{\int \exp\left(\varepsilon_{i,t-1}^{idio}\right) dF_{\varepsilon,t-1}^{idio}} \\ &= \eta_{i,t}^{island} + \Delta\varepsilon_{i,t}^{island} + \ln \frac{\int \exp\left(\eta_{i,t}^{idio}\right) dF_{\eta,t}^{idio} \int \exp\left(\varepsilon_{i,t}^{idio}\right) dF_{\varepsilon,t}^{idio}}{\int \exp\left(\varepsilon_{i,t-1}^{idio}\right) dF_{\varepsilon,t-1}^{idio}} \\ &= \eta_{i,t}^{island} + \Delta\varepsilon_{i,t}^{island},\end{aligned} \tag{9}$$

where $\Delta \ln c_{i,t} = \ln c_{i,t} - \ln c_{i,t-1}$ and $\Delta\varepsilon_{i,t}^{island} = \varepsilon_{i,t}^{island} - \varepsilon_{i,t-1}^{island}$.

The relevant model version of partial insurance against permanent and transitory shocks builds around the pass-through to the combined *island* and *idio*-shocks, i.e., to $\varepsilon_{i,t} = \varepsilon_{i,t}^{idio} + \varepsilon_{i,t}^{island}$ and $\eta_{i,t} = \eta_{i,t}^{idio} + \eta_{i,t}^{island}$. As is clear from (9), the *island*-shock translates one-for-one to consumption—the pass-through of shock to consumption is one—and the *idio*-shock does not translate into consumption—the pass-through of shock to consumption is zero. The pass-through of the combined shock is then a convex combination of these two. Directly applying (1), we obtain

$$\begin{aligned}1 - \lambda_{trans} &= \frac{\operatorname{cov}(\Delta \ln c_{i,t}, \varepsilon_{i,t})}{\operatorname{var}(\varepsilon_{i,t})} = \frac{\operatorname{cov}(\Delta\varepsilon_{i,t}^{island}, \varepsilon_{i,t})}{\operatorname{var}(\varepsilon_{i,t})} = \frac{\operatorname{cov}(\varepsilon_{i,t}^{island} - \varepsilon_{i,t-1}^{island}, \varepsilon_{i,t}^{island} + \varepsilon_{i,t}^{idio})}{\operatorname{var}(\varepsilon_{i,t})} \\ &= \frac{\operatorname{var}(\varepsilon_{i,t}^{island})}{\operatorname{var}(\varepsilon_{i,t}^{island} + \varepsilon_{i,t}^{idio})} = \frac{\operatorname{var}(\varepsilon_{i,t}^{island})}{\operatorname{var}(\varepsilon_{i,t}^{island}) + \operatorname{var}(\varepsilon_{i,t}^{idio})},\end{aligned} \tag{10}$$

[6] For example, Kaplan and Violante (2010) make the same assumption when studying partial insurance within a standard incomplete markets model.

and

$$1-\lambda_{perm} = \frac{\operatorname{cov}(\Delta \ln c_{i,t}, \eta_{i,t})}{\operatorname{var}(\eta_{i,t})} = \frac{\operatorname{cov}(\eta_{i,t}^{island}, \eta_{i,t})}{\operatorname{var}(\eta_{i,t})} = \frac{\operatorname{cov}(\eta_{i,t}^{island}, \eta_{i,t}^{island} + \eta_{i,t}^{idio})}{\operatorname{var}(\eta_{i,t})} \tag{11}$$
$$= \frac{\operatorname{var}(\eta_{i,t}^{island})}{\operatorname{var}(\eta_{i,t}^{island} + \eta_{i,t}^{idio})} = \frac{\operatorname{var}(\eta_{i,t}^{island})}{\operatorname{var}(\eta_{i,t}^{island}) + \operatorname{var}(\eta_{i,t}^{idio})},$$

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, λ_{trans}, is given by the fraction of the variance of transitory shocks attributable to the *idio*-component. This way, it becomes clear that the *island* and *idio* shocks serve as an abstraction that allows to capture partial insurance.

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 (7) 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 $x \in \{\varepsilon, \eta\}$. This means that we consider a tax and transfer system that cross-sectionally redistributes endowments, and rules out wasteful government consumption or debt-financed transfer payments.

3.2 Measurement of the Insurance Value of Taxes and Transfers

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 (9).

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). As there are two types of shocks, in principle multiple combinations of $\{\lambda^{pre}_{trans}, \lambda^{pre}_{perm}\}$ can exist that make agents indifferent. We assume that $\lambda^{pre}_{trans} = \lambda^{post}_{trans} = 1$, i.e., 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; instead all transitory shocks happen purely within islands, and thus can be insured away fully by agents trading state-contingent claims. Note that incomplete market models typically find very high insurance against transitory shocks through private savings alone (Busch and Ludwig, 2024; De Nardi *et al.*, 2020; Kaplan and Violante, 2010), which makes overall full insurance a plausible assumption.

This leaves partial insurance against permanent risk as the relevant margin of the model. Consider agents born in period τ. When they face the stochastic income stream $y^{pre} = \{y^{pre}_{i,t}\}^{\infty}_{t=\tau}$ with a degree of partial insurance λ^{pre}_{perm} this translates into stochastic streams of *idio* and *island* components $\{y^{pre,island}_{i,t}, y^{pre,idio}_{i,t}\}^{\infty}_{t=\tau}$. The two components are such that the implied distribution of their sum, $(y^{pre,island}_{i,t} + y^{pre,idio}_{i,t})$, corresponds to the distribution of the total income $y^{pre}_{i,t}$. Likewise, in the alternative scenario they face income streams $\{y^{post,island}_{i,t}, y^{post,idio}_{i,t}\}^{\infty}_{t=\tau}$ which are consistent with income stream $y^{post} = \{y^{post}_{i,t}\}^{\infty}_{t=\tau}$ and partial insurance λ^{post}_{perm}. With the two income streams, y^{pre} and y^{post}, as well as insurance λ^{post}_{perm} at hand, we then find the level of partial insurance λ^{pre}_{perm} that makes agents ex ante indifferent:

$$\mathbb{E}_{\tau} \sum_{t=\tau}^{\infty} (\beta\delta)^{t-\tau} U\left(c_t(a_{i,t}, y^{pre,island}_{i,t}, y^{pre,idio}_{i,t})\right) = \mathbb{E}_{\tau} \sum_{t=\tau}^{\infty} (\beta\delta)^{t-\tau} U\left(c_t(a_{i,t}, y^{post,island}_{i,t}, y^{post,idio}_{i,t})\right). \tag{12}$$

4 The Insurance Value of Taxes and Transfers in Sweden

4.1 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, which we take from Busch *et al.* (2018). The data moments are calculated using longitudinal data on household earnings changes from LINDA for the period 1979-2010. LINDA is compiled from administrative sources (the Income Register) and tracks a representative sample with approximately 300,000 individuals per year. Gross (pre-government) income at the household level includes earnings from labor and capital income, while net (post-government) income adds transfers and taxes. The measure of net (post-government) earnings lumps together four main groups of public programs that are consistently measured over time. The groups are (1) labor-market-related policies, (2) aid to low-income families, (3) pension payments, and (4) taxes.

Labor-market-related policies mainly consist of unemployment benefit payments. Busch *et al.* (2022) show that this component of social insurance policy is particularly important for mitigating cyclical variation of downside household earnings risk. Aid to low-income families encompasses family support, housing assistance, and direct cash transfers from the public sector. These transfers are particularly important to stabilize the earnings of low-income households, who are more likely to meet the criteria for receiving such aid during recessions. Pension payments, although not directly influenced by business cycles, can impact households with members close to or at retirement age. These individuals might opt for pension benefits instead of unemployment benefits if they choose to retire after losing their job. Taxes include income taxes on both labor and capital income, but taxes paid on capital income constitute a small part of total tax payments.

Let y_t^{pre} and y_t^{post} denote log of pre- and post-government household income, respectively. For each of the two income measures, we separately fit the following permanent-transitory process (where we drop the explicit reference to pre or post-government income):

$$\begin{aligned} y_t &= z_t + \varepsilon_t \\ z_t &= z_{t-1} + \eta_t \end{aligned} \tag{13}$$

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:

$$\varepsilon_t \sim \begin{cases} \mathcal{N}(\bar{\mu}_\varepsilon, \sigma^2_{\varepsilon,1}) & \text{with prob. } p_{\varepsilon,1} \\ \mathcal{N}(\bar{\mu}_\varepsilon, \sigma^2_{\varepsilon,2}) & \text{with prob. } 1 - p_{\varepsilon,1} \end{cases} \tag{14}$$

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

$$\eta_t \sim \begin{cases} \mathcal{N}(\bar{\mu}_{\eta,t} + \mu_{\eta,1} + \phi_1 x_t, \sigma^2_{\eta,1}) & \text{with prob. } p_{\eta,1} \\ \mathcal{N}(\bar{\mu}_{\eta,t} + \mu_{\eta,2} + \phi_2 x_t, \sigma^2_{\eta,2}) & \text{with prob. } p_{\eta,2} \\ \mathcal{N}(\bar{\mu}_{\eta,t} + \mu_{\eta,3} + \phi_3 x_t, \sigma^2_{\eta,3}) & \text{with prob. } p_{\eta,3} \end{cases} \tag{15}$$

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 $\mathbb{E}[\exp(\eta_t)] = 1$. 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

$$\chi_{trans} = \{\sigma_{\varepsilon,1}, \sigma_{\varepsilon,2}, p_{\varepsilon,1}\} \tag{16}$$

$$\chi_{perm} = \{\mu_{\eta,2}, \mu_{\eta,3}, \sigma_{\eta,1}, \sigma_{\eta,2}, p_{\eta,1}, p_{\eta,2}, \phi_2, \phi_3\} \tag{17}$$

by the simulated method of moments (SMM).[7] We target the time series of L9050 and L5010[8] 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 pre-government income, and 213 moments for post-government income, which we use as targets in the estimation of the two income processes for Sweden.

To construct the simulated time series of income growth moments, we write earnings growth as a function of the shocks, using equation (13):

$$y_t - y_{t-s} = \varepsilon_t - \varepsilon_{t-s} + \sum_{j=0}^{s-1} \eta_{t-j}, \tag{18}$$

for different horizons $s = 1, 3, 5$, and then calculate the relevant statistical moments of these distributions. To construct the simulated life-cycle variance profile, we use a time-invariant distribution of shocks by imposing $x_t = 0\ \forall t$. We then normalize the series and rescale it such that the resulting simulated variance profile exhibits the same mean as its empirical counterpart.

[7]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 $p_{\eta,2} = p_{\eta,3}, \sigma_{\eta,2} = \sigma_{\eta,3}$.

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

We simulate individual 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 life cycle 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 $m \in \{L5010(\Delta^1 y_t), L5010(\Delta^3 y_t), L5010(\Delta^5 y_t), \ldots, var(y_{age=25}), \ldots, var(y_{age=60})\}$. In each simulation, we draw a matrix of random variables $X_r = \{\varepsilon_1^i, \varepsilon_2^i, \ldots, \varepsilon_T^i, \eta_1^i, \ldots, \eta_T^i\}_{i=1}^{I}$ 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 a histogram of the simulated distribution for the estimated mixture of Normals corresponding to pre-government (blue, filled) and post-government (red, solid border). In the plots, we use the normalization such that $E[\exp(\eta_t)] = 1$. For completeness, Figures 3 and 4 in Appendix A.2 show the simulated moments at these parameters together with the empirical moments over time.

Table 1: Estimated Parameter Values

Parameter	Description	Pre-Gov.	Post-Gov.
$p_{\varepsilon,1}$	Mixture prob. of ε distribution	0.860	0.864
$\sigma_{\varepsilon,1}$	Std. dev. of ε distribution mix. comp. 1	0.039	0.042
$\sigma_{\varepsilon,2}$	Std. dev. of ε distribution mix. comp. 2	0.489	0.384
$p_{\eta,1}$	Mixture prob. of η distribution mix. comp. 1	0.983	0.985
$p_{\eta,2}$	Mixture prob. of η distribution mix. comp. 2	0.009	0.007
$p_{\eta,3}$	Mixture prob. of η distribution mix. comp. 3	0.009	0.007
$\sigma_{\eta,1}$	Std. dev. of η distribution mix. comp. 1	0.084	0.057
$\sigma_{\eta,2}$	Std. dev. of η distribution mix. comp. 2	0.034	0.084
$\sigma_{\eta,3}$	Std. dev. of η distribution mix. comp. 3	0.034	0.084
$\mu_{\eta,2}$	Mean of mixt. comp. 2 of η distribution	0.022	0.015
$\mu_{\eta,3}$	Mean of mixt. comp. 3 of η distribution	-0.214	-0.070
ϕ_2	Aggregate risk transmission mixt. comp. 2	1.922	1.972
ϕ_3	Aggregate risk transmission mixt. comp. 3	0.534	0.341
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 figure, the distribution of permanent income changes varies over the cycle in an asymmetric way for both measures of income (pre- and post-government). In about average growth times (as from 2010 to 2011), 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, and the difference between the 90th and 10th percentiles ($P90 - P10$) decreases by 30%. Strong negative GDP growth (as from 2008 to 2009) goes hand-in-hand with a left-skewed distribution, while strong positive GDP growth (as from 2009 to 2010) comes with a right-skewed distribution.

The right skewness in an expansionary year (2009–10) is captured by a positive coefficient of skewness (the third standardized moment); and the mirror image holds true for a contractionary year (2008–9). This sign difference also shows in measures of Kelley's skewness, which is based on the 10th, 50th, and 90th percentiles of the distribution:

Figure 1: Cross-Sectional Distribution of Permanent Income Changes

(a) 2008–9 (*GDP growth:* −5.04%)

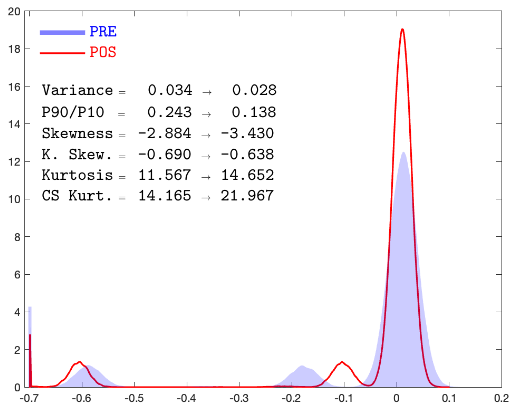

(b) 2009–10 (*GDP growth:* 6.59%)

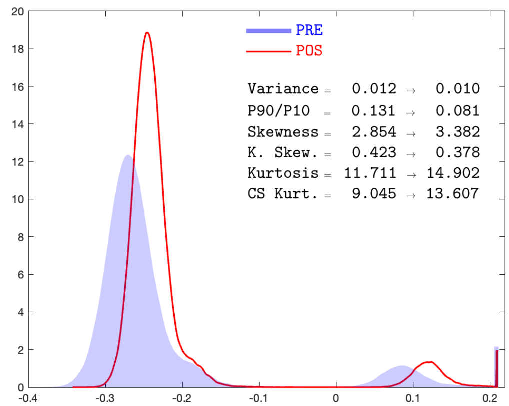

(c) 2010–11 (*GDP growth:* 2.49%)

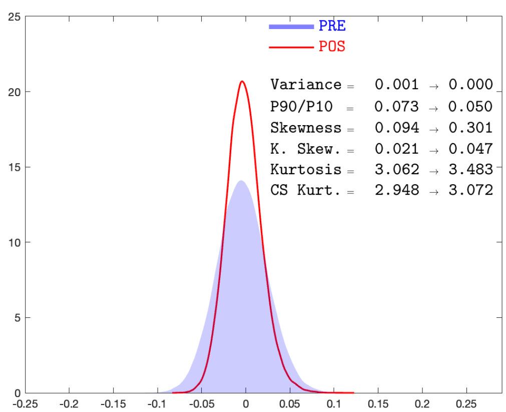

Note: Each figure shows the distribution of simulated pre-government permanent income changes η in blue (shaded without border) and post-government in red (solid outer border). The different figures correspond to three years that are representative of different states of the business cycle. The second, third, and fourth standardized moments of the distribution are reported alongside robust percentile-based measures to capture dispersion, asymmetry, and concentration. For readability, we truncate the density plots below -0.7 (panel (a)) and above 0.21 (panel (b)); we then add histogram bars that report the cumulative density below or above the thresholds, respectively.

$\mathcal{KS} = ((P90 - P50) - (P50 - P10))/(P90 - P10)$. $\mathcal{KS}$ 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 $\mathcal{KS} = -0.690$ for 2008–9 indicates that $(P90 - P50)$

accounts for 15.3% of the $(P90 - P10)$ dispersion.[9] On the other hand, in the growth period from 2009–10, the value of $\mathcal{KS} = 0.423$ indicates that $(P90 - P50)$ accounts for about 71% of the $(P90 - P10)$ dispersion.

The tax and transfer system dampens the pass-through of the business cycle to the distribution, 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 2008–9 and 2009–10. Also for post-government income, $\mathcal{KS}$ changes from negative in 2008–9 to positive in 2009–10. However, the difference is less pronounced than for pre-government income. In 2008–9, $\mathcal{KS} = -0.638$ indicates that $(P90 - P50)$ accounts for about 18% of the $(P90 - P10)$ dispersion. In 2009, $\mathcal{KS} = 0.378$ indicates that $(P90 - P50)$ accounts for 69%. 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, and (iii.), increase concentration of income changes in both contractionary and expansionary years. The question we turn to now is: how do households value this?

4.2 Measures of Insurance

We now make use of the structure outlined in Section 3 and feed it with the estimated income processes for the two income measures. In particular, we consider a range of possible values of insurance against permanent income risk after taxes and transfers, λ_{perm}^{post}, and use Equation (12) to back out λ_{perm}^{pre}. We start with the case $\lambda_{perm}^{post} = 0$, and return to the implications of this later. When $\lambda_{perm}^{post} = 0$, the obtained λ_{perm}^{pre} measures the degree of partial insurance provided by the government under the assumption that there is no additional partial insurance. This assumption can be 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.

[9] Note that $(P90 - P50)/(P90 - P10) = 0.5 + \mathcal{KS}/2$.

From an ex-ante perspective, the distribution of possible consumption 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 income shocks faced by agents translate into this consumption distribution, and thus matter for welfare. These shocks accumulate and generate a distribution that widens as a cohort ages. In addition, the idiosyncratic shock distributions are estimated to vary with the aggregate state of the economy, which itself is risky. Thus, we consider two complementary measures. The first measure builds around a cohort of agents that lives through the Swedish macroeconomic history captured by the time series of x_t that is used in the estimation of the income process. In this sense, it takes an ex-post perspective. For this cohort, we construct a set of simulated income and consumption profiles, which we translate into average life-cycle utility. The second measure takes an ex-ante perspective also regarding the aggregate state, for which we ad a stochastic process. In particular, we fit an AR(1) process, and then use the estimated process when constructing the insurance measure, which considers ex-ante expected life-cycle utility.

Let us first turn to a cohort that enters the Swedish economy in year 1979 (the first year for which the micro data for the estimation is available), and then lives through the macroeconomic history experienced until 2011. The income process (13) with the parameters reported in Table 1 implies a distribution of possible paths of the permanent income component. 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

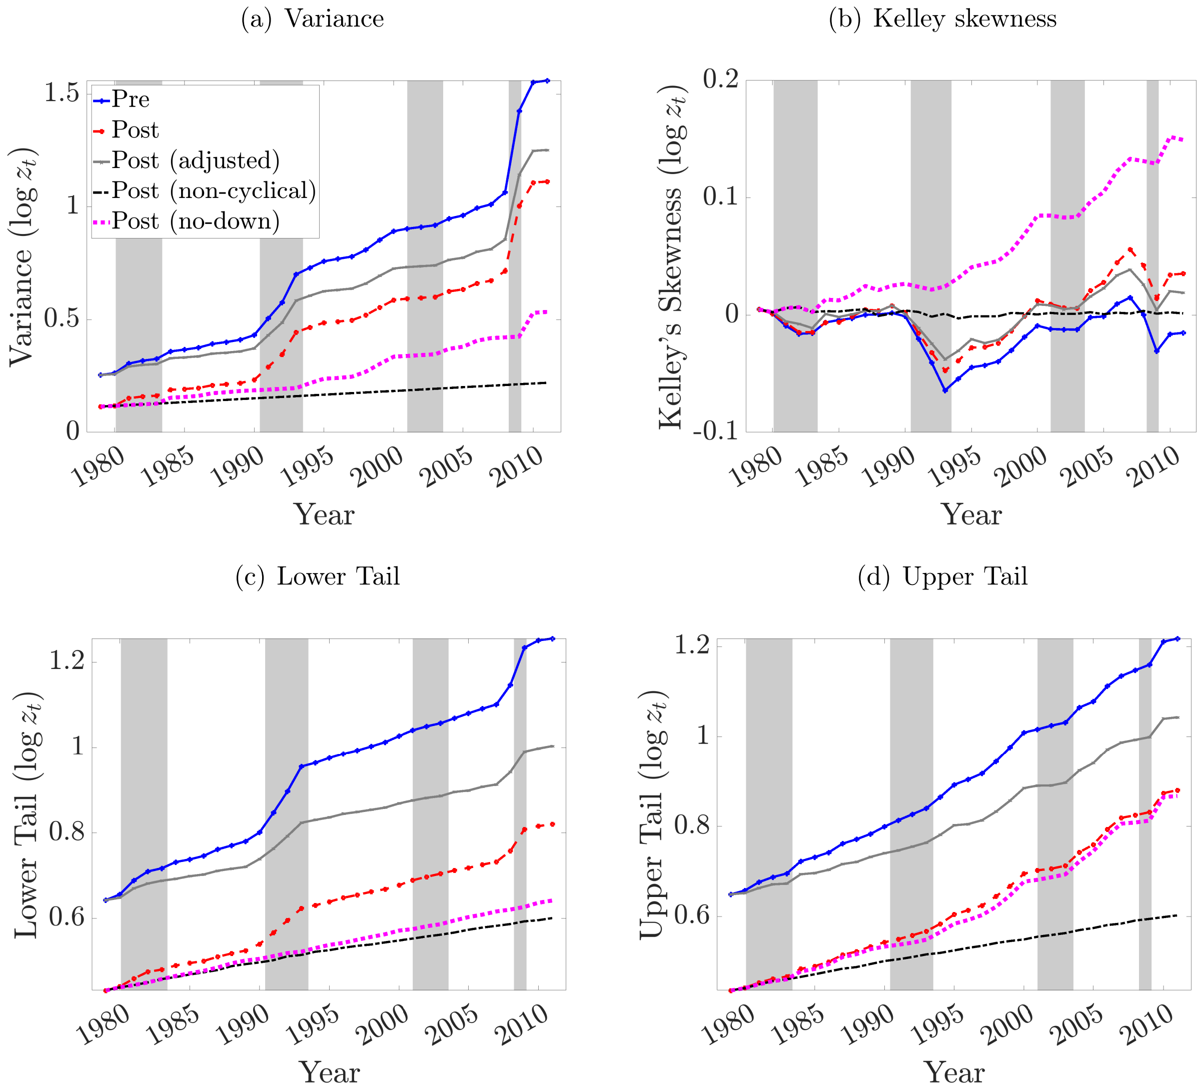

Note: Each figure shows a moment of the simulated cross-sectional distribution of permanent income for a cohort that lives through the Swedish macroeconomic history and faces, (i), the estimated pre-government income process; (ii), the estimated post-government income process; (iii), the post-government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

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.

We now employ the model measure derived in Section 3.2 to derive the degree of partial insurance against permanent income risk, λ^{pre}_{perm}, implied by the tax and transfer system. Thus, in line with the description in Section 3.2, the goal is to find the λ^{pre}_{perm}, 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^{pre}_{perm}, we scale the estimated parameters of the permanent shocks such that the variance of the resulting distribution for η^{idio}_t is equal to fraction λ^{post}_{perm} 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 $E\left[exp\left(\eta^{island}\right)\right] = E\left[exp\left(\eta^{idio}\right)\right] = 1$.

[NOW: evaluation] Under log utility we find $\lambda^{pre}_{perm} = 0.43$, 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.

4.3 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.

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 4.1, 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 and 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. This implies that the dispersion evolves as displayed in panel (a) of Figure 2, but Kelley's skewness is zero throughout.

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

		Higher-Order Risk Present		Gaussian with same Variance	
	Scenario:	λ^{pre}	CEV	λ^{pre}	CEV
	From pre-government income to...	*ln* utility			
(I)	...post-government income (post)	38.49%	11.31%	39.69%	11.83%
(II)	...post adjusted for initial dispersion	13.15%	3.78%	14.84%	4.27%
(III)	...post w/o cyclicality	71.40%	21.59%	72.97%	22.82%
(IV)	...post w/o cyclicality and adjusted	45.23%	13.37%	48.10%	14.51%
(V)	...post w/o reaction to negative x_t	50.16%	14.89%	62.41%	19.22%
(VI)	...post w/o neg. reaction & adjusted	24.50%	7.12%	37.54%	11.16%
		CRRA w/ Risk Aversion = 2			
(I)	...post-government income (post)	18.34%	32.48%	37.83%	28.13%
(II)	...post adjusted for initial dispersion	8.71%	15.46%	16.06%	11.34%
(III)	...post w/o cyclicality	74.79%	123.16%	73.52%	59.33%
(IV)	...post w/o cyclicality and adjusted	55.50%	93.93%	50.27%	38.46%
(V)	...post w/o reaction to negative x_t	62.89%	105.28%	62.68%	49.35%
(VI)	...post w/o neg. reaction & adjusted	45.66%	78.39%	39.88%	29.79%

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}.

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 the CEV 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, 2024). 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

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%**	**7%**	**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}.

of (log) permanent shocks—see equation (11). 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

$$\lambda^{gov} = 1 - \frac{1 - \lambda^{pre}_{perm}}{1 - \lambda^{post}_{perm}} = \frac{\lambda^{pre}_{perm} - \lambda^{post}_{perm}}{1 - \lambda^{post}_{perm}} \tag{19}$$

Note that for $\lambda^{post}_{perm} = 0$ (our benchmark case), equation (19) 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. (2000). Unequal societies: Income distribution and the social contract. *American Economic Review*, **90** (1), 96–129.

— (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. (2018). *Asymmetric Business-Cycle Risk and Social Insurance*. NBER Working Papers 24569, National Bureau of Economic Research, Inc.

—, —, — and — (2022). Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance. *American Economic Journal: Macroeconomics*, **14** (2), 207–42.

— and Ludwig, A. (2024). Higher-Order Income Risk over the Business Cycle. *International Economic Review*, **65** (3), 1105–31.

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*.

—, — and — (2017). Optimal Tax Progressivity: An Analytical Framework*. *The Quarterly Journal of Economics*, **132** (4), 1693–1754.

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. (2010). 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.

Yaari, M. E. (1965). Uncertain Lifetime, Life Insurance, and the Theory of the Consumer. *The Review of Economic Studies*, **32 (2)**, 137–150.

A Estimation Details

A.1 Global Optimization Details

A.2 Data Fit of 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

(a) P9050$(y_t - y_{t-1})$

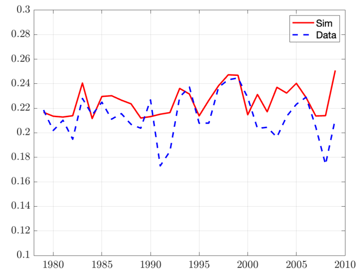

(b) P9050$(y_t - y_{t-3})$

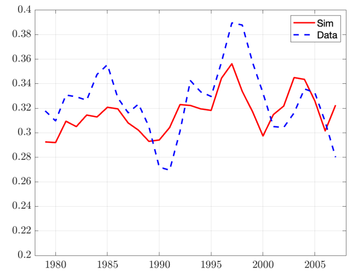

(c) P9050$(y_t - y_{t-5})$

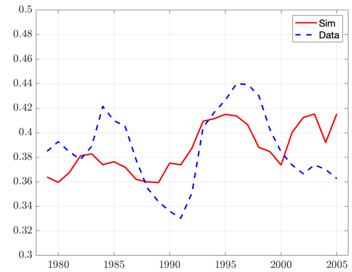

(d) P5010$(y_t - y_{t-1})$

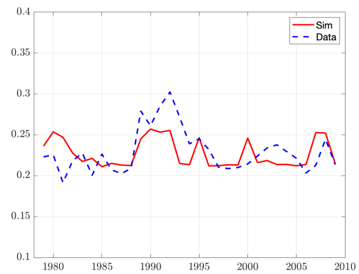

(e) P5010$(y_t - y_{t-3})$

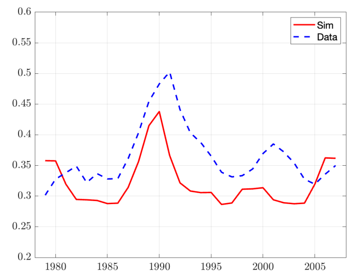

(f) P5010$(y_t - y_{t-5})$

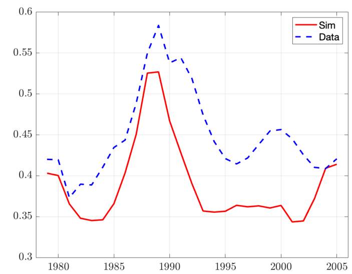

(g) KS$(y_t - y_{t-1})$

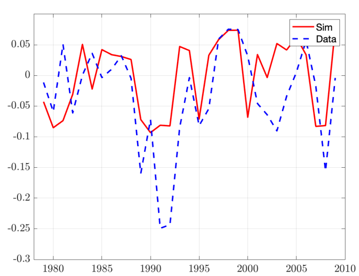

(h) KS$(y_t - y_{t-3})$

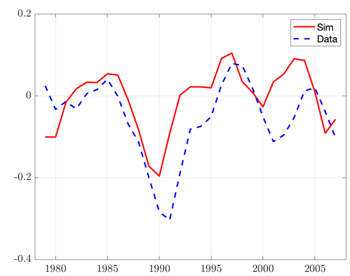

(i) KS$(y_t - y_{t-5})$

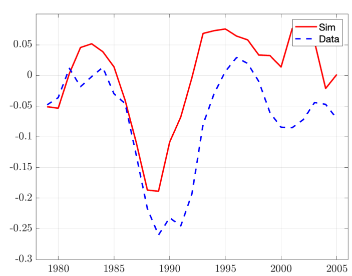

Note: Each panel shows the time series of a moment of short-run, medium-run, or long-run income changes together with the corresponding moment implied by the estimated income process.

Figure 4: Post-Government Income Fit

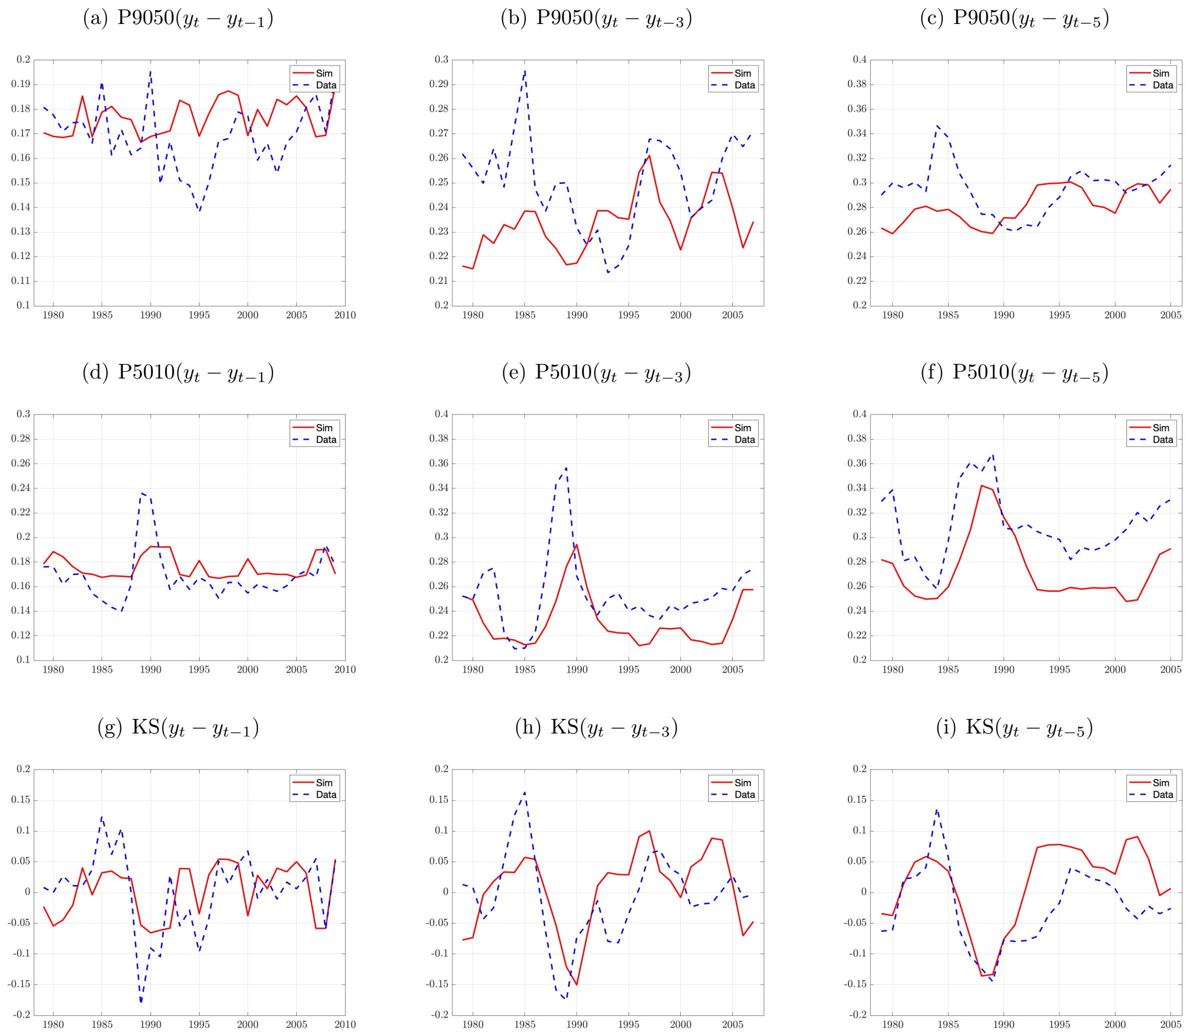

Note: See notes to Figure 3.

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 $1 - \lambda$.

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