Aggregate Shocks and Cross-Section Dynamics: Quantifying Redistribution and Insurance in US Household Data

Thomas J. Sargent, Yatheesan J. Selvakumar, and Ziyue Yang*

November 17, 2025

Abstract

This paper uses additive functionals and dynamic mode decompositions to analyze the co-evolution from 1990 to 2023 of cross-sections of private earned income, post-tax-and-transfer income, and consumption in the Consumer Expenditure Survey (CEX). We quantify how cross-sectional inequality and redistribution interact with aggregate income. We construct value functions for quantiles of synthetic consumers who are exposed to both i.i.d. and serially correlated risks in income and consumption growth rates. For the median household, welfare costs from serially correlated risk are orders of magnitude larger than welfare costs from i.i.d. risk. For each quantile, we also compare the benefits of eliminating risks in consumption growth with the benefits from participating in the US tax and transfer system. In absolute values, benefits from the latter, which are positive (negative) for low (high) quantile consumers, far exceed those from the former.

Keywords: Long-run risk, welfare, heterogeneity, common trends, business cycles.

JEL Classification: E3, E6, C7

*We thank Fernando Alvarez, Greg Kaplan, Lars Peter Hansen, and Jia Li for helpful comments at preliminary presentations of our work.

1 Introduction

This paper uses a parametric statistical model of cointegrated CEX cross-section quantiles to organize evidence about how aggregate income interacts with insurance and redistribution.[1] We infer substantial heterogeneity in how aggregate shocks affect different quantiles, and our welfare calculations quantify how the existing US tax and transfer system provides significant insurance and redistribution, particularly for lower-income households, with benefits that vastly exceed those from further stabilizing the business cycle. We confirm patterns that rationalize representative agent macroeconomic models and an associated "neo-classical synthesis" that recommends separating monetary-fiscal policies to moderate aggregate business cycles from microeconomic policies that redistribute income and protect against idiosyncratic shocks.[2]

We construct what Koopmans (1947) called a "Kepler stage" descriptive model of business cycles to distinguish it from the "Newton stage" structural models that Koopmans and his Cowles Commission colleagues invented to study consequences of historically unprecedented government policies.[3] Koopmans interpreted Burns and Mitchell's (1946) approach as an ad hoc algorithm designed to infer a single business cycle "factor" that intermediates co-movements of a diverse set of quantities and nominal prices.[4] Borrowing ideas about additive and multiplicative functionals from Hansen (2012), we use trend, martingale and stationary components of additive functionals to represent how aggregate income interacted dynamically with CEX cross-sections for private income, private income net of taxes and transfers, and consumption.[5] We synthesize an aggregate (earned) income series as a cross-section Chisini mean of private earned income from CEX quantiles.[6] We posit that this

[1]These CEX consumption and post-tax-and-transfer income cross-sections measure outcomes after the operation of monetary-fiscal policies aimed at stabilizing business cycles as well as the redistribution and insurance already present in the US tax and transfer system, meaning our data measure outcomes after the operation of these systems. This is important for interpreting our welfare comparisons in Section 5.

[2]For accounts of the "neoclassical synthesis" and its origins, see Lucas (1987) and Sargent (2015, 2024, In Press).

[3]See Koopmans (1950), Hood and Koopmans (1953), and Marschak (1953).

[4]As emphasized by Koopmans (1947) and Sargent and Sims (1977), the role of a purely descriptive parametric statistical model like ours is to detect and organize patterns that should be matched by a "structural" statistical model, cast in terms of parameters that can be interpreted as describing purposes and constraints of economic decision makers who live inside the model.

[5]By basing our statistical representations on Hansen (2012), we avoid pitfalls described by Hamilton (2018) that are associated with the Hodrick-Prescott pre-filtering method that is still widely used in macroeconomics.

[6]Thus, our measure of this "aggregated" variable emerges from measurements of diverse disaggregated variables, in the tradition of Burns and Mitchell.

aggregate income series is an additive functional that is driven by a first-order vector autoregression of aggregate income growth and quantiles of the three cross-sections scaled by aggregate income. Our specification implies that quantiles for the three CEX cross sections are also additive functionals that we can decompose into their trend, martingale, and stationary components.[7] We use these representations to quantify interactions between aggregate income and the three CEX cross-sections. We also use them as inputs into our section 5 welfare calculations. These calculations quantify how thoroughly the US tax and transfer system redistributes income across quantiles and insures different quantiles against shocks to income growth rates. Section 2 describes how we construct quantiles for the CEX data. Section 3 describes how we construct and represent additive functionals. Section 4 describes dynamic mode decomposition and their connection to vector autoregressions. It also presents findings about how aggregate income influences cross-sections. Section 5 constructs welfare comparisons like those of Lucas (1987, Sec. III) and Lucas (2003), who computed measures of the welfare of a representative consumer exposed to i.i.d. risk in consumption growth. We extend Lucas's analysis to study a cross section of consumers who are exposed to serially correlated risks in consumption growth. We then study welfare consequences from participating in the existing US tax and transfer system. We also study how different quantiles gain from increasing the rate of growth of log consumption and compare them with gains from eliminating exposures to consumption growth risk. Section 6 concludes. In appendices, we provide technical details and possible extensions.

2 Data Description and Compression

The Consumer Expenditure Survey (CEX) is a nationally representative survey of US households conducted by the Bureau of Labor Statistics. We study quarterly CEX waves from 1990 to 2023. Our statistical analysis focuses on three variables:[8] (1) Private income – labor income plus financial income, (2) Post-tax income – private income plus transfers minus taxes, and (3) Consumption.

For each time t, we rank all households in the sample by their real consumption level,

[7] Satisfaction of technical conditions described by Sargent et al. (2025) justify our use of a first-order VAR rather than the higher-order VAR that appears in the more general additive functional of Hansen (2012).

[8] See Appendix A for details on the data construction and preparation.

deflated by the PCE chain-type price index. Let $\hat{c}_{i,t}$ denote the real consumption of household $i = 1, \ldots, N$ at time t, where N is the total number of households.

We partition the ranked households into 100 percentiles. For each percentile $p = 1, \ldots, 100$, let $\mathcal{P}_p$ denote the set of households in that percentile. If N is divisible by 100, each set contains exactly $N/100$ households. Otherwise, the first 99 sets each contain $\lfloor N/100 \rfloor$ households, and the 100th set contains the remainder. We define the log real consumption of percentile p as

$$c_{p,t} = \log\left(\frac{1}{N_p}\sum_{i\in\mathcal{P}_p}\hat{c}_{i,t}\right), \quad p = 1, \ldots, 100,$$

where N_p denotes the number of households in percentile p.

We define the corresponding log private-income and post-tax income percentiles analogously and denote them by $m_{p,t}$ and $d_{p,t}$.[9] We construct these income measures in real terms by deflating with the PCE chain-type price index, as for consumption. To diminish the influence of outliers, we keep $M = 97$ quantiles indexed by $p \in \{2, 3, \ldots, 98\} =: \mathcal{P}$ for each variable. The vector $\mathbf{c}_t = [c_{2,t}, \ldots, c_{98,t}]^\top$ collects the remaining quantiles in an $(M \times 1)$ column vector for consumption.

Analogously, $\mathbf{m}_t = [m_{2,t}, \ldots, m_{98,t}]^\top$ and $\mathbf{d}_t = [d_{2,t}, \ldots, d_{98,t}]^\top$ collect the private- and post-tax income quantiles as $(M \times 1)$ column vectors. For each time period t, we define a Chisini mean (Chisini, 1929) of the cross-section of $\mathbf{u}_t \in \{\mathbf{m}_t, \mathbf{d}_t, \mathbf{c}_t\}$ as:

$$Y_t^{(u)} = \frac{1}{M}\sum_{p\in\mathcal{P}} u_{p,t} \tag{1}$$

Figure 1 compares nominal versions of our Chisini means from the CEX quantiles with corresponding aggregates from the National Income and Product Accounts (NIPA). Here $Y_t^{(m)}$, $Y_t^{(d)}$, and $Y_t^{(c)}$ denote Chisini means of private income, post-tax income, and consumption, respectively.[10] Levels of the series differ in panel (1a), especially for consumption, but

[9] For each of consumption, private income, and post-tax income we rank households separately, so a household's consumption percentile need not coincide with its percentiles in the income distributions.

[10] We compute private income using NIPA Table 2.1 on Personal Income and its Disposition. NIPA private income = Personal income (1) - (Employer) supplements to wages and salaries (6) - Government social benefits (17) + Contributions to government social insurance (25). The numbers in parentheses indicate the line item in NIPA Table 2.1.

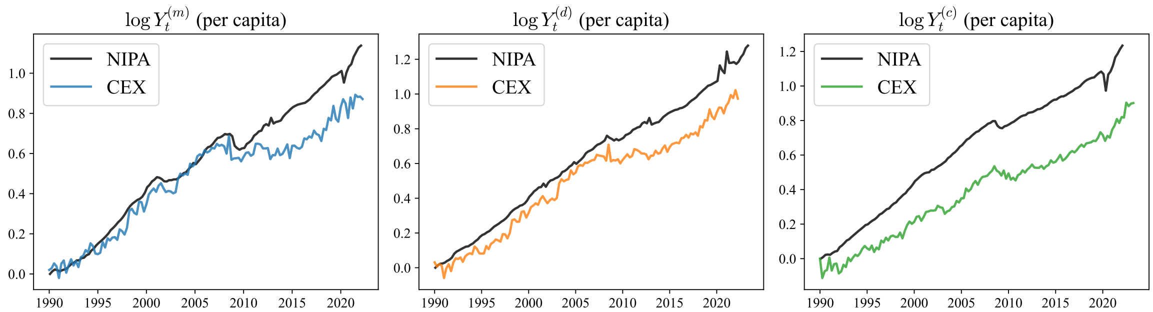

(a) Levels

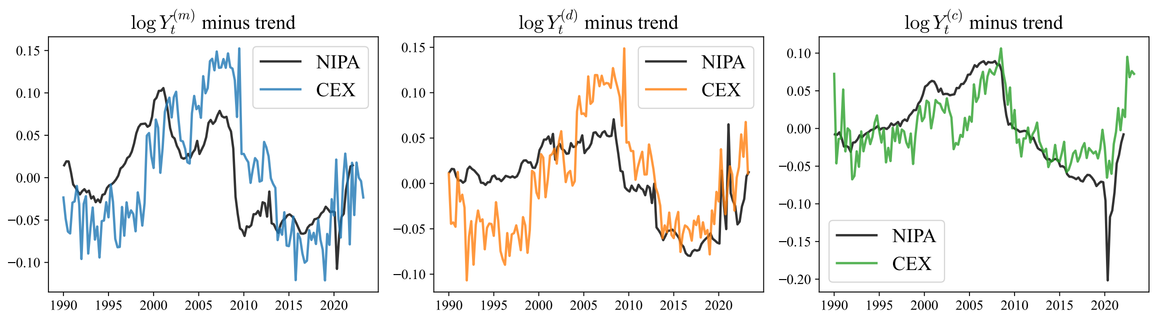

(b) Deviations from linear trends

Figure 1: NIPA and Chisini CEX means

their cyclical patterns in panel (1b) are broadly similar. Differences in consumption levels can be attributed to coverage and definition differences. NIPA consumption includes services to nonprofit institutions, government expenditures (Medicare/Medicaid), employer expenditures, owner-occupied rent, and financial services and insurance, which are not included in the CEX measure. According to Carroll et al. (2015), these differences account for somewhat more than a quarter of the gap between the two series from 1992 to 2010.

3 Descriptive Statistical Model

We construct our statistical model by first defining the key objects. Recall that $\mathbf{m}_t, \mathbf{d}_t, \mathbf{c}_t \in \mathbb{R}^M$ are vectors of quantiles for private income, post-tax income, and consumption at time t. For each $t \in \{1, \ldots, T\}$, we subtract the Chisini mean $Y_t := Y_t^{(m)}$ from all cross-section quantiles[11]:

$$\widetilde{\mathbf{m}}_t = \mathbf{m}_t - Y_t, \tag{2}$$

$$\widetilde{\mathbf{d}}_t = \mathbf{d}_t - Y_t, \tag{3}$$

$$\widetilde{\mathbf{c}}_t = \mathbf{c}_t - Y_t. \tag{4}$$

For each of the three variables (private income, post-tax income, and consumption), let $\boldsymbol{\nu}^{(i)} \in \mathbb{R}^M$ denote the vector of time-series means across quantiles, where $i \in \{m, d, c\}$ indexes the variable type. Let $\nu := \frac{1}{T-1}\sum_{t=2}^{T}(Y_t - Y_{t-1}) \in \mathbb{R}$ denote the time series mean of aggregate income growth. We define the $\mathbf{y}_t \in \mathbb{R}^{3M+1}$ state vector as

$$\mathbf{y}_t = \begin{bmatrix} Y_t - Y_{t-1} - \nu \\ \widetilde{\mathbf{m}}_t - \boldsymbol{\nu}^{(m)} \\ \widetilde{\mathbf{d}}_t - \boldsymbol{\nu}^{(d)} \\ \widetilde{\mathbf{c}}_t - \boldsymbol{\nu}^{(c)} \end{bmatrix}. \tag{5}$$

We assume that $\{Y_t\}$ is an additive functional driven by the following first order vector autoregression:

$$\mathbf{y}_{t+1} = \mathbf{B}\,\mathbf{y}_t + \mathbf{a}_{t+1} \tag{6}$$

$$Y_{t+1} - Y_t - \nu = e_1\,\mathbf{B}\,\mathbf{y}_t + e_1\,\mathbf{a}_{t+1}, \tag{7}$$

where $\mathbf{B} \in \mathbb{R}^{(3M+1)\times(3M+1)}$ is a transition matrix with spectral radius $\rho(\mathbf{B}) < 1$, $\mathbf{a}_{t+1} \in \mathbb{R}^{3M+1}$ are innovations satisfying $\mathbb{E}[\mathbf{a}_{t+1}] = \mathbf{0}$, $\mathbb{E}[\mathbf{a}_{t+1}\,\mathbf{y}_{t-j}^\top] = \mathbf{0}$ for all $j \geq 0$, $\mathbb{E}[\mathbf{a}_{t+1}\,\mathbf{a}_{t+1}^\top] = \boldsymbol{\Omega}$, and $e_1 = [1\ 0\ \cdots\ 0] \in \mathbb{R}^{1\times(3M+1)}$ is a row vector that selects the first entry of $\mathbf{y}_t$. Later in Section 4, we describe a dynamic mode decomposition algorithm for estimating $\mathbf{B}$ outlined in Sargent et al. (2025).

Following Hansen (2012)[12], we can decompose the additive functional Y_t into trend, mar-

[11] By construction, the cross-sectional mean $\frac{1}{M}\sum_{p\in\mathcal{P}} \tilde{m}_{p,t} = 0$ for all t. We compute cross quantile means and standard deviations for each variable. For example, for private income we compute $\mu_t^{(m)} := \frac{1}{M}\sum_{p\in\mathcal{P}} \tilde{m}_{p,t}$ and $\sigma_t^{(m)} := \sqrt{\frac{1}{M}\sum_{p\in\mathcal{P}}(\tilde{m}_{p,t} - \mu_t^{(m)})^2}$. We report these moments in Figure 11 in Appendix A.

[12] Also see the QuantEcon lecture on additive and multiplicative functionals available here https://python-advanced.quantecon.org/additive_functionals.html.

tingale, stationary, and constant components.

Proposition 1 (Additive Functional Decomposition). Suppose $\mathbf{y}_t$ satisfies (6) with $\rho(\mathbf{B}) < 1$. Then Y_t admits the decomposition

$$Y_t = \underbrace{t\nu}_{\text{Trend}} + \underbrace{\sum_{j=1}^{t} \mathbf{H}\,\mathbf{a}_j}_{\text{Martingale}} + \underbrace{(-\mathbf{g}\,\mathbf{y}_t)}_{\text{Stationary}} + \underbrace{(\mathbf{g}\,\mathbf{y}_0 + Y_0)}_{\text{Constant}}. \tag{8}$$

where $\mathbf{H}, \mathbf{g} \in \mathbb{R}^{1\times(3M+1)}$ are defined by

$$\mathbf{H} = e_1 + e_1\,\mathbf{B}(\mathbf{I} - \mathbf{B})^{-1}, \tag{9}$$

$$\mathbf{g} = e_1\,\mathbf{B}(\mathbf{I} - \mathbf{B})^{-1}. \tag{10}$$

Proof. Define $Z_t \coloneqq Y_t + \mathbf{g}\,\mathbf{y}_t$. Using (6) and (7),

$$Z_{t+1} - Z_t - \nu = (Y_{t+1} - Y_t - \nu) + \mathbf{g}(\mathbf{y}_{t+1} - \mathbf{y}_t) = \left[e_1\,\mathbf{B} + \mathbf{g}(\mathbf{B} - \mathbf{I})\right]\mathbf{y}_t + (e_1 + \mathbf{g})\,\mathbf{a}_{t+1}\,.$$

Choose $\mathbf{g}$ such that $\mathbf{g}(\mathbf{B} - \mathbf{I}) + e_1\,\mathbf{B} = \mathbf{0}$ which yields $\mathbf{g} = e_1\,\mathbf{B}(\mathbf{I} - \mathbf{B})^{-1}$ as in (10). With this choice, we have

$$Z_{t+1} - Z_t - \nu = \mathbf{H}\,\mathbf{a}_{t+1}, \qquad \mathbf{H} \coloneqq e_1 + \mathbf{g},$$

and substituting the expression for $\mathbf{g}$ from (10) gives $\mathbf{H} = e_1 + e_1\,\mathbf{B}(\mathbf{I} - \mathbf{B})^{-1}$, i.e., (9). By construction of $\mathbf{a}_{t+1}$, it also follows that $\{Z_t - t\nu\}$ is a martingale with increments $\mathbf{H}\,\mathbf{a}_{t+1}$, so $Z_t = Z_0 + t\nu + \sum_{j=1}^{t} \mathbf{H}\,\mathbf{a}_j$. Substituting $Z_t = Y_t + \mathbf{g}\,\mathbf{y}_t$ and rearranging yields (8) as desired. □

The four components in (8) are: (i) a deterministic trend $t\nu$, (ii) a martingale $\sum_{j=1}^{t} \mathbf{H}\,\mathbf{a}_j$, (iii) an asymptotically stationary component $-\mathbf{g}\,\mathbf{y}_t$, and (iv) a constant $(\mathbf{g}\,\mathbf{y}_0 + Y_0)$.[13] Figure 2 presents decomposition (8). We estimate all model parameters ν, $\boldsymbol{\nu}^{(m)}$, $\boldsymbol{\nu}^{(d)}$, $\boldsymbol{\nu}^{(c)}$, $\mathbf{B}$, $\boldsymbol{\Omega}$ with CEX data through Q4 2008. Thus, $t\nu$ is our estimate of "pre-crisis trends". Shortly before the 2008 financial crisis, the martingale component was close to zero and Y_t was above its deterministic trend. After 2008, Y_t fell and remained persistently below trend until the end

[13] The martingale and stationary components are typically correlated.

of the sample. Initially, the martingale component dropped to minus 15%. The stationary component declined gradually and became negative in 2015, while the martingale component increased until 2020. The fall in Y_t during 2021 came mostly from the stationary component. At the end of the sample, Y_t was around 15% below trend. Below we shall describe substantial heterogeneity in quantile-specific stationary components after 2008.

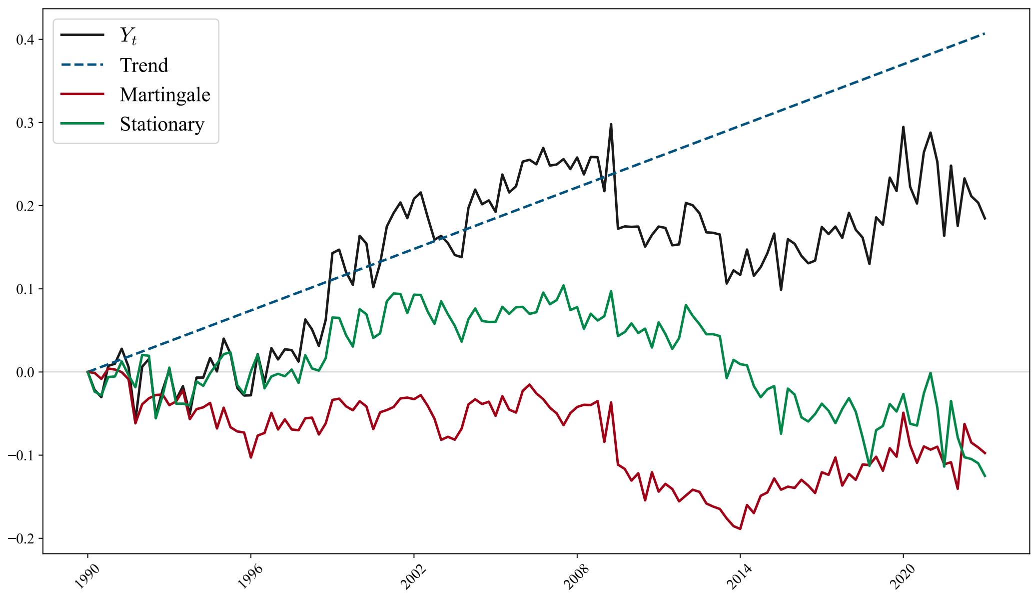

Figure 2: Decomposition of $Y_t - Y_0$

Corollary 1 (Quantile Decompositions). Under the conditions of Proposition 1, each quantile $m_{p,t}$, $d_{p,t}$, and $c_{p,t}$ for $p \in \mathcal{P}$ admits a decomposition into trend, martingale, stationary, and constant components.

Proof. We show this for the pth percentile of private income. Define $e_p^{(m)} \in \mathbb{R}^{1\times(3M+1)}$ to select the component of $\mathbf{y}_t$ corresponding to $\tilde{m}_{p,t} - \nu_p^{(m)}$ where $\nu_p^{(m)}$ is the time series mean of $\tilde{m}_{p,t}$. By definition, $\tilde{m}_{p,t} = m_{p,t} - Y_t$ and, from (5), $e_p^{(m)}\, \mathbf{y}_t = \tilde{m}_{p,t} - \nu_p^{(m)}$. Hence $m_{p,t} = Y_t + e_p^{(m)}\, \mathbf{y}_t + \nu_p^{(m)}$. Substituting (8) for Y_t yields

$$m_{p,t} = \underbrace{t\nu}_{\text{Trend}} + \underbrace{\sum_{j=1}^{t} \mathbf{H}\,\mathbf{a}_j}_{\text{Martingale}} \underbrace{-(\mathbf{g} - e_p^{(m)})\,\mathbf{y}_t}_{\text{Stationary}} + \underbrace{(\mathbf{g}\,\mathbf{y}_0 + Y_0 + \nu_p^{(m)})}_{\text{Constant}}. \tag{11}$$

Define selection vectors $e_p^{(d)}$ and $e_p^{(c)}$ to pick the components $\tilde{d}_{p,t} - \nu_p^{(d)}$ and $\tilde{c}_{p,t} - \nu_p^{(c)}$ from

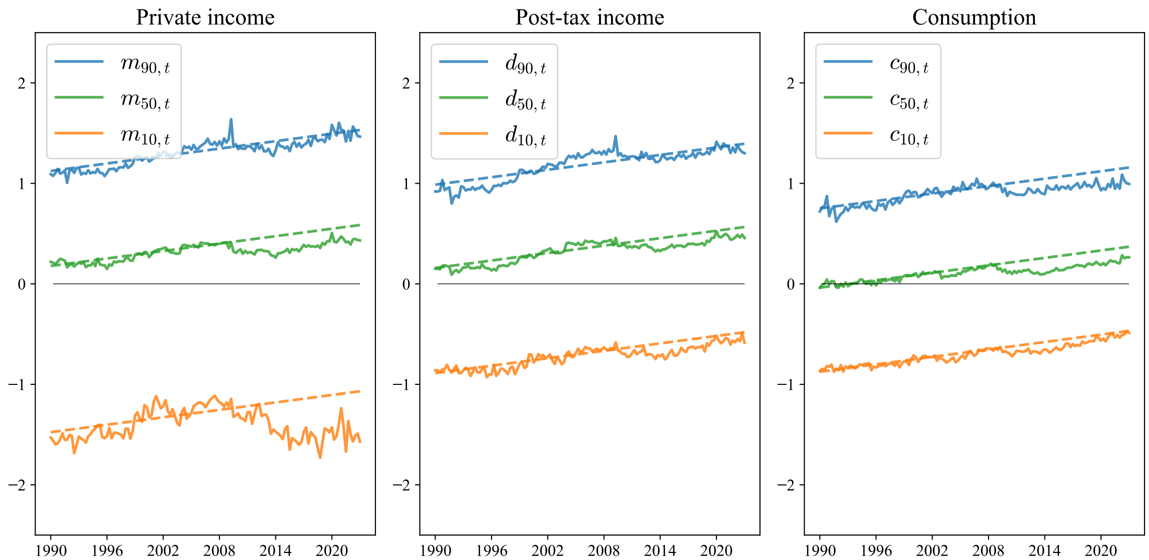

Figure 3: Evolution of quantiles

$\mathbf{y}_t$. Then the pth percentiles of post-tax income and consumption admit analogous decompositions. □

The decompositions (8) and (11) share common deterministic trend $t\nu$ and martingale $\sum_{j=1}^{t} \mathbf{H}\,\mathbf{a}_j$ components across all quantiles and variables, but have quantile-specific stationary and constant components. This structure is built into our statistical model specification (6)–(7).

In Figure 3, we use appropriate versions of (11) to decompose income and consumption quantiles. For the 10th percentile, private income is much lower than both post-tax income and consumption. For 90th percentile households, private incomes are higher than post-tax incomes, indicating substantial transfers from higher parts of the distribution. Differences across percentiles are largest for private income and smallest for consumption.

The decompositions in Figure 3 indicate interesting patterns in quantiles' deviations from the deterministic trend. After the 2008 financial crisis, private income of the 90th percentile fell 10 percent below trend, but by 2021 the deviation from trend was zero and it has remained there since then. Similar patterns apply to 90th percentile post-tax income. But after 2008, 90th percentile consumption fell more – around 15 percent below deterministic trend, and it stayed below deterministic trend through the end of our sample. As for lower quantiles, after 2008, private incomes at the 10th percentile fell around 50 percent below deterministic trend. Private income gradually recovered until 2020, but then fell a lot. By the end of the sample, 10th percentile private income was 55 percent below deterministic trend. But 10th percentile post-tax income and consumption fell much less, falling at most 15-20 percent below deterministic trend; by 2021, both of these 10th percentile variables had returned to trend, and they stayed there through the end of our sample.

A striking aspect of Figure 3 is that, at the 10th percentile, post-tax income is systematically higher and less volatile than private income, while post-tax income and consumption are very similar. This indicates substantial redistribution and insurance at the 10th percentile, but little consumption smoothing through management of private savings. Things look very different at the 90th percentile, where post-tax income is systematically lower than private income, and consumption is systematically lower than post-tax income. This indicates the presence of substantial redistribution and social insurance and also of consumption smoothing through management of private savings. These differences across quantiles shape differences in discounted expected utilities that we shall report in Section 5.

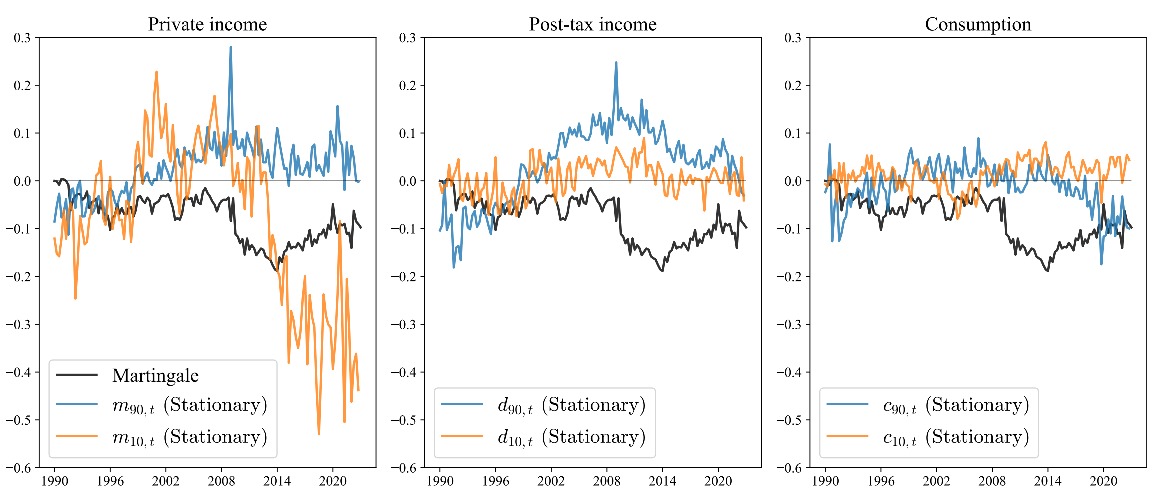

Figure 4: Components of quantiles

Figure 4 provides another perspective by using equation (11) to decompose deviations of quantiles from deterministic trend into two parts: the common martingale component of Y_t (black), which makes the same contributions to all quantiles of all three variables, and idiosyncratic stationary components (blue for 90th percentiles, orange for 10th percentiles).

The common martingale component (which is the same martingale that appears in Figure 2) fell 15% during the 2008 crisis and gradually recovered to about minus 4% at the end of our sample.

The left panel of Figure 4 shows that for the 10th percentile, martingale and stationary components both contribute to the large drop in private income relative to its deterministic trend. While the martingale component dominates initially, outcomes after 2014 are mostly driven by the stationary component, which accounts for the entire 50 percent gap from deterministic trend by 2023. On the other hand, after 2008 the stationary component for the 90th percentile was slightly positive, and by the end of the sample mostly offset the negative martingale component.

Turning to the middle panel of Figure 4, post-tax income stationary components for the 10th percentile were close to zero, making the common martingale component the predominant driver of the fall relative to trend. The stationary component for the 90th percentile that was positive after 2008, and so partially offset the negative martingale component, became slightly negative at the end of the sample.

Turning now to consumption displayed in the right panel of Figure 4, stationary components of both 10th and 90th percentiles were close to zero, so that the common martingale component makes the most important contribution to consumption. Toward the end of the sample, the stationary component started to push consumption below deterministic trend for the 90th percentile. That the idiosyncratic stationary components are, on the whole, close to zero for 10th and 90th consumption percentiles indicates pervasive effective risk-sharing.[14]

4 Dynamic Mode Decomposition

In this section, we discuss the dynamic mode decomposition (DMD) and associated reduced rank first-order VAR that we used to estimate parameters of our statistical model (6)-(7). In Section 4.1 we describe how a data reduction technique based on a singular value decomposition underlies the method. In Section 4.2 we briefly describe estimated parameters and plot

[14]This feature is consistent with the bottom right panel of Figure 13 that plots time-series of demeaned $\tilde{c}_{p,t}$ for different percentiles. It shows that consumption paths are highly correlated across quantiles, exhibiting very similar fluctuations in our sample. In macro models with preferences and complete market structures like those presented in Ljungqvist and Sargent (2018, ch. 8), highly correlated consumption paths like these indicate the presence of widespread sharing of aggregate risks.

four dominant DMD modes against components of our data set. In Section 4.3 we extract quantitative inferences about interactions between aggregate income and our cross-section quantiles.

4.1 Estimation

We estimated parameters of the VAR (6) by constructing a dynamic mode decomposition (DMD), a "machine-learning" technique used to study fluid dynamics[15] that Sargent et al. (2025) link to a vector autoregression and a linear Gaussian state-space system. Our analysis begins with a data set formed by quarterly observations of the vector $\mathbf{y}_t \in \mathbb{R}^{(3M+1)\times 1}$ defined in (5). We use these data to estimate the first-order VAR

$$\begin{aligned} \mathbf{y}_{t+1} &= \mathbf{B}\,\mathbf{y}_t + \mathbf{a}_{t+1} \\ \mathbb{E}[\mathbf{a}_{t+1}\,\mathbf{a}_{t+1}^\top] &= \mathbf{\Omega}, \quad \mathbf{a}_{t+1} \perp \mathbf{y}_t \end{aligned} \tag{12}$$

The coefficient matrix $\mathbf{B}$ has dimension $(3M+1)\times(3M+1)$ which, since $M = 97$, is 292×292 in our application. Our 33 years of quarterly observations mean that we have $T = 133$ observations with which to estimate the 292^2 coefficients in $\mathbf{B}$. Thus, specification (6) confronts us with an underdetermined least squares problem that we solve by using ideas connected to DMD. We set an integer $N \ll T$ and seek a rank-N first-order VAR representation

$$\begin{aligned} \mathbf{y}_{t+1} &= \widehat{\mathbf{B}}\,\mathbf{y}_t + \widehat{\mathbf{a}}_{t+1} \\ \mathbb{E}[\widehat{\mathbf{a}}_{t+1}\widehat{\mathbf{a}}_{t+1}^\top] &= \widehat{\mathbf{\Omega}}, \quad \widehat{\mathbf{a}}_{t+1} \perp \mathbf{y}_t\,. \end{aligned} \tag{13}$$

where $\widehat{\mathbf{B}}$ is a rank N matrix and

$$\widehat{\mathbf{a}}_{t+1} \coloneqq \mathbf{y}_{t+1} - \widehat{\mathbf{B}}\,\mathbf{y}_t, \qquad \widehat{\mathbf{\Omega}} \coloneqq \frac{1}{T-1}\sum_{t=1}^{T-1}\widehat{\mathbf{a}}_{t+1}\widehat{\mathbf{a}}_{t+1}^\top.$$

We can use the following algorithm to estimate such $\widehat{\mathbf{B}}$:

[15] See, for example, Tu et al. (2014) and Brunton and Kutz (2022).

1. Construct

$$\mathbf{Y} = \begin{bmatrix} \mathbf{y}_1 & \mathbf{y}_2 & \dots & \mathbf{y}_{T-1} \end{bmatrix}, \quad \mathbf{Y}' = \begin{bmatrix} \mathbf{y}_2 & \mathbf{y}_3 & \dots & \mathbf{y}_T \end{bmatrix}$$

where $\mathbf{Y}, \mathbf{Y}' \in \mathbb{R}^{(3M+1)\times(T-1)}$.

2. Compute the singular value decomposition (SVD)

$$\mathbf{Y} = \mathbf{U}\,\mathbf{\Sigma}\,\mathbf{V}^\top .$$

3. Retain the N leading singular values and vectors to form $\mathbf{U}_N\,\mathbf{\Sigma}_N\,\mathbf{V}_N$, so that

$$\mathbf{Y} \approx \mathbf{U}_N\,\mathbf{\Sigma}_N\,\mathbf{V}_N^\top .$$

4. Compute a $(3M+1) \times N$ matrix $\mathbf{\Phi}$ and an $N \times N$ diagonal matrix $\mathbf{\Lambda}$ with the following formulas:

$$\widetilde{\mathbf{B}} = \mathbf{U}_N^\top\,\mathbf{Y}'\,\mathbf{V}_N\,\mathbf{\Sigma}_N^{-1}, \qquad \widetilde{\mathbf{B}}\,\mathbf{W} = \mathbf{W}\,\mathbf{\Lambda}, \qquad \mathbf{\Phi} = \mathbf{Y}'\,\mathbf{V}_N\,\mathbf{\Sigma}_N^{-1}\,\mathbf{W},$$

with $\widetilde{\mathbf{x}}_t = \mathbf{\Phi}^+\,\mathbf{y}_t$ where $\mathbf{\Phi}^+$ is the Moore-Penrose pseudoinverse of $\mathbf{\Phi}$.

The rank N least squares estimator and its modal identities are

$$\widehat{\mathbf{B}} = \mathbf{Y}'\,\mathbf{V}_N\,\mathbf{\Sigma}_N^{-1}\,\mathbf{U}_N^\top, \quad \widehat{\mathbf{B}}\mathbf{\Phi} = \mathbf{\Phi}\mathbf{\Lambda}, \quad \widehat{\mathbf{B}} = \mathbf{\Phi}\mathbf{\Lambda}\mathbf{\Phi}^+. \tag{14}$$

Here $\mathbf{\Lambda}$ contains the DMD eigenvalues and $\mathbf{\Phi}$ the associated exact DMD modes for the rank N estimator $\widehat{\mathbf{B}}$.

4.2 Parameters and DMD Modes

Setting $N = 4$, we estimate the following parameters:[16]

[16]Visualizing the singular value of the data matrix $\mathbf{Y}$ provides an informal guide to choosing a suitable low-rank approximation of the data. The SVD has two dominant singular values associated with $\mathbf{Y}$ (see Figure 14 in Appendix B). Appendix B plots 50 singular values for our $\mathbf{Y}$ data matrix and associated dynamic modes for the largest three singular values. Using the elbow method for picking the number of modes that matches the largest descent in singular values, we choose $N = 4$ to capture the dominant dynamics.

$$
\mathbf{\Lambda} = \begin{bmatrix} 0.984 & 0 & 0 & 0 \\ 0 & 0.916 & 0 & 0 \\ 0 & 0 & 0.494 & 0 \\ 0 & 0 & 0 & 0.351 \end{bmatrix} \tag{15}
$$

Figure 5 plots the estimated DMD modes. We associate each mode $\hat{x}_{n,t}$ with $n \in \{1,2,3,4\}$ with time series of first and second moments of post-tax income and consumption with the highest correlations.

The first mode is highly persistent ($\lambda_1 = 0.984$) and strongly correlated with the mean of post-tax income ($\text{Corr}(\mu_t^{(d)}, \hat{x}_{1,t}) = 0.94$). The second mode is less persistent ($\lambda_2 = 0.916$) and correlated with the standard deviation of post-tax income ($\text{Corr}(\sigma_t^{(d)}, \hat{x}_{2,t}) = 0.87$). Modes 3 ($\lambda_3 = 0.494$) and 4 ($\lambda_4 = 0.351$) are further down the persistence spectrum, and they are correlated with the mean and standard deviation of consumption, respectively. Thus, the four dominant DMD modes we estimated capture the dynamics of first and second moments of post-tax income and consumption.[17]

Since the first two modes are far more persistent than 3 and 4, we focus on them in what follows and relegate discussion of modes 3 and 4 to Appendix C. Figure 6 shows loadings $\mathbf{\Phi}$ of quantiles on the two dominant modes. The left panel plots quantile-specific loadings on private income. Since $\mu_t^{(m)} = 0$ for all t by construction, the loadings of private income quantiles must average to zero, thus shedding light on increases in income inequality associated with an increase in mode 1.

Loadings of post-tax quantiles (middle upper panel) on mode 1 are all positive. This suggests that mode 1 serves as an aggregate income factor. It also shows that low and high quantile incomes are more sensitive to mode 1, confirming the "U"-shape of income sensitivities across the earnings distribution detected by Guvenen et al. (2017). The right panel plots loadings of consumption quantiles on mode 1. All loadings are positive or close to 0, suggesting that a rise in mode 1 coincides with a rise in aggregate consumption. Low quantiles are more sensitive to mode 1 than high quantiles and their sensitivity is similar across both post-tax income and consumption. This suggests limited consumption smoothing

[17] Modes are linked to the data moments by picking the time series of moments that have the highest absolute value of correlation with each mode. The absolute value is used because DMD modes and their amplitudes are only defined up to an arbitrary nonzero scaling.

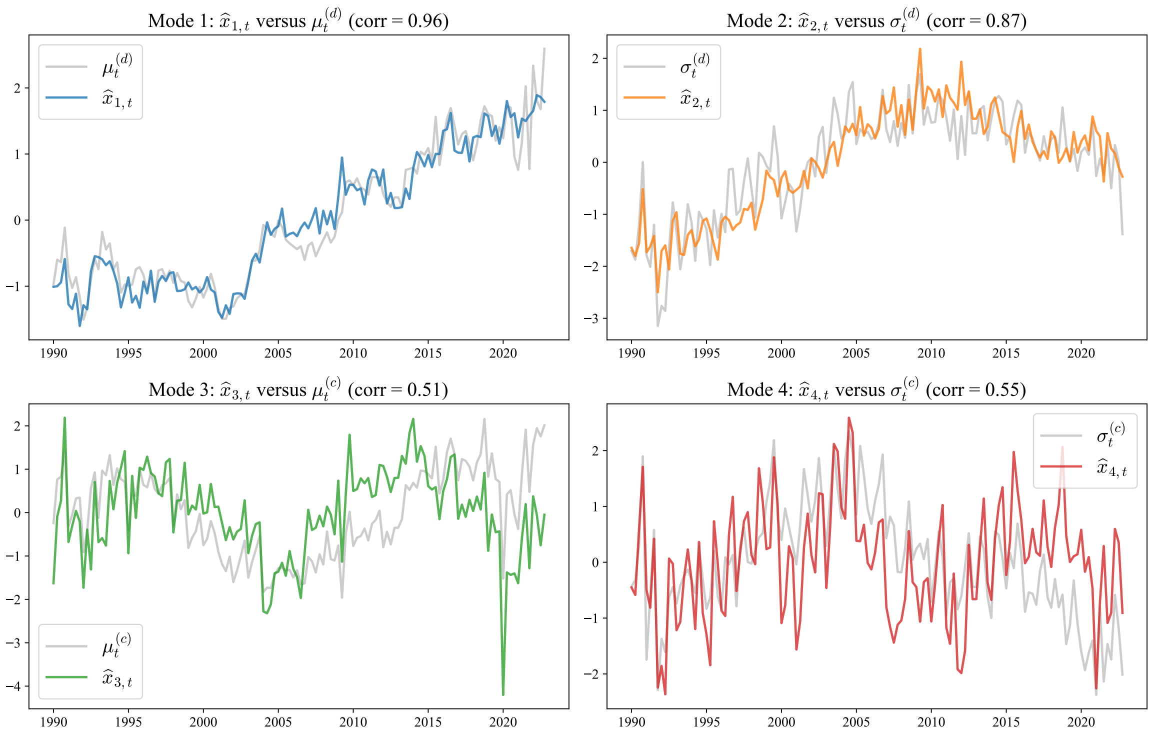

Figure 5: DMD modes and data moments

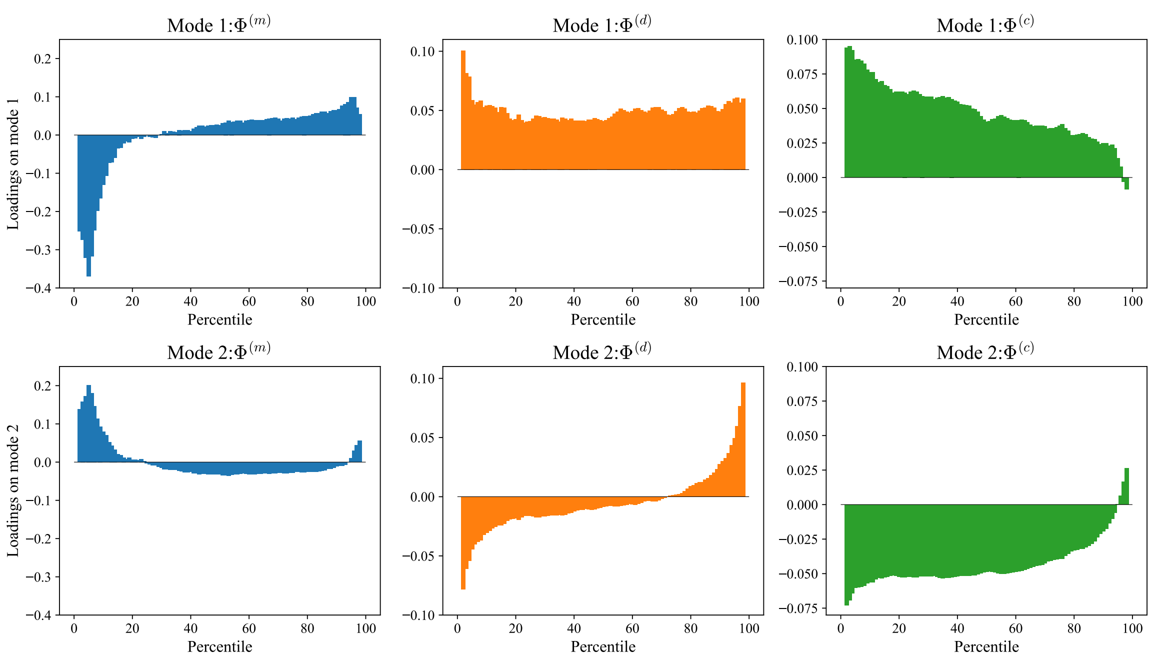

Figure 6: DMD loadings on modes 1 and 2

through management of their private savings. For low-quantile households, most smoothing of consumption relative to private income comes through the tax and transfer system embedded in post-tax income, rather than through self-insurance via private savings, so it is not fully captured in these loadings. High consumption quantiles are less sensitive to mode 1 than high post-tax income quantile counterparts, indicating more consumption smoothing through management of private savings than accomplished by low quantile consumers.

Turning to loadings on mode 2 in the lower three panels of Figure 6, post-tax income loadings are negative at the bottom of the distribution and positive at the top (middle lower panel), consistent with interpreting mode 2 as a variance factor for post-tax income. An increase in mode 2 therefore raises post-tax income inequality: high quantile incomes rise while low quantile incomes fall. Consumption loadings are negative for low and middle quantiles and become positive only in the top tail (middle lower panel), so a rise in mode 2 reduces consumption for most households while modestly increasing it for the richest, implying greater dispersion in consumption as the variance of post-tax income increases.

4.3 Micro-Macro Interactions

To investigate the influences of aggregate income growth $Y_t - Y_{t-1} - \nu$ on inequality dynamics, we compute population least squares regressions of innovations to various percentile differences on innovations to $Y_t - Y_{t-1} - \nu$. We can define a vector $\mathbf{b}$ to a quantile difference $z_{t+1} = \mathbf{b}\, \mathbf{y}_{t+1}$, (e.g., a 90–10 quantile difference for private income). By a using different vector $\mathbf{b}$, we create different quantile differences. We can then project the innovation $\mathbf{a}_{z,t+1}$ to z_{t+1} on the innovation $\mathbf{a}_{1,t+1}$ to aggregate income growth $y_{1,t+1} = e_1\, \mathbf{y}_{t+1}$:

$$\mathbf{a}_{z,t+1} = \vartheta\, \mathbf{a}_{1,t+1} + \xi_{t+1} \tag{16}$$

The variance of $\mathbf{a}_{z,t+1}$ is $\mathbf{b}\boldsymbol{\Omega}\mathbf{b}^\top$ and the covariance $\text{Cov}(\mathbf{a}_{z,t+1}, \mathbf{a}_{1,t+1}) = \mathbf{b}\boldsymbol{\Omega} e_1^\top$, so the least squares regression coefficient is

$$\vartheta := \frac{\text{Cov}(\mathbf{a}_{z,t+1}, \mathbf{a}_{1,t+1})}{\text{Var}(\mathbf{a}_{1,t+1})} = \frac{\mathbf{b}\boldsymbol{\Omega} e_1^\top}{e_1 \boldsymbol{\Omega} e_1^\top}. \tag{17}$$

Hence, the fraction of variance of the innovation $\mathbf{a}_{z,t+1}$ in the quantile difference attributable to the innovation to aggregate income growth $\mathbf{a}_{1,t+1}$ is

$$\frac{\left(\mathbf{b}\boldsymbol{\Omega}e_1^\top\right)^2}{\left(\mathbf{b}\boldsymbol{\Omega}\mathbf{b}^\top\right)\left(e_1\boldsymbol{\Omega}e_1^\top\right)} = \frac{\vartheta^2\left(e_1\boldsymbol{\Omega}e_1^\top\right)}{\mathbf{b}\boldsymbol{\Omega}\mathbf{b}^\top} \tag{18}$$

The three panels of Figure 7 plot variance ratio (18) for several $p_{\text{high}} - p_{\text{low}}$ quantile differences of private income, post-tax income, and consumption. The panels also report quantile differences parameterized by $\mathbf{b}\,\mathbf{y}_{t+1}$ ranging from 90-10 percentile gaps to 82-18 percentile gaps. For private income, fractions (18) range between 6% and 15%, with the highest values for the mid-range percentile differences (around 85-15). For post-tax income and consumption cross-sections, variance ratios (18) are smaller. Figure 7 shows that aggregate income is a larger source of variations in inequality in private incomes than of inequalities in either post-tax incomes or consumption. Figure 7 thus quantifies pervasive insurance and redistribution apparent in our CEX cross-sections from 1990-2023.[18]

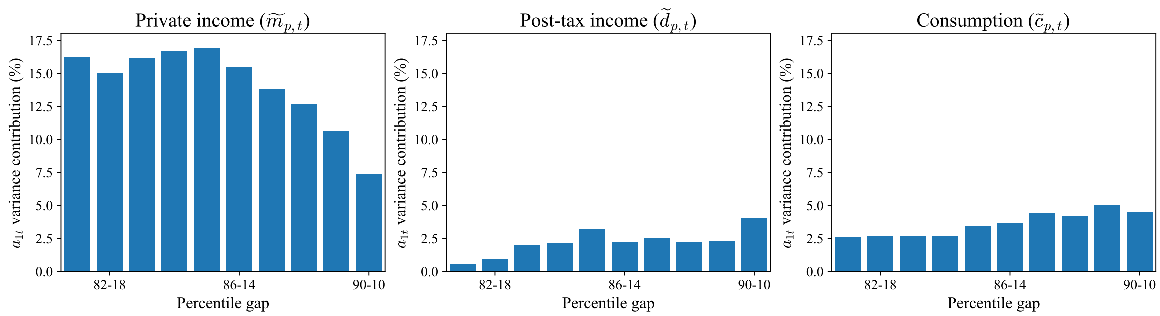

Figure 7: Contributions of aggregate income to inequality

[18]See Gramm et al. (2024, Fig. 2.1) for graphical evidence of substantial US redistribution of income via taxes and transfers.

5 Welfare Comparisons

> The gains from removing all existing variability in aggregate consumption ... are surely well below 1 percent of national income. Policies that deal with the very real problems of society's less fortunate – wealth redistribution and social insurance – can be designed in total ignorance of the nature of business-cycle dynamics. Lucas (1987, p. 105)

This section uses our statistical model of the co-evolution of quantile-specific private income, private income net of taxes and transfers, and consumption to extend welfare calculations of Lucas (1987, 2003). These calculations confirm a claim of Lucas (1987, p. 105) that fiscal redistribution and insurance promote individual consumers' discounted utilities substantially more than further reductions of aggregate fluctuations around trend growth rates. Subsection 5.1 studies cross-quantile welfare benefits from removing exposures to risks in consumption growth rates. We capitalize on our statistical model's representation of serially correlated risks in quantile-specific consumption growth rates.[19] Subsection 5.2 studies cross-quantile welfare consequences of participating in the US tax and transfer system. Subsection 5.3 studies cross-quantile welfare consequences of alterations in the trend rate of growth.

5.1 Welfare Costs of Martingale and Stationary Shocks

Lucas (1987, Sec. III), Lucas (2003), and Tallarini (2000) used the following model to calculate welfare benefits from removing a representative consumer's exposure to i.i.d. risk in US post WWII consumption growth per capita.[20] With c_t as the log of per capita consumption C_t, a time-0 representative agent orders consumption streams according to the discounted expected value

$$v(c_0) = (1-\beta)\,\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t c_t, \tag{19}$$

[19] In Appendix F, we describe how to include a Markov process of transitions across quantiles.

[20] Obstfeld (1994), Dolmas (1998), Hansen et al. (1999), Alvarez and Jermann (2004), and Barillas et al. (2009) studied extensions including some that were based on discrete-time versions of risk-sensitive value functions and other recursive utility functionals.

where $\beta \in (0,1)$ is a discount factor, and c_t evolves according to

$$c_{t+1} - c_t = \mu + \sigma_\epsilon \epsilon_{t+1}, \quad \epsilon_{t+1} \sim \mathcal{N}(0,1). \tag{20}$$

Under this specification, the conditional expectation of $C_{t+1} = \exp(c_{t+1})$ is

$$\mathbb{E}_t\, C_{t+1} = \exp\left(\mu + \frac{1}{2}\sigma_\epsilon^2\right) C_t, \tag{21}$$

and the expected discounted value (19) has the form

$$v(c_0) = c_0 + \frac{\beta}{(1-\beta)}\mu. \tag{22}$$

A value of a risk-free path that starts at $\bar{c}_0$ at $t = 0$ has the same conditional means $\mathbb{E}_0\, C_{t+j}, j \geq 0$ as the risky sequence generated by (20) is

$$v(\bar{c}_0) = \bar{c}_0 + \frac{\beta}{(1-\beta)}\left(\mu + \frac{1}{2}\sigma_\epsilon^2\right). \tag{23}$$

Setting $v(c_0) = v(\bar{c}_0)$, we derive the increment

$$c_0 - \bar{c}_0 = \frac{\beta \sigma_\epsilon^2}{2(1-\beta)}. \tag{24}$$

Hence, equation (24) is a "compensating difference" that equates (22) to (23). The gap $c_0 - \bar{c}_0$ in equation (24) is a percentage of consumption that a representative consumer would be willing to sacrifice if he could instead live in a parallel economy in which σ_ϵ has been set to zero as a result of ideal countercyclical monetary and fiscal policies. Lucas (1987, Sec. III) and Lucas (2003) interpreted gap $c_0 - \bar{c}_0$ as an upper bound on welfare gains that could be gathered by post WWII countercyclical US monetary and fiscal policies.

The logarithmic random walk-with-drift specification (20) on which formula (24) is based exposes a representative consumer to i.i.d risk in consumption growth. Our specification

$$c_{p,t} = \underbrace{t\nu}_{\text{Trend}} + \underbrace{\sum_{j=1}^{t} \mathbf{H}\,\mathbf{a}_j}_{\text{Martingale}} \underbrace{-(\mathbf{g} - e_p^{(c)})\,\mathbf{y}_t}_{\text{Stationary}} + \underbrace{(\mathbf{g}\,\mathbf{y}_0 + Y_0 + \nu_p^{(c)})}_{\text{Constant}} \tag{25}$$

exposes a consumer in the pth quantile to i.i.d. risk in consumption growth through $\mathbf{a}_{t+1}$ and to serially correlated risk in consumption growth through the martingale $\sum_{j=1}^{t} \mathbf{H}\,\mathbf{a}_j$ and the serially correlated state variables $\mathbf{y}_t$.[21]

We shall calculate quantile-specific equalizing differences for hypothetical consumers who are also exposed to quantile-specific serially correlated risks in their consumption growth rates. We compute percentages of current consumption that someone stuck forever in the pth consumption quantile would be willing to sacrifice in order to eliminate both the stationary and the martingale random components on the right hand side of representation (25).

We model a household permanently in the pth consumption quantile who values consumption streams at time t according to[22]

$$v_p(\boldsymbol{\zeta}_t, c_{p,t}) = (1-\beta)\,\mathbb{E}_t\left[\sum_{j=0}^{\infty} \beta^j c_{p,t+j}\right], \qquad \beta \in (0,1), \tag{26}$$

To link this to (6), define the augmented state $\boldsymbol{\zeta}_t = \begin{bmatrix} 1 \\ \mathbf{y}_t \end{bmatrix}$ and the transition $\mathbf{A} = \begin{bmatrix} 1 & \mathbf{0} \\ \mathbf{0} & \mathbf{B} \end{bmatrix}$. Let $\mathbf{a}_{t+1} = \mathbf{J}\boldsymbol{\epsilon}_{t+1}$ where $\boldsymbol{\epsilon}_{t+1} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $\boldsymbol{\Omega} = \mathbf{J}\mathbf{J}^\top$. Choose $\mathbf{C} = \begin{bmatrix} \mathbf{0} \\ \mathbf{J} \end{bmatrix}$ so that the lower block of $\boldsymbol{\zeta}_{t+1} = \mathbf{A}\,\boldsymbol{\zeta}_t + \mathbf{C}\boldsymbol{\epsilon}_{t+1}$ reproduces $\mathbf{y}_{t+1} = \mathbf{B}\,\mathbf{y}_t + \mathbf{J}\boldsymbol{\epsilon}_{t+1}$ in (6).

Define the loadings

$$\mathbf{D}_p^{(c)} := \begin{bmatrix} \nu & e_1\,\mathbf{B} + e_p^{(c)}(\mathbf{B} - \mathbf{I}) \end{bmatrix}, \tag{27}$$

$$\mathbf{F}_p^{(c)} := \left(e_1 + e_p^{(c)}\right)\mathbf{J}. \tag{28}$$

These expressions follow from differencing (25), substituting $\mathbf{y}_{t+1} = \mathbf{B}\,\mathbf{y}_t + \mathbf{a}_{t+1}$ with $\mathbf{a}_{t+1} = \mathbf{J}\boldsymbol{\epsilon}_{t+1}$, and collecting terms in $\boldsymbol{\zeta}_t$ and $\boldsymbol{\epsilon}_{t+1}$. The state-space system for consumption growth

[21] Bansal and Yaron (2004) imposed cross-equation restrictions from consumption Euler equations for US asset prices to infer the presence of such serially correlated risks in US aggregate consumption. Hansen et al. (2008) and Hansen and Sargent (2010) also studied long-run risks but intentionally imposed those cross-equation restrictions at a different point in their analysis.

[22] In our exercises, we set $\beta = 0.99$ for quarterly data.

is

$$\boldsymbol{\zeta}_{t+1} = \mathbf{A}\,\boldsymbol{\zeta}_t + \mathbf{C}\boldsymbol{\epsilon}_{t+1}, \tag{29}$$

$$c_{p,t+1} - c_{p,t} = \mathbf{D}_p^{(c)}\boldsymbol{\zeta}_t + \mathbf{F}_p^{(c)}\,\boldsymbol{\epsilon}_{t+1}. \tag{30}$$

In this representation, $\mathbf{F}_p^{(c)}\,\boldsymbol{\epsilon}_{t+1}$ captures exposure to i.i.d. risk in quantile-p consumption growth, while $\mathbf{D}_p^{(c)}\boldsymbol{\zeta}_t$ captures exposure to persistent risk.

We seek a value function $v_p(\boldsymbol{\zeta}, c_{p,t})$ that satisfies the recursion

$$v_p(\boldsymbol{\zeta}, c_p) = (1-\beta)c_p + \beta\,\mathbb{E}[v_p(\boldsymbol{\zeta}', c_p')]$$

where $\boldsymbol{\zeta}, c_p$ are current period's state and consumption, and $\boldsymbol{\zeta}'$, c_p' are next period's state and consumption satisfying (29)–(30). We use this simplified notation for recursive representations in the remainder of this section. The following proposition provides a closed-form solution and follows the logic of Hansen et al. (2008) and Hansen and Sargent (2010) for computing value functions under long-run risk.

Proposition 2. The value function for a consumer permanently in the pth consumption quantile is linear in the state $\boldsymbol{\zeta}$ and current log consumption c_p:

$$v_p(\boldsymbol{\zeta}, c_p) = c_p + \boldsymbol{\lambda}_p^{(c)\top}\boldsymbol{\zeta}, \tag{31}$$

where

$$\boldsymbol{\lambda}_p^{(c)\top} = \beta\mathbf{D}_p^{(c)}(\mathbf{I} - \beta\,\mathbf{A})^{-1}. \tag{32}$$

Proof. Define $\bar{v}_p(\boldsymbol{\zeta}, c_p) \coloneqq v_p(\boldsymbol{\zeta}, c_p) - c_p$. The recursion implies

$$\begin{aligned}\bar{v}_p(\boldsymbol{\zeta}, c_p) &= \beta\,\mathbb{E}\left[\bar{v}_p(\boldsymbol{\zeta}', c_p') + (c_p' - c_p)\right] \\ &= \beta\,\mathbb{E}\,\bar{v}_p(\boldsymbol{\zeta}', c_p') + \beta\mathbf{D}_p^{(c)}\boldsymbol{\zeta},\end{aligned}$$

where the second equality uses (30). Conjecture $\bar{v}_p(\boldsymbol{\zeta}, c_p) = \boldsymbol{\lambda}_p^{(c)\top}\boldsymbol{\zeta}$ for some row vector $\boldsymbol{\lambda}_p^{(c)\top}$.

Using $\mathbb{E}\,\boldsymbol{\zeta}' = \mathbf{A}\,\boldsymbol{\zeta}$ from (29), we obtain

$$\boldsymbol{\lambda}_p^{(c)\top}\boldsymbol{\zeta} = \beta\boldsymbol{\lambda}_p^{(c)\top}\mathbf{A}\,\boldsymbol{\zeta} + \beta\mathbf{D}_p^{(c)}\boldsymbol{\zeta}.$$

Since this must hold for all $\boldsymbol{\zeta}$, we have $\boldsymbol{\lambda}_p^{(c)\top}(\mathbf{I} - \beta\,\mathbf{A}) = \beta\mathbf{D}_p^{(c)}$. Because $\mathbf{A} = \mathrm{diag}(1, \mathbf{B})$ with $\rho(\mathbf{B}) < 1$ and $\beta \in (0, 1)$, we have $\rho(\beta\,\mathbf{A}) = \beta < 1$, so $(\mathbf{I} - \beta\,\mathbf{A})$ is invertible, yielding (32). Substituting back into $v_p(\boldsymbol{\zeta}, c_p) = c_p + \bar{v}_p(\boldsymbol{\zeta}, c_p)$ gives (31). □

To compare the quantile-p value function (31) to one associated with a risk-free certainty equivalent consumption path (in the spirit of Lucas (1987, 2003); Tallarini (2000)), we iterate (29)–(30) and derive by induction, for $j \geq 1$,

$$c_{p,t+j} = c_{p,t} + \boldsymbol{\psi}_{p,j}^{(c)\top}\boldsymbol{\zeta}_t + \sum_{i=0}^{j-1}\mathbf{J}_{p,i}^{(c)}\boldsymbol{\epsilon}_{t+j-i}, \tag{33}$$

where

$$\boldsymbol{\psi}_{p,j}^{(c)\top} = \sum_{m=0}^{j-1}\mathbf{D}_p^{(c)}\,\mathbf{A}^m, \tag{34}$$

$$\mathbf{J}_{p,i}^{(c)} = \mathbf{F}_p^{(c)} + \sum_{\ell=0}^{i-1}\mathbf{D}_p^{(c)}\,\mathbf{A}^\ell\,\mathbf{C}, \quad i \geq 0. \tag{35}$$

The coefficient $\mathbf{J}_{p,i}^{(c)}$ in (35) represents the cumulative loading of $c_{p,t+j}$ on the shock $\boldsymbol{\epsilon}_{t+j-i}$ that arrives i steps earlier: a direct effect via $\mathbf{F}_p^{(c)}$ plus indirect effects where the shock first moves the state (through $\mathbf{C}$), the state propagates (through $\mathbf{A}^\ell$), and then affects consumption growth (through $\mathbf{D}_p^{(c)}$).

For the certainty equivalent comparison, we define the certainty equivalent value function

$$w_p(\boldsymbol{\zeta}, c_p^{CE}) = c_p^{CE} + \boldsymbol{\lambda}_p^{(c)\top}\boldsymbol{\zeta} + \alpha_p^{(c)}, \tag{36}$$

where c_p^{CE} is a certainty equivalent (log) consumption at time t for a consumer in the pth quantile, and $\alpha_p^{(c)} := c_p - c_p^{CE}$ is a compensating differential that adjusts for the risk in the consumption growth process defined by

$$\alpha_p^{(c)} = \frac{\beta}{2}\sum_{i=0}^{\infty}\beta^i \mathbf{J}_{p,i}^{(c)}\mathbf{J}_{p,i}^{(c)\top}, \tag{37}$$

where $\mathbf{J}_{p,i}^{(c)}$ is defined in (35). Equation (37) follows by equating the discounted value of the risky path to that of a risk-free path with the same conditional expectations for levels $C_{p,t+j} = \exp(c_{p,t+j})$, using the log-normal moment formula applied to (33) in the spirit of Lucas (1987).

The compensation $\alpha_p^{(c)}$ depends on both i.i.d. risk (through $\mathbf{F}_p^{(c)}$ in (28)) and persistent risk (through the state loadings $\mathbf{D}_p^{(c)}$ in (27)). This compensation differs across quantiles but is constant across time, a pattern built into our specification of the additive process (8) for Y_t.

To measure the allocation of benefits from redistribution and insurance, we temporarily pretend that pth quantile households must consume either total private income $m_{p,t}$ or income net of taxes and transfers $d_{p,t}$ instead of the $c_{p,t}$ that they actually consume. Comparing value functions associated with consuming private income provides us with measures of welfare gains to a pth quantile consumer from participating in the US tax and transfer system. Comparing value functions for consuming income net of transfers and for consumption allows us to measure benefits that are presumably achieved through management of personal savings.[23]

The left panel of Figure 8 plots quantile-specific compensations $\alpha_p^{(c)}$ for eliminating all risks, both serially correlated and i.i.d., from growth rates in consumption, income after tax and transfer income, and private income. The right panel plots compensations $\mathbf{F}_p^{(c)}\mathbf{F}_p^{(c)\top}$ associated with the i.i.d. component $\mathbf{a}_{t+1}$. Compensations are highest for lower quantiles of private income, reflecting the substantial income risk borne by lower quantile households. This is largely driven by their loadings $(e_1 + e_p^{(c)})\mathbf{J}$ in formula (28) for $\mathbf{F}_p^{(c)}\mathbf{F}_p^{(c)\top}$ (right chart), approximately half of which is driven by i.i.d. risk for low-quantile households, while i.i.d. risk is a relatively small driver of compensations for higher quantiles. Agents in the 50th percentile of the consumption distribution would be willing to pay about 0.8 percent of consumption

[23]Our three CEX data series make it possible for us to entertain mental experiments like this. Lucas (1987, Sec. III), Lucas (2003), and Tallarini (2000) based their calculations on their calibrated versions of the exogenous consumption endowment, representative-agent economy of Lucas (1978).

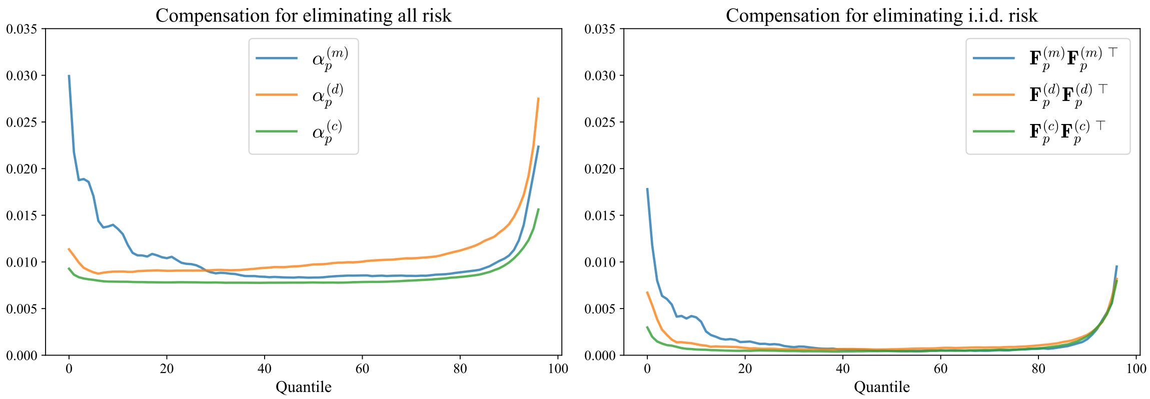

Figure 8: Compensations for eliminating risks in income or consumption growth

to remove all risk in consumption growth. This is approximately 16-times larger than the 0.05% of consumption for Lucas's (2003) representative consumer. Lucas's estimate aligns with our estimate for i.i.d risk, which is about 0.05% of consumption. The remaining 0.95 percent comes from exposure to serially correlated risk in consumption growth.[24]

Second, for all quantiles, proportional consumption compensations are smaller than proportional compensations for private and post-tax incomes, while these proportional compensations are similar for all quantiles. These patterns indicate that decrements in values due to exposures to consumption growth risks are lower than for exposures to income risks. They also indicate that consumption growth risks are ultimately borne relatively evenly across quantiles.

Third, differences in proportional compensations between post-tax income and consumption increase as we move toward higher quantiles. This seems to indicate more consumption smoothing by higher quantile households through management of private savings. Notice that the corresponding difference is negligible for $\mathbf{F}_p^{(c)}\mathbf{F}_p^{(c)\top}$ in the right panel. This suggests that the differences in the left chart mainly reflect how households use private savings to insure themselves against serially correlated risks in consumption growth.

[24] Bansal and Yaron (2004), Hansen et al. (2008), and Hansen and Sargent (2010) explore implications for market prices of risk of exposing representative consumers who dislike it to very persistent risk in consumption growth. Bansal and Yaron call it "long-run risk".

5.2 Welfare consequences of US tax-transfer system

We interpret $v_p(\boldsymbol{\zeta}_t, c_{p,t}) - v_p(\boldsymbol{\zeta}_t, m_{p,t})$ as measuring welfare benefits to a pth quantile household from managing personal savings and from participating in the US tax and transfer system. Figure 9 plots $v_p(\boldsymbol{\zeta}_t, c_{p,t}) - v_p(\boldsymbol{\zeta}_t, m_{p,t})$ for all $p = 10$ and $p = 90$. For the 10th percentile it ranges from 0.66 to 0.71, which is orders of magnitude bigger than the gains from removing exposure to growth rate risks shown in Figure 8. Notice how for the 10th percentile the value of consumption relative to private income increases after 2008, a consequence of consumption having fallen much less than income. The negative numbers for 90th percentile households in the right panel indicate that for them, participation in the US tax-and-transfer system is harmful and consequential, dwarfing the preceding hypothetical benefits from removing exposure to growth rate risk.

For the 10th and 90th percentiles, Figure 10 plots the following decomposition of $v_p(\boldsymbol{\zeta}_t, c_{p,t}) - v_p(\boldsymbol{\zeta}_t, m_{p,t})$:

$$\begin{aligned} v_p(\boldsymbol{\zeta}_t, c_{p,t}) - v_p(\boldsymbol{\zeta}_t, m_{p,t}) &= [v_p(\boldsymbol{\zeta}_t, c_{p,t}) - v_p(\boldsymbol{\zeta}_t, d_{p,t})] \\ &\quad + [v_p(\boldsymbol{\zeta}_t, d_{p,t}) - v_p(\boldsymbol{\zeta}_t, m_{p,t})] \end{aligned} \tag{38}$$

For the 10th percentile, most of the left-hand side (green line in Figure 10) is contributed by the second term in brackets on the right side of equation (38) (orange line). The first term is small (blue line), confirming that 10th percentile households live virtually hand-to-mouth, so

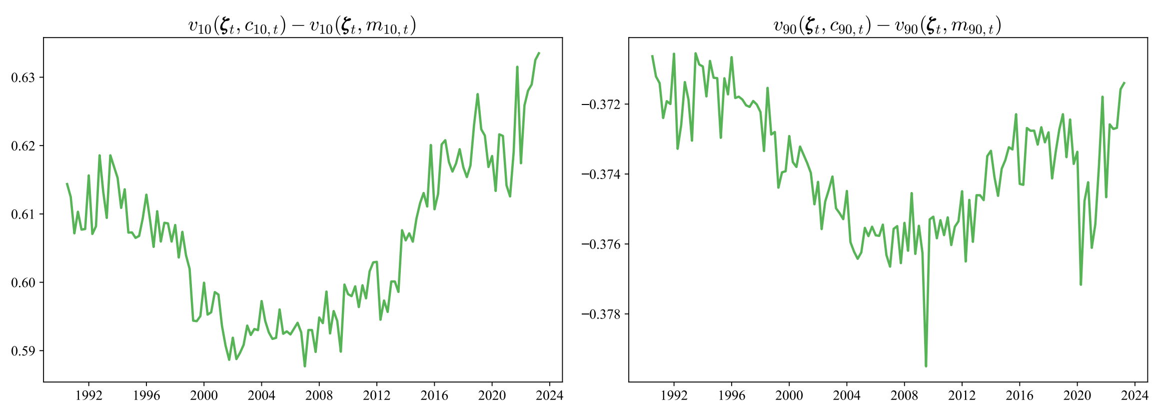

Figure 9: $v_p(\boldsymbol{\zeta}_t, c_{p,t}) - v_p(\boldsymbol{\zeta}_t, m_{p,t})$ for $p = 10$ (left panel) and $p = 90$ (right panel).

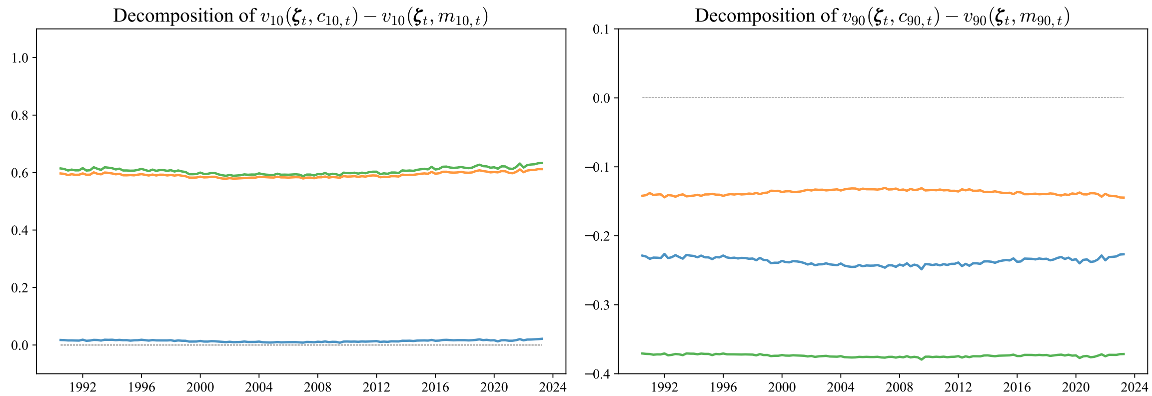

Figure 10: Where $p = 10$ in the left panels and $p = 90$ on the right panels, the green lines plot $v(\boldsymbol{\zeta}_t, c_{p,t}) - v(\boldsymbol{\zeta}_t, m_{p,t})$; the blue lines plot $v(\boldsymbol{\zeta}_t, c_{p,t}) - v(\boldsymbol{\zeta}_t, d_{p,t})$; the orange line plots $v(\boldsymbol{\zeta}_t, d_{p,t}) - v(\boldsymbol{\zeta}_t, m_{p,t})$. See equation (38).

for 10th percentile households the majority of welfare gains on the left of (38) come from their participation in the US tax and transfer system. In contrast, for 90th percentile consumers, the first term in brackets on the right side of equation (38) (blue line) contributes much more, indicating consequences of substantial private savings, while the second term (orange line) plays a smaller role.

5.3 Welfare consequences of growth

We now calculate reductions in consumption that households would be willing to accept in exchange for a once-and-for-all increase in the trend growth rate ν. The discounted expected utility of someone who is permanently in the pth consumption quantile of an economy with aggregate income trend growth $\bar{\nu} > \nu$ is

$$v_p(\boldsymbol{\zeta}, \bar{c}_p) = \bar{c}_p + \left(\bar{\boldsymbol{\lambda}}_p^{(c)}\right)^\top \boldsymbol{\zeta}, \tag{39}$$

where $\left(\bar{\boldsymbol{\lambda}}_p^{(c)}\right)^\top = \beta \bar{\mathbf{D}}_p^{(c)} (\mathbf{I} - \beta\,\mathbf{A})^{-1}$ and $\bar{\mathbf{D}}_p^{(c)} = \begin{bmatrix} \bar{\nu} & e_1\,\mathbf{B} + e_p^{(c)}(\mathbf{B} - \mathbf{I}) \end{bmatrix}$. A proportional decrease in the consumption that leaves a pth quantile consumer indifferent between ν and $\bar{\nu}$ is

$$\begin{aligned} c_p - \bar{c}_p &= (\bar{\boldsymbol{\lambda}}_p^{(c)\top} - \boldsymbol{\lambda}_p^{(c)\top})\boldsymbol{\zeta} && (40)\\ &= \beta(\bar{\mathbf{D}}_p^{(c)} - \mathbf{D}_p^{(c)})(\mathbf{I} - \beta\,\mathbf{A})^{-1}\boldsymbol{\zeta}. && (41) \end{aligned}$$

Using definitions of $\mathbf{D}_p^{(c)}, \bar{\mathbf{D}}_p^{(c)}$ and $\boldsymbol{\zeta}$, we can simplify this to

$$c_p - \bar{c}_p = \beta(\bar{\nu} - \nu)\phi_1(\boldsymbol{\zeta}), \tag{42}$$

where $\phi_1(\boldsymbol{\zeta})$ is the first element of the vector $(\mathbf{I} - \beta\,\mathbf{A})^{-1}\boldsymbol{\zeta}$. The compensating difference is constant across time and quantiles, another feature that our additive process (8) builds in.

Formula (42) indicates that households in all quantiles would be willing to sacrifice roughly 25 percent of current consumption to increase aggregate trend growth ν by 1 percentage point, i.e., from 1.2% to 2.2%. This accords with Lucas's (p. 1, 2003) conclusion that "the potential for welfare gains from better long-run, supply-side policies exceeds by far the potential from further improvements in short-run demand management."

To compare the value of growth with the value of eliminating consumption risk, we transform (42) to compute the increment in trend growth that would leave a pth quantile consumer as well off as if all consumption risk were removed:

$$\bar{\nu} - \nu = \frac{(c_p - \bar{c}_p)}{\beta\phi_1(\boldsymbol{\zeta})} = \frac{\alpha_p^{(c)}}{\beta\phi_1(\boldsymbol{\zeta})} \tag{43}$$

We also compute counterpart comparisons for households constrained to consume total private income and for households consuming income net of transfers and taxes. To decompose the sources of these welfare differences, we write

$$v_p(\boldsymbol{\zeta}, c_p) - v_p(\boldsymbol{\zeta}, m_p) = (c_p - m_p) + ({\boldsymbol{\lambda}_p^{(c)}}^\top - {\boldsymbol{\lambda}_p^{(m)}}^\top)\boldsymbol{\zeta} \tag{44}$$

where the first term measures the increment in value attributable to redistribution, while the second term captures insurance against persistent risk. More explicitly, the welfare difference at time t and state $\boldsymbol{\zeta}_t$ is

$$\Delta v_{p,t} \coloneqq v_p(\boldsymbol{\zeta}_t, c_{p,t}) - v_p(\boldsymbol{\zeta}_t, m_{p,t}) = (c_{p,t} - m_{p,t}) + \big({\boldsymbol{\lambda}_p^{(c)}}^\top - {\boldsymbol{\lambda}_p^{(m)}}^\top\big)\boldsymbol{\zeta}_t.$$

Applying the same linear mapping as in (42)–(43), we convert this welfare difference into a trend growth increment:

$$\tilde{\nu} - \nu = \frac{\Delta v_{p,t}}{\beta \phi_1(\boldsymbol{\zeta}_t)}, \qquad \phi_1(\boldsymbol{\zeta}_t) \coloneqq e_1^\top (\mathbf{I} - \beta\, \mathbf{A})^{-1} \boldsymbol{\zeta}_t.$$

Table 1 reports these compensations for eliminating all risks in consumption growth by setting $\bar{c}_p = c_p^{CE}$, and the value of insurance for the 10th, 50th and 90th percentiles.[25] For all quantiles, the utility benefits of eliminating all risks to consumption growth are equivalent to just a 0.03–0.04 percentage-point increase in annual trend growth (see the first two columns of Table 1). In contrast, for a 10th percentile household, access to the US tax, transfer, and financial system is welfare-equivalent to an increase in annual trend growth of aggregate income by about 2.6 percentage points, whereas the corresponding increments are about -0.8 and -1.5 percentage points for the 50th and 90th percentiles, respectively (rightmost columns of the table).

	Value of eliminating all risk to consumption growth		Value of social and private insurance	
	$(c_p - c_p^{CE})$	$(\bar{\nu} - \nu)$	$v_p(\boldsymbol{\zeta}, c_p) - v_p(\boldsymbol{\zeta}, m_p)$	$\tilde{\nu} - \nu$
Quantile	%	%, annual.	%	%, annual.
$p = 10$	0.79	0.032	63.35	2.56
$p = 50$	0.78	0.031	-21.00	-0.85
$p = 90$	0.93	0.038	-37.14	-1.50

Table 1: Compensating differentials for eliminating risks to consumption growth

6 Concluding Remarks

We have used additive functionals and dynamic mode decompositions to describe the coevolution of CEX cross-sections of private earned income, post-tax income, and consumption from 1990 to 2023. We decomposed a Chisini mean of private earned income and cross-section CEX quantiles of private earned incomes, incomes after taxes and transfers, and consumption levels into constant, deterministic trend, martingale, and stationary components. We use our statistical model to analyze how aggregate income has evolved along with those cross sections.

[25] For "eliminating all risk" we use (43) with $\alpha_p^{(c)}$. For "insurance" we map the value difference $\Delta v_{p,t} = v_p(\boldsymbol{\zeta}_t, c_{p,t}) - v_p(\boldsymbol{\zeta}_t, m_{p,t})$ into a growth increment via $\tilde{\nu} - \nu = \Delta v_{p,t}/(\beta \phi_1(\boldsymbol{\zeta}_t))$.

Our decomposition of aggregate income indicates that aggregate income was briefly above its deterministic trend before the 2008 financial crisis. After 2008, the martingale component fell suddenly and the stationary component fell gradually. By the end of the sample, the martingale component was close to zero and the stationary component was negative, dragging aggregate income below trend. There is substantial heterogeneity in stationary components of quantiles. For 10th percentile households, private income fell substantially after the financial crisis, stayed below deterministic trend and fell further during and after the COVID. Nevertheless, for 10th percentile households, both post-tax income and consumption fell less and by 2023 had returned close to trend. For the 90th percentile, private income returned to trend by 2023, as did post-tax income, but consumption remained below trend.

For private income inequality measures (e.g., 90-10 percentile differences), the fraction of variance attributable to aggregate income innovations ranges from around 7% to 15%. These fractions are smaller for post-tax income and consumption, indicating that redistribution and insurance arrangements have shaped the transmission of aggregate shocks to inequality.

Our extension of Lucas's welfare calculations (Lucas, 1987, 2003) to heterogeneous consumers exposed to serially correlated consumption growth risks yields several findings. Consumers in the 50th percentile of the consumption distribution would sacrifice approximately 0.8% of current consumption to eliminate all consumption growth risk—16 times larger than Lucas's representative agent estimate of 0.05%. But when we eliminate only i.i.d risks, our estimates are similar to Lucas's, indicating that persistent components of risks account for 0.95 of the compensating differences. Even so, benefits of eliminating the consumption growth risk are small when compared to benefits from increasing trend growth. We use our additive functional model to compute that eliminating all consumption growth risk is welfare equivalent to just a 0.03–0.04 percentage point increase in annual trend growth.

Benefits from participating in the US tax and transfer system dwarf the benefits of eliminating the risk of consumption growth. For 10th percentile consumers, the value increment from the tax and transfer system is orders of magnitude larger than the welfare cost of consumption growth risk. For this group, the benefit is equivalent to roughly a 2.5–2.6 percentage point increase in annual trend growth. For 90th percentile consumers, the redistribution through the tax and transfer system imposes welfare costs that exceed the welfare costs of growth rate risk. Our decomposition of these welfare effects shows that for 10th percentile households, most gains derive from tax and transfer redistribution rather than consumption smoothing through private savings. For 90th percentile consumers, private savings contribute substantially more to consumption smoothing.

The statistical framework that we have used to study the joint dynamics of aggregate income and cross sections of incomes and consumption can be extended to include household mobility (Appendix F), model uncertainty (Appendix G), or serve as an empirical target for evaluating structural HANK models.[26]

[26] Appendix A describes data sources. Appendix B plots 50 singular values for our $\mathbf{Y}$ data matrix and associated dynamic modes for the largest three singular values. Appendix D explores consequences of imposing long-run restrictions like those of Blanchard and Quah (1993) on our section 3 additive functional (6)- (7) of aggregate income $\{Y_t\}$. Appendix E provides another decomposition of the section 5 value differences $v_p(\boldsymbol{\zeta}_t, c_{p,t}) - v_p(\boldsymbol{\zeta}_t, m_{p,t})$. Appendix F describes how to extend our analysis to include Markov transitions across consumption and income quantiles. Using methods presented by Barillas et al. (2009), Appendix G describes how to extend the section 5 welfare calculations to include agents' concerns about misspecification of their subjective statistical models.

A Data

Figure 11 presents the evolution of cross-quantile moments. The levels, the standard deviation of private income ($\sigma_t^{(m)}$) is higher than post-tax income ($\sigma_t^{(d)}$) or consumption ($\sigma_t^{(c)}$). Moreover, $\sigma_t^{(d)}$ appears to be larger than $\sigma_t^{(c)}$ in most periods. $\sigma_t^{(m)}$ exhibits substantial variation over time, increasing during the 2008 financial crisis and the COVID-19 period. In contrast, $\sigma_t^{(d)}$ and $\sigma_t^{(c)}$ are more stable, reflecting the smoothing effects of taxes, transfers, and consumption behavior. Figure 12 shows the 90-10 percentile differences for all three variables.

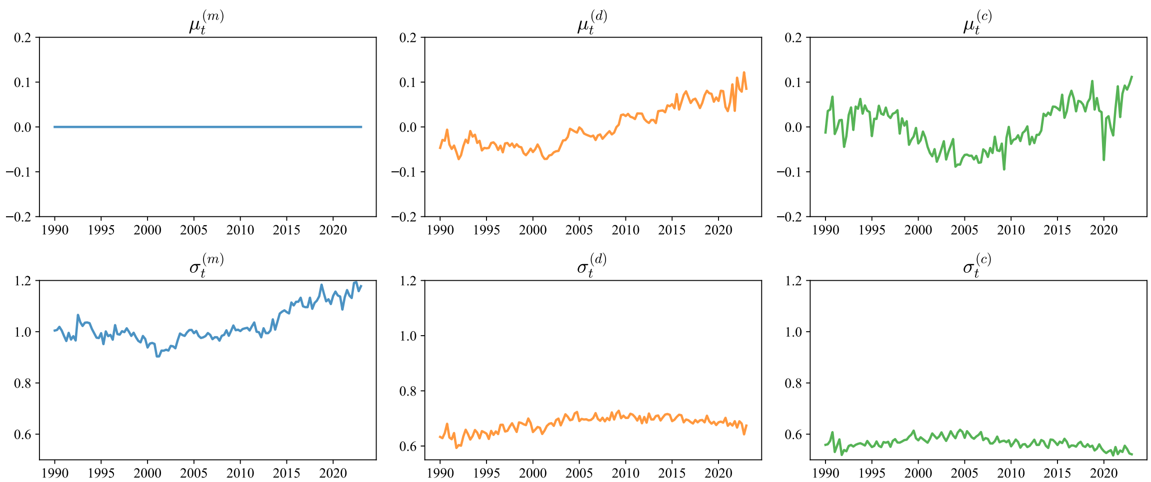

Figure 11: Cross-quantile moments

The difference is largest and most volatile for private income, smaller for post-tax income, and smallest for consumption. These patterns are consistent with consumption smoothing, whereby households adjust their consumption less than their income fluctuations, and with the redistributive effects of taxes and transfers. The figure also shows that this measure of inequality has increased for private income since 2008, while it has gradually declined for both post-tax income and consumption over the same period.

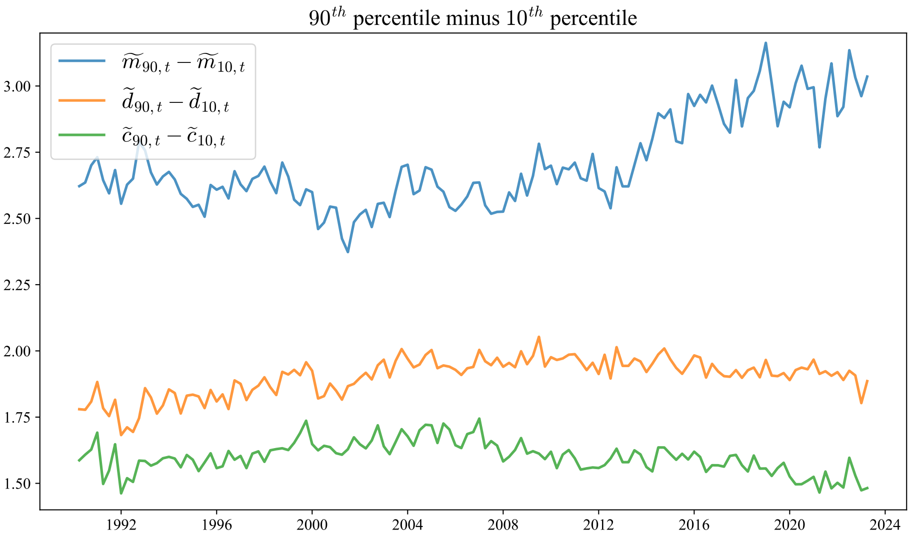

Figure 12: 90–10 percentile gap

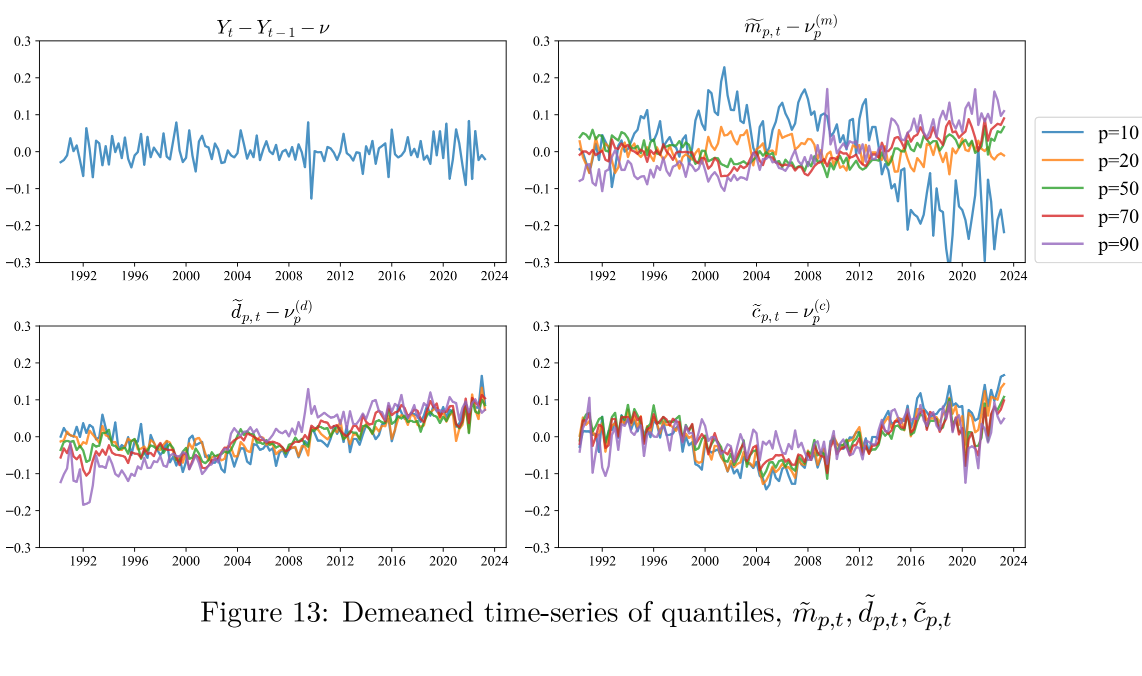

Figure 13: Demeaned time-series of quantiles, $\tilde{m}_{p,t}, \tilde{d}_{p,t}, \tilde{c}_{p,t}$

	Code Mnemonic		
	1990-2004	2004-2013	2013-2022
Private income			
Income from salary or wages	FSALARYX	FSALARYM	FSALARYM
Income from non-farm business	FNONFRMX	FNONFRMM	FSMPFRXM
Income from own farm	FFRMINCX	FFRMINCM	
Income from interest on savings accounts or bonds	INTEARNX	INTEARNM	INTRDVXM
Regular income earned from dividends, royalties, estates	FININCX	FININCXM	ROYESTXM
Income from pensions or annuities	PENSIONX	PENSIONM	RETSURVM
Net income or loss received from roomers or boarders	INCLOSSA	INCLOSAM	
Net income or loss received other rental properties	INCLOSSB	INCLOSBM	NETRENTM
Income from regular contributions from alimony and other	ALIOTHX	ALIOTHXM	
Income from care of foster children, cash scholarships	OTHRINCX	OTHRINCM	OTHRINCM
Transfer income			
Income from Social Security benefits and Railroad Benefit checks	FRRETIRX	FRRETIRM	FRRETIRM
Supplemental Security Income from all sources	FSSIX	FSSIXM	FSSIXM
Income from unemployment compensation	UNEMPLX	UNEMPLXM	
Income from workmen's compensation and veteran's payments	COMPENSX	COMPENSM	OTHREGXM
Income from public assistance including job training	WELFAREX	WELFAREM	WELFAREM
Income from other child support	CHDOTHX	CHDOTHXM	
Food stamps	JFDSTMPA		
Food stamps and electronic benefits	FOODSMPX	FOODSMPM	JFSAMTM

Table 2: Categorizing CEX income into private and transfers

B Singular Values and DMD Modes

Figure 14 plots singular values of the $\mathbf{Y}$ matrix. As noted in Section 4, we retain only the four largest singular values when we estimate $\widehat{\mathbf{B}}$ via DMD. This choice is guided by a standard "elbow rule": the singular values decline rapidly for the first few components and then flatten out, with a visible kink at the fourth singular value, suggesting that additional components mainly capture noise rather than systematic cross-sectional dynamics.

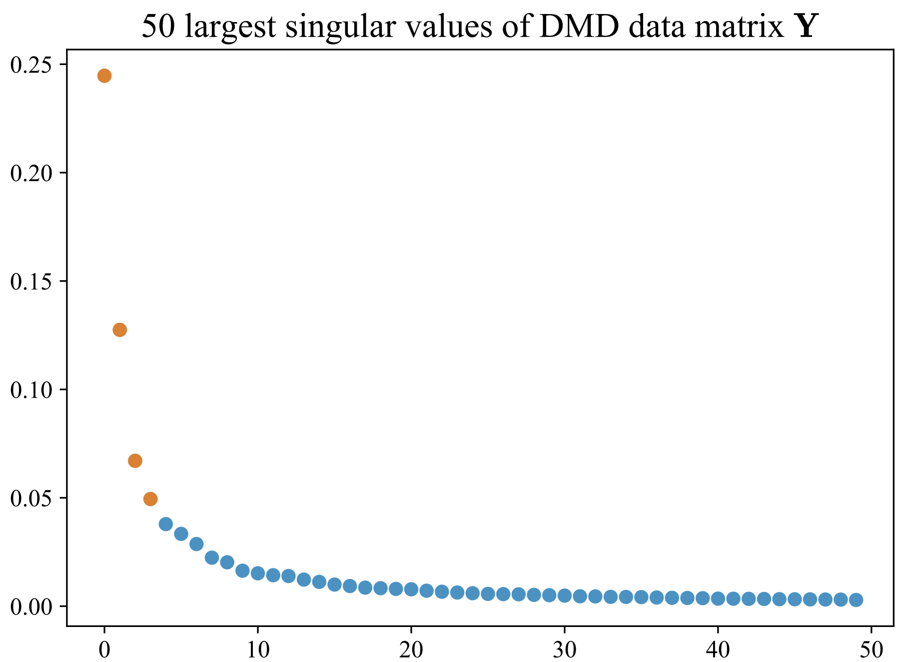

Figure 14: 50 largest singular values of DMD data matrix $\mathbf{Y}$

C Loadings Associated with Modes 3 and 4

Figure 15 plots the loadings associated with DMD modes 3 and 4. One noteworthy feature is that the loadings are substantially smaller than those associated with modes 1 and 2 in Figure 6.

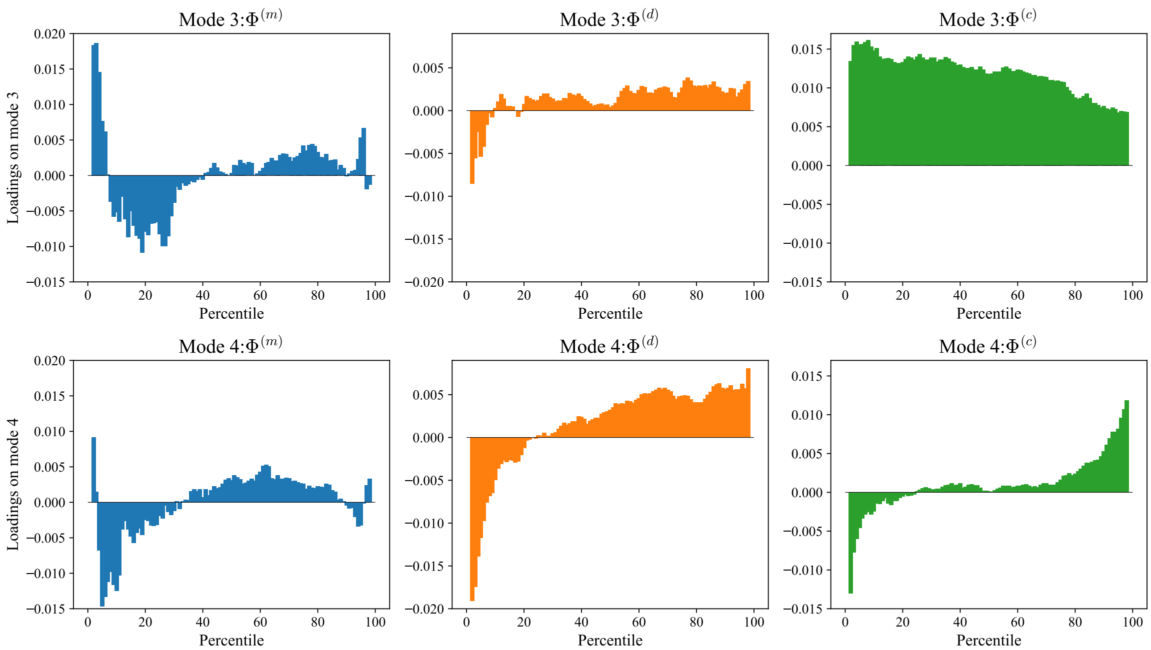

Figure 15: Loadings associated with DMD modes 3 and 4

D A Blanchard-Quah Analysis

To impose the long-run restrictions of Blanchard and Quah (1993), we represent the additive functional (6)-(7) for $\{Y_t\}$ as

$$\mathbf{y}_{t+1} = \mathbf{B}\,\mathbf{y}_t + \mathbf{J}\boldsymbol{\epsilon}_{t+1} \tag{45}$$

$$Y_{t+1} - Y_t - \nu = e_1\,\mathbf{B}\,\mathbf{y}_t + e_1\mathbf{J}\boldsymbol{\epsilon}_{t+1} \tag{46}$$

by setting $\mathbf{a}_{t+1} = \mathbf{J}\boldsymbol{\epsilon}_{t+1}$, where $\mathbf{J}\mathbf{J}^\top = \boldsymbol{\Omega}$ and $\boldsymbol{\epsilon}_{t+1} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. We want to impose that $\boldsymbol{\epsilon}_{1,t+1}$ is the only shock that has a permanent effect on aggregate income, i.e. for $j \geq 1$,

$$\lim_{j\to\infty} \mathbb{E}[Y_{t+j}|\boldsymbol{\epsilon}_{1,t+1}] \neq 0$$

$$\lim_{j\to\infty} \mathbb{E}[Y_{t+j}|\boldsymbol{\epsilon}_{i,t+1}] = 0 \quad \forall i \geq 2$$

To implement this, we compute $\mathbf{J}$ according to

$$\mathbf{J} = (\mathbf{I} - \mathbf{B})\widetilde{\mathbf{J}} \tag{47}$$

$$\mathbf{S}_y(0) = \widetilde{\mathbf{J}}\widetilde{\mathbf{J}}^\top \tag{48}$$

where $\mathbf{S}_y(0) = (\mathbf{I} - \mathbf{B})^{-1}\boldsymbol{\Omega}((\mathbf{I} - \mathbf{B})^{-1})^\top$ is the spectral density matrix of $\mathbf{y}_t$ at frequency zero.[27]

Figure 16 plots the impulse response functions of 10th and 90th percentiles of income and consumption computed using our Blanchard and Quah identified system (45)-(46). The impulse response functions show that a positive permanent shock to aggregate income initially increases inequality in private income, with the 90–10 percentile gap widening for about 50 periods before gradually returning to its steady state. For post-tax income, both the 90th and 10th percentiles rise, with the 90th percentile rising more. Inequality in post-tax income thus rises, but much less than private income. This highlights the smoothing consequences of fiscal policy. For consumption, the 90th percentile responds very little, while consumption of the 10th percentile rises. These muted responses combine to produce a small fall in consumption inequality in response to a permanent aggregate income shock.

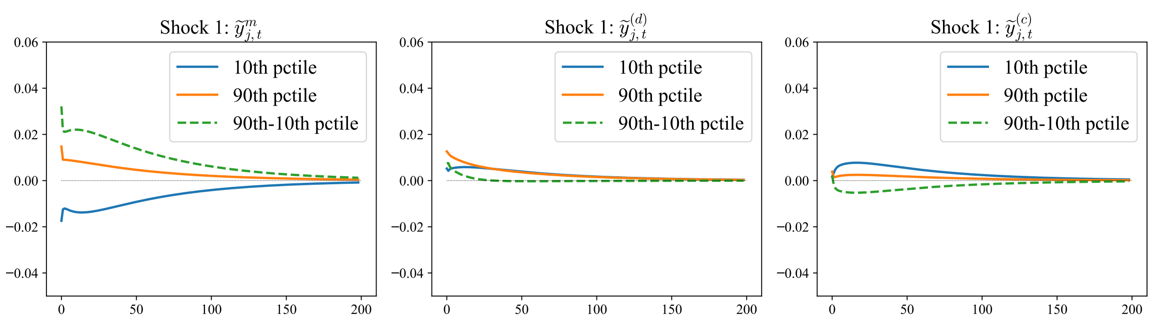

Figure 16: Impulse response

[27]Due to the reduced-rank $\mathbf{B}$ and thus the approximate reduced-rank nature of $\boldsymbol{\Omega}$, we encountered numerical issues when computing the Cholesky factorization of $\mathbf{S}_y(0)$ to obtain $\widetilde{\mathbf{J}}$. Instead we compute the $\mathbf{LDL}$ decomposition of $\mathbf{S}_y(0)$, and inspect the diagonal entries in $\mathbf{D}$. We found that all negative entries were within numerical precision of zero, and replaced them with zero. Calling the new matrix $\widetilde{\mathbf{D}}$, we compute $\widetilde{\mathbf{J}} = \mathbf{L}\widetilde{\mathbf{D}}$.

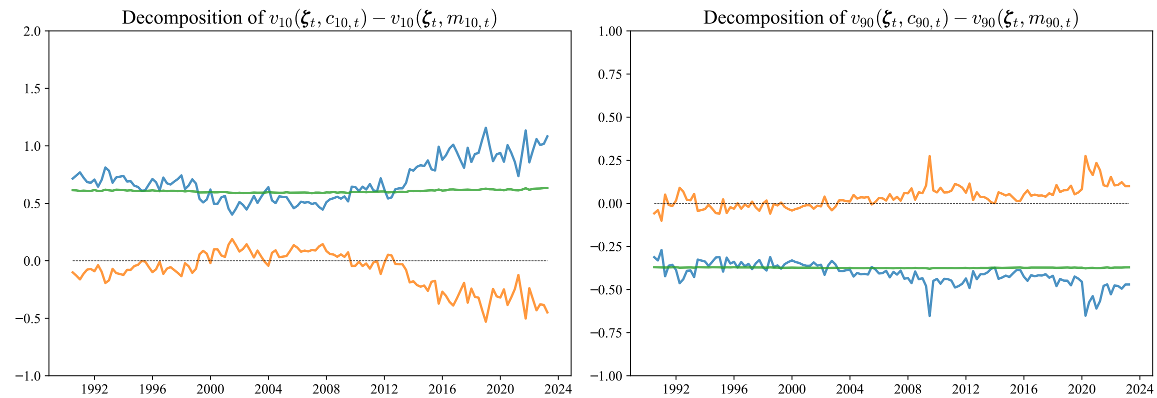

Figure 17: Where $p = 10$ (left) and $p = 90$ (right), the green lines plot $v_p(\boldsymbol{\zeta}_t, c_{p,t}) - v_p(\boldsymbol{\zeta}_t, m_{p,t})$; the orange lines plot $(\boldsymbol{\lambda}_p^{(c)\top} - \boldsymbol{\lambda}_p^{(m)\top})\boldsymbol{\zeta}_t$; the blue lines plot $c_{p,t} - m_{p,t}$. See equation (49).

E Another Decomposition

In Section 5.3, we construct an alternative decomposition of $v_p(\boldsymbol{\zeta}, c_p) - v_p(\boldsymbol{\zeta}, m_p)$:

$$v_p(\boldsymbol{\zeta}, c_p) - v_p(\boldsymbol{\zeta}, m_p) = (c_p - m_p) + (\boldsymbol{\lambda}_p^{(c)\top} - \boldsymbol{\lambda}_p^{(m)\top})\boldsymbol{\zeta} \tag{49}$$

We interpret the second term on the right as a measure of the increment in value attributable to social and private insurance against serially correlated risk in growth rates, while the first term measures consequences of redistribution.

Figure 17 plots components of this decomposition at each time t. In the case of the 10th percentile, the second term has become larger since 2012, while the first term has become smaller. For the 90th percentile, the second term has fallen since 2020, while the first term has risen. We can further decompose each term on the right hand side of value decomposition (49),

$$(\boldsymbol{\lambda}_p^{(c)\top} - \boldsymbol{\lambda}_p^{(m)\top})\boldsymbol{\zeta} = (\boldsymbol{\lambda}_p^{(c)\top} - \boldsymbol{\lambda}_p^{(d)\top})\boldsymbol{\zeta} + (\boldsymbol{\lambda}_p^{(d)\top} - \boldsymbol{\lambda}_p^{(m)\top})\boldsymbol{\zeta} \tag{50}$$

$$c_p - m_p = (c_p - d_p) + (d_p - m_p) \tag{51}$$

Equation (50) decomposes the value attributable to insurance against serially correlated risk in consumption growth into private insurance (first term) versus social insurance (second term). Similarly, equation (51) decomposes the consequences of redistribution into private savings (first term) versus social (second term).

The panels in the top row of Figure 18 plot the components of (50) for the 10th (left) and the 90th (right) percentiles. For the 10th percentile, the value of social insurance (orange) is the predominant driver of the value of total insurance (green), while the value of private insurance (blue) is negligible. For the 90th percentile, the value of private insurance (blue) is the predominant driver of overall insurance (green), while social insurance (orange) is close to zero across the sample. These findings are again consistent with the hand-to-mouth nature of low income households, and that the welfare system plays a small role for upper quantile households.

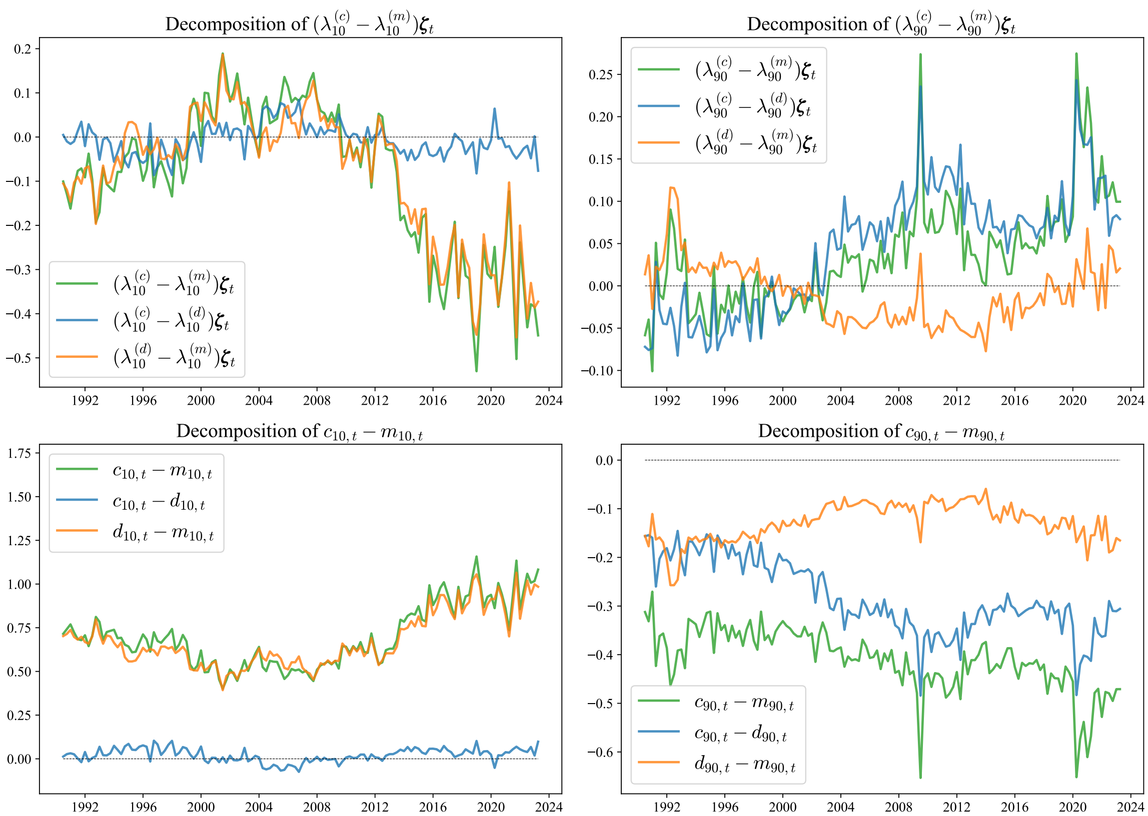

Figure 18: Where $p = 10$ (left panels) and $p = 90$ (right panels) the top panels plot the decomposition of $(\boldsymbol{\lambda}_p^{(c)\top} - \boldsymbol{\lambda}_p^{(m)\top})\boldsymbol{\zeta}_t$; see equation (50). The bottom panels plot the decomposition of $c_{p,t} - m_{p,t}$; see equation (51).

The panels in the bottom row of Figure 18 plot the components of (51) for the 10th and 90th percentile. The left-hand panel shows that the majority of $c_{10,t} - m_{10,t}$ is driven by $d_{10,t} - m_{10,t}$, suggesting again that redistribution plays a large role for these households. The right-hand panel shows that, for the 90th percentile, $c_{90,t} - d_{90,t}$ is the dominant driver, suggesting that private savings is more important than social redistribution.

F Mobility across quantiles

We now allow mobility across consumption quantiles. Let $\mathbf{\Pi} \in \mathbb{R}^{M \times M}$ be a row stochastic transition matrix with element $\Pi_{pp'}$ representing the probability of moving from quantile p to quantile p'. At time t, a consumer currently in quantile p receives flow payoff $c_{p,t}$ and observes the consumption vector $\mathbf{c}_t \in \mathbb{R}^M$. The value function satisfies

$$v_p(\boldsymbol{\zeta}, \mathbf{c}) = (1 - \beta)c_p + \beta \sum_{p' \in \mathcal{P}} \Pi_{pp'} \, \mathbb{E}\left[v_{p'}(\boldsymbol{\zeta}', \mathbf{c}')\right]. \tag{52}$$

Here the state and consumption dynamics obey (29)–(30) for each quantile p.

Proposition 3 (Linear value with mobility). Fix $\beta \in (0, 1)$. Assume the dynamics (29)–(30) and a row stochastic matrix $\mathbf{\Pi}$. There exists a unique solution to (52) of the form

$$v_p(\boldsymbol{\zeta}, \mathbf{c}) = \boldsymbol{\lambda}_p^{(c)\top} \boldsymbol{\zeta} + \sum_{q \in \mathcal{P}} b_{pq} c_q. \tag{53}$$

Let $n = \dim(\boldsymbol{\zeta})$. Define the $M \times M$ matrix $\mathbf{B}^{(\Pi)} \in \mathbb{R}^{M \times M}$ with entries $[\mathbf{B}^{(\Pi)}]_{pq} = b_{pq}$, where the superscript reminds us that this matrix depends on the mobility matrix $\mathbf{\Pi}$. Define $\mathbf{D}^{(c)} \in \mathbb{R}^{M \times n}$ with rows $[\mathbf{D}^{(c)}]_{p\cdot} = \mathbf{D}_p^{(c)}$, and $\mathbf{\Lambda}^{(c)} \in \mathbb{R}^{M \times n}$ with rows $[\mathbf{\Lambda}^{(c)}]_{p\cdot} = \boldsymbol{\lambda}_p^{(c)\top}$. Then

$$\mathbf{B}^{(\Pi)} = (1 - \beta)(\mathbf{I} - \beta \mathbf{\Pi})^{-1}, \tag{54}$$

$$\mathbf{\Lambda}^{(c)} = \beta \mathbf{\Pi} \mathbf{\Lambda}^{(c)} \, \mathbf{A} + \beta \mathbf{\Pi} \mathbf{B}^{(\Pi)} \mathbf{D}^{(c)}. \tag{55}$$

Proof. Conjecture that the value function has the form (53) and substitute into the value equation (52). From the dynamics (29)–(30), we have $\mathbb{E}_t \, \boldsymbol{\zeta}_{t+1} = \mathbf{A} \, \boldsymbol{\zeta}_t$ and $\mathbb{E}_t \, c_{q,t+1} = c_{q,t} +$

$\mathbf{D}_q^{(c)}\boldsymbol{\zeta}_t$ since $\mathbb{E}_t\,\boldsymbol{\epsilon}_{t+1} = \mathbf{0}$. Thus

$$\begin{aligned}\mathbb{E}_t\, v_{p'}(\boldsymbol{\zeta}_{t+1}, \mathbf{c}_{t+1}) &= \boldsymbol{\lambda}_{p'}^{(c)\top}\,\mathbf{A}\,\boldsymbol{\zeta}_t + \sum_q b_{p'q}\big(c_{q,t} + \mathbf{D}_q^{(c)}\boldsymbol{\zeta}_t\big)\\ &= \Big[\boldsymbol{\lambda}_{p'}^{(c)\top}\,\mathbf{A} + \sum_q b_{p'q}\mathbf{D}_q^{(c)}\Big]\boldsymbol{\zeta}_t + \sum_q b_{p'q}c_{q,t}.\end{aligned}$$

Substituting into (52) gives

$$v_p(\boldsymbol{\zeta}_t, \mathbf{c}_t) = (1-\beta)c_{p,t} + \beta\sum_{p'}\Pi_{pp'}\Big[\Big(\boldsymbol{\lambda}_{p'}^{(c)\top}\,\mathbf{A} + \sum_q b_{p'q}\mathbf{D}_q^{(c)}\Big)\boldsymbol{\zeta}_t + \sum_q b_{p'q}c_{q,t}\Big].$$

For this to equal $\boldsymbol{\lambda}_p^{(c)\top}\boldsymbol{\zeta}_t + \sum_q b_{pq}c_{q,t}$ for all $\boldsymbol{\zeta}_t$ and $\mathbf{c}_t$, we require

$$b_{pq} = (1-\beta)\delta_{pq} + \beta\sum_{p'}\Pi_{pp'}b_{p'q}, \tag{56}$$

$$\boldsymbol{\lambda}_p^{(c)\top} = \beta\sum_{p'}\Pi_{pp'}\Big[\boldsymbol{\lambda}_{p'}^{(c)\top}\,\mathbf{A} + \sum_q b_{p'q}\mathbf{D}_q^{(c)}\Big], \tag{57}$$

where δ_{pq} denotes the Kronecker delta. Rewriting (56) in matrix form yields $\mathbf{B}^{(\Pi)} = (1-\beta)\mathbf{I} + \beta\boldsymbol{\Pi}\mathbf{B}^{(\Pi)}$, so $\mathbf{B}^{(\Pi)} = (1-\beta)(\mathbf{I}-\beta\boldsymbol{\Pi})^{-1}$ as stated in (54). The inverse exists because $\rho(\beta\boldsymbol{\Pi}) \le \beta < 1$.

Stacking (57) across all p yields

$$\boldsymbol{\Lambda}^{(c)} = \beta\boldsymbol{\Pi}\boldsymbol{\Lambda}^{(c)}\,\mathbf{A} + \beta\boldsymbol{\Pi}\mathbf{B}^{(\Pi)}\mathbf{D}^{(c)},$$

which is equation (55). Rearranging: $\boldsymbol{\Lambda}^{(c)}(\mathbf{I} - \beta\,\mathbf{A}) = \beta\boldsymbol{\Pi}\mathbf{B}^{(\Pi)}\mathbf{D}^{(c)}$. Since $\rho(\beta\,\mathbf{A}) \le \beta < 1$, the matrix $\mathbf{I} - \beta\,\mathbf{A}$ is invertible, giving

$$\boldsymbol{\Lambda}^{(c)} = \beta\boldsymbol{\Pi}\mathbf{B}^{(\Pi)}\mathbf{D}^{(c)}(\mathbf{I} - \beta\,\mathbf{A})^{-1}.$$

This can also be written as a linear system for the rows $\boldsymbol{\lambda}_p^{(c)\top}$ via (57). Uniqueness follows from the invertibility of $\mathbf{I} - \beta\,\mathbf{A}$. □

Remark 1 (Connection to no-mobility case). Setting $\boldsymbol{\Pi} = \mathbf{I}$ (no mobility across quantiles) recovers Proposition 2. In this case, $\mathbf{B}^{(\Pi)} = \mathbf{I}$ and the value function simplifies to

$$v_p(\boldsymbol{\zeta}, c_p) = c_p + \boldsymbol{\lambda}_p^{(c)\top} \boldsymbol{\zeta}$$

where $\boldsymbol{\lambda}_p^{(c)\top} = \beta \mathbf{D}_p^{(c)} (\mathbf{I} - \beta\,\mathbf{A})^{-1}$, matching equation (32).

G Fear of Misspecification

We extend the analysis to agents who fear model misspecification, following Hansen et al. (2008), Hansen and Sargent (2010), and the discrete version of Hansen and Sargent (2010). Consider a consumer permanently in quantile p who orders consumption streams according to a robust value function $v_p(\boldsymbol{\zeta}, c_p)$ satisfying

$$v_p(\boldsymbol{\zeta}, c_p) = (1-\beta) c_p + \mathcal{T}_\theta[\beta v_p(\boldsymbol{\zeta}', c_p')], \tag{58}$$

where the robust operator $\mathcal{T}_\theta$ is defined by

$$\mathcal{T}_\theta[\beta v] = \min_{\substack{m(\boldsymbol{\epsilon}) \geq 0 \\ \mathbb{E}[m(\boldsymbol{\epsilon})]=1}} \mathbb{E}\left[m(\boldsymbol{\epsilon}) \big(\beta v + \theta \log m(\boldsymbol{\epsilon}) \big) \right] = -\theta \log \mathbb{E}[\exp(-\beta v/\theta)]. \tag{59}$$

Here $\theta > 0$ is a robustness parameter penalizing deviations from the reference model, and the expectation is taken over $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The minimization in (59) implements the multiplier preferences of Hansen et al. (1999).

Proposition 4 (Robust value function). Under the dynamics (29)–(30) and robustness parameter $\theta > 0$, the value function (58) has the form

$$v_p(\boldsymbol{\zeta}, c_p) = c_p + \boldsymbol{\lambda}_p^{(c)\top} \boldsymbol{\zeta} + \kappa_p, \tag{60}$$

where $\boldsymbol{\lambda}_p^{(c)}$ is given by (32) and

$$\kappa_p = -\frac{\beta^2}{2(1-\beta)\theta} \left\| \mathbf{C}^\top \boldsymbol{\lambda}_p^{(c)} + \mathbf{F}_p^{(c)} \right\|^2. \tag{61}$$

Proof. See Proposition 2.1 of the discrete version of Hansen and Sargent (2010)[28]. □

Remark 2. As in Appendix F, we can incorporate mobility across quantiles by introducing a transition matrix $\mathbf{\Pi}$ and extending the robust value function to

$$v_p(\boldsymbol{\zeta}, \mathbf{c}) = (1-\beta)c_p + \beta \sum_{p'} \Pi_{pp'} \mathcal{T}_\theta[v_{p'}(\boldsymbol{\zeta}', \mathbf{c}')]. \tag{62}$$

[28] Note that the discrete version of Hansen and Sargent (2010) can be found at http://www.tomsargent.com/research/longrunrisk_tom_14.pdf.

References

Alvarez, Fernando and Urban J Jermann, "Using asset prices to measure the cost of business cycles," Journal of Political economy, 2004, 112 (6), 1223–1256.

Bansal, Ravi and Amir Yaron, "Risks for the long run: A potential resolution of asset pricing puzzles," The journal of Finance, 2004, 59 (4), 1481–1509.

Barillas, Francisco, Lars Peter Hansen, and Thomas J Sargent, "Doubts or variability?," Journal of Economic Theory, 2009, 144 (6), 2388–2418.

Blanchard, Olivier Jean and Danny Quah, "The dynamic effects of aggregate demand and supply disturbances: Reply," The American Economic Review, 1993, 83 (3), 653–658.

Brunton, Steven L. and J. Nathan Kutz, Data-Driven Science and Engineering: Machine Learnings, Dynamical Systems, and Control, second edition, Cambridge University Press, 2022.

Burns, Arthur F and Wesley C Mitchell, Measuring business cycles, National bureau of economic research, 1946.

Carroll, C. D., T. F. Crossley, and J. Sabelhaus, Improving the Measurement of Consumer Expenditures, University of Chicago Press, 2015.

Chisini, O., "Sul concetto di media," Periodico di Matematiche, 1929.

Dolmas, Jim, "Risk preferences and the welfare cost of business cycles," Review of Economic Dynamics, 1998, 1 (3), 646–676.

Gramm, Phil, Robert Ekelund, and John Early, The Myth of American Inequality: How Government Biases Policy Debate (With a New Preface), Rowman & Littlefield, 2024.

Guvenen, Fatih, Sam Schulhofer-Wohl, Jae Song, and Motohiro Yogo, "Worker Betas: Five Facts about Systematic Earnings Risk," American Economic Review: Papers & Proceedings, 2017, 107 (5), 398–403.

Hamilton, James D, “Why you should never use the Hodrick-Prescott filter,” Review of Economics and Statistics, 2018, 100 (5), 831–843.

Hansen, Lars Peter, “Dynamic valuation decomposition within stochastic economies,” Econometrica, 2012, 80 (3), 911–967.

_ and Thomas J Sargent, “Fragile beliefs and the price of uncertainty,” Quantitative Economics, 2010, 1 (1), 129–162.

_ , John C Heaton, and Nan Li, “Consumption Strikes Back? Measuring Long-Run Risk,” Journal of Political economy, 2008, 116 (2), 260–302.

_ , Thomas J Sargent, and Thomas D Jr. Tallarini, “Robust permanent income and pricing,” The Review of Economic Studies, 1999, 66 (4), 873–907.

Hood, William C and T C Koopmans, Studies in econometric method, Wiley, 1953.

Koopmans, Tjalling C, “Measurement without theory,” The Review of Economics and Statistics, 1947, 29 (3), 161–172.

_ , Statistical inference in dynamic economic models, New York: Wiley, 1950.

Ljungqvist, L and T J Sargent, Recursive Macroeconomic Theory, 4 ed., MIT Press, 2018.

Lucas, Robert E Jr., “Asset prices in an exchange economy,” Econometrica, 1978, pp. 1429–1445.

_ , Models of business cycles, Vol. 26, Oxford Blackwell, 1987.

_ , “Macroeconomic priorities,” American Economic Review, 2003, 93 (1), 1–14.

Marschak, Jacob, “Economic Measurment for Policy and Prediction,” in William C Hood and T C Koopmans, eds., Studies in econometric method, Wiley, 1953, pp. 1–26.

Obstfeld, Maurice, “Evaluating risky consumption paths: The role of intertemporal substitutability,” European economic review, 1994, 38 (7), 1471–1486.

Sargent, Thomas J., “Robert E. Lucas Jr.’s collected papers on monetary theory,” Journal of Economic Literature, 2015, 53 (1), 43–64.

—, "Haok and Hank Models," in Sofía Bauducco, Andrés Fernández, and Giovanni L. Violante, eds., Heterogeneity in Macroeconomics: Implications for Monetary Policy, Series on Central Banking, Analysis, and Economic Policies, Santiago, Chile: Central Bank of Chile, 2024, pp. 13–38.

—, "Macroeconomics after Lucas," Journal of Political Economy, In Press.

— and Christopher A. Sims, "Business cycle modeling without pretending to have too much a priori economic theory," in "New methods in business cycle research," Minneapolis, Minnesota: Federal Reserve Bank of Minneapolis, 1977, pp. 145–168.

—, Yatheesan J. Selvakumar, and Ziyue Yang, "Dynamic Mode Decompositions and Vector Autoregressions," Technical Report, New York University 2025.

Tallarini, Thomas D Jr., "Risk-sensitive real business cycles," Journal of monetary Economics, 2000, 45 (3), 507–532.

Tu, Jonathan H, Clarence W Rowley, Dirk M Luchtenburg, Steven L Brunton, and J Nathan Kutz, "On dynamic mode decomposition: Theory and applications," Journal of Computational Dynamics, 2014, 1 (2), 391–421.