Aggregate Shocks and Cross-Section Dynamics: Quantifying Redistribution and Insurance in U.S. Household Data

Thomas J. Sargent
New York University

Yatheesan J. Selvakumar*
New York University

June 29, 2025

Abstract

This paper uses additive functionals and dynamic mode decompositions to represent and analyze the co-evolution of cross-sections of private earned income, post-tax-and-transfer income, and consumption in the Consumer Expenditure Survey (CEX) from 1990 to 2021. We quantify how cross-sectional inequality and redistribution dynamically interact with aggregate income. Cross-section dynamics contribute only modestly to the innovation variance of aggregate income growth, while an innovation to aggregate income affects cross-section dynamics more. We construct value functions for heterogeneous synthetic consumers who are exposed to both serially correlated and i.i.d risks in their income and consumption growth rates. For the median household, we find that the cost associated with serially correlated risk is orders of magnitude larger than costs associated with i.i.d risk. We then compare, for each quantile, benefits of eliminating risks in their consumption growth with benefits from their participation 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 estimates a parametric statistical model of cointegrated CEX cross-section quantiles and uses it to organize evidence about how aggregate income dynamically interacts with redistribution and social insurance across households.[1] We detect patterns like those that motivated 20th century representative agent macroeconomic models and the associated "neo-classical synthesis" policy perspective that recommends separating macroeconomic monetary-fiscal policies for moderating business cycles from microeconomic policies that redistribute income across individuals and insure individuals against idiosyncratic shocks.[2]

We construct an example of what Koopmans (1947) called a purely descriptive "Kepler stage" model of business cycles to distinguish it from the "Newton stage" structural models that Koopmans and his colleagues at the Cowles Commission constructed to study consequences of historically unprecedented government policies. Burns and Mitchell's (1946) book as providing a set of averaging procedures for detecting a single business cycle "factor" that drove co-movements of a collection of variables influenced by "the" business cycle.[3] 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 our CEX cross-sections.[4] We use only the CEX data and synthesize aggregate (earned) income as the cross-section Chisini mean of private earned income from the CEX quantiles.[5] We posit that aggregate income is an additive functional, 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 CEX quantiles are also additive functionals, which we decompose into their trend, martingale, and stationary components.[6] We use these representations to quantify the dynamic interaction

[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. We urge the reader to keep this in mind when looking at our welfare comparisons in Section 6.

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

[3]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.

[4]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.

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

[6]Satisfaction of technical conditions described by Sargent and Selvakumar (2024) justify our use of a first-order VAR rather

between cross-section and aggregate income, and study the heterogeneous welfare benefits of social and private insurance as well as insurance against both serially correlated and i.i.d risk to consumption growth.

Section 2 describes our data and how we construct quantiles. Section 3 presents our approach to constructing and decomposing the additive functionals. Section 4 describes dynamic mode decompositions and their connection to vector autoregressions. Section 5 presents our findings on the influence of aggregate income on cross-sections and vice versa. Section 6 constructs welfare comparisons in the spirit of Lucas (1987, Sec. III) and Lucas (2003), who studied welfare of a representative consumer exposed to i.i.d. risk in consumption growth. We extend their framework to study diverse consumers who are also exposed to serially correlated risk in consumption growth. Finally, we compare hypothetical welfare gains of eliminating risks in individual consumption growth with gains from participating in the existing US tax and transfer system, or from increasing trend rates of consumption growth. Section 7 concludes. Seven appendices provide technical details and describe feasible and worthwhile extensions.

2 Data Description and Compression

The Consumer Expenditure Survey (CEX) is a nationally representative survey of U.S. households conducted by the Bureau of Labor Statistics. We study quarterly waves of the Consumer Expenditure Survey (CEX) from 1990 to 2021. Our statistical analysis focuses on three variables:[7] (1) Private income – labor income plus financial income, (2) Post-tax income – private income plus transfers minus taxes, and (3) Consumption. For each time period t, we rank all households in the sample by their real consumption level $c_{1,t}, c_{2,t}, \ldots, c_{\tilde{I},t}$, where $\tilde{I}$ is the number of households, deflated by the PCE chain-type price index. Then we split them into one hundred equally sized bins that we call percentiles. If $\tilde{I}$ is divisible by 100, then each percentile contains $I = \frac{\tilde{I}}{100}$ households. If not, the bottom 99 bins each contain I households and the 100th bin contains the remainder. Except for this detail, define $y_{p,t}^{cons}$ as mean consumption

than the higher-order VAR that appears in the more general additive functional of Hansen (2012).

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

of households in percentile p:

$$y_{p,t}^{cons} = \log \frac{1}{I} \sum_{j=1}^{I} c_{j+I(p-1),t} \quad \text{for } p = 1, \ldots, 100 \tag{1}$$

and call it the log real consumption of the p-th percentile of the consumption distribution. We transform the private income and post-tax income observations in a similar way. To diminish influences of outliers, we remove the 1st, 2nd, and 99th percentile bins, leaving $M = 97$ quantiles for each variable. The vector $\mathbf{y}_t^{cons} = [y_{3,t}^{cons}, \ldots, y_{97,t}^{cons}]$ collects all the remaining quantiles in a $(1 \times M)$ vector for consumption, which we repeat for private and post-tax income. For each time period t, we define a Chisini mean (Chisini, 1929) of the cross-section of private income[8]:

$$Y_t = \frac{1}{M} \sum_{p=1}^{M} y_{p,t}^{priv} \tag{2}$$

Figure 1 compares nominal versions of our Chisini means from the CEX quantiles with corresponding aggregates from the National Income and Product Accounts (NIPA).[9] Levels of the series differ in panel (1a), especially for consumption, but their cyclical patterns in panel (1b) are broadly similar. The differences in consumption levels can be attributed to coverage and definitional 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.

[8] In the case where $\tilde{I}$ is divisible by 100, the Chisini mean is equal to the mean of acros $\tilde{I}$ households.

[9] 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.

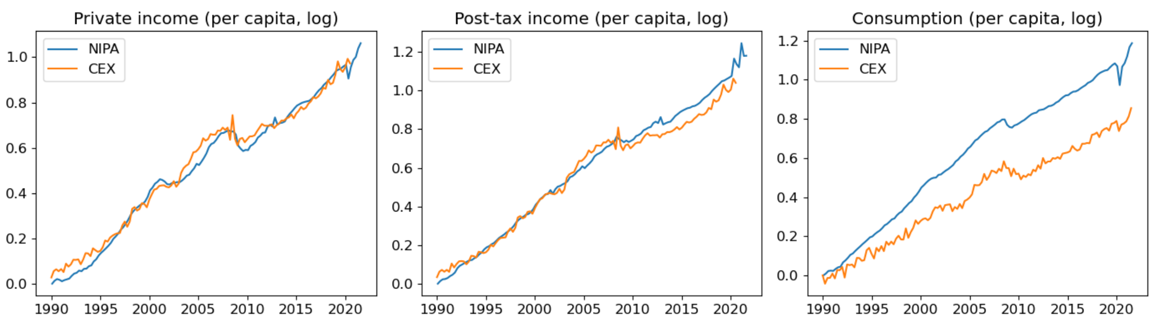

(a) Levels

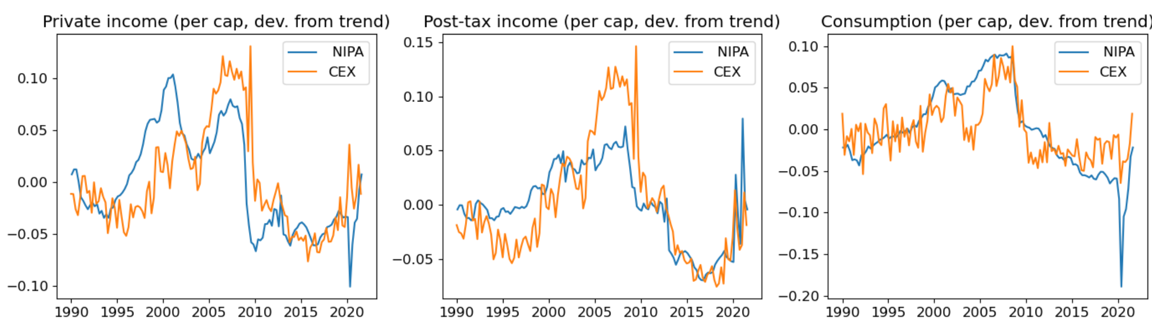

(b) Deviations from deterministic trends

Figure 1: NIPA and Chisini CEX means

3 A Descriptive Statistical Model

To prepare our data, we first subtract Y_t from all cross section quantiles:[10]

$$\widetilde{\mathbf{y}}_t^{priv} = \mathbf{y}_t^{priv} - Y_t \tag{5}$$

$$\widetilde{\mathbf{y}}_t^{post} = \mathbf{y}_t^{post} - Y_t \tag{6}$$

$$\widetilde{\mathbf{y}}_t^{cons} = \mathbf{y}_t^{cons} - Y_t \tag{7}$$

Next, define $\widetilde{\mathbf{y}}_t$ to be the $(1 \times 3M)$ vector

$$\widetilde{\mathbf{y}}_t = \left[\widetilde{\mathbf{y}}_t^{priv} - \boldsymbol{\nu}^{priv}, \quad \widetilde{\mathbf{y}}_t^{post} - \boldsymbol{\nu}^{post}, \quad \widetilde{\mathbf{y}}_t^{cons} - \boldsymbol{\nu}^{cons} \right]$$

where $\boldsymbol{\nu}^i$ is a $1 \times M$ vector of time-series means for each quantile of $i = \{priv, post, cons\}$. Then define $\mathbf{y}_t$ to be the $(3M+1) \times 1$ vector of aggregate income growth and our three scaled CEX cross-sections:

$$\mathbf{y}_t = \begin{bmatrix} Y_t - Y_{t-1} - \nu \\ \widetilde{\mathbf{y}}_t^\top \end{bmatrix}, \tag{8}$$

where ν is the in-sample mean of $Y_t - Y_{t-1}$. We assume that $\{Y_t\}$ is an **additive functional** defined by

$$\mathbf{y}_{t+1} = \mathbf{B}\,\mathbf{y}_t + \mathbf{a}_{t+1} \tag{9}$$

$$Y_{t+1} - Y_t - \nu = e_1\,\mathbf{B}\,\mathbf{y}_t + e_1\,\mathbf{a}_{t+1}, \tag{10}$$

[10] By construction, the cross-sectional mean

$$\frac{1}{M}\sum_{j=1}^{M} \tilde{y}_{j,t}^{priv} = 0 \quad \forall t \tag{3}$$

We compute cross-quantile means and standard deviations for each variable. For example, for private income we compute

$$\mu_t^{priv} := \frac{1}{M}\sum_{i=1}^{M} \tilde{y}_{i,t}^{priv}, \quad \sigma_t^{priv} := \sqrt{\frac{1}{M}\sum_{i=1}^{M} (\tilde{y}_{it}^{priv} - \mu_t^{priv})^2}. \tag{4}$$

We report these moments in Figure 11 in appendix A.

where $E[\mathbf{a}_t\, \mathbf{y}_{t-j}^\top] = 0 \quad \forall j \geq 1$, $E[\mathbf{a}_t\, \mathbf{a}_t^\top] = \mathbf{\Omega}$, $e_1 = \begin{bmatrix} 1 & 0 & \cdots & 0 \end{bmatrix}$ is a $(3M+1) \times 1$ selection vector and all eigenvalues of the $(3M+1) \times (3M+1)$ matrix $\mathbf{B}$ are strictly less than 1 in modulus so that $\{\mathbf{y}_t\}$ is an asymptotically stationary stochastic process[11]. Following Hansen (2012), decompose Y_t as[12]:

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

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

$$\mathbf{g} = e_1 \mathbf{B}(\mathbf{I} - \mathbf{B})^{-1}$$

The resolvent operator $(\mathbf{I} - \mathbf{B})^{-1}$ appears in the formulas for $\mathbf{H}$ and $\mathbf{g}$; $t\nu$ is a deterministic trend, $\sum_{j=1}^{t} \mathbf{H}\, \mathbf{a}_j$ is a martingale that we regard as a "stochastic trend", $\mathbf{g}\, \mathbf{y}_t$ is stationary, and $(\mathbf{g}\, \mathbf{y}_0 + Y_0)$ is a constant.[13]

Figure 2 presents decomposition (11) of Y_t.[14] Before the 2008 financial crisis, Y_t stayed above the deterministic trend with the martingale component close to zero and the stationary component the main driver. After 2008, Y_t fell and remained persistently below trend until the end of the sample. Initially, the martingale component fell sharply to minus 15% and was the main driver of Y_t. The stationary component fell gradually and became negative in 2015, while the martingale component increased. By the end of 2021, Y_t was 10% below the deterministic trend, driven equally by both components. We show below that this decomposition of Y_t conceals substantial heterogeneity in the quantile-specific stationary components after 2008.

We also decompose cross-section quantiles for individual incomes and consumptions. Since $\tilde{y}_{p,t}^{priv} = y_{p,t}^{priv} - Y_t - \nu_p^{priv}$, we represent the p-th percentile of private income as the

[11] We estimate $\mathbf{B}$ using the Dynamic Mode Decomposition outlined in our related paper Sargent and Selvakumar (2024). We present a brief overview in Section 4

[12] Also see the QuantEcon lecture on additive and multiplicative functionals available here https://python-advanced.quantecon.org/additive_functionals.html. Hansen (2012) specifies an equivalent decomposition with $\mathbf{a}_t = \mathbf{J}\mathbf{z}_t$, where $\mathbf{z}_t$ is a standard multivariate Gaussian random vector and $\mathbf{\Omega} = \mathbf{J}\mathbf{J}^\top$. Our representation circumvents numerical issues associated with the Cholesky decomposition of $\mathbf{\Omega}$.

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

[14] We estimate all model parameters – ν, ν^{priv}, ν^{post}, ν^{cons}, $\mathbf{B}$, $\mathbf{\Omega}$ – on data until Q4 2008. Thus, the trend $t\nu$ can be considered the "pre-crisis trend". It is important to keep in mind that Y_t is a measure of aggregate private income, capturing a subset of US GDP.

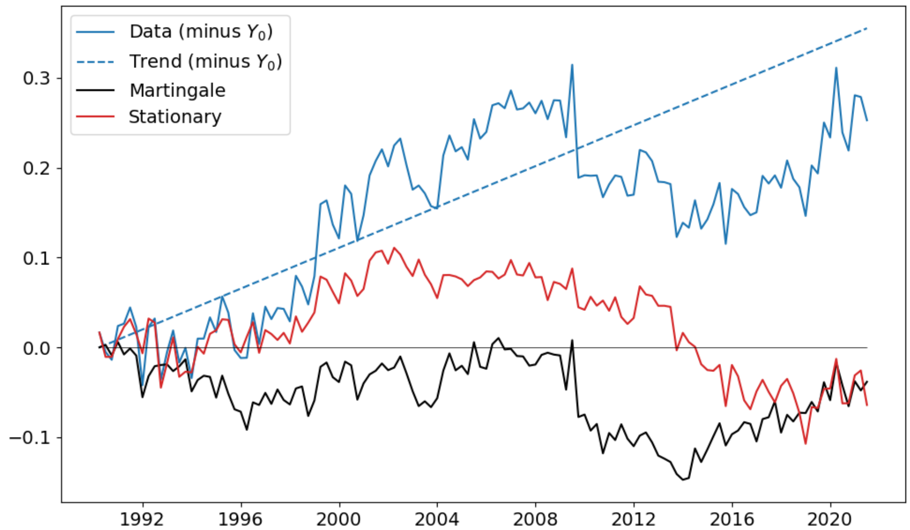

Figure 2: Decomposition of aggregate income Y_t

additive functional

$$y_{p,t}^{priv} = \underbrace{t\nu}_{\text{Trend}} + \underbrace{\sum_{j=1}^{t} \mathbf{H}\,\mathbf{a}_j}_{\text{Martingale}} + \underbrace{e_p^{priv}\,\mathbf{B}\,\mathbf{y}_{t-1} + e_p^{priv}\,\mathbf{a}_t - \mathbf{g}\,\mathbf{y}_t}_{\text{Stationary}} + \underbrace{(\mathbf{g}\,\mathbf{y}_0 + Y_0 + \nu_p^{priv})}_{\text{Constant}} \tag{12}$$

where e_p^{priv} is a $(3M+1)\times 1$ vector that selects the p-th percentile of private income. We form counterparts of the decomposition for all quantiles of $\mathbf{y}_t^{post}$ and $\mathbf{y}_t^{cons}$. Evidently, the deterministic trend $t\nu$ and martingale $\sum_{j=1}^{t}\mathbf{H}\,\mathbf{a}_j$ components are common across quantiles and variables, but the stationary and constant parts vary. We built in this structure when we specified our statistical model (9)–(10).

Figure 3 decomposes incomes and consumption quantiles using (12). For the 10th percentile, the level of 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. Consequently, differences across percentiles are largest for private income and smallest for consumption.

Keeping in mind that these deviations are sums of the common martingale and the quantile-specific stationary components in representation (12), the Figure 3 decompositions indicate interesting patterns in quantiles' deviations from their own deterministic trends.

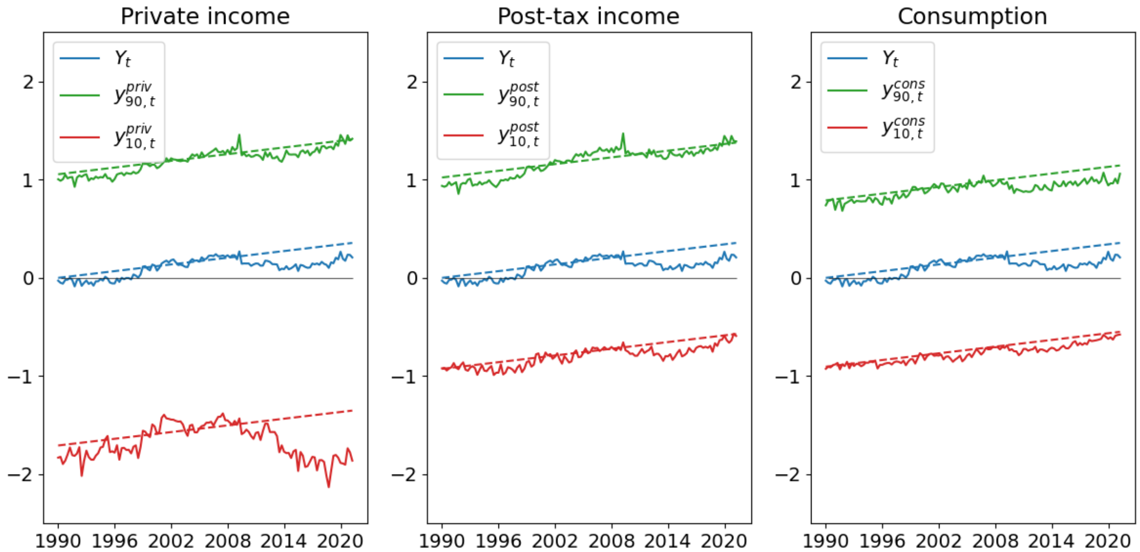

Figure 3: Evolution of quantiles

After the 2008 financial crisis, private income of the 90th percentile fell 10 percent below trend, and by 2021 the deviation from trend was zero. Similar numbers apply to post-tax income. Consumption for the 90th percentile, however, fell further – around 15 percent below trend and remained low until the end of our sample. As for lower quantiles, private incomes at the 10th percentile fell around 50 percent below trend following 2008. By 2021, private income was still 40 percent below its trend. Post-tax income and consumption were much less sensitive, falling at most 15-20 percent below trend. Both variables subsequently increased and were back to trend by 2021.

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 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. Similar differences across quantiles will also appear in discounted expected utility functions that we report in section 6.

To provide another perspective, Figure 4 reports results of using equation (12) to decom-

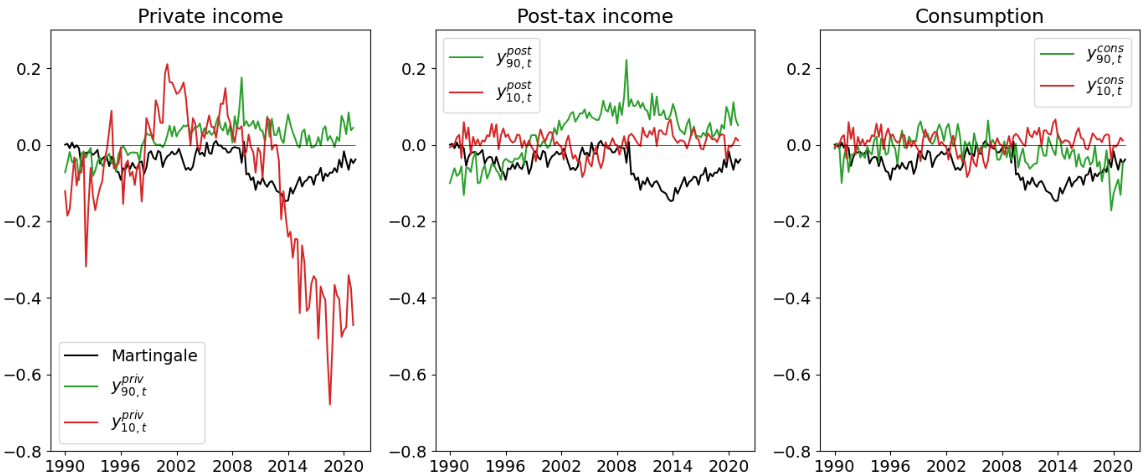

Figure 4: Components of quantiles

pose deviations of quantiles from their deterministic trends 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 (green for 90th percentiles, red for 10th percentiles). The common martingale component (which is the same as the martingale in Figure 2) fell 15% during the 2008 crisis to below and recovered to around minus 5% 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, the dynamics after 2014 are largely driven by the stationary component, which accounts for nearly the entire 40 percent gap from trend by 2021. On the other hand, after 2008 the stationary component for the 90th percentile was slightly positive, and largely offset the negative martingale component by the end of the sample.

Post-tax income stationary components for the 10th percentile was close to zero, making the common martingale component the predominant driver of the fall relative to trend. The stationary component for the 90th percentile was slightly positive (5 percent), partially offsetting the negative effect of the martingale component.

For consumption, the stationary components of both the 10th and 90th percentile were close to zero, making the common martingale component the major driver of consumption over our period. Toward the end of the sample, the stationary component started to drive consumption below trend for the 90th percentile. That the idiosyncratic stationary components are, on the whole, close to zero for the 10th and 90th percentiles suggest that dynamics of consumption are fairly similar across the distribution with the risks being evenly distributed.[15]

Remarkably, the 90th percentile of consumption is around 15 percent below trend at the end of our sample, compared to only 5 percent below trend for the 10th percentile, despite the 10th percentile of private income remaining close to 40 percent below trend at the end of the sample.

4 Role of Dynamic Mode Decomposition

We estimated the parameters of our model by constructing a Dynamic Mode Decomposition (DMD), a "machine-learning" technique used to study fluid dynamics (see Brunton and Kutz (2022)) that Sargent and Selvakumar (2024) link to vector autoregressions and linear Gaussian state-space systems. Our analysis begins with a data set formed by quarterly observations of the $(3M + 1) \times 1$ vector $\mathbf{y}_t$ defined in (9). We use these data to estimate the first-order VAR

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

The coefficient matrix $\mathbf{B}$ is $(3M+1) \times (3M+1)$ which, since $M = 98$, equals 295×295 in our application. Our 32 years of quarterly observations means that we have $T = 128$ observations with which to estimate the 295^2 coefficients in $\mathbf{B}$. Consequently, estimating (9) confronts us with an underdetermined least squares problem that we solve by using ideas connected to Dynamic Mode Decompositions. We set a integer $N << T$ and seek a rank-N first-order

[15]This feature is consistent with the bottom right panel of Figure 13 that plots time-series of demeaned $\widetilde{y}_{p,t}^{cons}$ 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.

VAR representation:

$$\begin{aligned} \mathbf{y}_t &= \widehat{\mathbf{B}}\,\mathbf{y}_{t-1} + \widehat{\mathbf{a}}_t \\ \mathbb{E}[\widehat{\mathbf{a}}_t \widehat{\mathbf{a}}_t^\top] &= \widehat{\boldsymbol{\Omega}}, \quad \widehat{\mathbf{a}}_t \perp \mathbf{y}_{t-1}\,. \end{aligned} \tag{14}$$

where $\widehat{\mathbf{B}}$ is a $(3M+1) \times (3M+1)$ matrix with rank N. We use the following algorithm to estimate $\widehat{\mathbf{B}}$ in (14):

1. Construct

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

where $\mathbf{Y}$ and $\mathbf{Y}'$ are both $(3M+1) \times T$ matrices.

2. Compute singular value decomposition (SVD) of data matrix $\mathbf{Y} = \mathbf{U}\,\boldsymbol{\Sigma}\,\mathbf{V}^\top$

3. Compress $\mathbf{Y}$ by retaining only its N largest singular values and corresponding vectors, then redefine $\mathbf{Y}$ as the compressed data matrix

$$\mathbf{Y} \approx \mathbf{U}_{[:,:N]}\,\boldsymbol{\Sigma}_{[:N,:N]}\,\mathbf{V}_{[:N,:]}^\top$$

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

$$\boldsymbol{\Phi} = \mathbf{Y}'\,\mathbf{V}_{[:N,:]}\,\boldsymbol{\Sigma}_{[:N,:N]}^{-1}\,\mathbf{W}, \quad \boldsymbol{\Lambda} = \mathbf{W}^{-1}(\mathbf{U}_{[:,:N]}^\top\,\mathbf{B}\,\mathbf{U}_{[:,:N]})\,\mathbf{W}$$

where $\mathbf{W}$ is a matrix of eigenvectors.

5. Compute

$$\widehat{\mathbf{B}} = \mathbf{Y}'\,\mathbf{Y}^+ = \mathbf{Y}'\,\mathbf{V}\,\boldsymbol{\Sigma}^{-1}\,\mathbf{U}^\top = \boldsymbol{\Phi}\boldsymbol{\Lambda}\boldsymbol{\Phi}^+ \tag{15}$$

where $\mathbf{Y}^+$ is the pseudoinverse of $\mathbf{Y}$, $\boldsymbol{\Lambda}$ is a diagonal matrix of eigenvalues of $\widehat{\mathbf{B}}$, and $\boldsymbol{\Phi}$ contains corresponding eigenvectors.

Formula (15) opens the door to constructing a Dynamic Mode Decomposition because it implies that our reduced-rank first-order VAR (14) can be represented as

$$\widehat{\mathbf{x}}_{t+1} = \mathbf{\Lambda}\widehat{\mathbf{x}}_t + \mathbf{\Lambda}\mathbf{\Phi}^+\widehat{\mathbf{a}}_t \tag{16}$$

$$\mathbf{y}_t = \mathbf{\Phi}\widehat{\mathbf{x}}_t + \widehat{\mathbf{a}}_t \tag{17}$$

where $\widehat{\mathbf{x}}_t = \mathbb{E}[\mathbf{x}_t \,|\, \mathbf{y}_{t-1}] = \mathbf{\Lambda}\mathbf{\Phi}^+ \mathbf{y}_{t-1}$. Because of how it is connected to a linear Gaussian state-space representation that rationalizes the reduce rank specification of $\widehat{\mathbf{B}}$, we call representation (16)–(17) a "pseudo-innovations representation" that is associated with the linear Gaussian state-space representation (see Appendix B).

$$\mathbf{x}_{t+1} = \mathbf{\Lambda}\,\mathbf{x}_t + \mathbf{C}\,\mathbf{w}_{t+1} \tag{18}$$

$$\mathbf{y}_t = \mathbf{\Phi}\,\mathbf{x}_t + \mathbf{v}_t, \quad \mathbb{E}\,\mathbf{v}_t\,\mathbf{v}_t^\top = \mathbf{R}, \tag{19}$$

$$\tag{20}$$

where $\mathbf{\Sigma}_\infty = (\mathbf{\Phi}^\top\mathbf{\Omega}^{-1}\mathbf{\Phi})^{-1}, \mathbf{R} = \mathbf{\Omega} - \mathbf{\Phi}\mathbf{\Sigma}_\infty\mathbf{\Phi}^\top, \mathbf{C}\,\mathbf{C}^\top = \mathbf{\Sigma}_\infty - \mathbf{\Lambda}\mathbf{\Phi}^+\,\mathbf{R}(\mathbf{\Lambda}\mathbf{\Phi}^+)^\top$.

5 Additional Findings

In section 5.1 we briefly describe estimated parameters and plot the dominant two DMD modes against components of our data set. In section 5.2 we extract from our statistical model some quantitative inferences about interactions between aggregate income and our cross-section quantiles.

5.1 Parameters and DMD Modes

Setting $N = 3$, we estimate the following parameters:[16]

$$\mathbf{\Lambda} = \begin{bmatrix} 0.98 & 0 & 0 \\ 0 & 0.85 & 0 \\ 0 & 0 & 0.40 \end{bmatrix}, \quad \mathbf{C}\,\mathbf{C}^\top = \begin{bmatrix} 0.27 & 0.19 & -0.04 \\ 0.19 & 0.27 & -0.01 \\ -0.04 & -0.01 & 0.4 \end{bmatrix} \tag{21}$$

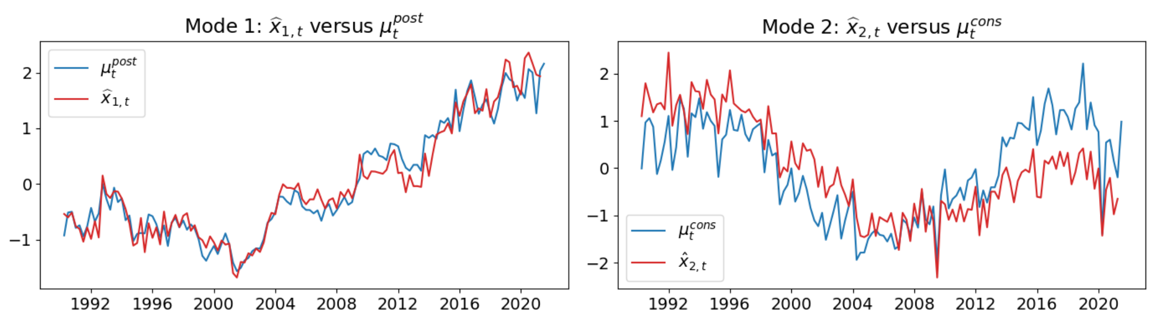

Figure 5: First and second DMD modes

Figure 5 plots the estimated DMD modes. The first mode is highly persistent ($\lambda_1 = 0.98$) and strongly correlated with the mean of post-tax income ($\text{Corr}(\mu_t^{post}, \hat{x}_{1t}) = 0.95$). The second mode is less persistent ($\lambda_2 = 0.85$) and moderately correlated with the mean of consumption ($\text{Corr}(\mu_t^{cons}, \hat{x}_{2t}) = 0.58$).

Figure 6 shows the loadings of quantiles on the two dominant modes. The left panel plots quantile-specific loadings on private income. Since $\mu_t^{priv} = 0$ for all t by construction, the loadings of private income quantiles must average to zero, thus shedding light on the increased income inequality associated with an increase in mode 1.

The loadings of post-tax quantiles (middle top panel) on mode 1 are all positive, which is consistent with mode 1 serving 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 Guvenan et al. (2017). The

[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 C).Appendix C plots 50 singular values for our $\mathbf{Y}$ data matrix and associated dynamic modes for the largest three singular values.Sargent and Selvakumar (2024) present simulations that indicate how setting N too small adversely affects inference. We err on the side of caution and estimate the model with $N = 3$ and verify that mode 3 indeed resembles noise (Figure 15). We proceed by focusing on two dominant modes.

right panel plots loadings of consumption quantiles on mode 1. All loadings are positive, suggesting that a rise in mode 1 coincides with a rise 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 through management of their private savings. High consumption quantiles are less sensitive to mode 1 than high post-tax income quantile counterparts, indicating more consumption smoothing on their part.

Turning to loadings on mode 2 in the lower three panels of Figure 6, consumption quantile loadings are all positive, so an increase in mode 2 raises all consumption quantiles. In contrast, an increase in mode 2 is associated with a reduction in post-tax income inequality, since incomes for low quantiles rise and incomes for high quantiles fall.

5.2 Micro-Macro Interactions

To measure how macro and micro dynamics interact statistically, we compute an orthogonal decomposition of the innovation covariance matrix:

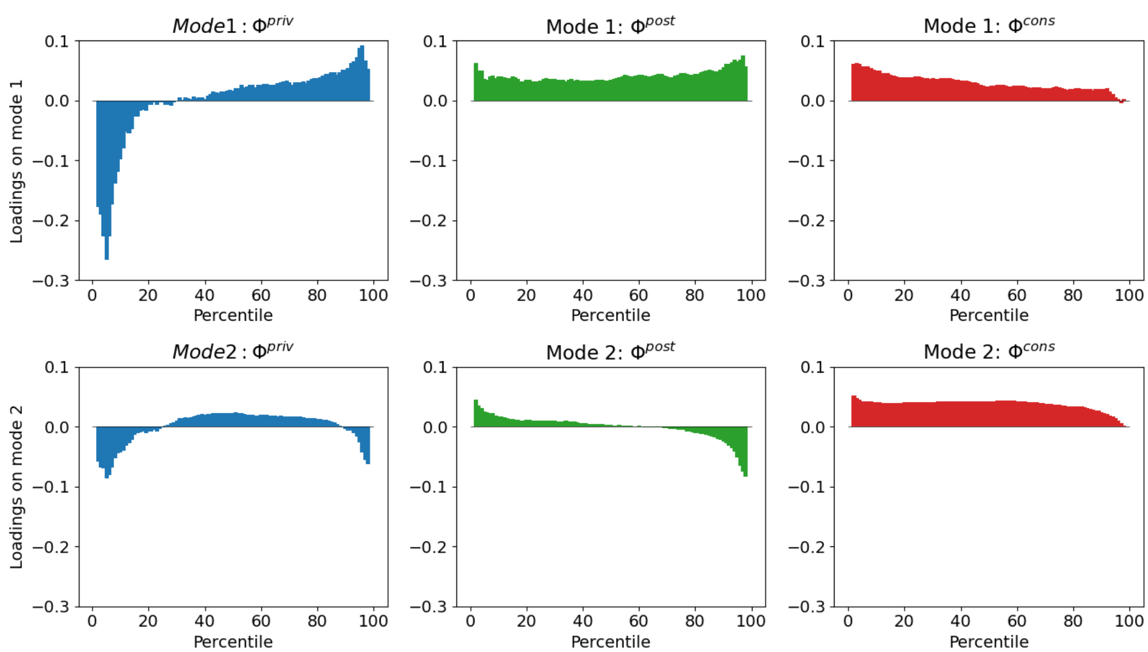

Figure 6: DMD loadings on modes 1 and 2

	$\mathbf{\Omega}_{(1,1)}$	$\mathbf{R}_{(1,1)}$	$(\mathbf{\Phi\Sigma}_\infty\mathbf{\Phi}^\top)_{(1,1)}$
Variance	0.00103	0.0010	$2.8e^{-5}$
Fraction of $\mathbf{\Omega}_{(1,1)}$ (%)	-	97%	3%

Table 1: Decomposition of variance of aggregate income growth innovation

$$\widehat{\mathbf{\Omega}} = \mathbf{R} + \mathbf{\Phi\Sigma}_\infty\mathbf{\Phi}^\top \tag{22}$$

Since the (1,1) element of $\mathbf{y}_t$ corresponds to aggregate income growth $Y_t - Y_{t-1} - \nu$, we can use the (1,1) elements of $\mathbf{\Omega}, \mathbf{R}$ and $\mathbf{\Phi\Sigma}_\infty\mathbf{\Phi}^\top$ to compute the fraction of the variance of innovation to $Y_t - Y_{t-1} - \nu$ that can be attributed to the DMD modes $\widehat{\mathbf{x}}_t$, represented by the (1,1) element of $\mathbf{\Phi\Sigma}_\infty\mathbf{\Phi}^\top$. Since $\widehat{\mathbf{x}}_t$ was computed using data on cross-section quantiles, $(\mathbf{\Phi\Sigma}_\infty\mathbf{\Phi}^\top)_{(1,1)}$ reveals the importance of cross-section dynamics on the dynamics of aggregate income growth. The residual fraction, computed using the (1,1) element of $\mathbf{R}$, represents the fraction of variance attributable to the component that is orthogonal to cross-section dynamics. Table 1 shows that only 3% of the innovation variance can be attributed to the common modes $\widehat{\mathbf{x}}_t$, while 97% can be attributed to idiosyncratic component. The table suggests that, as intermediated through the dynamic modes, cross-section dynamics have limited consequences for the dynamics of aggregate income growth.[17]

To investigate possible reverse influences of the dynamics of $Y_t - Y_{t-1} - \nu$ on inequality dynamics, we compute population least squares regressions of innovations to several percentile differences on innovations to $Y_t - Y_{t-1} - \nu$. We use an appropriate vector $\mathbf{b}$ to define a measure of inequality $z_t = \mathbf{b}\,\mathbf{y}_t$, (e.g. a 90-10 quantile difference for private income), such that by choosing different vectors $\mathbf{b}$, we create different measures of inequality. We want to compute a population regression of the innovation $a_{z,t}$ to z_t on the innovation $a_{1,t}$ to aggregate income growth $y_{1,t} = \mathbf{c}\,\mathbf{y}_t$:

$$a_{z,t} = \beta a_{1,t} + \xi_t \tag{23}$$

[17]See similar findings by Chang et al. (2024) and a discussion of them by Sargent (2024), who recommended interpreting them as outcomes of the comprehensive redistribution and social insurance present in the US tax and transfer system, as well as countercyclical monetary-fiscal policies that Lucas (1987, 2003) concluded had mostly attenuated post WWII business cycles. In this context, it is tempting to wonder how much HANK model-inspired Ramsey policies, like those studied by Bhandari et al. (2021) that balance price stability against providing insurance to less well endowed consumers, could have contributed to those outcomes.

The innovation variance of z_t is $\mathbf{b}\boldsymbol{\Omega}\mathbf{b}^\top$; and $\texttt{Cov}(a_{z,t}, a_{1,t}) = \mathbf{b}\boldsymbol{\Omega}\mathbf{c}^\top$. Thus, the fraction of variance in $a_{z,t}$ attributable to $a_{1,t}$ is

$$\frac{\beta\mathbf{c}\boldsymbol{\Omega}\mathbf{c}^\top\beta^\top}{\mathbf{b}\boldsymbol{\Omega}\mathbf{b}^\top} \tag{24}$$

The left panel of Figure 7 plots contributions to variances of inequality of private income, post-tax income and consumption attributable to the innovation to aggregate income growth. We use our $\mathbf{b}\,\mathbf{y}_t$ measures of inequality, ranging from 90-10 percentile gaps to 82-18 percentile gaps. For private income, fractions of attributable variance are between 6% and 15%, with the highest values for the mid-range percentile differences (around 86-14). This is between two and five times as large as the 3% variance attribution to the common modes of aggregate income in Table 1. For post-tax income and consumption cross-sections, variance ratios are smaller. These variance contributions indicate that aggregate income is a larger source of inequality in private income than of post-tax income and consumption. Thus, Figure 7 provides another perspective on the pervasive insurance and redistribution present in our CEX cross-sections from 1990-2021.[18]

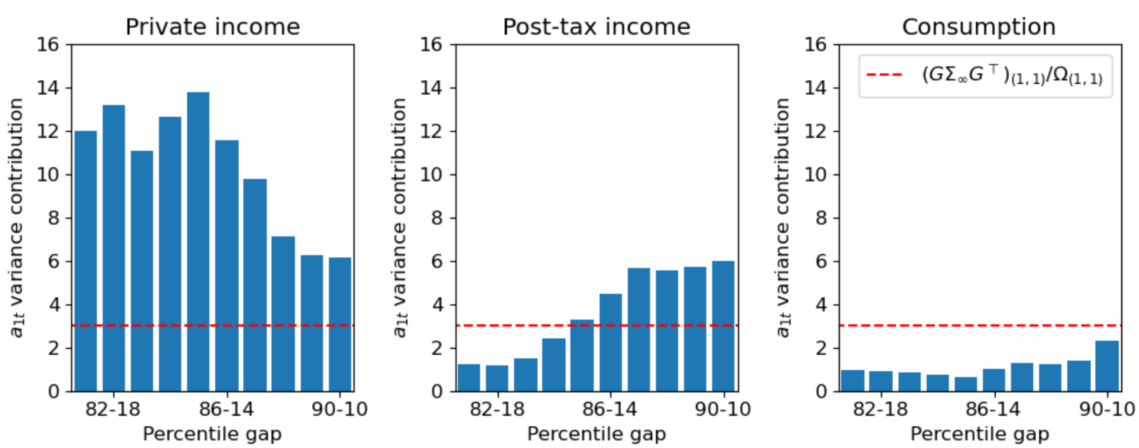

Figure 7: Contributions of aggregate income to inequality

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

6 Welfare Costs of Martingale and Stationary Shocks

> 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)

By capitalizing on our statistical model's representation of serially correlated risks in quantile-specific consumption growth rates for our CEX consumers, this section extends welfare calculations of Lucas (1987, 2003). We compute welfare consequences of eliminating risks to consumption growth for hypothetical individuals who are permanently stuck in the same quantiles of income or consumption.[19] We also use the CEX income and consumption data to estimate quantile-specific values from participation in the US tax and transfer system. These calculations allow us to appraise a statement of Lucas (1987, p. 105) that fiscal redistribution and insurance promote individual consumers' discounted utilities much more than further moderations of aggregate fluctuations around trend growth rates. Our estimates confirm Lucas's statement.

Lucas (1987, Sec. III), Lucas (2003), and Tallarini (2000) used versions of the following model to calculate welfare benefits of completely eliminating a representative consumer's exposure to i.i.d. risk in U.S. post WWII consumption growth per capita.[20] Where c_t is the log of per capita consumption C_t, a representative agent orders consumption streams according to the discounted expected value

$$v = (1-\beta)E_0 \sum_{t=0}^{\infty} \beta^t c_t, \quad \beta \in (0,1) \tag{25}$$

When c_t is generated by the following random walk with drift

$$c_{t+1} - c_t = \mu + \sigma_\epsilon \epsilon_{t+1}, \quad \epsilon_{t+1} \sim \mathcal{N}(0,1), \tag{26}$$

[19] See Appendix F where we incorporate the possibility of social mobility to our calculations.

[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.

the expected discounted value (25) is

$$v = v(c_0) = \frac{\beta}{(1-\beta)}\mu + c_0 \tag{27}$$

and the conditional expectation of C_{t+1} is

$$E_t C_{t+1} = \exp\left(\mu + \frac{1}{2}\sigma_\epsilon^2\right) C_t. \tag{28}$$

A value function $v(c_0)$ for a risk-free path that starts at c_t^0 at $t = 0$ and that has the same conditional means $E_0 C_{t+j}$ as the risky sequence generated by (26) is

$$v(c_0^n) = \frac{\beta}{(1-\beta)}\left(\mu + \frac{\sigma_\epsilon^2}{2}\right) + c_0^n. \tag{29}$$

An increment

$$c_0 - c_0^n = \frac{\beta \sigma_\epsilon^2}{2(1-\beta)} \tag{30}$$

is a "compensating difference" that equates (27) to (29). The gap $c_0 - c_0^n$ in equation (30) is a percentage of consumption that a representative consumer would 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 this as an upper bound on the welfare gains that could be gathered from further improving post WWII countercyclical U.S. monetary and fiscal policies.

We now provide an extension of formula (30) that applies to our setting. The logarithmic random walk-with-drift specification (26) on which formula (30) is based exposes a representative consumer to i.i.d risk in consumption growth. By way of contrast, our specification[21]

$$y_{p,t}^{cons} = \underbrace{t\nu}_{\text{Trend}} + \underbrace{\sum_{j=1}^{t} \mathbf{H}\,\mathbf{a}_j}_{\text{Martingale}} + \underbrace{e_p^{cons}\,\mathbf{B}\,\mathbf{y}_{t-1} + e_p^{cons}\,\mathbf{a}_t - \mathbf{g}\,\mathbf{y}_t}_{\text{Stationary}} + \underbrace{(\mathbf{g}\,\mathbf{y}_0 + Y_0 + \nu_p^{cons})}_{\text{Constant}} \tag{31}$$

exposes a consumer in the pth quantile to both i.i.d. risk in consumption growth through $\mathbf{a}_t$ and serially correlated risk in consumption growth through the martingale $\sum_{j=1}^{t} \mathbf{H}\,\mathbf{a}_j$ and

[21] Here the e_p^{cons} is a $(3M+1) \times 1$ selects $y_{p,t}^{cons}$, the p-th percentile of consumption.

the serially correlated state variables $\mathbf{y}_t$.[22]

We extend Lucas's equalizing difference calculation (30) for his representative agent's consumption process (26) to quantile-specific equalizing differences for hypothetical consumers who are also exposed to quantile-specific serially correlated risks in their consumption growth rates. We shall compute percentages of current consumption that someone stuck forever in the p-th 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 (31).

To this end, our person of interest is someone permanently in the p-th consumption quantile who values consumption streams according to the discounted expected value[23]

$$V_p(\boldsymbol{\zeta}_t, y_{p,t}^{cons}) = (1-\beta)E_0 \sum_{t=0}^{\infty} \beta^t y_{p,t}^{cons}, \quad \beta \in (0,1) \tag{32}$$

where $y_{p,t}^{cons}$ is the log of their consumption. Similar to (12), additive functional (10) implies that $y_{p,t}^{cons} = e_p^{cons}\widetilde{\mathbf{y}}_t + Y_t + \nu_p^{cons}$. Consequently,

$$\boldsymbol{\zeta}_{t+1} = \mathbf{A}\,\boldsymbol{\zeta}_t + \mathbf{C}\,\boldsymbol{\varepsilon}_{t+1} \tag{33}$$

$$y_{p,t+1}^{cons} - y_{p,t}^{cons} = \mathbf{D}_p^{cons}\,\boldsymbol{\zeta}_t + \mathbf{F}_p^{cons}\boldsymbol{\varepsilon}_{t+1}, \tag{34}$$

where $\boldsymbol{\zeta}_t = \begin{bmatrix} 1 & \mathbf{y}_t \end{bmatrix}^\top$; $\mathbf{D}_p^{cons} = \begin{bmatrix} \nu & e_1\,\mathbf{B} + e_p^{cons}(\mathbf{B} - \mathbf{I}) \end{bmatrix}$ and $\mathbf{F}_p^{cons} = e_1 + e_p^{cons}\mathbf{C}$. In representation (33)-(34), $\mathbf{F}_p\boldsymbol{\varepsilon}_{t+1}$ expresses the the consumer's exposure to an i.i.d. component of risk in quantile-p consumption growth, while $\mathbf{D}_p^{cons}\,\boldsymbol{\zeta}_t$ expresses the consumer's exposure to persistent risk in quantile-p consumption growth.

We seek a value function $V_p(\boldsymbol{\zeta}_t, y_{p,t}^{cons})$ in (32) that satisfies a recursion

$$V_p(\boldsymbol{\zeta}_t, y_{p,t}^{cons}) = (1-\beta)y_t^{cons,p} + \beta V_p(\boldsymbol{\zeta}_{t+1}, y_{p,t+1}^{cons}).$$

[22] 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.

[23] In our exercises below, we will set $\beta = 0.99$ for our quarterly dataset.

Using logic of Hansen et al. (2008) and Hansen and Sargent (2010), we verify that[24]

$$V_p(\boldsymbol{\zeta}_t, y_{p,t}^{cons}) = \boldsymbol{\lambda}_p^{cons\top} \boldsymbol{\zeta}_t + y_{p,t}^{cons} \tag{35}$$

where

$$\boldsymbol{\lambda}_p^{cons\top} = \beta\, \mathbf{D}_p^{cons} (\mathbf{I} - \beta\, \mathbf{A})^{-1}. \tag{36}$$

We want to compare a quantile-p value function (35) with one associated with a risk-free certainty-equivalent consumption path. In the spirit of Lucas (1987, 2003) and Tallarini (2000), we seek an adjustment to consumption at time t that renders the p-th quantile consumer indifferent between the certainty equivalent plan and the original risky consumption plan given in (33) - (34). To calculate such a certainty-equivalent, risk-free consumption plan, we first deduce from (33)-(34) that at time $t+j$, p-th quantile consumption satisfies

$$y_{p,t+j}^{cons} = y_{p,t}^{cons} + \boldsymbol{\psi}_p^{cons} \boldsymbol{\zeta}_t + \sum_{i=0}^{j-1} \mathbf{J}_{p,i}^{cons} \boldsymbol{\varepsilon}_{t+j-i} \tag{37}$$

where

$$\boldsymbol{\psi}_p^{cons\top} = \boldsymbol{\lambda}_p^{cons\top} = \beta\, \mathbf{D}_p^{cons} (\mathbf{I} - \beta\, \mathbf{A})^{-1} \tag{38}$$

$$\mathbf{J}_{p,i}^{cons} = \mathbf{F}_p^{cons} + \sum_{\ell=0}^{i-1} \mathbf{D}_p^{cons}\, \mathbf{A}^l\, \mathbf{C} \tag{39}$$

$$\alpha_p^{cons} = \frac{\beta}{2} \sum_{i=0}^{\infty} \beta^i \mathbf{J}_{p,i}^{cons} \mathbf{J}_{p,i}^{cons\top}, \tag{40}$$

and the scalar α_p^{cons} appears in the following value function $W_p(\boldsymbol{\zeta}_t, y_{p,t}^{cons}) = (1-\beta) \sum_{j=0}^{\infty} \beta^j y_{p,t+j}^{cons}$ for the p-th quantile for a certainty-equivalent path that starts at $(\boldsymbol{\zeta}_t, y_{p,t}^{cons})$:

$$W^p(\boldsymbol{\zeta}_t, y_{p,t}^{cons}) = y_{p,t}^{cons} + \boldsymbol{\lambda}_p^{cons\top} \boldsymbol{\zeta}_t + \alpha_p^{cons} \tag{41}$$

The presence of serially correlated risk in consumption growth in representation (31) shapes

[24] Hansen and Sargent (2010) deployed continuous-time versions of this value function. The discrete-time version is available here http://www.tomsargent.com/research/longrunrisk_tom_14.pdf.

how α_p^{cons} depends on consumption quantile p's loading onto $\boldsymbol{\zeta_t}$ through $\mathbf{D}_p^{cons}$ and also on how the p-th quantile loads onto ε_{t+1} through $\mathbf{F}_p^{cons}$. The inclusion of risk in value function (41) enters through α_p^{cons} and $\boldsymbol{\lambda}_p^{cons\top}$. Within α_p^{cons}, the i.i.d. part enters through $\mathbf{F}_p^{cons}$, while the persistent part enters in two ways, one through the discounted sum of previous shocks ε_{t+1}, and the other through $\mathbf{D}_p^{cons}$ as a constituent of $\boldsymbol{\lambda}_p^{cons\top}$.

Let $y_{p,t}^{cons,CE}$ be log consumption at time t in the certainty-equivalent plan. We seek a proportional decrease in the certainty-equivalent path that leaves a p-th quantile consumer indifferent between the risky and risk-free paths, i.e. $V_p(\boldsymbol{\zeta}_t, y_{p,t}^{cons}) = W_p(\boldsymbol{\zeta}_t, y_{p,t}^{cons,CE})$, which implies that

$$y_{p,t}^{cons} - y_{p,t}^{cons,CE} = \alpha_p^{cons} \tag{42}$$

Thus, $y_{p,t}^{cons} - y_{p,t}^{cons,CE}$ differs across quantiles but is constant across time. This feature is a direct implication of our specification of the additive process (11) for Y_t, where we built in time-invariance of the compensating differences.

To study evidence of redistribution and social insurance in our CEX data set, we also compute discounted expected values, in which we temporarily pretend that, rather than consumption $y_{p,t}^{cons}$, p-th quantile households consume either total private income $y_{p,t}^{priv}$ or income net of taxes and transfers $y_{p,t}^{post}$. Comparing these objects for private income and income net of transfers allows us to estimate the value to the p-th quantile consumer of participating in the US tax and transfer system. Comparing them for income net of transfers and for consumption allows us to estimate value increments presumably achieved with consumption smoothing through management of individual savings.[25]

The left panel of Figure 8 plots certainty-equivalent compensations for eliminating all risk (long-run and i.i.d) from quantile-specific growth (α_p^{cons}) for all quantiles of consumption, after tax and transfer income, and private income. The right panel plots objects associated with the i.i.d. component of these parameterized by $\mathbf{F}_p^{cons}\mathbf{F}_p^{cons\top}$. First, compensation is highest for the lower quantiles of private income, reflecting the substantial income risk they face. This is largely driven by their loadings on ε_{t+1}, encoded by $\mathbf{F}_p^{cons}\mathbf{F}_p^{cons\top}$ (right chart). Agents in the 50th percentile of consumption distribution would be willing to pay around

[25]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).

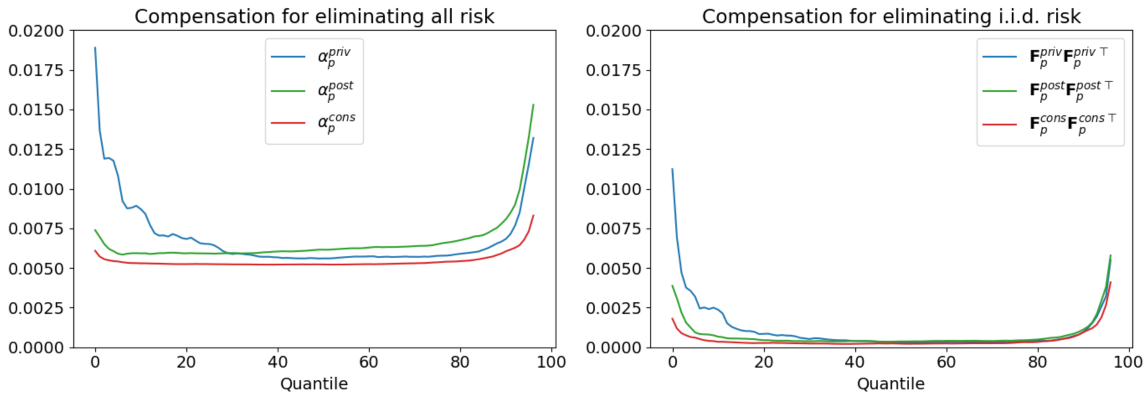

Figure 8: Compensations for eliminating risks in income or consumption growth

0.5 percent of their current level of consumption to remove all risk in consumption growth. This is approximately 10-times larger than Lucas's estimate of 0.05% of consumption for his representative consumer. Lucas's estimate, however, is inline with our estimate for i.i.d risk for the same consumer, which is about 0.05% of consumption. This implies that the remaining 0.45 percent comes from exposure to serially correlated risk in consumption growth.[26]

Second, for all quantiles, proportional consumption compensations are smaller than analogous 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, the difference in proportional compensations between post-tax income and consumption increase as we move toward higher quantiles. To us, this indicates more consumption smoothing by higher quantile households through management of private savings. Notice that the same object is negligible for $\mathbf{F}_p^{cons}\mathbf{F}_p^{cons\top}$ in the right panel. This suggests that the differences in the left chart reflect the ability of private savings to insure households from serially correlated risks in consumption growth, rather than i.i.d risk.

[26] 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".

6.1 Welfare benefits of social and private insurance

To put compensating consumption level differences ranging from .05% (Lucas's) to .5% (ours) into perspective, we find it useful to compare these numbers with increments or decrements of our computed value functions for p-th quantile consumers that were delivered by US tax and transfer system. To accomplish this, we compute $V_p(\boldsymbol{\zeta}_t, y_{p,t}^{cons}) - V_p(\boldsymbol{\zeta}_t, y_{p,t}^{priv})$. We interpret this difference as value of actions of the US tax and transfer system and/or self-insurance for a p-th quantile household. Figure 9, which plots this difference in values for all quantiles, shows for the 10th percentile that it is orders of magnitude larger (0.77) than the value of eliminating exposure to growth rate risk for any of the three variables plotted in Figure 8. The value difference of consumption over private income also increases after 2008, albeit small in magnitude, reflecting the fact that consumption fell much less than private income for this group. For the 90th percentile, the difference in values is negative, reflecting in part that their private income is higher than their consumption as a result of how the US tax and transfer system treats consumers in the 90th percentile; for 90th percentile persons, the absolute value of the difference plotted in the right panel of Figure 9 indicates that to these consumers, the redistribution and insurance dwarf the welfare consequences of eliminating exposure to risks in their income and consumption growth rates.

For these two percentiles, Figure 10 plots the following decomposition of $V_p(\boldsymbol{\zeta}_t, y_{p,t}^{cons}) -$

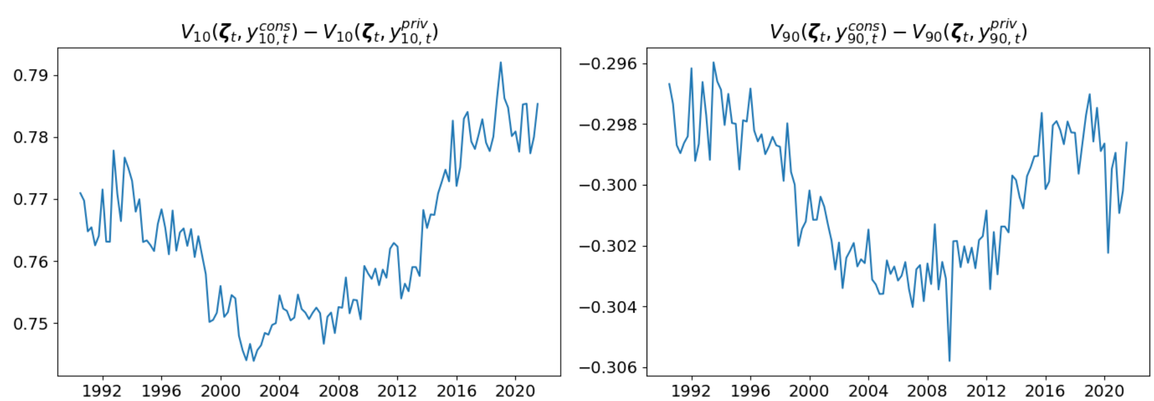

Figure 9: $V_p(\boldsymbol{\zeta}_t, y_{p,t}^{cons}) - V_p(\boldsymbol{\zeta}_t, y_{p,t}^{priv})$ for $p = 10$ (left panel) and $p = 90$ (right panel).

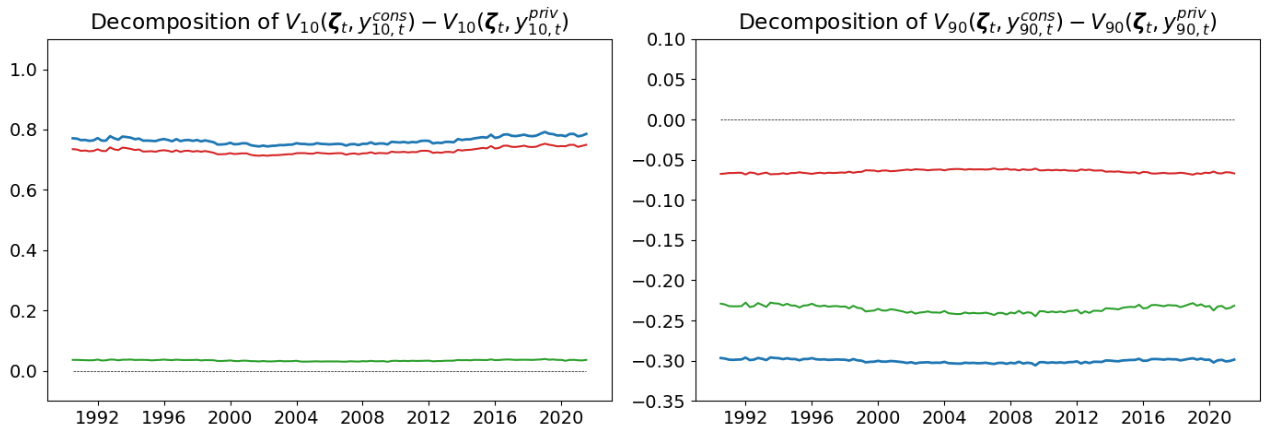

Figure 10: Where $p = 10$ in the left panels and $p = 90$ on the right panels, the blue lines plot $V(\zeta_t, y^{cons}_{p,t}) - V(\zeta_t, y^{priv}_{p,t})$; the green lines plot $V(\zeta_t, y^{cons}_{p,t}) - V(\zeta_t, y^{post}_{p,t})$; the red line plots $V(\zeta_t, y^{post}_{p,t}) - V(\zeta_t, y^{priv}_{p,t})$. See equation (43).

$V_p(\zeta_t, y^{priv}_{p,t})$:

$$V_p(\zeta_t, y^{cons}_{p,t}) - V_p(\zeta_t, y^{priv}_{p,t}) = \left[V_p(\zeta_t, y^{cons}_{p,t}) - V_p(\zeta_t, y^{post}_{p,t})\right] + \left[V_p(\zeta_t, y^{post}_{p,t}) - V_p(\zeta_t, y^{priv}_{p,t})\right] \tag{43}$$

For the 10th percentile, most of the left hand side (blue) is contributed by the second term in brackets on the right side of equation (43) (red). That the first term is small (green) confirms that 10th percentile households live virtually hand-to-mouth, so for 10th precentile households the majority of welfare gains on the left of (43) 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 (43) (the green line) contributes much more, indicating consequences of substantial private savings.

6.2 Compensating differences in deterministic trends

We now calculate equalizing differences in today's level of consumption that our hypothetical households would be willing to sacrifice for an increase in the trend growth rate ν. Let $\widehat{V}_p$ be discounted expected utility of someone who is permanently in the p-th consumption quantile in a fictitious economy with aggregate income trend growth $\widehat{\nu} > \nu$. The value function $\widehat{V}_p$ can be represented as

$$\widehat{V}_p(\boldsymbol{\zeta}_t, \widehat{y}^{cons}_{p,t}) = \widehat{\boldsymbol{\lambda}}^{cons\top}_p \boldsymbol{\zeta}_t + \widehat{y}^{cons}_{p,t}, \tag{44}$$

where $\widehat{\boldsymbol{\lambda}}^{cons\top}_p = \beta \widehat{\mathbf{D}}^{cons}_p (\mathbf{I} - \beta\,\mathbf{A})^{-1}$ and $\widehat{\mathbf{D}}^{cons}_p = \begin{bmatrix} \widehat{\nu} & e_1\,\mathbf{B} + e^{cons}_p(\mathbf{B} - \mathbf{I}) \end{bmatrix}$. A proportional decrease in the high-growth rate $\widehat{\nu}$ path that leaves the p-th quantile of consumption indifferent between ν and $\widehat{\nu}$ is

$$y^{cons}_{p,t} - \widehat{y}^{cons}_{p,t} = (\widehat{\boldsymbol{\lambda}}^{cons\top}_p - \boldsymbol{\lambda}^{cons\top}_p)\boldsymbol{\zeta}_t \tag{45}$$

$$= \beta(\widehat{\mathbf{D}}^{cons}_p - \mathbf{D}^{cons}_p)(\mathbf{I} - \beta\,\mathbf{A})^{-1}\boldsymbol{\zeta}_t. \tag{46}$$

Using definitions of $\mathbf{D}^{cons}_p, \widehat{\mathbf{D}}^{cons}_p$ and $\boldsymbol{\zeta}_t$, we can write the proportional decrease as

$$y^{cons}_{p,t} - \widehat{y}^{cons}_{p,t} = \beta(\widehat{\nu} - \nu)\phi_1, \tag{47}$$

where ϕ_1 is the first element of the vector $(\mathbf{I} - \beta\,\mathbf{A})^{-1}\boldsymbol{\zeta}_t$. The compensating difference is constant across time and quantiles, a feature that our additive process (11) builds in.

Formula (47) computes that households in all quantiles would be willing to sacrifice 24.75 percent of current consumption to increase aggregate trend growth ν by 1 percent (from 1.2% to 2.2%) per annum. This is orders of magnitude larger than the Figure 8 that any household would be willing to sacrifice to remove all risks to consumption growth. This finding, that households in all quantiles care much more about increasing trend growth than further reducing risk in their consumption growth rates, extends 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."*

It might be enlightening to compute increments in trend growth that would make a hypothetical household in the p-th quantile of consumption indifferent between increasing trend growth and eliminating all consumption risk. To do this, we transform (47) to obtain

$$\widehat{\nu}_p - \nu = \frac{(y^{cons}_{p,t} - y^{cons,CE}_{p,t})}{\beta\phi_1} = \frac{\alpha^{cons}_p}{\beta\phi_1} \tag{48}$$

We can compute analogous measures for value functions for total (social and private) insurance. Table 2 reports these compensations for eliminating all risks in consumption growth and the value of insurance for the 10th, 50th and 90th percentiles. We use equation (48) to compute utility-equivalent trend growth increments. The table shows that the value of eliminating all risks to consumption growth is equivalent to around a 0.022 percentage point increase in annual trend growth, and this is similar across all quantiles. These are much smaller than our estimates of values of redistribution and social and private insurance. For the 10th percentile, access to the redistribution provided by the US tax, transfer, and financial system is welfare-equivalent to an increase in annual trend growth of aggregate income by more than a 3 percentage points. This number is negative for the median and the 90th percentile – minus 0.75 percentage points and minus 1.21 percentage points respectively. These quantities further highlight the role the US tax and transfer system in redistributing welfare across low and high income quantiles.

Quantile	Value of eliminating all risk to consumption growth		Value of social and private insurance	
	$(y_{p,t}^{cons} - y_{p,t}^{cons,CE})$ %	$(\nu^{CE} - \nu)$ %, annual.	$V_p(\zeta_t, y_{p,t}^{cons}) - V_p(\zeta_t, y_{p,t}^{priv})$ %	$\nu^{insurance} - \nu$ %, annual.
$p = 10$	0.54	0.022	78.53	3.17
$p = 50$	0.52	0.021	-18.70	-0.75
$p = 90$	0.56	0.023	-29.86	-1.21

Table 2: Compensating differentials for eliminating risks to consumption growth

7 Concluding Remarks

We have used additive functionals and dynamic mode decompositions to analyze the co-evolution of Consumer Expenditure Survey cross-sections of private earned income, post-tax income, and consumption from 1990 to 2021. We decomposed a Chisini mean of private earned income and cross-section quantiles into constant, deterministic trend, martingale, and stationary components following Hansen (2012). Our statistical model reveals patterns that illuminate how aggregate income dynamically interacts with redistribution and social insurance across households.

Our decomposition of aggregate income shows that before the 2008 financial crisis, aggregate income was briefly above its deterministic trend with the stationary component the main driver. After 2008, the martingale component fell suddenly, while the stationary component fell gradually. By the end of the sample, they contributed equally to the trend gap. This aggregate pattern conceals substantial heterogeneity in the stationary components of quantiles. For the 10th percentile, private income fell substantially after the financial crisis and remained below trend through 2021. However, both post-tax income and consumption fell less and returned close to trend by 2021. For the 90th percentile, private income returned to trend by 2021, as did post-tax income, but consumption remained below trend.

Cross-section dynamics contribute only 3% to the innovation variance of aggregate income growth, while innovations to aggregate income affect cross-section dynamics more substantially. For private income inequality measures (e.g. 90-10 percentile difference), the fraction of variance attributable to aggregate income innovations ranges from 6% to 15%. These fractions are smaller for post-tax income and consumption, indicating that redistribution and insurance mechanisms attenuate the transmission of aggregate shocks to inequality.

Our extension of Lucas's (1987, 2003) welfare calculations 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 0.5% of current consumption to eliminate all consumption growth risk—ten times larger than Lucas's representative agent estimate of 0.05%. Our estimates recover similar numbers to Lucas's when we eliminate only i.i.d risk, suggesting that the persistent component of risk accounts for 0.45 percentage points of the 0.5% compensating difference. Nevertheless, the further benefits of eliminating consumption growth risk is quantitatively small. We capitalize on our additive functional assumption to compute that eliminating all consumption growth risk is equivalent to an 0.025 percentage point increase in trend growth.

The welfare benefits from participating in the US tax and transfer system dwarf those from eliminating consumption growth risk. For 10th percentile consumers, the value increment from the tax and transfer system equals 0.77, orders of magnitude larger than the 0.5% welfare cost of consumption growth risk. For this group, the benefit is equivalent to around a 3 percentage point increase in trend growth. For 90th percentile consumers, the redistribution through the tax and transfer system imposes welfare costs that similarly exceed the welfare costs of growth rate risk. 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.

Seven appendices conclude our paper. Appendix A describes data sources. Appendix B reviews connections between linear-quadratic-Gaussian state space models and Dynamic Mode Decompositions on which our section 4 analysis is based. Appendix C 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 in the tradition of Blanchard and Quah (1993) on our section 3 additive functional (9)- (10) of aggregate income $\{Y_t\}$. Appendix E provides another decomposition of the section 6 value difference $V_p(\boldsymbol{\zeta}_t, y_{p,t}^{cons}) - V_p(\boldsymbol{\zeta}_t, y_{p,t}^{priv})$. Appendix F describes how to extend our analysis to include transitions across consumption and income quantiles. Appendix G describes how to extend our section 6 welfare calculations to incorporate concerns about model misspecification using method described by Barillas et al. (2009).

A Data

Figure 11 presents the evolution of cross-quantile moments. The levels, the standard deviation of private income (σ_t^{priv}) is higher than post-tax income (σ_t^{post}) or consumption (σ_t^{cons}). Moreover, σ_t^{post} appears to be larger than σ_t^{cons} in most periods. σ_t^{priv} exhibits substantial variation over time, increasing during the 2008 financial crisis and the COVID-19 period. In contrast, the standard deviations of post-tax income (σ_t^{post}) and consumption (σ_t^{cons}) are more stable, reflecting the smoothing effects of taxes, transfers, and consumption behavior.

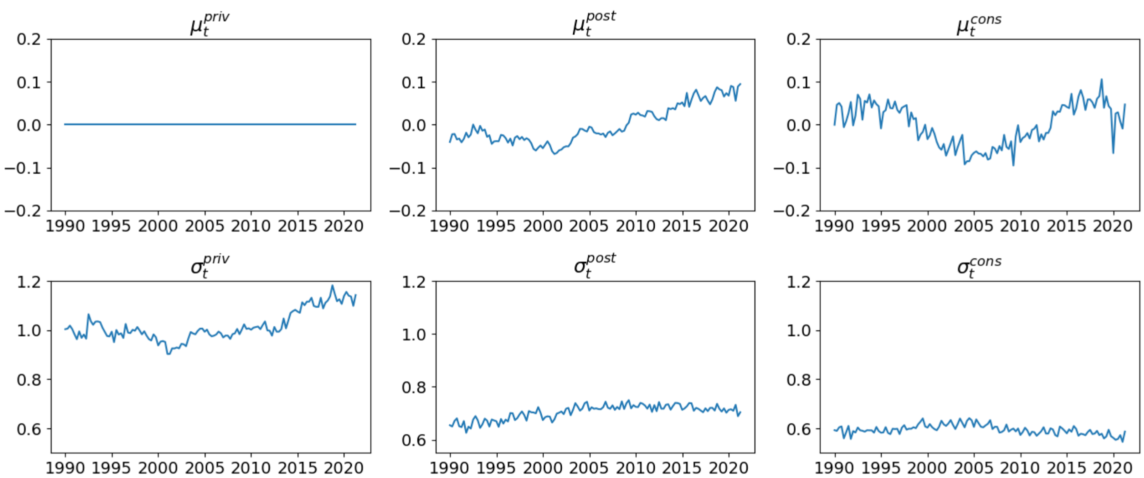

Figure 11: Cross-quantile moments

Figure 12 shows the 90-10 percentile differences for all three variables. 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.

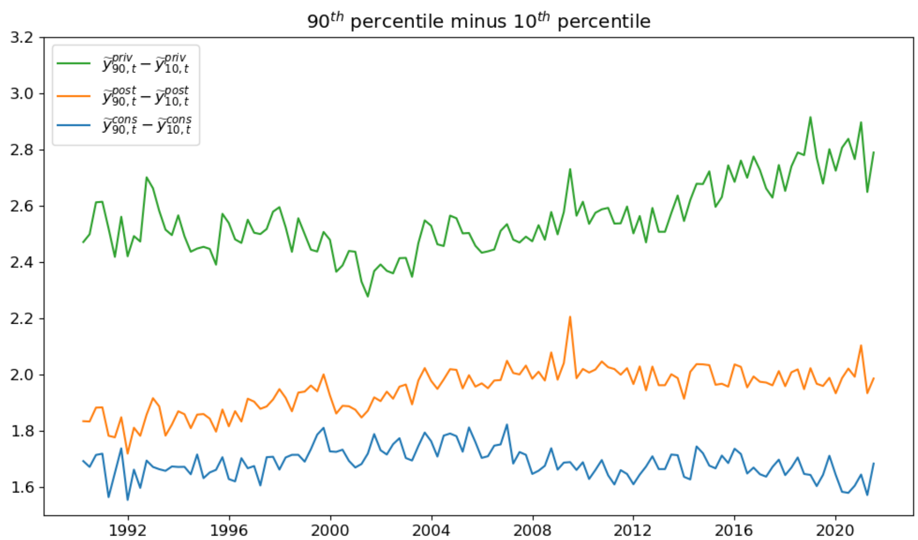

Figure 12: 90th percentile minus 10th percentile

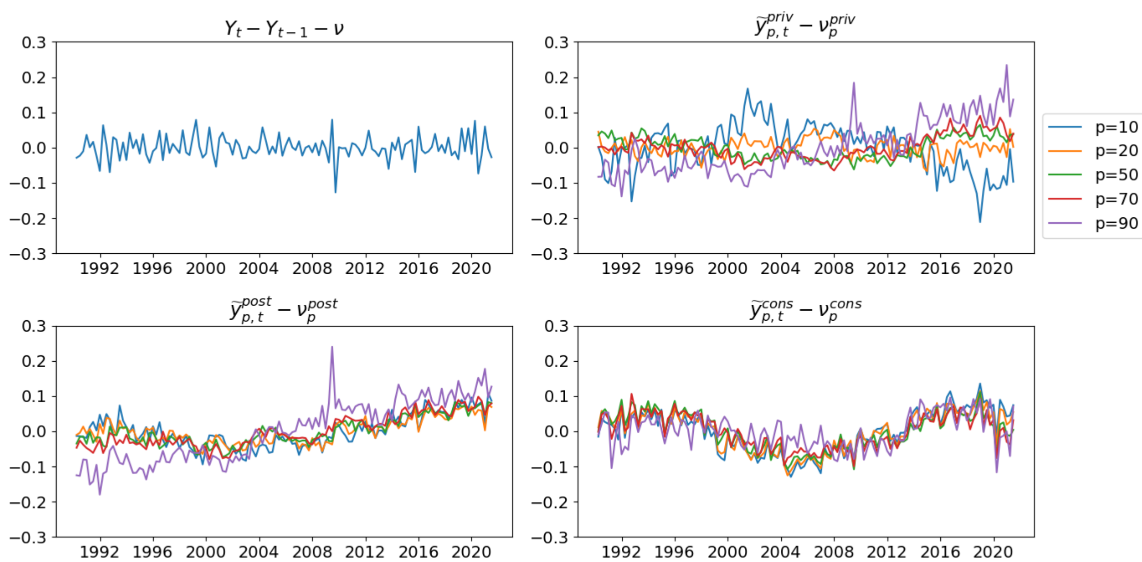

Figure 13: Demeaned time-series of quantiles, $\tilde{y}^{priv}_{p,t}$, $\tilde{y}^{post}_{p,t}$, $\tilde{y}^{cons}_{p,t}$

	Code Mneumonic		
	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 3: Categorizing CEX income into private and transfers

B Connections to LQG State-Space Models

Sargent and Selvakumar (2024) describe conditions under which approximation (14) is consistent with an underlying linear-quadratic-Gaussian (LQG) state-space model:

$$\mathbf{x}_{t+1} = \mathbf{A}\,\mathbf{x}_t + \mathbf{C}\,\mathbf{w}_{t+1} \tag{49}$$

$$\mathbf{y}_t = \mathbf{G}\,\mathbf{x}_t + \mathbf{v}_t \tag{50}$$

where $\mathbf{y}_t$ is our observed $(3M+1)\times 1$ vector, $\mathbf{x}_t$ is an $N \times 1$ vector of latent states, $\mathbf{w}_t \sim N(0, \mathbf{I}_{N\times N})$ is the state innovation, $\mathbf{v}_t \sim N(0, \mathbf{R}_{(3M+1)\times(3M+1)})$ is the observation noise. The conditions are[27]

- $M \gg N$
- $\mathbf{A}$ is diagonal so that $\mathbf{x}_{i,t+1} = \mathbf{A}_{ii}\,\mathbf{x}_{i,t} + \sum_j \mathbf{C}_{ij}\,\mathbf{w}_{j,t+1}$
- $\mathbf{G}$ has full column rank
- $\|\,\mathbf{G}^\top\,\mathbf{G}\,\| = O(M)$ and $\|\,\mathbf{R}\,\| = o(M)$

The state-space model in equations (49)-(50) implies an infinite-order VAR for y_t:

$$\mathbf{y}_t = \sum_{j=1}^{\infty} \mathbf{B}_j^{\infty}\,\mathbf{y}_{t-j} + \mathbf{a}_t \tag{51}$$

with:

$$\mathbb{E}[\mathbf{a}_t\,\mathbf{a}_t^\top] = \boldsymbol{\Omega} = \mathbf{G}\,\boldsymbol{\Sigma}_\infty\,\mathbf{G}^\top + \mathbf{R} \tag{52}$$

$$\mathbb{E}[\mathbf{a}_t\,\mathbf{y}_{t-j}^\top] = \mathbf{0} \quad \forall j \geq 1 \tag{53}$$

$$\mathbf{B}_j^{\infty} = \mathbf{G}(\mathbf{A} - \mathbf{K}\,\mathbf{G})^{j-1}\,\mathbf{K} \quad \forall j \geq 1 \tag{54}$$

where $\mathbf{K}$ is the steady-state Kalman gain and $\boldsymbol{\Sigma}_\infty$ is the steady-state covariance of the Kalman filter state estimate error. The innovations representation associated with this infinite-order

[27] See Sargent and Selvakumar (2024) and Iao and Selvakumar (2024).

VAR is:

$$\widehat{\mathbf{x}}_{t+1} = \mathbf{A}\,\widehat{\mathbf{x}}_t + \mathbf{K}\,\mathbf{a}_t \tag{55}$$

$$\mathbf{y}_t = \mathbf{G}\,\widehat{\mathbf{x}}_t + \mathbf{a}_t \tag{56}$$

where $\widehat{\mathbf{x}}_t = \mathbb{E}[\mathbf{x}_t \,|\, \mathbf{y}^{t-1}]$ is the Kalman filter estimate of the state based on past observations, and $\mathbf{a}_t = \mathbf{y}_t - \mathbb{E}[\mathbf{y}_t \,|\, \mathbf{y}^{t-1}]$ is the innovation or prediction error. Sargent and Selvakumar (2024) show that $(\mathbf{A} - \mathbf{K}\,\mathbf{G}) \approx \mathbf{0}$ under their restrictions, which implies that

$$\mathbf{B}_j^{\infty} \approx \begin{cases} \mathbf{G}\,\mathbf{K} & j = 1 \\ \mathbf{0} & j > 1, \end{cases} \tag{57}$$

so that the infinite-order VAR in equation (51) is well-approximated by a first-order VAR:

$$\mathbf{y}_t = \mathbf{B}\,\mathbf{y}_{t-1} + \mathbf{a}_t \tag{58}$$

where $\mathbf{B} \approx \mathbf{B}_1^{\infty} = \mathbf{G}\,\mathbf{K}$. Moreover, $(\mathbf{A} - \mathbf{K}\,\mathbf{G}) \approx \mathbf{0}$ implies that $\mathbb{E}[\mathbf{x}_t \,|\, \mathbf{y}^{t-1}] \approx \mathbb{E}[\mathbf{x}_t \,|\, \mathbf{y}_{t-1}]$, which we exploit to align the innovations representation with the reduced-rank VAR estimated by DMD.

As stated in the text, estimating reduced-rank VAR (14) generates a pseudo-innovations representation:

$$\begin{aligned} \widehat{\mathbf{x}}_{t+1} &= \mathbf{\Lambda}\widehat{\mathbf{x}}_t + \mathbf{\Lambda}\Phi^{+}\widehat{\mathbf{a}}_t \\ \mathbf{y}_t &= \mathbf{\Phi}\widehat{\mathbf{x}}_t + \widehat{\mathbf{a}}_t, \end{aligned}$$

which aligns with the authentic innovations representation in equations (55)-(56) after we set $\mathbf{A} = \mathbf{\Lambda}, \mathbf{K} = \mathbf{\Lambda}\mathbf{\Phi}^{+}, \mathbf{G} = \mathbf{\Phi}, \mathbf{\Sigma}_\infty = (\mathbf{G}^\top\,\mathbf{\Omega}^{-1}\,\mathbf{G})^{-1}, \mathbf{R} = \mathbf{\Omega} - \mathbf{G}\,\mathbf{\Sigma}_\infty\,\mathbf{G}^\top, \mathbf{C}\,\mathbf{C}^\top = \mathbf{\Sigma}_\infty - \mathbf{K}\,\mathbf{R}\,\mathbf{K}^\top$.

C Singular Values and DMD Modes

Figure 14 plots singular values of the $\mathbf{Y}$ matrix. Figure 15 plots time series of three DMD modes $\hat{\mathbf{x}}_t$ associated with the largest three singular values. As noted in the text, we retain only the two largest singular values when we estimate $\mathbf{B}$ via reduced-rank least squares.

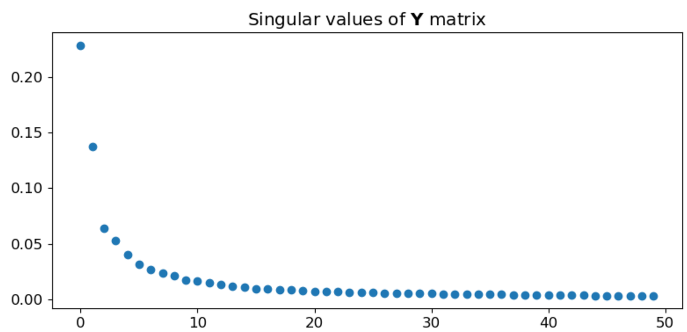

Figure 14: 50 largest singular values of DMD data matrix $\mathbf{Y}$

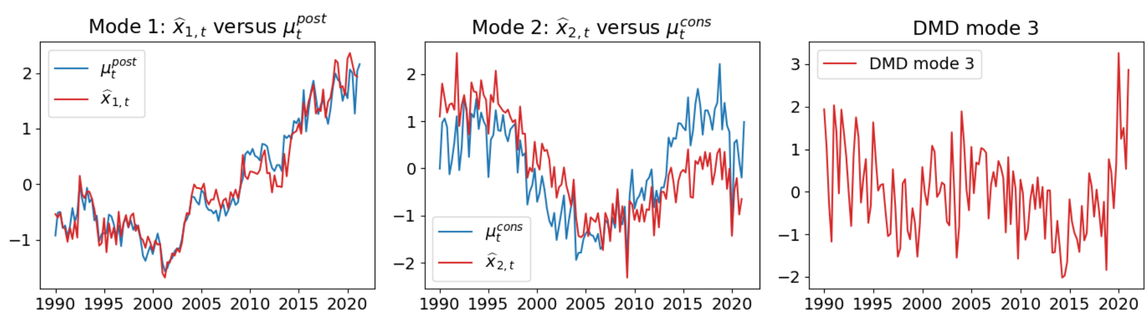

Figure 15: Time-series of estimated DMD modes $\widehat{\mathbf{x}}_t$

D A Blanchard-Quah Analysis

To impose the long-run restrictions of Blanchard and Quah (1993), we represent the additive functional (9)- (10) for $\{Y_t\}$ as

$$\mathbf{y}_{t+1} = \mathbf{B}\,\mathbf{y}_t + \mathbf{Q}\varepsilon_{t+1} \tag{59}$$

$$Y_{t+1} - Y_t - \nu = e_1\,\mathbf{B}\,\mathbf{y}_t + e_1\mathbf{Q}\varepsilon_{t+1} \tag{60}$$

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

$$\lim_{j\to\infty} \mathbb{E}[Y_{t+j}|\varepsilon_{1t}] \neq 0$$

$$\lim_{j\to\infty} \mathbb{E}[Y_{t+j}|\varepsilon_{it}] = 0 \quad \forall i \geq 1$$

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

$$\mathbf{Q} = (\mathbf{I} - \mathbf{B})\widetilde{\mathbf{Q}} \tag{61}$$

$$\mathbf{S}_y(0) = \widetilde{\mathbf{Q}}\widetilde{\mathbf{Q}}^\top \tag{62}$$

where $\mathbf{S}_y(0) = (\mathbf{I} - \mathbf{B})\mathbf{\Omega}(\mathbf{I} - \mathbf{B})^\top$ is the spectrial density matrix of $\mathbf{y}_t$ at frequency zero.[28]

Figure 16 plots a prediction error variance decomposition of Y_t attributable to $\boldsymbol{\varepsilon_{1,t}}$) and to all other shocks. It shows that indeed first shock accounts for an increasing share of the variance of aggregate income Y_t as the horizon increases, reaching about 95% at very long horizons.

Figure 17 plots the impulse response functions of 10th and 90th percentiles of income and consumption computed using our Blanchard and Quah identified system (59)-(60). The impulse response functions show that a positive permanent shock to aggregate income initially increases inequality in private income as the 90-10 percentile gap widening for about 50 periods before gradually returning to its steady state. For post-tax income, 90 and 10 percentiles both rise, the 90th percentile albeit 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.

[28]Due to the reduced-rank $\mathbf{B}$ and thus the approximate reduced-rank nature of $\mathbf{\Omega}$, we encountered numerical issues when computing the Cholesky factorization of $\mathbf{S}_y(0)$ to obtain $\widetilde{\mathbf{Q}}$. 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 then with zero. Calling the new matrix $\widetilde{\mathbf{D}}$, we compute $\widetilde{\mathbf{Q}} = \mathbf{L}\widetilde{\mathbf{D}}$.

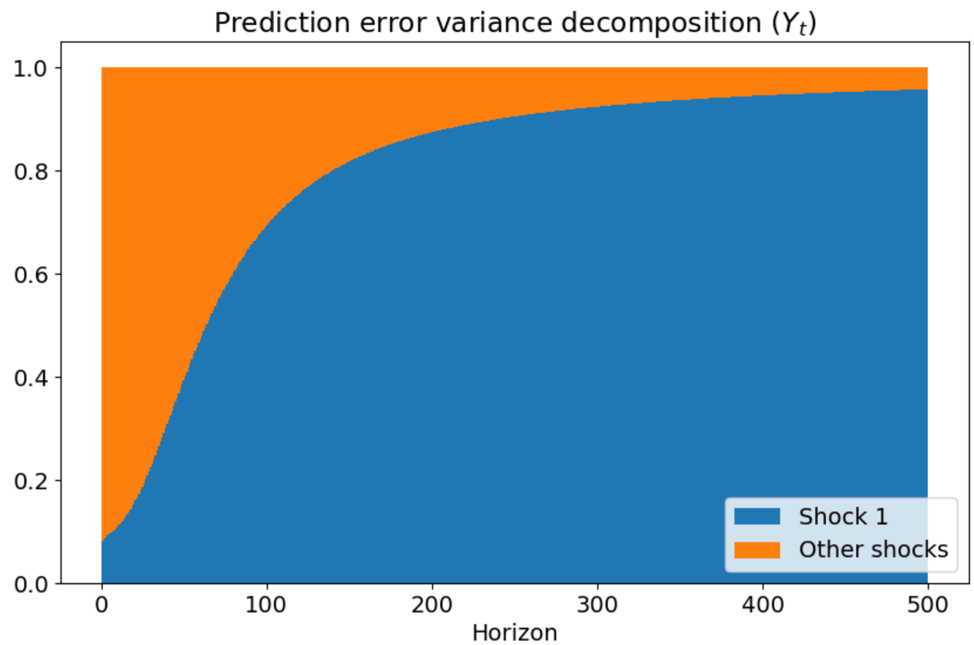

Figure 16: Variance decomposition of structural shocks

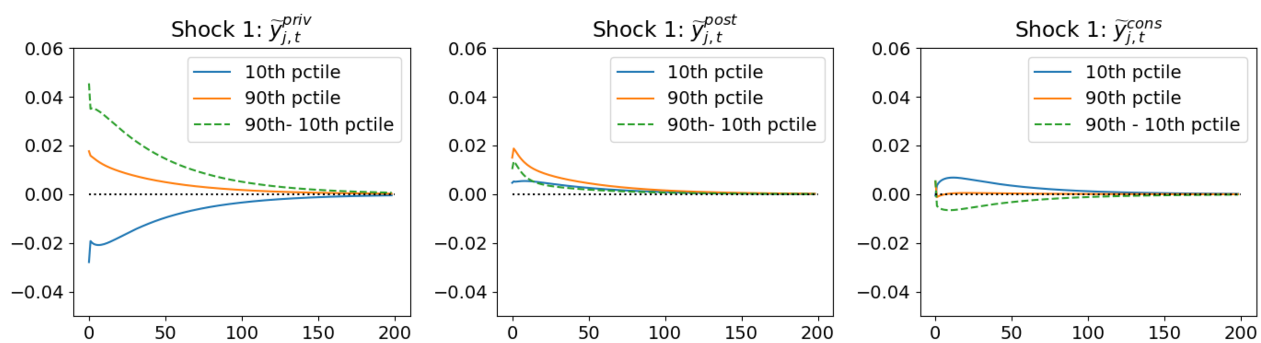

Figure 17: Impulse response

E Another Decomposition

We construct an alternative decomposition of $V_p(\zeta_t, y_{p,t}^{cons}) - V_p(\zeta_t, y_{p,t}^{priv})$:

$$V_p(\zeta_t, y_{p,t}^{cons}) - V_p(\zeta_t, y_{p,t}^{priv}) = (\boldsymbol{\lambda}_p^{cons\top} - \boldsymbol{\lambda}_p^{priv\top})\zeta_t + (y_{p,t}^{cons} - y_{p,t}^{priv}) \tag{63}$$

We interpret the first 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

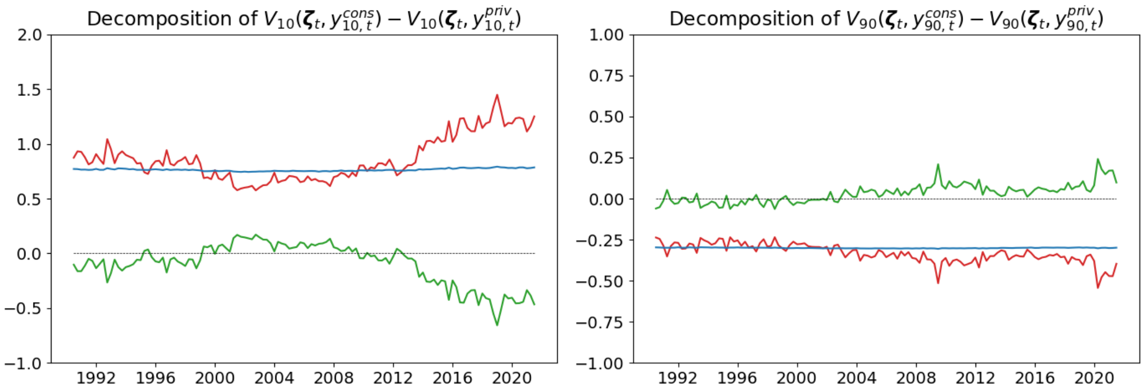

Figure 18: Where $p = 10$ in the left panels and $p = 90$ on the right panels, the blue lines plot $V_p(\boldsymbol{\zeta}_t, y^{cons}_{p,t}) - V_p(\boldsymbol{\zeta}_t, y^{priv}_{p,t})$; the green line plots $(\boldsymbol{\lambda}^{cons\top}_p - \boldsymbol{\lambda}^{priv\top}_p)\zeta_t$; the red line plots $y^{cons}_{p,t} - y^{priv}_{p,t}$. See equation (63).

second term measures consequences of redistribution. Figure 18 plots components of this decomposition. For both 10th and 90th percentile, the second term (red) is the predominant actor. In the case of the 10th percentile, the second term has become larger since 2012, while the first term becomes 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 (63),

$$(\boldsymbol{\lambda}^{cons\top}_p - \boldsymbol{\lambda}^{priv\top}_p)\zeta_t = (\boldsymbol{\lambda}^{cons\top}_p - \boldsymbol{\lambda}^{post\top}_p)\zeta_t + (\boldsymbol{\lambda}^{post\top}_p - \boldsymbol{\lambda}^{priv\top}_p)\zeta_t \tag{64}$$

$$y^{cons}_{p,t} - y^{priv}_{p,t} = (y^{cons}_{p,t} - y^{post}_{p,t}) + (y^{post}_{p,t} - y^{priv}_{p,t}) \tag{65}$$

Equation (64) 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 (65) decomposes the consequences of redistribution into private savings (first term) versus social (second term).

The panels in the top row of Figure 19 plot the components of (64) for the 10th (left) and the 90th (right) percentiles. For the 10th percentile, the value of social insurance (red) is the predominant driver of the value of total insurance (blue), while the value of private insurance (green) is negligible. For the 0th percentile, the value of private insurance (green) is the predominant driver of overall insurance, while social insurance is close to zero across the

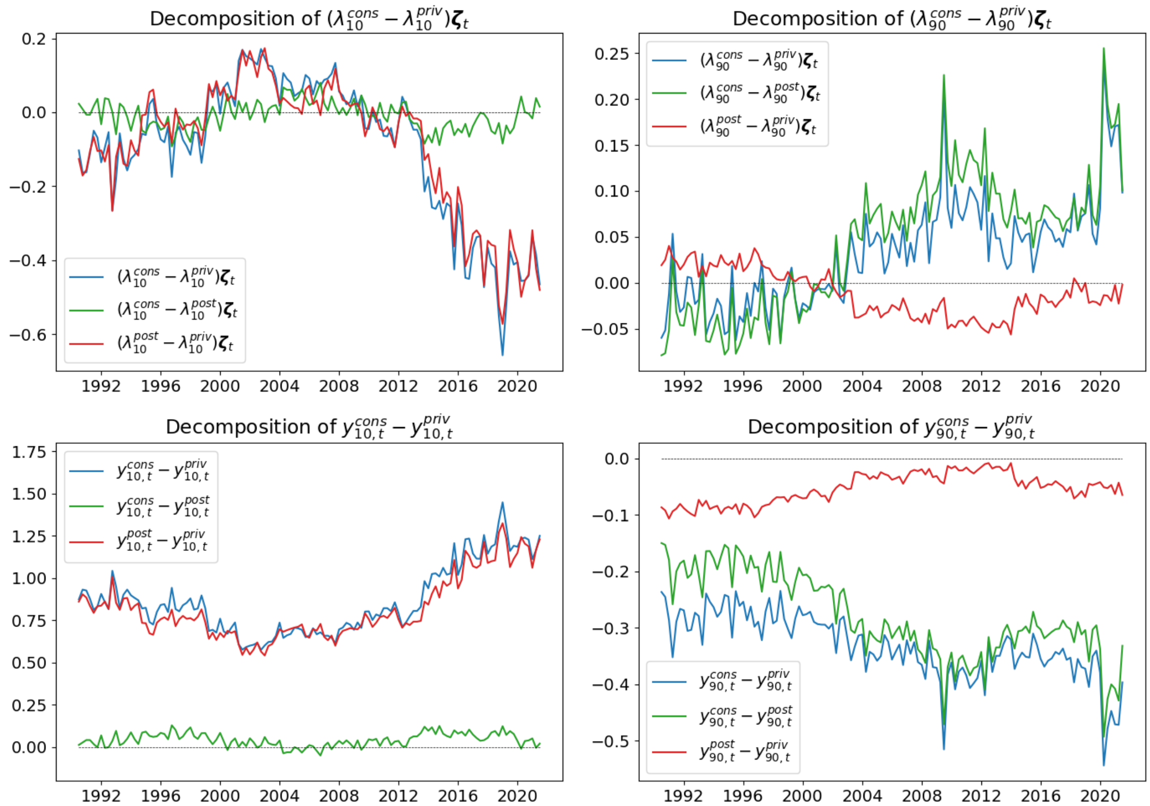

Figure 19: Where $p = 10$ (left panels) and $p = 90$ (right panels) the top panels plot the decomposition of $(\boldsymbol{\lambda}_p^{cons\top} - \boldsymbol{\lambda}_p^{priv\top})\zeta_t$; see equation (64). The bottom panels plots the decomposition of $y_{p,t}^{cons} - y_{p,t}^{priv}$; see equation (65).

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 for upper quantile households.

The panels in the bottom row of Figure 19 plot the components of (65) for the 10th and 90th percentile. The left hand side panel shows that the majority of $y_{10,t}^{cons} - y_{10,t}^{priv}$ is driven by $y_{10,t}^{post} - y_{10,t}^{priv}$, suggesting again that redistribution plays a large role for these households. The right hand side panel shows that, for the 90th percentile $y_{10,t}^{cons} - y_{10,t}^{post}$ is the dominant driver, suggesting that private savings is more important than social redistribution.

F Mobility across quantiles

We can construct an hypothetical consumer residing now in consumption quantile p who is exposed to possible transitions to other consumption quantiles with one-quarter transition probabilities described by a time-invariant $M \times M$ stochastic matrix $\mathbf{P}$. The scaled expected discounted utility $(1-\beta)E_0 \sum_{t=0}^{\infty} \beta^t y_{p,t}^{cons}$ is now given by value function

$$V(\boldsymbol{\zeta}_t, y_{p,t}^{cons}) = (1-\beta) y_{p,t}^{cons} + \beta \sum_{i=1}^{M} \mathbf{P}_{p,i} V(\boldsymbol{\zeta}_{t+1}, y_{i,t+1}^{cons}) \tag{66}$$

If we guess that this value function takes a form

$$V(\boldsymbol{\zeta}_t, y_{p,t}^{cons}) = \boldsymbol{\lambda}_p^{cons\top} \boldsymbol{\zeta}_t + \sum_{i=1}^{M} b_{p,i} y_{i,t}^{cons}$$

we can verify that $\boldsymbol{\lambda}_p^{cons}, b_{i,1}, \ldots, b_{i,M}$ for $i = 1, \ldots, M$ satisfy

$$\lambda_i^{cons\top} = \beta \sum_j P_{ij} \left(\lambda_j^{cons\top} \mathbf{A} + \sum_k a_{j,k} \mathbf{D}_k^{cons\top} \right) \tag{67}$$

$$b_{i,k} = (1-\beta)\delta_{ik} + \beta \sum_j \mathbf{P}_{ij} b_{j,k} \tag{68}$$

where $\delta_{i,k}$ is the Kronecker delta that obeys $\delta_{i,k} = 1$ if $i = k$ and 0 otherwise.

G Fear of Misspecification

We can study agents who fear model misspecification. We follow Hansen et al. (2008) and Hansen and Sargent (2010) who order a representative household's preferences with a value function that satisfies the recursion

$$W(\boldsymbol{\zeta}_t, y_{p,t}^{cons}) = (1-\beta)\, \mathbf{y}_t^{cons,p} + \mathtt{T}_1 \left[\beta W(\boldsymbol{\zeta}_{t+1}, y_{p,t+1}^{cons}) \right] \tag{69}$$

$$\mathtt{T}_1 \left[\beta W(\boldsymbol{\zeta}_{t+1}, y_{p,t+1}^{cons}) \right] = \min_{m(\boldsymbol{\varepsilon}_{t+1}) \geq 0,\, \mathbb{E}[m(\boldsymbol{\varepsilon}_{t+1})]=1} \mathbb{E} \left[m(\boldsymbol{\varepsilon}_{t+1}) \left(\beta W(\boldsymbol{\zeta}_{t+1}, y_{p,t+1}^{cons}) + \theta_1 \log m(\boldsymbol{\varepsilon}_{t+1}) \right) \middle| \boldsymbol{\zeta}_t, y_{p,t}^{cons} \right]$$

where we can interpret $\mathtt{T}_1$ as implementing the multiplier preferences of Hansen et al. (1999). The worst-case model is given by

$$\zeta_{t+1} = \mathbf{A}\,\zeta_t + \mathbf{C}\,\varepsilon^*_{t+1} \tag{70}$$

$$y^{cons}_{p,t+1} - y^{cons}_{p,t} = \mathbf{D}^{cons}_p\,\zeta_t + \mathbf{G}^{cons}_p\,\varepsilon^*_{t+1} \tag{71}$$

where $\varepsilon^{p*}_{t+1} = \varepsilon_{t+1} + \mathbf{w}^*_p$ where $\mathbf{w}^*_p = \frac{-\beta}{\theta}(\mathbf{C}^\top\,\boldsymbol{\lambda}^{cons}_p + \mathbf{G}^{cons\top}_p)$. Hansen and Sargent (2010) show that the value function has representation

$$W(\zeta_t, y^{cons}_{p,t}) = y^{cons}_{p,t} + \boldsymbol{\lambda}^{cons\top}_p \zeta_t + \boldsymbol{\kappa}^{cons}_p \tag{72}$$

with $\boldsymbol{\kappa}^{cons}_p = \frac{-\beta^2}{2(1-\beta)\theta}||\boldsymbol{\lambda}^{cons\top}_p\,\mathbf{C} + \mathbf{G}^{cons}_p\,||^2$, and $\boldsymbol{\lambda}^{cons}_p$ as given in (67). As in appendix F, we could also allow transitions across quantiles by appropriatedly introducing a stochastic matrix $\mathbf{P}$ governing them.

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.

Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J Sargent, "Inequality, Business Cycles, and Monetary-Fiscal Policy," *Econometrica*, 2021, *89* (6), 2559–2599.

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* number burn46-1, 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.

Chang, Minsu, Xiaohong Chen, and Frank Schorfheide, "Heterogeneity and aggregate fluctuations," *Journal of Political Economy*, 2024, *132* (12), 4021–4067.

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.

Guvenan, 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.

Iao, Man Chon and Yatheesan J. Selvakumar, "Estimating HANK with Micro Data," Technical Report, New York University 2024.

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

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.

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.

_ **and Yatheesan J. Selvakumar**, "Dynamic Mode Decompositions and Vector Autoregressions," Technical Report, New York University 2024.

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