

The Value of Public Insurance Against Idiosyncratic Income Risk: A Variance-Adjustment Statistic*

Christopher Busch

Rocio Madera

LMU Munich, CESifo and CEPR

Southern Methodist University and CESifo

April 2, 2026

Abstract

We introduce a tractable method built around an incomplete-markets model to assess the value of public insurance against permanent idiosyncratic income risk. Our approach translates statistical differences between gross and disposable incomes into a welfare-equivalent variance adjustment. Under homoskedastic Gaussian shocks, the variance ratio of permanent shocks to gross and disposable incomes provides a sufficient statistic for the size of insurance. More generally, with cyclical non-Gaussian shocks, public insurance amounts to a variance reduction by 38–49% in Sweden (tax registers) and 24–31% in the United States (PSID), depending on risk attitudes. Consumption-based measures in the PSID confirm our model-based measure.

Keywords: Tax and transfer system, partial insurance, incomplete markets.

*We thank David Domeij and Fatih Guvenen for detailed discussions in early stages of the project. For helpful comments we thank Jeanne Commault, Bulent Guler, Marcus Hagedorn, Burhan Kuruscu, Jonas Loebbing, Alex Ludwig, Ellen McGrattan, Sergio Ocampo, Gonzalo Paz-Pardo (discussant), Luis Pérez, Kjetil Storesletten, Mirko Wiederholt, and Peter Zorn, as well as participants at various conferences and seminars.

Busch: LMU Munich, Geschwister-Scholl-Platz 1, 80539 Munich, Germany;
christopher.busch@econ.lmu.de.

Madera: Southern Methodist University, Umphrey Lee 301, 3300 Dyer Street, Dallas, TX 75205, USA;
romadera@smu.edu.

1 Introduction

Household incomes are characterized by risky fluctuations. When households are risk-averse, riskiness of their income trajectories typically has welfare consequences. Among its many functions, public policy has an insurance role: it cushions disposable income against this volatility. Some policies are explicitly designed as a buffer for specific sources of downside risk—most notably, unemployment insurance dampens temporary income losses due to job loss. Others, like progressivity of the income tax system, compress the distribution of possible income changes without targeting explicit reasons for income losses (or gains). Existing evidence for various countries documents that the overall tax and transfer schemes are successful in buffering households’ disposable income against gross income fluctuations (see, e.g., [Blundell *et al.*, 2014](#); [De Nardi *et al.*, 2021](#)). This smoothing implies insurance from an ex ante perspective.

The goal of this paper is to assess the value of this insurance. For that purpose, it is crucial to gauge the extent to which household consumption expenditures are shielded from household income risk. Thus, assessing the degree of (consumption) insurance provided by the tax and transfer system requires information on income before taxes, income after taxes, as well as on consumption expenditures. Many rich (administrative) data sources allow for the detailed exploration of household income trajectories, but do not cover equally reliable panel data on consumption expenditures. We introduce a tractable method to assess the insurance value of the existing tax and transfer system in such a (typical) data setting: Our approach works with data on gross and net measures of household-level incomes, and translates statistical differences induced by taxes and transfers into a welfare-equivalent variance adjustment. We illustrate it using moments from extensive panel data based on Swedish tax registers and from the Panel Study of Income Dynamics for the United States.

At the core of our approach is a simple observation: while taxes and transfers affect the distributions of incomes and of income changes in various ways, one can describe it by means of a hypothetical adjustment of the variance that has the same insurance consequences. This way we translate potentially asymmetric (or more generally: non-Gaussian) features of the distribution into a symmetric measure (the variance). If households care more about

non-Gaussian features, then these translate stronger into the welfare-equivalent variance adjustment. Our focus on this variance adjustment builds on the empirical benchmark going back to [Blundell *et al.* \(2008\)](#), who introduce an insurance coefficient that exploits the covariance structure of income and consumption changes. By introducing structural assumptions, namely the specification of preferences over uncertain outcome streams and an incomplete asset market structure, we extend the capabilities of the framework to capture differences in statistical moments beyond the variance.

We propose an analytical model framework, in which households receive income shocks against which consumption can be partially insured on incomplete asset markets. We achieve analytical tractability by the abstraction of an island structure in the spirit of [Heathcote *et al.* \(2014\)](#) and more recently [Boerma and Karabarbounis \(2021\)](#): There are two types of shocks, some purely idiosyncratic and some common to all individuals that form an island. A full set of claims contingent on the purely idiosyncratic state is available, while no claims contingent on island-level outcomes can be traded across islands. This dichotomy gives rise to an equilibrium in which no claims are traded across islands, while all purely idiosyncratic shocks are perfectly insured against within islands. In equilibrium, consumption can be expressed in closed form as a function of the sequence of received income shocks.

Note that the ratio of the variance of island-level shocks to the total variance of the combined island-level and idiosyncratic shocks is exogenous. This property of the model allows us to parameterize the degree of partial insurance: We translate the shocks to an income process that captures total risk into two components, which capture (non-insurable) island-level risk and (fully insurable) within-island risk. For a given parameterized degree of partial insurance, the model thus maps an (exogenous) income process 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 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 expressed as a welfare-equivalent variance reduction. It 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.

We treat household income as the fundamental source of risk. Another possibility is to pose a process for wages (or for individual productivity) and to explicitly model labor supply endogenously (as, e.g., [Heathcote *et al.*, 2014](#)). Our choice is data-driven: similar to consumption, high-quality tax register data on incomes rarely comes with reliable measures of hours worked. We exploit the systematic differences between the two income measures in a flexible manner that does not require the specification of a parametric function to link gross and net income. Note that throughout we focus on the value of insurance against household-level risk. Formally, the welfare-equivalent variance reduction is computed holding household decisions fixed. This is a deliberate choice: We abstract from the costs of providing this insurance—distortions to labor supply and the financing of transfers—, as well as from redistribution and general equilibrium effects. In other words, we take household income risk as given and measure the share the government absorbs, rather than calculating optimal policy.

Sufficient statistic under Gaussian shocks. In line with basic intuition, we show that the ratio of the variance of permanent net income shocks to the variance of permanent gross income shocks is a sufficient statistic for the insurance value of taxes and transfers when the income process is Gaussian and shocks are homoskedastic. Our general framework, which accommodates non-Gaussian and heteroskedastic shocks, complements this analytical result with a richer measurement device, which we bring to the data in two applications.

Application to Sweden. We first apply our approach to Sweden based on statistical moments of tax register data that covers the period 1976–2011. Sweden provides a benchmark case for our approach because of the availability of high-quality data on relevant income measures, without the corresponding data on consumption expenditures. In an early study of the tax and transfer system in Sweden, [Floden and Linde \(2001\)](#) found large welfare gains from

insurance against (privately) 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 the tax and transfer system for mitigating the pass-through of aggregate fluctuations to individual income risk.

The first ingredient for the analysis is the formulation of a statistical income process that captures three main features of household-level income trajectories. We estimate two sets of parameters of the income 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. Our estimates corroborate results from the literature for a large set of diverse countries (see, e.g., Blundell *et al.*, 2014; Busch *et al.*, 2022; Guvenen *et al.*, 2014): 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.

We find that the degree of partial insurance provided by the tax and transfer system amounts to a welfare-equivalent variance reduction by about 48.6%, which translates into a welfare gain, expressed as a consumption equivalent variation (CEV), of about 11.8% under log utility. We then focus on the part of that gain that is attributable to smoothing typical shocks received over the life cycle. Taxes and transfers insure about 17.6% of these (CEV: 4.2%). However, remaining risk (in post-government household-level income) is still substantial: The variance reduction that is welfare-equivalent to moving to an acyclical income stream adjusted for initial dispersion amounts to 32.2%, i.e., an additional 14.6%—under log utility this corresponds to a 3.5% CEV to completely eliminate cyclical fluctuations of risk (in line with explorations of the welfare costs of business cycles in incomplete markets settings a la Imrohoroglu, 1989a; Storesletten *et al.*, 2001; Krusell *et al.*, 2009).

Application to the United States. We also apply our approach to a country for which consumption data exists. Specifically, we consider the Panel Study of Income Dynamics, which allows us to compare our model-based measure of insurance provided by the tax and transfer

system to an alternative based on pass-through estimates from the two income measures (pre- and post-government). In this illustration we find that the model-based measure aligns well with direct estimates: the various approaches yield a measure of government-provided insurance against permanent income risk that amounts to about 24–25%—markedly lower than what we find in the application to Sweden.

When parameterizing the model with stronger risk attitudes, we find that higher-order moments—and how taxes and transfers affect those—do translate into the welfare-equivalent variance adjustments in both Sweden and the United States. This delivers additional insights about the relative strength of tax and transfer schemes in reducing downside risk (typically through transfers) vis a vis upside chances (typically through progressive taxation). In the United States, the reduction of downside risk (which agents do like) dominates the overall measure. In Sweden, the dampening of upside chances dominates (which agents do not like).

Literature. Our model-based framework fits in between two distinct approaches in the literature. One strand of the literature resorts to panel data on income (before and after taxes), and consumption. This is the path taken in [Blundell *et al.* \(2008\)](#) (BPP), who fit a statistical model of income and consumption to data from the PSID. It is noteworthy that the consumption measure is imputed based on CEX data. In the absence of consumption data that matches the richness and reliability of income data, we instead impose a model structure that allows us to interpret the available income data. Note that the imposition of an income process is shared between the approach taken in [Blundell *et al.* \(2008\)](#) and ours. As the BPP coefficients exploit the covariance structure of income measures and consumption, they do not explicitly account for non-Gaussian features—our approach, on the other hand, translates non-Gaussian features into welfare-equivalent variance adjustments to the extent implied by risk attitudes of model agents. This *ex ante* perspective also makes our argument complementary to a recent discussion in [Ghosh and Theloudis \(2025\)](#), who explore limitations of the BPP approach by considering a quadratic approximation of the Euler equation, or [Busch and Ludwig \(2024\)](#) who analyze a non-linear consumption function in a structural model.

Another strand in the literature builds around fully structural models and interprets the data through the lens of standard incomplete markets (SIM) frameworks (going back to [Aiya-gari, 1995](#); [Huggett, 1993](#); [Imrohoroglu, 1989b](#)). These models explicitly limit the available means of private insurance and explicitly model specific mechanisms—most crucially, agents can build up savings in a riskfree asset ([Kaplan and Violante, 2010](#), were the first to show that in a calibrated version of the one-asset SIM model, insurance amounts to about two thirds of the empirical estimate from BPP). Instead, by using the abstraction of an island structure, our model is agnostic with respect to the exact source of insurance, and as a measurement tool captures multiple possible mechanisms. Further, within these model frameworks, the link between gross and net income is typically restricted to a specific functional form, in particular a tax function a la [Bénabou \(2002\)](#). Our approach does not require such a parametric restriction, and instead builds around separately estimated income processes for pre- and post-government income, which capture empirical regularities of these income measures.

The rest of the paper is organized as follows. [Section 2](#) builds up the measurement of partial insurance from established definitions and introduces the notion of a welfare-equivalent variance adjustment. [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. [Section 5](#) compares the model-based measure of government-provided insurance to a direct measure based on estimated consumption pass-through coefficients for pre- and post-government income. [Section 6](#) concludes.

2 Partial Insurance by the Public Insurance System

2.1 Pass-Through Measures of Insurance

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, individual risk is partially insured against. Focussing attention on permanent shocks, if shocks received by individual i at all time periods t were directly observable in the data, one could pin down a *pass-through* coefficient β using an OLS regression of consumption changes on permanent shocks (η):

$$\beta = \frac{\text{cov}(\Delta \ln c_{i,t}, \eta_{i,t})}{\text{var}(\eta_{i,t})}. \quad (1)$$

Given this coefficient, the parameter of partial insurance against permanent shocks is given by

$$\lambda = 1 - \beta. \quad (2)$$

Given that transitory and permanent 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 coefficient β (and its counterpart for transitory shocks) is 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 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 $T(y) = y - \phi y^{1-\pi}$, so that disposable income of individual i in period t is given by:

$$y_{i,t}^{disp} = \phi y_{i,t}^{1-\pi}. \quad (3)$$

The progressivity parameter, π ,¹ directly translates into the elasticity of disposable income with respect to gross income, $1 - \pi$: $\Delta \ln y_{i,t}^{disp} = (1 - \pi)\Delta \ln y_{i,t}$, and thus

$$\frac{\text{cov}(\Delta \ln c_{i,t}, \Delta \ln y_{i,t}^{disp})}{\text{var}(\Delta \ln y_{i,t}^{disp})} = \frac{(1 - \pi)\text{cov}(\Delta \ln c_{i,t}, \Delta \ln y_{i,t})}{(1 - \pi)^2\text{var}(\Delta \ln y_{i,t})}. \quad (4)$$

Let $\lambda^{disp} = 1 - \text{cov}(\Delta \ln c_{i,t}, \Delta \ln y_{i,t}^{disp})/\text{var}(\Delta \ln y_{i,t}^{disp})$ denote the degree of partial consumption insurance against shocks to disposable income, and likewise λ for gross income $y_{i,t}$. The relationship in (4) implies that $\lambda = 1 - (1 - \pi)(1 - \lambda^{disp})$. It is useful to consider two bounds for common reference values:

$$\lambda = \begin{cases} \pi & \text{if } \lambda^{disp} = 0 \text{ (no self-insurance)} \\ 1 & \text{if } \lambda^{disp} = 1 \text{ (full self-insurance)} \end{cases}. \quad (5)$$

This implies that, if agents are able to fully self-insure, then the degree of public insurance is irrelevant. If agents have no ability to self-insure, then total insurance is equal to public insurance, which is exactly equal to the degree of progressivity.² Previous studies show that most agents are somewhere in between (e.g., [Blundell *et al.*, 2008](#), estimate a $\lambda^{disp} = 0.36$ 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.³ 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

¹If $\pi > 0$, marginal tax rates exceed average rates and hence the tax and transfer system is considered progressive. Conversely, when $\pi = 0$ the tax and transfer scheme is flat.

²Note that λ^{disp} captures the amount of self-insurance against *disposable* income, and that generally the degree of self-insurance is endogenous with respect to existing policies that determine the degree of public insurance.

³In survey data, measurement error and low frequency pose challenges. In administrative data, imputations are required. In bank records, samples are rarely representative. In all cases, a pervasive measurement issue regardless of the source is the disconnect between expenditures and consumption, particularly serious for durable consumption.

terms of a lower data requirement comes with strong parametric assumptions regarding exactly this link: The variance of changes is scaled by $(1 - \pi)^2$, see (4), 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.

2.2 The Insurance Value as Variance Adjustment

To build intuition, consider a simple static framework to assess the insurance value of taxes and transfers. Agents have preferences over consumption and receive an income drawn from some (known) distribution, which they consume. Taxes and transfers change this distribution. We consider preferences being captured by a utility function, whose first k derivatives exist and have alternating signs (positive odd derivatives and negative even derivatives).

In order to isolate the insurance channel of taxes and transfers, consider mean-preserving changes of the distribution. For any change of the distribution, there is a welfare-equivalent adjustment of the variance. If taxes and transfers alter other distributional features, then this will be reflected in the welfare-equivalent variance adjustment to the extent that agents care about these distributional features—Likewise, if agents only care about distributional features up to the variance, then changes of other distributional moments will not be reflected in the welfare-equivalent adjustment of the variance.

Denote pre-government and post-government incomes by Y^{pre} and Y^{post} , respectively. Let $\mu_k(\cdot)$ denote the k 'th central moment of a distribution. Let λ^{gov} denote a proportional reduction of the variance of Y^{pre} that makes agents indifferent between the resulting compressed distribution and the exposure to the distribution of Y^{pos} . Focussing on the variance, a second-order Taylor series expansion yields:

$$\begin{aligned} U(\bar{Y}) + U'(\bar{Y})\mathbb{E}(Y^{pre} - \bar{Y}) + \frac{1}{2}U''(\bar{Y})(1 - \lambda^{gov})\mu_2(Y^{pre}) \\ = U(\bar{Y}) + U'(\bar{Y})\mathbb{E}(Y^{post} - \bar{Y}) + \frac{1}{2}U''(\bar{Y})\mu_2(Y^{post}), \end{aligned}$$

where $\bar{Y} = \mathbb{E}[Y^{pre}] = \mathbb{E}[Y^{post}]$. As we consider a mean-preserving change of the distribution, the expression simplifies to

$$\lambda^{gov} = 1 - \frac{\mu_2(Y^{post})}{\mu_2(Y^{pre})}. \quad (6)$$

When the only difference between the distributions of Y^{pre} and Y^{post} is a reduction of the variance, the expression delivers the tautological statement that the welfare-equivalent adjustment of the variance is exactly equal to this reduction. Importantly, the strength of risk attitudes, as captured by derivatives of the utility function, does not matter for this as we are considering a mean-preserving change, and thus evaluate the derivatives at the same level of consumption.

Now consider a setting where the tax and transfer system alters a symmetric distribution of pre-government income such that post-government income is not only less dispersed, but also features right-skewness. This is captured by a third-order Taylor approximation to expected utility under Y^{post} :

$$\frac{1}{2}U''(\bar{Y})(1 - \lambda^{gov})\mu_2(Y^{pre}) = \frac{1}{2}U''(\bar{Y})\mu_2(Y^{post}) + \frac{1}{6}U'''(\bar{Y})\mu_3(Y^{post}),$$

and thus

$$\lambda^{gov} = 1 - \frac{\mu_2(Y^{post})}{\mu_2(Y^{pre})} - \frac{1}{3} \frac{U'''(\bar{Y})}{U''(\bar{Y})} \frac{\mu_3(Y^{post})}{\mu_2(Y^{pre})}. \quad (7)$$

Given $U''(.) < 0$ and $U'''(.) > 0$, a positively skewed distribution of post-government income reduces the right-hand side, and thus yields a higher λ^{gov} . This is intuitive: taxes and transfers reduce the variance and induce right-skewness. If agents like right-skewness (which they do with the given preferences), then the welfare-equivalent reduction of the variance which preserves the symmetry of the distribution needs to be larger. Note that the welfare-equivalent variance reduction is not higher due to risk aversion, but due to prudence (captured by $U'''(.) / U''(.)$). The same logic extends directly to starting out with non-symmetric distributions of Y^{pre} , as well as to higher-order moments.

3 A Framework for the Measurement of Partial Insurance

3.1 The Model Economy

Overview. 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 sequence of island-level shocks.

The abstraction of islands allows the model to capture partial insurance in a way that will become clear shortly. Importantly for the quantitative analysis, there is no need to define empirical counterparts to the model islands. A possible interpretation is that an island represents an extended network of family members, who perfectly share the purely idiosyncratic risks faced by each member. Some shocks will hit every member equally and hence it cannot be insured within the family network—for example, regional and sectoral shocks.

Population and endowment structure. Each period a mass $(1 - \delta)$ of newborns enters the economy at age 0 and replaces workers which stochastically leave the economy: at any age, the probability of survival to the next period is constant at $\delta \in (0, 1)$ (a la [Yaari, 1965](#)). Individual income (endowment) of a period- τ -born agent i in period t is given by a transitory component $\varepsilon_{i,t}^{\tau,idio}$ and both an island-level and an idiosyncratic permanent component of income $z_{i,t}^{\tau,l}$ for $l \in \{island, idio\}$. We assume a time-invariant distribution of the transitory shock component, which we denote by $F_{\varepsilon,t}^{\tau} = F_{\varepsilon}$.⁴ We denote the permanent shocks by $\eta_{i,t}^{\tau,l}$ for $l \in \{island, idio\}$, and assume that their distribution varies systematically with the aggregate state of the economy, which we refer to by x_t : $\eta_{i,t}^{\tau,l} \sim F_{\eta,t}^{\tau,l} = F_{\eta,x_t}^l$. All stochastic components of income are independent and normalized such that, in every period t , $\int \exp\left(\eta_{i,t}^{\tau,l}\right) dF_{\eta,x_t}^l = 1$ for $l \in \{island, idio\}$ and likewise for ε . The aggregate state is drawn from a distribution

⁴Pre-empting what will be a model result shown below, notice that the stationarity of the ε distribution is innocuous for equilibrium allocations as transitory shocks are perfectly insurable in the model.

that depends on the last realization, so that we have $x_{t+1} \sim H(x_t)$. Age 0 agents entering in year τ hold zero wealth and are allocated to an island j of agents which share the same initial realizations of the permanent components $\{z_{i,\tau}^{\tau, island}, z_{i,\tau}^{\tau, idio}\}_i$, followed by the same sequence of island-level shocks $\{\eta_{i,t}^{\tau, island}\}_{i,t=\tau+1}^\infty$. The following endowment process captures the described features:

$$\begin{aligned} y_{i,t}^\tau &= z_{i,t}^{\tau, island} + z_{i,t}^{\tau, idio} + \varepsilon_{i,t}^{\tau, idio}, & \varepsilon_{i,t}^{\tau, idio} &\sim F_\varepsilon \\ z_{i,t}^{\tau, l} &= z_{i,t-1}^{\tau, l} + \eta_{i,t}^{\tau, l}, & \eta_{i,t}^{\tau, l} &\sim F_{\eta, x_t}^l, \quad \text{for } l \in \{island, idio\}; \quad x_t \sim H(x_{t-1}). \end{aligned} \quad (8)$$

Preferences. Agents maximize expected discounted lifetime utility, whereby we assume time- and state-separable preferences with a CRRA per-period utility function $U(c) = (c^{1-\gamma} - 1) / (1 - \gamma)$. We use log utility as the benchmark, and also inspect the role of stronger risk attitudes implied by a parameter of relative risk aversion, γ , larger than 1. The discount factor β is constant across the population. Expected lifetime utility of an agent i born in period τ derived from some stochastic consumption sequence is given by

$$\mathbb{E}_\tau \sum_{t=\tau}^{\infty} (\beta\delta)^{t-\tau} U(c_{i,t}), \quad (9)$$

where \mathbb{E}_τ is the expectation operator that uses all information available in period τ .

Asset markets. Every period t agents engage in asset trade. Within islands, agents can trade one-period bonds $b^{within}(s_{i,t+1}^\tau, x_{t+1}; s_i^{\tau, t}, x^t)$ at price $q^{within}(s_{i,t+1}^\tau, x_{t+1}; s_i^{\tau, t}, x^t)$, which pay out one unit of consumption contingent on the realization of next period's individual endowment state $s_{i,t+1}^\tau \equiv \{z_{i,t+1}^{\tau, idio}, \varepsilon_{i,t+1}^{\tau, idio}, z_{i,t+1}^{\tau, island}\}$ and x_{t+1} .⁵

Across islands, agents cannot trade claims contingent on the island-level shock, but only on the purely idiosyncratic components $s_{i,t+1}^{\tau, idio} \equiv \{z_{i,t+1}^{\tau, idio}, \varepsilon_{i,t+1}^{\tau, idio}\}$ and x_{t+1} , denoted by $b^{across}(s_{i,t+1}^{\tau, idio}, x_{t+1}; s_i^{\tau, t}, x^t)$, at price $q^{across}(s_{i,t+1}^{\tau, idio}, x_{t+1}; s_i^{\tau, t}, x^t)$. Claims are in zero net supply.

⁵We follow notation conventions and denote period- t realizations of the endowment by $s_{i,t}^\tau$ and the history from τ to t by $s_{i,\tau}^{\tau, t} = \{s_{i,\tau}^\tau, \dots, s_{i,t}^\tau\}$, and similarly, x_t denotes the aggregate state in period- t , while x^t denotes the history from period 0 onwards.

The per-period budget constraint of a generation- τ household with realized log income $y_{i,t}^\tau$ (composed of the components $s_{i,t}^\tau$) in aggregate state of the economy x_t is given by

$$\begin{aligned} c_{i,t}^\tau + \int_x \int_s q^{within}(s, x; s_i^{\tau,t}, x^t) b^{within}(s, x; s_i^{\tau,t}, x^t) dF(s_{i,t+1}^\tau; x_{t+1}) dF(x_{t+1}) + \\ \int_x \int_s q^{across}(s, x; s_i^{\tau,t}, x_t) b^{across}(s, x; s_i^{\tau,t}, x_t) dF(s_{i,t+1}^{\tau,idio}; x_{t+1}) dF(x_{t+1}) \\ = \exp(y_{i,t}^\tau) + b^{within}(s_{i,t}^\tau, x_t; s_i^{\tau,t-1}, x^{t-1}) + b^{across}(s_{i,t}^{\tau,idio}, x_t; s_i^{\tau,t-1}, x^{t-1}), \end{aligned} \quad (10)$$

where $F(x_{t+1})$ is the distribution of the aggregate state x_{t+1} , which in turn simply affects the distribution of the individual state $s_{i,t+1}^\tau$, which is given by $F(s_{i,t+1}^\tau; x_{t+1})$ and $F(s_{i,t+1}^{\tau,idio}; x_{t+1})$.

Information and equilibrium. Agents of a generation τ observe their initial realization of the permanent component $\left\{ z_{i,\tau}^{\tau, island}, z_{i,\tau}^{\tau, idio} \right\}_i$; they do not know the sequence of island-level shocks $\left\{ \eta_{i,t}^{\tau, island} \right\}_{i,t=\tau+1}^\infty$ that subsequently defines their island. At the beginning of a period, agents observe the aggregate state of the economy x_t and their individual draws from the shock distributions. While they do not know the sequence of these shock distributions ex ante, agents do know the distributions conditional on the aggregate state, and they know the process that governs the evolution of this aggregate state, i.e., $x_{t+1} \sim H(x_t)$. Based on this, agents form expectations about possible trajectories of x_t and the implied possible trajectories of the distributions of $s_{i,t}^\tau$.

A sequential markets equilibrium, given an initial distribution of households, is a sequence of prices $q^{within}(s, x; s_i^{\tau,t}, x^t)$, $q^{across}(s, x; s_i^{\tau,t}, x^t)$ and allocations of consumption $c_{i,t}^\tau$, and assets $b^{within}(s, x; s_i^{\tau,t}, x^t)$ and $b^{across}(s, x; s_i^{\tau,t}, x^t)$, such that at given prices, the allocations solve the household problems and markets clear in every period and state of the world.

There is an equilibrium with no trade across islands. This mimics the result in [Heathcote et al. \(2014\)](#) and [Boerma and Karabarbounis \(2021\)](#), whose models feature similar asset market structures. While in a no-trade equilibrium in the spirit of [Constantinides and Duffie \(1996\)](#) idiosyncratic endowment shocks remain uninsured, in the no-trade equilibrium of the island economy there is partial insurance: island-level shocks remain uninsured while idiosyncratic shocks are insured against perfectly through the state-contingent claims.

Consider a static island-planner who optimally allocates available resources within an island. By equally distributing within the island, the planner equates the expected marginal rate of substitution between consumption today and consumption tomorrow across all agents. The resulting allocation is supported as an equilibrium where prices of claims to all possible future realizations reflect the expected marginal rate of substitution. With the given preferences, the expected MRS is a function of expected consumption growth. At the candidate allocation from the planner problem, expected consumption growth is identical everywhere. This implies that on every island, individuals face the same price for within-island claims. There are no gains from trade across islands and all trade happens within islands.

In this equilibrium with no trade across islands, the period t log consumption of an agent i of generation τ , with income components $(z_{i,t}^{island}, z_{i,t}^{idio}, \varepsilon_{i,t}^{idio})$ is given by

$$\ln c_{i,t}^\tau \left(z_{i,t}^{\tau, island}, z_{i,t}^{\tau, idio}, \varepsilon_{i,t}^{idio} \right) = z_{i,t}^{\tau, island} + \ln \int \exp(y^{\tau, idio}) dF_{y^{\tau, idio}, t}^\tau, \quad (11)$$

where $F_{y^{\tau, idio}, t}^\tau$ is the period t distribution of idiosyncratic income of individuals from generation τ ($y_{i,t}^{\tau, idio} = z_{i,t}^{\tau, idio} + \varepsilon_{i,t}^{\tau, idio}$). 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. The distribution of this idiosyncratic component depends on both time and age: 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 $(t - \tau)$, because the permanent shocks $\eta_{i,t}^{idio}$ accumulate over age, resulting in a widening distribution of the permanent component $z_{i,t}^{idio}$. Note that we assume transitory shocks at the purely idiosyncratic level, identical to both [Heathcote *et al.* \(2014\)](#) and [Boerma and Karabarbounis \(2021\)](#). This assumption yields perfect insurance against transitory income risk. It rests on insights from calibrated incomplete market models, which 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](#)).

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 analytical 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. This also allows us to directly obtain 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.

Degree of partial insurance. Within the model, we make the common assumption that agents can observe transitory and permanent shocks directly.⁶ The consumption function implies the following consumption change from $t - 1$ to t :

$$\begin{aligned} \Delta \ln c_{i,t} \left(z_{i,t}^{\tau, island}, z_{i,t}^{\tau, idio}, \varepsilon_{i,t}^{idio} \right) &= \eta_{i,t}^{island} + \ln \frac{\int \exp(\eta_{i,t}^{idio}) dF_{\eta,t}^{idio} \int \exp(\varepsilon_{i,t}^{idio}) dF_{\varepsilon,t}^{idio}}{\int \exp(\varepsilon_{i,t-1}^{idio}) dF_{\varepsilon,t-1}^{idio}} \\ &= \eta_{i,t}^{island}. \end{aligned} \quad (12)$$

The relevant model version of partial insurance against permanent shocks builds around the pass-through to the combined *island* and *idio*-shocks, i.e., to $\eta_{i,t} = \eta_{i,t}^{idio} + \eta_{i,t}^{island}$. As is clear from (12), 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 overall pass-through of the combined shock is then a convex combination of these two measures. Directly applying (1), we obtain

$$\begin{aligned} 1 - \lambda &= \frac{\text{cov}(\Delta \ln c_{i,t}, \eta_{i,t})}{\text{var}(\eta_{i,t})} = \frac{\text{cov}(\eta_{i,t}^{island}, \eta_{i,t})}{\text{var}(\eta_{i,t})} = \frac{\text{cov}(\eta_{i,t}^{island}, \eta_{i,t}^{island} + \eta_{i,t}^{idio})}{\text{var}(\eta_{i,t})} \\ &= \frac{\text{var}(\eta_{i,t}^{island})}{\text{var}(\eta_{i,t}^{island} + \eta_{i,t}^{idio})} = \frac{\text{var}(\eta_{i,t}^{island})}{\text{var}(\eta_{i,t}^{island}) + \text{var}(\eta_{i,t}^{idio})}, \end{aligned} \quad (13)$$

such that the degree of partial insurance against permanent shocks, λ , is given by the fraction of the variance of permanent 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.

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

3.2 Measurement of the Insurance Value of Taxes and Transfers

We now use the model structure outlined above in order to measure the degree of partial insurance provided by the tax and transfer system. 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, taxes and transfers alter the endowment stream (8) faced by agents. Importantly, we maintain the normalization that $\int \exp(\eta_{i,t}^x) dF_{\eta,t}^x = 1$ for $x \in \{island, idio\}$ and $\int \exp(\varepsilon_{i,t}^{idio}) dF_{\varepsilon,t}^{idio} = 1$. This means that we consider cross-sectional redistribution of endowments, and rule out wasteful government consumption or debt-financed transfer payments.

We then consider the following experiment. Agents live in one of two possible scenarios. In the first—pre-government—the endowment stream describes household level gross income. In the second—post-government—the endowment stream describes disposable income. In each scenario, an exogenous degree of partial insurance governs the split between the *island* and *idio* components of the total idiosyncratic shocks. Given the amount of partial insurance, we obtain stochastic consumption streams per equation (12).

We first assume a degree of partial insurance against (total) individual shocks in the *post-government* scenario—i.e., we assume a value for λ^{post} . Next, we 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). Consider agents born in period τ . When they face the stochastic income stream $y^{pre} = \{y_{i,t}^{pre}\}_{i,t=\tau}^{\infty}$ with a degree of partial insurance λ^{pre} against permanent shocks, this translates into stochastic streams of *idio* and *island* components $\{z_{i,t}^{pre,island}, z_{i,t}^{pre,idio}, \varepsilon_{i,t}^{pre,idio}\}_{i,t=\tau}^{\infty}$. The components are such that the implied distribution of their sum in period t , $(z_{i,t}^{pre,island} + z_{i,t}^{pre,idio} + \varepsilon_{i,t}^{pre,idio})$, corresponds to the distribution of the total income $y_{i,t}^{pre}$. Likewise, in the other scenario they face income streams $\{z_{i,t}^{post,island}, z_{i,t}^{post,idio}, \varepsilon_{i,t}^{post,idio}\}_{i,t=\tau}^{\infty}$ which are consistent with income stream $y^{post} = \{y_{i,t}^{post}\}_{i,t=\tau}^{\infty}$ and partial insurance λ^{post} .

We now denote the consumption function of a generation τ that results from optimal behavior when facing some income stream y and a degree of partial insurance λ by $c_t^\tau(y, \lambda)$. With the two income streams, y^{pre} and y^{post} , as well as insurance λ^{post} at hand, we find the

level of partial insurance λ^{pre} that yields the island- and idiosyncratic income sequences that makes agents ex ante indifferent:

$$\mathbb{E}_{i|\tau} \sum_{t=\tau}^{\infty} (\beta\delta)^{t-\tau} U(c_t^\tau(y_{i,t}^{pre}, \lambda^{pre})) = \mathbb{E}_{i|\tau} \sum_{t=\tau}^{\infty} (\beta\delta)^{t-\tau} U(c_t^\tau(y_{i,t}^{post}, \lambda^{post})), \quad (14)$$

where $\mathbb{E}_{i|\tau}$ denotes the expectation operator taken over possible individual realizations, conditional on birth in period τ .

The λ^{pre} that solves (14) combines the overall insurance provided by the government and additional private insurance captured by λ^{post} . We can therefore define the overall pass-through from gross income to consumption as the product of the pass-through from disposable income to consumption, $1 - \lambda^{post}$, and the pass-through from gross income to disposable income, $1 - \lambda^{gov}$. This second term captures the insurance provided by the government, which mimics the discussion of the tax function in Section 2. So the implied measure of government-provided insurance is given by the following relationship:

$$(1 - \lambda^{pre}) = (1 - \lambda^{post})(1 - \lambda^{gov}). \quad (15)$$

Note again that imposing a different λ^{post} yields a different λ^{pre} from solving (14). This gives rise to λ^{gov} not depending on the parameterization of λ^{post} in the analytical case discussed next—nor in our numerical implementations below.

3.3 A Sufficient Statistic of Public Insurance

We now consider a special case of the model framework and assume that overall permanent shocks $\eta_{i,t}^\tau$ are distributed Normally with a time-constant variance σ_η^2 . The crucial insight is that the measure of insurance is not sensitive with respect to the (exogenous) insurance λ^{post} , the survival probability δ , nor the imposed preference parameters, i.e., the degree of risk aversion (pinned down by parameter γ) and the time discount factor β .

The degree of insurance λ yields a Normally distributed island-level component with scaled variance $(1 - \lambda)\sigma_\eta^2$: $F_{\eta,t}^{\tau, island} = \mathcal{N}(- (1 - \lambda)\sigma_\eta^2/2, (1 - \lambda)\sigma_\eta^2)$. As permanent shocks accumulate from birth onwards, the consumption distribution of generation τ in period t is

log-Normal:

$$\ln c_{i,t}^\tau \sim \mathcal{N} \left(-\frac{(t+1-\tau)(1-\lambda)\sigma_\eta^2}{2}, (t+1-\tau)(1-\lambda)\sigma_\eta^2 \right). \quad (16)$$

This allows for a fully analytical solution of the degree of insurance in (14), which with CRRA per-period utility functions becomes

$$\mathbb{E}_{i|\tau} \sum_{i,t=\tau}^{\infty} (\beta\delta)^{t-\tau} \frac{((c_{i,t}^{\tau,pre})^{1-\gamma} - 1)}{1-\gamma} = \mathbb{E}_{i|\tau} \sum_{i,t=\tau}^{\infty} (\beta\delta)^{t-\tau} \frac{((c_{i,t}^{\tau,post})^{1-\gamma} - 1)}{1-\gamma}. \quad (17)$$

The log-Normal distributions of $(c_{i,t}^{\tau,pre})^{1-\gamma}$ and $(c_{i,t}^{\tau,post})^{1-\gamma}$ with means and variances as implied by (16) give analytical expressions for the period- τ expectations of all elements in the discounted sum of (17), which simplifies to

$$\begin{aligned} & \sum_{i,t=\tau}^{\infty} (\beta\delta)^{t-\tau} \exp \left(-(t+1-\tau) \frac{1}{2} \gamma(1-\gamma)(1-\lambda^{pre})\sigma_\eta^{2,pre} \right) = \\ & \sum_{i,t=\tau}^{\infty} (\beta\delta)^{t-\tau} \exp \left(-(t+1-\tau) \frac{1}{2} \gamma(1-\gamma)(1-\lambda^{post})\sigma_\eta^{2,post} \right). \end{aligned}$$

Equality of the two geometric series requires that $(1-\lambda^{pre})\sigma_\eta^{2,pre} = (1-\lambda^{post})\sigma_\eta^{2,post}$, and thus we obtain $(1-\lambda^{pre}) = (1-\lambda^{post}) \frac{\sigma_\eta^{2,post}}{\sigma_\eta^{2,pre}}$, which implies by (15) that the government provided insurance is

$$\lambda^{gov} = 1 - \frac{\sigma_\eta^{2,post}}{\sigma_\eta^{2,pre}}. \quad (18)$$

A sufficient statistic of insurance. In summary, the ratio of variances of disposable income shocks and gross income shocks directly reflects the measure of government provided insurance. This measure is robust with respect to the exogenous parameters. In particular, it implies that the pass-through of post-government income to consumption, captured by the imposed degree of insurance λ^{post} does not affect the measure of government provided insurance. Thus, equation (18) captures that data on gross and net incomes alone suffices to evaluate the degree of government provided insurance: the ratio of the variances of permanent shocks of the two income measures is a sufficient statistic for the degree of partial insurance

in the case of homoskedastic log-Normal distributions. Equation (18) should not come as a surprise given the earlier discussion that delivered the same expression in Equation (6). Again, in the context of our model framework, we evaluate the degree of insurance by taxes and transfers as a welfare-equivalent reduction of the variance—in this case the variance of log-Normally distributed permanent shocks. In the empirical applications, the estimated permanent log income shocks feature non-Gaussian distributions, which systematically vary over the cycle. Therefore, we can expect that risk attitudes of agents play a role for a welfare-equivalent variance adjustment.

4 Public Insurance in Sweden

4.1 Regularities of Household Income in Sweden

In this section, we evaluate the degree of partial insurance provided by the tax and transfer system in Sweden. The two empirical ingredients necessary to apply the measurement device derived above are two stochastic income streams: one that captures the regularities of household income before taxes and transfers (pre-government income), and one that captures the regularities of household income after taxes and transfers (post-government income). We estimate these using a set of data moments, which we take from [Busch *et al.* \(2018\)](#). The data moments are based on 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 wages and salaries, the labor part of business income, as well as the taxable compensation for sick leave and parental leave. Net (post-government) income adds transfers and taxes, taking into account various public programs, all of which are consistently measured over time: (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 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, whereby the former account for the bulk of total tax payments.

Let $y_{i,t}^{pre}$ and $y_{i,t}^{post}$ 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):

$$\begin{aligned} y_{i,t} &= z_{i,t} + \varepsilon_{i,t} \\ z_{i,t} &= z_{i,t-1} + \eta_{i,t} \end{aligned} \tag{19}$$

where $\varepsilon_{i,t}$ is an *iid* transitory shock, and $\eta_{i,t}$ 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:

$$\varepsilon_{i,t} \sim \begin{cases} \mathcal{N}(\bar{\mu}_\varepsilon, \sigma_{\varepsilon,1}^2) & \text{with prob. } p_{\varepsilon,1} \\ \mathcal{N}(\bar{\mu}_\varepsilon, \sigma_{\varepsilon,2}^2) & \text{with prob. } 1 - p_{\varepsilon,1} \end{cases} \tag{20}$$

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 $\eta_{i,t}$ follows a mixture of three normals:

$$\eta_{i,t} \sim \begin{cases} \mathcal{N}(\bar{\mu}_{\eta,t} + \mu_{\eta,1} + \phi_1 x_t, \sigma_{\eta,1}^2) & \text{with prob. } p_{\eta,1} \\ \mathcal{N}(\bar{\mu}_{\eta,t} + \mu_{\eta,2} + \phi_2 x_t, \sigma_{\eta,2}^2) & \text{with prob. } p_{\eta,2} \\ \mathcal{N}(\bar{\mu}_{\eta,t} + \mu_{\eta,3} + \phi_3 x_t, \sigma_{\eta,3}^2) & \text{with prob. } p_{\eta,3} \end{cases} \tag{21}$$

where $p_{\eta,j}$, $j = 1, 2, 3$, denotes the probability of drawing from component j , where $\sum_{j=1}^3 p_{\eta,j} = 1$. The parameters ϕ_j determine how strongly aggregate risk as captured by x_t translates into changes of the distribution of idiosyncratic earnings risk. We operationalize x_t by standardized log GDP growth. This choice turns out to allow the model to capture the time series behavior of the distribution of household income changes well, and allows for a direct comparison with the empirical specification in [Busch *et al.* \(2022\)](#). In every t , $\bar{\mu}_{\eta,t}$ is chosen such that $\mathbb{E}_i[\exp(\eta_{i,t})] = 1$. In the estimation, we then shift the distribution and 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 the process. We estimate the set of parameters $\chi = \{\chi_{trans}, \chi\}$ where

$$\chi_{trans} = \{\sigma_{\varepsilon,1}, \sigma_{\varepsilon,2}, p_{\varepsilon,1}\} \quad (22)$$

$$\chi = \{\mu_{\eta,2}, \mu_{\eta,3}, \sigma_{\eta,1}, \sigma_{\eta,2}, p_{\eta,1}, p_{\eta,2}, \phi_2, \phi_3\} \quad (23)$$

by the simulated method of moments (SMM).⁷ We target the time series of the variance, the right tail (P90-P50), and the left tail (P50-P10) 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 300 moments for pre-government income, and 300 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 (19):

$$y_{i,t} - y_{i,t-s} = \varepsilon_{i,t} - \varepsilon_{i,t-s} + \sum_{j=0}^{s-1} \eta_{i,t-j}, \quad (24)$$

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

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.

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 ($n = 1, \dots, N$) that corresponds to cross-sectional target $m \in \{var(\Delta^1 y_{i,t}), \dots, var(y_{age=60})\}$. In each simulation, we draw a matrix of random variables $X_r = \{\varepsilon_{i,1}, \varepsilon_{i,2}, \dots, \varepsilon_{i,T}, \eta_{i,1}, \dots, \eta_{i,T}\}_{i=1}^I$ 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 1990s, which Sweden entered in what turned out to be an extended contractionary period, followed by a recovery leading into the 2000s. Figure 1 shows the density functions implied by the estimates for the two subperiods. Over the four-year period from 1989 to 1993, GDP plummeted to

a negative GDP growth of -3.7% ; the average annual growth rate of -0.93% is about 1.3 standard deviations below the average growth. With an average annual growth of 4% (about 0.8 standard deviations above the average), from 1996–2000 GDP grew by about 17% .

Table 1: Estimated Parameter Values

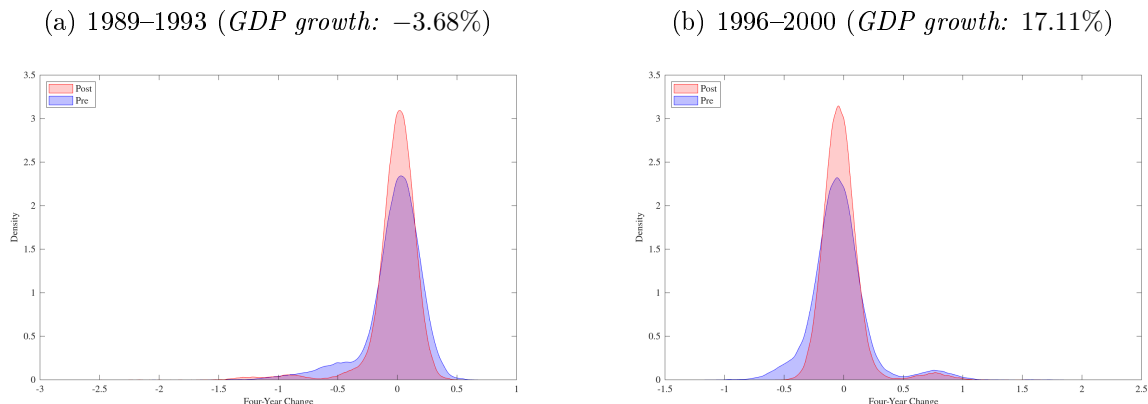
Parameter	Description	Pre-Gov.	Post-Gov.
$p_{\varepsilon,1}$	Mixture prob. of ε distribution	0.863	0.876
$\sigma_{\varepsilon,1}$	Std. dev. of ε mix. comp. 1	0.041	0.047
$\sigma_{\varepsilon,2}$	Std. dev. of ε mix. comp. 2	0.484	0.391
$p_{\eta,1}$	Mixture prob. of η mix. comp. 1	0.972	0.983
$p_{\eta,2}$	Mixture prob. of η mix. comp. 2	0.014	0.008
$p_{\eta,3}$	Mixture prob. of η mix. comp. 3	0.014	0.008
$\sigma_{\eta,1}$	Std. dev. of η mix. comp. 1	0.076	0.060
$\sigma_{\eta,2}$	Std. dev. of η mix. comp. 2	0.010	0.053
$\sigma_{\eta,3}$	Std. dev. of η mix. comp. 3	0.010	0.053
$\mu_{\eta,2}$	Mean of η mix. comp. 2	0.209	0.089
$\mu_{\eta,3}$	Mean of η mix. comp. 3	-0.446	-0.067
ϕ_2	Aggregate risk transmission mix. comp. 2	0.636	0.726
ϕ_3	Aggregate risk transmission mix. comp. 3	0.028	0.203
M	# moments targeted in estimation	300	300

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

Between these two periods, 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 four-year income changes. Each panel shows a histogram of the simulated distribution for the estimated mixture of Normals corresponding to pre-government (blue) and post-government (red) income. In the plots, we normalize the distribution such that $\mathbb{E}_i[\exp(\eta_{i,t})] = 1$. For completeness, Figures A.1 and A.2 in Appendix A.2 show the simulated moments of changes at different horizons at the estimated parameters together with the empirical moments over time. As captured in the figures, 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 from 1989 to 1993) goes hand-in-hand with

a left-skewed distribution, while strong positive GDP growth (as from 1996 to 2000) comes with a right-skewed distribution. Table 2 shows several statistical moments of distributions that summarize the change over the cycle and the difference between gross and net income.

Figure 1: Cross-Sectional Distribution of Permanent Income Changes



Note: Each figure shows the distribution of simulated pre-government permanent income changes η in blue and post-government in red. Distributions are normalized to a mean of 1 in levels.

The tax and transfer system compresses the distribution, as captured by a smaller estimated variance. In the downturn of 1989–93, the variance of net income is 0.0460 compared to 0.0602 for gross income. In the change from 1996–2000, these numbers are 0.0311 and 0.0625, respectively. The right-skewness in expansions is captured by a positive coefficient of skewness (the third standardized moment); and the mirror image holds true for contractions. This sign difference also shows in measures of Kelley’s skewness, which is based on the 10th, 50th, and 90th percentiles of the distribution: $\mathcal{KS} = ((P90 - P50) - (P50 - P10))/(P90 - P10)$. \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 allows for a direct interpretation of the magnitude of the change in the distribution over the cycle: For pre-government income, the value of $\mathcal{KS} = -0.2098$ for 1989–93 indicates that $(P90 - P50)$ accounts for 39.5% of the $(P90 - P10)$ dispersion.⁸ On the other hand, in the growth period from 1996–2000, the value of $\mathcal{KS} = 0.0111$ indicates that $(P90 - P50)$ accounts for about 50.6% of the $(P90 - P10)$ dispersion. This is to say that the distribution is relatively symmetric in growth times, while it is left-skewed in contractionary times.

⁸Note that $(P90 - P50)/(P90 - P10) = 0.5 + \mathcal{KS}/2$.

Table 2: Moments of Distribution of Four-Year Permanent Income Changes

Year	Variance		P90-P10		Skewness	
	<u>Pre</u>	<u>Post</u>	<u>Pre</u>	<u>Post</u>	<u>Pre</u>	<u>Post</u>
1989–1993	0.0602	0.0460	0.5206	0.3512	-1.6445	-3.0244
1996–2000	0.0625	0.0311	0.4814	0.3314	1.2223	2.2714
	Kelley		P50-P10		P90-P50	
	<u>Pre</u>	<u>Post</u>	<u>Pre</u>	<u>Post</u>	<u>Pre</u>	<u>Post</u>
1989–1993	-0.2098	-0.0984	0.3149	0.1929	0.2057	0.1583
1996–2000	0.0111	0.0516	0.2380	0.1572	0.2434	0.1743

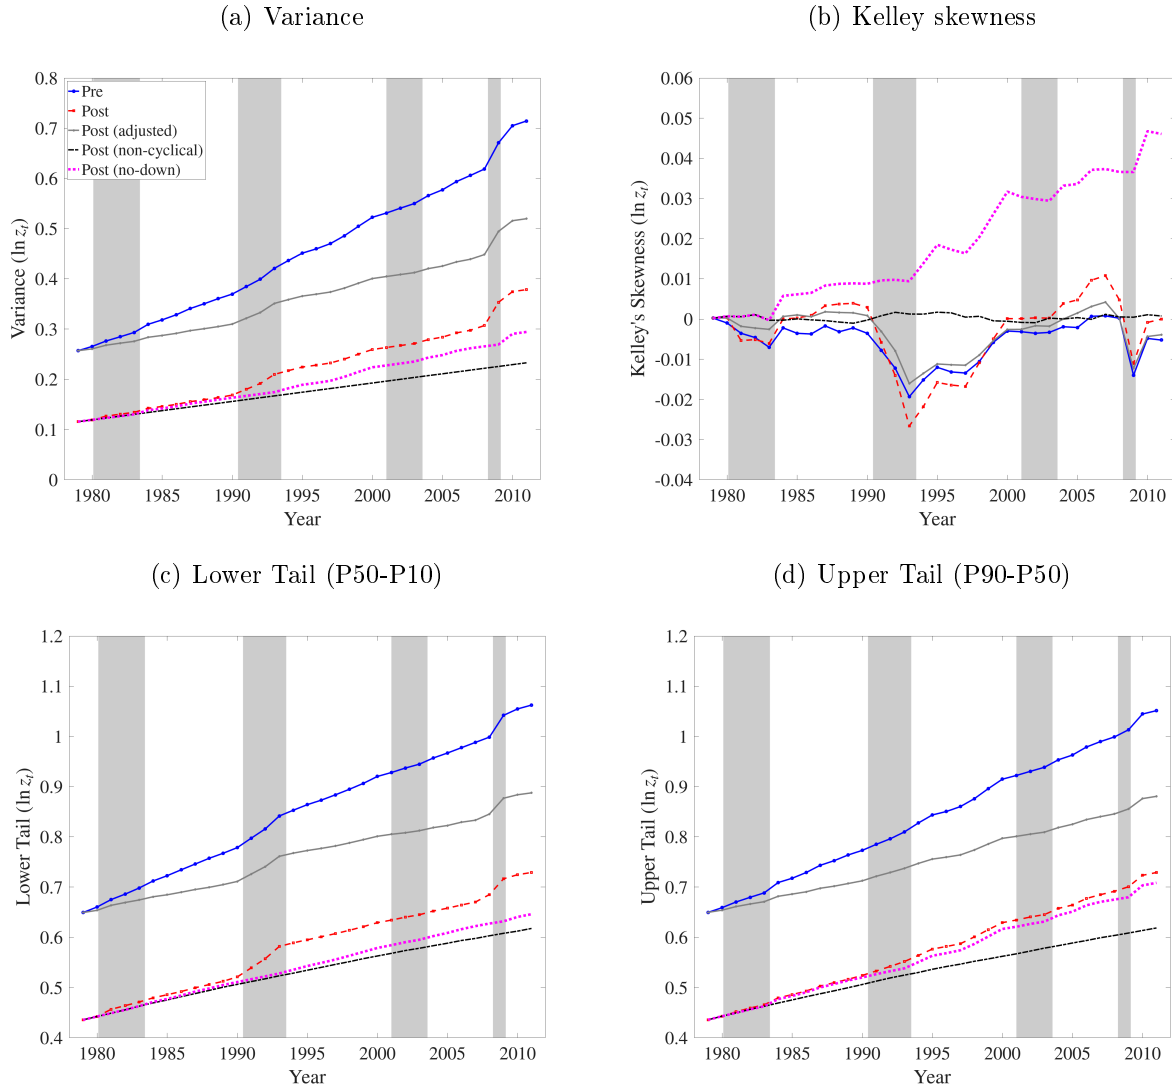
Note: Table shows moments of the distribution of four-year permanent income changes in a contractionary period (1989–1993) and an expansionary period (1996–2000).

Taxes and transfers dampen the pass-through of aggregate contractions to the distribution, which is captured in the parameter estimates for ϕ_2 and ϕ_3 in Table 1. Also for post-government income, \mathcal{KS} changes from negative 1989–93 to positive in 1996–2000. However, the difference is less pronounced than for pre-government income: In the contractionary period, $(P90 - P50)$ accounts for about 45% of the $(P90 - P10)$ dispersion ($\mathcal{KS} = -0.0984$), and from 1996–2000 $(P90 - P50)$ accounts for about 53% ($\mathcal{KS} = 0.0516$). Furthermore, the distribution is leptokurtic for both income measures in 2008 and 2009, with a somewhat higher kurtosis for post-government income, which captures that the tax and transfer system overall increases the concentration of the distribution.

Trajectories of a cohort. The income process (19) with the parameters reported in Table 1 implies a distribution of possible idiosyncratic paths of the permanent income component. One possible realization is the one of 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 realized until 2011 (i.e., the one on which the estimated income process is based). Given its relevance for the insurance measure we introduced above, we now consider the distribution of the permanent income component. First, consider the blue line in panel (a) in Figure 2: it shows the variance of the cross-sectional permanent income component of pre-government income. During the contractions of the early 1990s and around the Great Recession, 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 cyclicity of the distribution of shocks; or (v), a post-government income process that eliminates the reaction of of the distribution to downside changes.

Second, consider permanent component of post-government income for the same cohort. In each of the four panels of Figure 2, the red line reports the cross-sectional moments. The first key difference is that the overall dispersion at every age is smaller (see panel a): starting with an already less dispersed distribution, it also increases less over age, in line with the discussion of the estimated permanent income change component in the previous section. Second, in the years leading up to the recession of the early 1990s, the asymmetry as measured by Kelley’s skewness behaves very similar to what is observed for pre-government income; in the subsequent recovery Kelley’s skewness of post-government income gets less and less negative and after a dip in the early 2000’s contraction turns positive around the mid-2000s, before dropping again during the Great Recession. The remaining four series in each panel reflect moments of the distribution resulting from shutting down various features of the income process to which we turn below.

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

4.2 Measures of Insurance

We feed the measurement device outlined in Section 3 with the estimated income processes for the two income measures. A key parameter of the measurement device is the degree of insurance against permanent income risk after taxes and transfers, λ^{post} . We consider a range of possible values for λ^{post} and use Equation (14) to back out λ^{pre} , and return to the implications of this choice later. When $\lambda^{post} = 0$, the obtained λ^{pre} measures the degree of partial insurance provided by the government under the assumption that there is no additional partial insurance.

From an ex-ante perspective, the distribution of possible consumption streams that can realize over the life cycle is 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. In addition, the idiosyncratic shock distributions are estimated to vary with the aggregate state of the economy, which itself is risky. We capture this by adding a stochastic process for the aggregate state. In particular, we operationalize $x_{t+1} \sim H(x_t)$ by fitting an AR(1) process to x_t , and then use the estimated process when constructing the insurance measure, which considers ex-ante expected life-cycle utility. The choice of an AR(1) is not central; it mimics McKay (2017), who for a similar process of individual risk uses an AR(1) to capture the dynamics of an aggregate latent variable that shifts the distribution of individual shocks. The sequence of aggregate states experienced by the cohort entering the labor market in 1979 is one possible realization of this process.

In line with the description in Section 3.2, we now find the degree of partial insurance against permanent income risk, λ^{pre} , which yields a consumption stream that makes households indifferent to facing the post-government income stream. For a given λ^{pre} , we scale the estimated parameters of the permanent shocks such that the variance of the resulting distribution for $\eta_{i,t}^{idio}$ is equal to fraction λ^{pre} 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.⁹ We normalize such that $\mathbb{E}[\exp(\eta^{island})] = \mathbb{E}[\exp(\eta^{idio})] = 1$.

Overall insurance. Under log utility we find a welfare-equivalent variance adjustment of $\lambda^{pre} = 0.4856$, which means that the existing tax and transfer schedule in Sweden corresponds to insuring households against 49% of permanent shocks to household labor income, as shown in Table 3. 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}* . The 49% partial insurance translates into a CEV of 11.7%. Hence, the existing tax and transfer system provides sizable insurance. When setting the coefficient of relative risk aversion to 2, and thus imposing overall stronger risk attitudes, the implied degree of insurance is basically unchanged at $\lambda^{pre} = 0.4846$. Of course, the associated CEV is higher and roughly doubles to 25.0%.

⁹Appendix B gives expressions for the standardized moments of the scaled distribution.

Role of initial dispersion. To interpret the degree of insurance of 49% further, it is important to notice that government policy reduces the overall level of cross-sectional dispersion, the build-up over time as a cohort ages, as well as the cyclicity of shocks. In order to isolate the insurance value that stems from how the usual shocks received over time are buffered, 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 obtain $\lambda^{pre} = 0.1757$, i.e., moving from the pre- to the post-government income stream adjusted to the same initial variance amounts to partial insurance of 18%, which translates into a CEV of about 4.1% under log utility.

Gain of eliminating cyclicity. Given the already sizable insurance, what is the scope of additional government policy as a means of insurance against cyclical risk? We consider the same experiment for another counterfactual income process: Assume that on top of what the government already does, cyclicity 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. The implied degree of insurance is about 64% (or 32% when adjusting for initial variance at age 25). Considering the CEV connected to those insurance parameters, the scope of additional insurance is sizable: when adjusting for initial variance effects, the CEV is 7.7% (compared to the 4.1% for the cyclical process).

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., $\lambda^{post} = 0$. 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 parameterization. For this, we assume that instead 10% of permanent shocks to post-government income are insured. This delivers a slightly 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

Table 3: Partial Insurance and Welfare Gains of the Tax and Transfer System

Scenario:		Higher-Order Risk Present		Gaussian with same Variance	
		$\lambda^{pre} = \lambda^{gov}$	CEV	$\lambda^{pre} = \lambda^{gov}$	CEV
From pre-government income to...		\ln utility			
(I)	...post-government income (post)	48.56%	11.73%	49.74%	11.40%
(II)	...post adjusted for initial dispersion	17.57%	4.13%	17.27%	3.82%
(III)	...post w/o cyclicalities	63.50%	15.53%	62.88%	14.61%
(IV)	...post w/o cyclicalities and adjusted	32.24%	7.68%	30.40%	6.80%
(V)	...post w/o reaction to negative x_t	52.84%	12.81%	55.65%	12.83%
(VI)	...post w/o neg. reaction & adjusted	21.78%	5.14%	23.17%	5.15%
CRRA w/ Risk Aversion = 2					
(I)	...post-government income (post)	48.46%	24.96%	49.58%	24.86%
(II)	...post adjusted for initial dispersion	17.54%	8.51%	17.94%	8.42%
(III)	...post w/o cyclicalities	63.21%	33.50%	63.02%	32.49%
(IV)	...post w/o cyclicalities and adjusted	31.94%	15.93%	31.21%	15.00%
(V)	...post w/o reaction to negative x_t	55.04%	28.71%	55.64%	28.25%
(VI)	...post w/o neg. reaction & adjusted	24.96%	11.77%	23.93%	11.37%

Note: The term λ^{pre} denotes the degree of partial insurance against permanent shocks. See text for details on the scenarios. The CEV columns denote the 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} .

additional calculations we did for the benchmark case. Results are reported in Table 4. As we discussed in Section 3.2, the obtained partial insurance parameters λ^{pre} now combine both, the partial insurance provided by the tax and transfer system, and the additional partial insurance that comes from other insurance channels. We follow (15) to back out the implied government insurance. Up to rounding error the obtained measures for partial insurance provided by the tax and transfer system, λ^{gov} , are effectively identical to the ones obtained in the benchmark case. This mimics the analytical result in the case of homoskedastic Gaussian risk discussed in Section 3.3.

Role of Risk Attitudes. In the preceding discussion we considered relatively weak risk attitudes. For these, it turned out that the non-log-Gaussian features of the distribution of permanent shocks do not play a role. As alluded to in the simple exercise in Section 2.2, when agents assign a higher utility weight to features such as skewness and kurtosis of the

Table 4: Partial Insurance and Welfare Gains of the Tax and Transfer System ($\lambda^{post} = 0.1$)

Scenario:		λ^{pre}	λ^{gov}	λ^{pre}	λ^{gov}
From pre-government income to...		\ln utility		Risk Aversion = 2	
(I)	...post-government income (post)	53.80%	48.66%	53.73%	48.59%
(II)	...post adjusted for initial dispersion	25.80%	17.56%	25.76%	17.51%
(III)	...post w/o cyclicalities	67.11%	63.46%	66.85%	63.17%
(IV)	...post w/o cyclicalities and adjusted	38.89%	32.10%	38.60%	31.78%
(V)	...post w/o reaction to negative x_t	57.72%	53.02%	59.59%	55.10%
(VI)	...post w/o neg. reaction & adjusted	29.66%	21.84%	31.49%	23.88%

Note: The term λ^{pre} denotes the degree of partial insurance against permanent shocks. See text for details on the scenarios. λ^{gov} gives the part associated with the tax and transfer system.

distribution, we expect distributional differences along these features to be reflected in the welfare-equivalent variance adjustment captured by λ^{gov} . At the same time, as long as the distribution is log-Normal, we expect the model-based λ^{gov} to be robust with respect to the strength of risk attitudes. Thus, we now reparameterize the model with stronger risk attitudes and reconsider the benchmark exercise that compares gross and net incomes.

In Table 5, we show the extent of variance adjustment that makes households ex ante indifferent to facing the stream induced by the tax and transfer system for $\gamma = 3$ and $\gamma = 4$; in the first rows we repeat the results for $\gamma = 1$ and $\gamma = 2$ from Table 3 to ease comparison. The first main take away is that we can confirm that under a (counterfactual) log-Normal distribution the attributed degree of government-provided insurance is robust.

For the distributions with higher-order risk, we see differences in the welfare-equivalent variance adjustment. Precisely, when agents have stronger risk attitudes ($\gamma = 4$), the estimated deviations from log-Normality imply a smaller welfare-equivalent variance adjustment (insurance of about 37% instead of 49%). This delivers insights into the relevance of different distributional changes induced by taxes and transfers: both negative and positive shocks

are dampened. The fact that a smaller welfare-equivalent variance reduction is necessary is indirect evidence for the latter to dominate—and this is something that agents do not like as much.¹⁰

Table 5: The Role of Risk Attitudes for Partial Insurance

From pre- to post-government income	Higher-Order Risk		Gaussian	
	$\lambda^{pre} = \lambda^{gov}$	<i>CEV</i>	$\lambda^{pre} = \lambda^{gov}$	<i>CEV</i>
<i>ln</i> utility	48.56%	11.73%	49.74%	11.40%
Risk Aversion = 2	48.46%	24.96%	49.58%	24.86%
Risk Aversion = 3	45.44%	40.62%	49.26%	42.10%
Risk Aversion = 4	37.40%	56.59%	48.82%	65.53%

Note: The term λ^{gov} denotes the degree of partial insurance against permanent shocks associated with the tax and transfer system. The *CEV* columns denote the 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} .

5 Comparison to Direct Estimates of Insurance

Overview. In the analysis up to now, we showed the possibility to assess the insurance value of taxes and transfers by means of a measurement tool that builds around a model which translates fundamental income risk into consumption. It needs two empirical inputs—estimated income processes reflecting regularities of gross income and net income, respectively—and is parameterized by the degree of partial insurance against net income shocks. While we emphasized the robustness of the results with respect to this parameterization, we can go further, and inform the framework by direct estimates of the pass-through of income to consumption. The resulting model-based measure of government insurance can then be compared to the corresponding measure based on directly estimated pass-through coefficients—despite potential measurement error of consumption, which we emphasized in the motivating discussion.

¹⁰This is consistent with the interpretation given in [Busch *et al.* \(2022\)](#), who document that the reduction of cyclicity of skewness of income changes in Sweden is mostly driven by less pronounced upside chances in expansionary periods.

Thus, we now proceed in two steps. For both we turn to studying households in the United States using data from the Panel Study of Income Dynamics (PSID). First, we take an established estimate of the pass-through from permanent shocks to household net income to consumption, directly from [Blundell *et al.* \(2008\)](#). We then use it to parameterize the measurement framework for an analysis of government provided insurance in the United States, which we inform by estimated income processes based on the PSID. Second, we revisit the estimation of [Blundell *et al.* \(2008\)](#) using later waves of the PSID, which allow for a consistent measure of some consumption components. We directly estimate pass-through coefficients for both, gross and net household income. Together, these two estimates yield a direct measure of government insurance. This, again, we compare to the one based on the measurement framework using income data alone. The described exercise delivers consistent estimates of government provided insurance in the United States of about 24%.

Parameterizing with BPP estimates. We estimate two income processes for the PSID using the income sample from [Busch *et al.* \(2022\)](#). Precisely, we fit process (19) using the same set of moments as for our Swedish sample, whereby we take into account that the PSID is only biannual after 1997 and thus instead of one-, three-, and five-year changes, we target moments of two-, four-, and six-year changes instead. In [Appendix A.3](#), we report the parameter estimates and time series of moments. The estimated process features distributions of permanent shocks which vary over the cycle in a systematic way, qualitatively in line with our estimates for Sweden.

Recall from the discussion in [Section 2](#) that the pass-through coefficient from permanent income shocks to consumption changes captures the degree of partial insurance. The key challenge in estimating this pass-through coefficient is that permanent shocks η_t are not directly observed—only total income y_t and consumption c_t are measured. [Blundell *et al.* \(2008\)](#) address this identification problem by constructing an instrument from leads and lags of income growth:

$$\tilde{\eta}_t^{BPP} = \Delta y_{t-1} + \Delta y_t + \Delta y_{t+1}, \quad (25)$$

which, under the permanent-transitory decomposition of income, isolates the permanent component while being uncorrelated with transitory shocks and measurement error. This yields the IV/GMM estimator:

$$\hat{\beta}_{perm} = \frac{\text{cov}(\Delta \ln c_t, \Delta y_{t-1} + \Delta y_t + \Delta y_{t+1})}{\text{cov}(\Delta y_t, \Delta y_{t-1} + \Delta y_t + \Delta y_{t+1})}. \quad (26)$$

Under the model assumptions, the instrument is correlated with Δy_t through the permanent component η_t but uncorrelated with transitory shocks and measurement error, and thus this is a consistent estimator of the pass-through of permanent shocks. The benchmark estimate in [Blundell *et al.* \(2008\)](#) implies a degree of partial insurance against permanent shocks to net income of $\lambda^{pos} = 0.36$. When feeding our estimated income processes into the model parameterized accordingly, we obtain a model-based measure of government provided insurance amounting to $\lambda^{gov} = 0.2391$.

Comparison to direct estimates. In order to compare this to a direct estimate based on pass-through coefficients for both pre- and post-government income, we add the consumption measures from recent survey waves (the PSID began reporting consumption in 1999) to our income sample. We define consumption as the sum of three consistently measured non-durable components, listed in [Table 6](#).¹¹ Household consumption is adjusted to real values using the same price index applied to income data and equivalized using the square root of family size.

Table 6: Components of Consumption Measure

Component	Raw PSID variables	Description
Food	FDHM + FDOUT + FDDEL	Food at home, out, and delivered
Public transportation	BUS + CAB + OTRAN	Public transit (bus, cab, other)
Utilities	HEAT + ELECTR + + WATER + OUTIL	Gas/heat, electricity, water, other utilities

Notes: Components of consumption used in the analysis.

¹¹While, the PSID reports additional non-durable categories starting in 2005, we focus on the components available starting in 1999, because the instrument requires three consecutive observations of changes per household, which limits the effective estimation window.

Table 7 presents two-step GMM estimates of the pass-through coefficient β_{perm} for both post-government and pre-government income. Both coefficients are significantly different from zero, confirming that permanent income shocks do affect household consumption: The post-government pass-through of $\hat{\beta}_{perm}^{post} = 0.242$ indicates that about 24% of a permanent shock to disposable income passes through to consumption, implying an insurance coefficient of $\hat{\lambda}^{post} = 1 - 0.242 = 0.758$; The pre-government pass-through of $\hat{\beta}_{perm}^{pre} = 0.183$ indicates that about 18% of a permanent shock to gross income passes through, implying $\hat{\lambda}^{pre} = 0.817$.

Table 7: GMM Estimates of Pass-Through Coefficients in PSID

$\hat{\beta}_{perm}^{post}$ (Post-Government Income)	0.242	(0.067)
$\hat{\beta}_{perm}^{pre}$ (Pre-Government Income)	0.183	(0.046)

Notes: Two-step GMM estimates with robust standard errors clustered at the individual level.

These estimates differ substantially from the insurance coefficient in [Blundell *et al.* \(2008\)](#), which corroborates [Ghosh and Theloudis \(2025\)](#), who also estimate higher insurance in an analysis of the newer PSID waves. Several factors potentially account for this discrepancy. First, our sample covers a more recent period (1999–2010), and the tax-and-transfer system and other household insurance opportunities may have evolved since the 1980s (their analysis builds on PSID data from 1978–1992). Second, the exact steps revolving around sample selection, variable construction, and consumption measurement can affect estimates. Third, the biennial structure of the post-1997 PSID—in contrast to the annual data used in the original BPP study—changes the horizon over which income and consumption growth are measured in our data, potentially attenuating estimated pass-through coefficients if shocks are partially smoothed within two-year intervals.

However, we are not interested here in estimates of consumption insurance *per se*. Instead, we want to assess the amount of partial insurance attributable to the tax and transfer system, given estimates of pass-through coefficients. The key qualitative finding that $\hat{\beta}_{perm}^{pre} < \hat{\beta}_{perm}^{post}$ implies positive government insurance and is consistent with BPP’s original conclusion. Now consider the relationship in Equation (15), which we repeat here for convenience: $(1 - \lambda^{pre}) = (1 - \lambda^{post})(1 - \lambda^{gov})$. Thus, we can back out the partial insurance provided by the tax and

transfer system from the two pass-through estimates as $\hat{\lambda}_{BPP}^{gov} = 1 - \frac{\hat{\beta}_{perm}^{pre}}{\hat{\beta}_{perm}^{post}} = 1 - \frac{0.183}{0.242} = 0.245$: According to the estimated pass-through coefficients, the U.S. tax-and-transfer system insures about 24.5% of permanent income shocks.

This lines up with the model-based measure obtained under the parameterization with the BPP estimate of $\lambda^{pos} = 0.36$. When using our direct estimate of $\lambda^{pos} = 0.76$ instead, the measurement framework yields effectively the same insurance value of taxes and transfers of 0.2426. Thus, the measurement framework based on the income processes alone yields an insurance value in line with what is implied by direct estimates of consumption insurance. We summarize the three insurance coefficients in Table 8. Thus, in our PSID sample, we estimate a substantially lower degree of insurance than the 48.6% we estimate for Sweden in our benchmark analysis (Table 3). This cross-country difference aligns with the broader literature documenting more comprehensive social insurance in Scandinavian countries relative to the United States.

Table 8: Government Insurance in PSID—Model-Based vs. Direct Estimates

	Parameterized Model		Direct Estimates of Pass-through
	$\lambda^{post} = 0.36$	$\lambda^{post} = 0.76$	$\lambda^{pre} = 1 - 0.183$; $\lambda^{post} = 1 - 0.242$
λ^{gov}	0.239	0.243	0.245

Notes: Degrees of government insurance implied by direct estimates of pass-through coefficients and by the model-based measurement device under ln utility.

As before, we expect the model-based measure to change when moving to a parameterization with stronger risk attitudes. We summarize corresponding results in Table 9. Consider first the Gaussian case again, which is shown in the last column of the table. The welfare-equivalent variance reduction associated with the tax and transfer system is again quite robust regarding risk attitudes. When using the income process estimates that capture non-Gaussian features, our measurement framework parameterized with stronger risk attitudes attributes a larger welfare-equivalent variance reduction to the tax and transfer system (up to 31%). This is consistent with the mirror image to the explanation in the Swedish case: In the PSID, the estimates imply that the dampening of downside movements of permanent

income changes dominates. Again, this is also consistent with the interpretation given in [Busch *et al.* \(2022\)](#), who emphasize the role of transfers that buffer against income losses in the PSID for the cyclical dynamics of skewness of income changes.

Table 9: Government Insurance in PSID—Role of Risk Attitudes

From pre- to post-government income	Higher-Order Risk	Gaussian
	λ^{gov}	λ^{gov}
\ln utility	24.26%	25.30%
Risk Aversion = 2	25.76%	25.50%
Risk Aversion = 3	28.04%	25.88%
Risk Aversion = 4	31.17%	26.43%

Notes: Degrees of government insurance for different parameterizations of the model; $\lambda^{post} = 0.76$.

These insights on the role of risk attitudes together with the insights from [Table 8](#) are also informative about the limits of the BPP insurance coefficients. The insurance coefficients exploit the covariance structure of income changes and consumption changes. Thus, when comparing coefficients for gross and net incomes in order to draw conclusions about the insurance provided by taxes and transfers, one effectively compares the variances of changes for the two income measures. By construction, this misses systematic differences of higher-order risk. Thus, when higher-order risk is present and affected by taxes and transfers—and when agents care about this—the BPP insurance coefficients miss a relevant part of insurance. This complements similar arguments made in [Busch and Ludwig \(2024\)](#), [Guvenen *et al.* \(2024\)](#) and [Ghosh and Theloudis \(2025\)](#) regarding non-linear consumption insurance that is not captured by the original BPP coefficient.

6 Conclusion

The tax and transfer system partially insures households against individual income risk. We introduce a framework to assess the value of this insurance. It is built around an incomplete-markets model and translates differences between distributional regularities of household gross income and disposable income into a welfare-equivalent variance adjustment. In the standard Gaussian case with homoskedastic shocks, the ratio of the dispersion of permanent shocks to gross and disposable incomes provides a sufficient statistic for the degree of

government-provided insurance, independent of preference parameters. In the general case, when taxes and transfers are also allowed to reshape cyclical higher-order moments, risk attitudes determine how these features map into the welfare-equivalent variance adjustment.

The framework is applicable in any typical setting with administrative income data, and especially suited for settings in which consumption data is absent. In an application to such data from Swedish tax registers, we find that the tax and transfer system insures about 49% of permanent income shocks. Resorting to the Panel Study of Income Dynamics, we also document that the model-based measure aligns well with empirical estimates based on survey data on consumption, implying a degree of insurance of about 24% in the United States, which increases to 31% under stronger risk attitudes.

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.
- BOERMA, J. and KARABARBOUNIS, L. (2021). Inferring inequality with home production. *Econometrica*, **89** (5), 2517–2556.

- 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.
- 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.
- GHOSH, A. and THELOUDIS, A. (2025). Consumption Partial Insurance in the Presence of Tail Income Risk.
- GUVENEN, F. (2011). Macroeconomics with heterogeneity: A practical guide. *Federal Reserve Bank of Richmond Economic Quarterly*, **97** (3), 255–326.
- , OZKAN, S. and MADERA, R. (2024). Consumption dynamics and welfare under non-gaussian earnings risk. *Journal of Economic Dynamics and Control*, **169**, 104945, heterogeneity and Macroeconomics of Labor Markets.

- , — 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.
- 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.
- 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

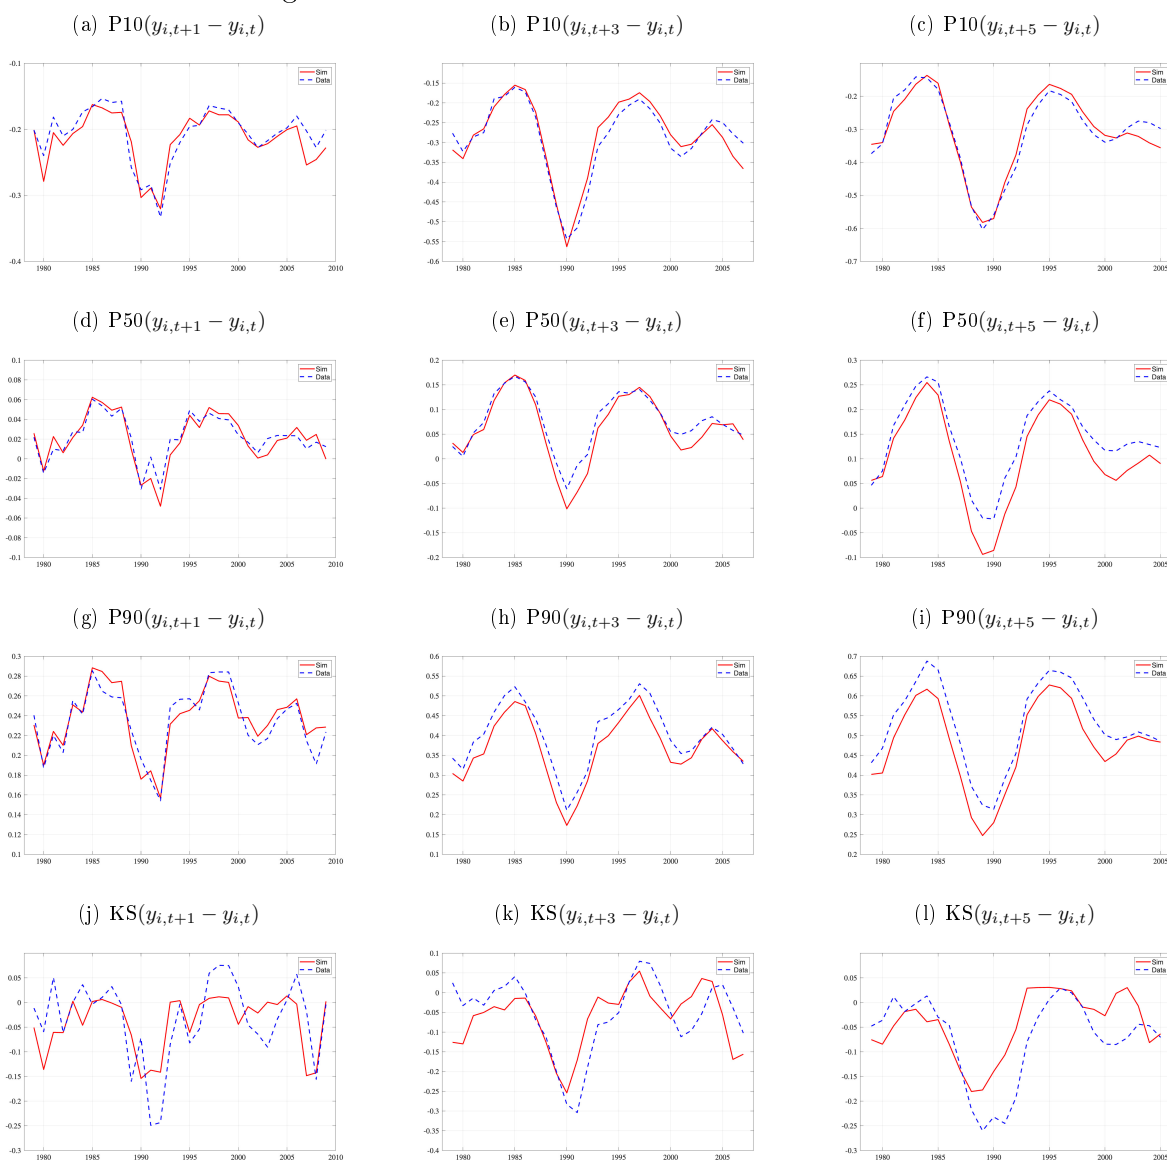
The global estimation procedure begins with a broad Global Search using the Differential Evolution algorithm (specifically the `de_rand_1_bin_radiuslimited` variant in the `BlackBoxOptim.jl` package), designed to scan the entire landscape of possible parameter values without getting trapped in false solutions. Once this algorithm identifies the neighborhood of the best fit, the code switches to a local step using the Nelder-Mead algorithm.

Intuitively, you can think of Differential Evolution as a process of natural selection. It starts with a diverse population of random guesses scattered across the map. In every generation, it creates new candidate solutions by mixing features from existing ones, taking a target solution and nudging it by the difference between two other random solutions. If the new mix performs better (fits the data closer), it survives to the next generation; if not, it is discarded. Over time, the entire population gravitates toward the optimal solution.

A.2 Data Fit of Estimated Income Processes for Sweden

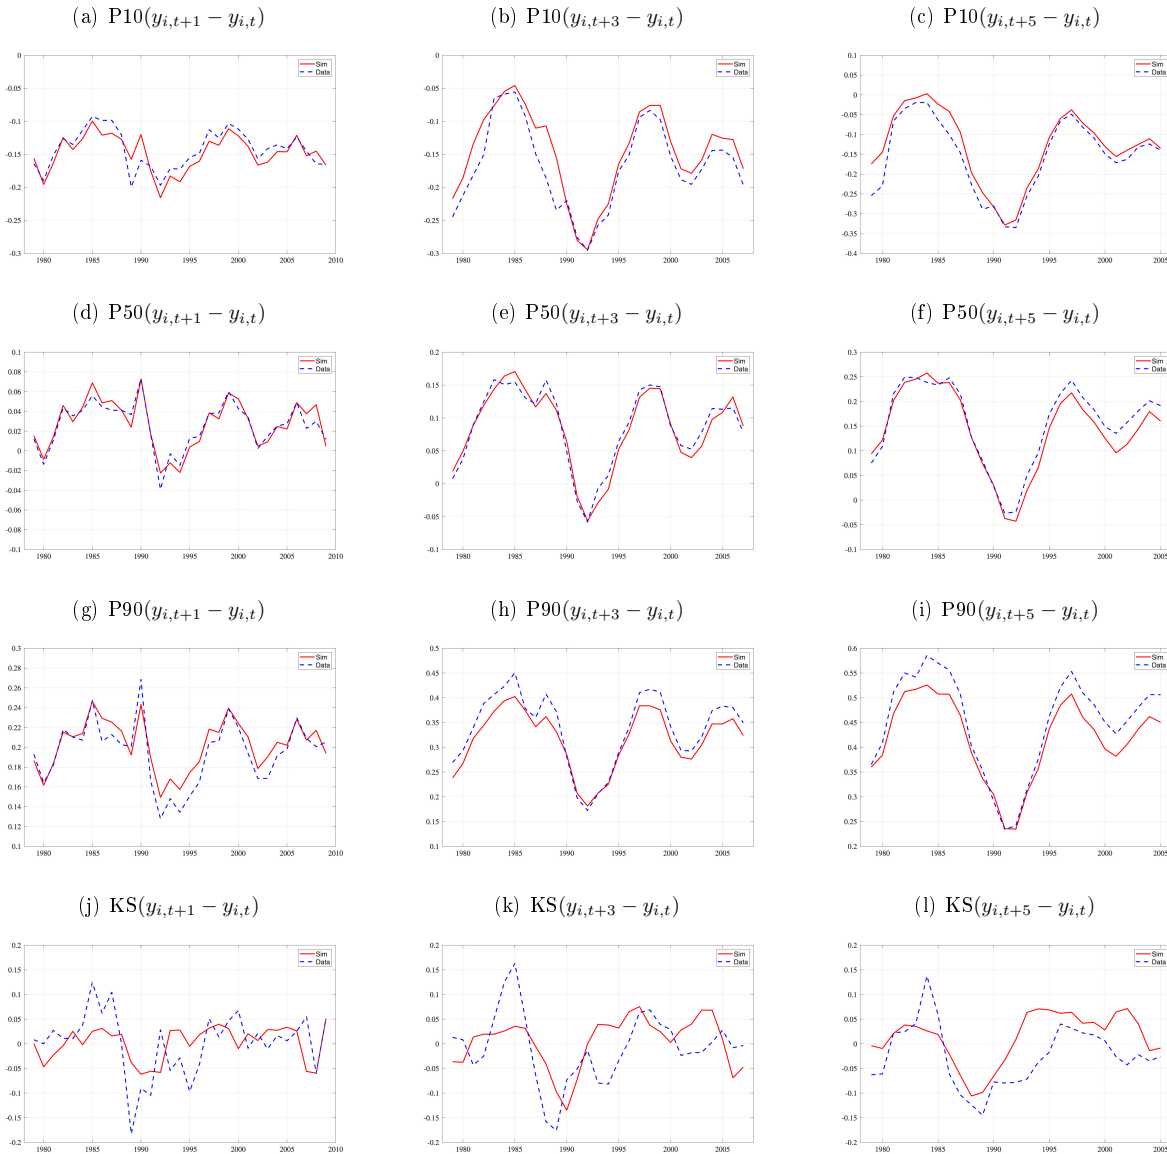
Figures [A.1](#) and [A.2](#) show moments implied by the estimated income processes for pre- and post-government household income in Sweden along with the data counterparts of the targeted set of moments.

Figure A.1: Sweden: 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 A.2: Sweden: Post-Government Income Fit



Note: See notes to Figure A.1.

A.3 PSID Estimation

Table A.1 shows parameter estimates of the processes for the PSID; Figures A.3 and A.4 show moments implied by the estimates. Figure A.5 illustrates four-year distributions of the permanent component for a period of strong aggregate growth: from 1982–86, GDP grew about 15.7%; while over the four year period from 2007–11 GDP initially dropped markedly during the onset of the Financial Crisis, before it started to mildly grow again between 2010

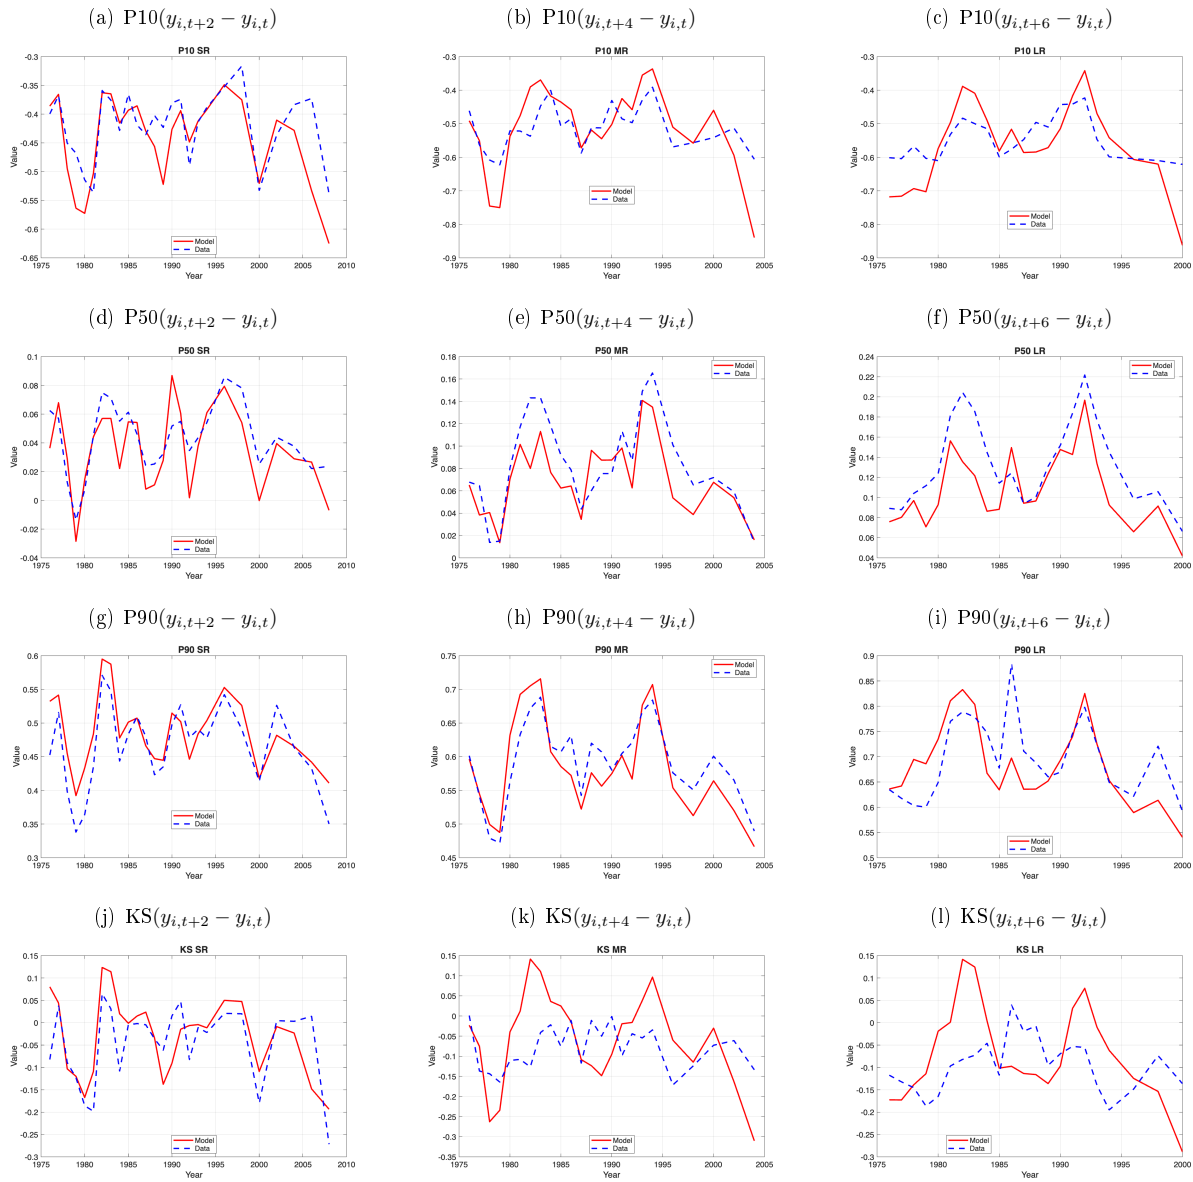
and 2011, making for an overall accumulated GDP growth of about 1% during the four years, far below the average growth trajectory. The implied distribution of permanent changes is right-skewed in 1982–86 and left-skewed in 2007–11.

Table A.1: Estimated Parameter Values (USA)

Parameter	Description	Pre-Gov.	Post-Gov.
$p_{\varepsilon,1}$	Mixture prob. of ε distribution	0.9010	0.9459
$\sigma_{\varepsilon,1}$	Std. dev. of ε mix. comp. 1	0.1566	0.1570
$\sigma_{\varepsilon,2}$	Std. dev. of ε mix. comp. 2	0.9989	0.8585
$p_{\eta,1}$	Mixture prob. of η mix. comp. 1 (fixed)	0.9500	0.9500
$p_{\eta,2}$	Mixture prob. of η mix. comp. 2 (fixed)	0.0250	0.0250
$p_{\eta,3}$	Mixture prob. of η mix. comp. 3 (fixed)	0.0250	0.0250
$\sigma_{\eta,1}$	Std. dev. of η mix. comp. 1	0.0443	0.0498
$\sigma_{\eta,2}$	Std. dev. of η mix. comp. 2	0.5000	0.2096
$\sigma_{\eta,3}$	Std. dev. of η mix. comp. 3	0.5000	0.2096
$\mu_{\eta,2}$	Mean of η mix. comp. 2	0.0839	0.2026
$\mu_{\eta,3}$	Mean of η mix. comp. 3	-0.3353	-0.2292
ϕ_2	Aggregate risk transmission mix. comp. 2	0.4000	0.4012
ϕ_3	Aggregate risk transmission mix. comp. 3	0.4003	0.4073
M	# moments targeted in estimation	255	255

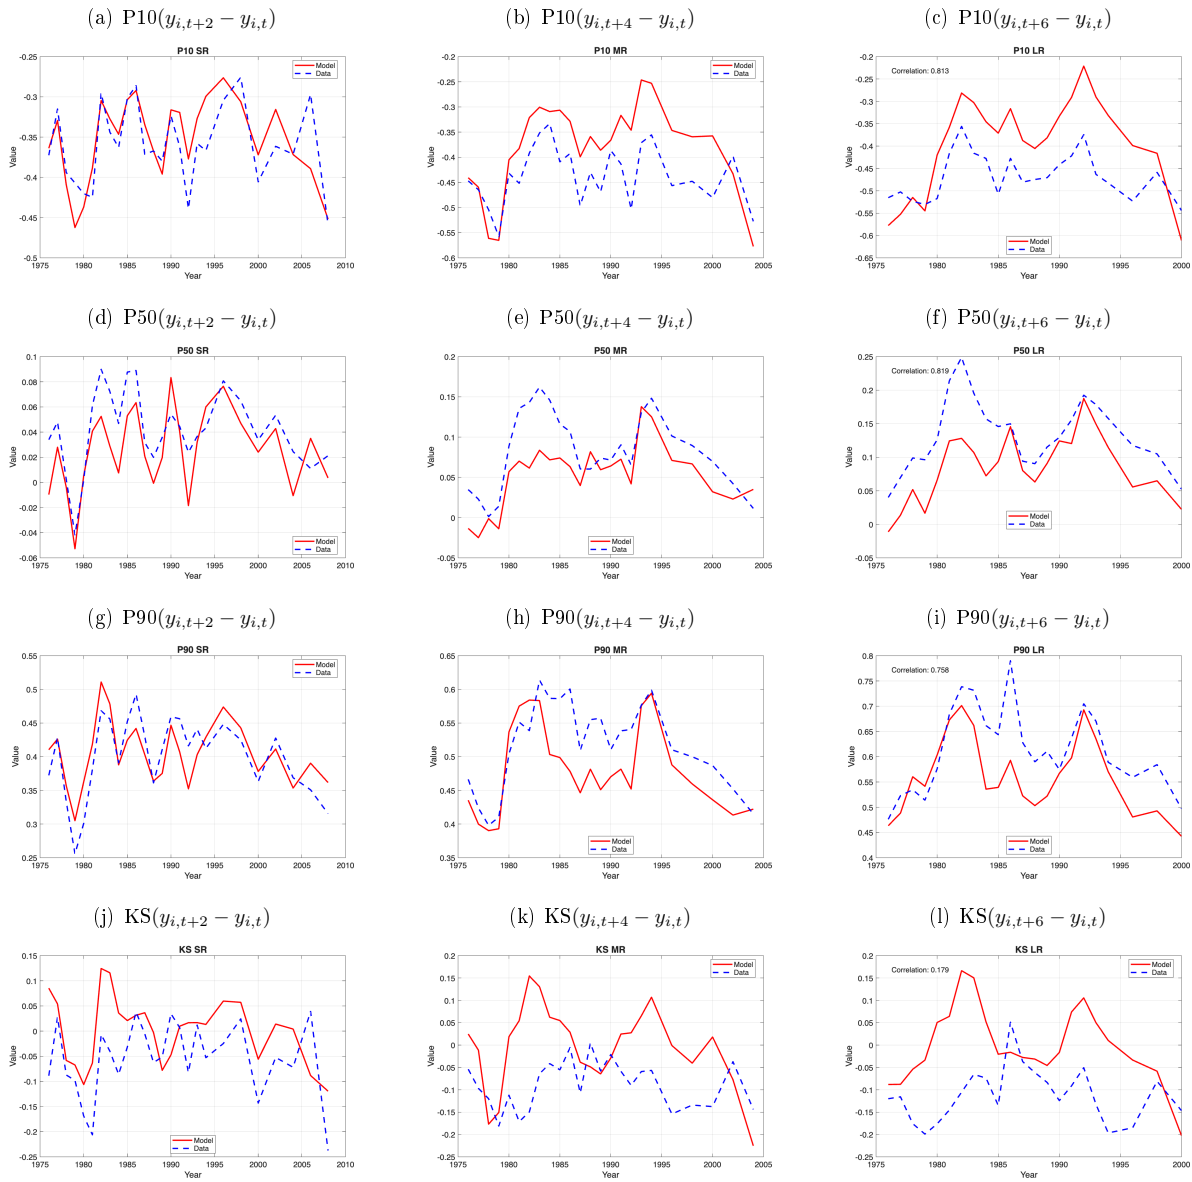
Note: Estimated parameters for gross household labor income (Pre-Gov.) and household income after taxes and transfers (Post-Gov.) in the PSID.

Figure A.3: USA: Pre-Government Income Fit



Note: See notes to Figure A.1.

Figure A.4: USA: Post-Government Income Fit




Note: See notes to Figure A.1.

Figure A.5: Cross-Sectional Distribution of Permanent Income Changes

(a) 1982–1986 (*GDP growth: 15.73%*)

(b) 2007–2011 (*GDP growth: 0.99%*)



Note: Each figure shows the distribution of simulated pre-government permanent income changes η in blue and post-government in red. Distributions are normalized to a mean of 1 in levels.

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 $1 - \lambda$.

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