
ECONtribute
Discussion Paper No. 351

Distributional Dynamics

Christian Bayer

Luis Calderon

Moritz Kuhn

January 2025

www.econtribute.de



**UNIVERSITÄT
ZU KÖLN**

Distributional Dynamics*

Christian Bayer[†]

*University of Bonn,
CEPR, IZA, and CESifo*

Luis Calderon[‡]

University of Bonn

Moritz Kuhn[§]

*University of Mannheim
CEPR, IZA, and CESifo*

January 5, 2025

Abstract

We develop a new method for deriving high-frequency synthetic distributions of consumption, income, and wealth. Modern theories of macroeconomic dynamics identify the joint distribution of consumption, income, and wealth as a key determinant of aggregate dynamics. Our novel method allows us to study their distributional dynamics over time. The method can incorporate different microdata sources, regardless of their frequency and coverage of variables, to generate high-frequency synthetic distributional data. We extend existing methods by allowing for more flexible data inputs. The core of the method is to treat the distributional data as a time series of functions whose underlying factor structure follows a state-space model, which we estimate using Bayesian techniques. We show that the novel method provides the high-frequency distributional data needed to understand better the dynamics of consumption and its distribution over the business cycle.

Keywords: *Consumption, income, and wealth inequality; Macroeconomic dynamics; Dynamic state-space model; Functional time-series data; Bayesian statistics*

JEL Classification: *E21, E32, E37, D31, C32, C55*

*We thank Nazarii Salish for discussion at early stages of this project and Lisa Dähne for research assistance. We also thank seminar participants at the Banca d'Italia, Hausdorff Center for Mathematics, KU Leuven, National University of Singapore, Paris School of Economics, and summer SED 2024 for helpful comments. Christian Bayer gratefully acknowledges funding through the ERC-CoG project Liquid-House-Cycle funded by the European Union's Horizon 2020 Program under grant agreement No. 724204. Bayer and Kuhn gratefully acknowledge support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy (EXC 2126/1 – 390838866) and through CRC TR 224 (Projects A03 and C05).

[†]Institute for Macroeconomics and Econometrics—University of Bonn, christian.bayer@uni-bonn.de.

[‡]Institute for Macroeconomics and Econometrics—University of Bonn, luis.calderon@uni-bonn.de.

[§]Department of Economics—University of Mannheim, mokuhn@uni-mannheim.de.

1 Introduction

Understanding the dynamics of the joint distribution of consumption, income, and wealth is central to understanding macroeconomic dynamics, the transmission of monetary and fiscal policy, and their cross-sectional effects (Andersen et al., 2021; Bayer et al., 2019; Bhandari et al., 2021; Holm, Paul, & Tischbirek, 2021; Mian, Straub, & Sufi, 2020). The limited availability of high-frequency information on the joint distribution is a significant limitation in this endeavor. We propose a novel and general technique based on functional data analysis and Bayesian time series methods to obtain high-frequency estimates of this (or other) joint distribution(s). The proposed method is flexible enough to combine distributional data from different microdata sources with aggregate data, even when these data are of mixed frequency and only one micro-dataset contains all variables of interest, while others contain only a subset.

The challenge is that the joint distributions of consumption, income, and wealth are functional data and hence, in principle, infinite-dimensional objects. Our novel method, however, exploits statistical dimensionality reduction techniques by assuming that the distributional dynamics can be captured by a factor model in which the factors themselves have a state-space representation. This assumption is based on insights from the heterogeneous agent macroeconomics literature, suggesting that a few factors should be sufficient to approximate the distributional dynamics, given that a small set of aggregate prices/shocks shape the distribution of consumption, income, and wealth in the short and medium-run (Auclert, Bardóczy, Rognlie, & Straub, 2021; Bayer, Born, & Luetticke, 2024). These prices also closely track movements in the aggregate economy. In fact, the empirical evidence generated by inequality research has so far found much support for this (Chodorow-Reich, Nenov, & Simsek, 2021; Di Maggio, Kermani, & Majlesi, 2020; Kuhn, Schularick, & Steins, 2020).

This set of findings from the macroeconomic literature has three implications for the joint evolution of aggregate and distributional data: First, the dynamics of the distributional data can be represented by a medium-size state-space model. Second, the states of this model are driven by a small set of aggregate factors and unobserved distributional shocks. The combination of these two facts is the key innovation to overcome the challenge of dealing with multidimensional functional data. In practice, we use factor analysis to uncover the lower-dimensional state-space representation of the distributional dynamics and its aggregate drivers. Third, given these factor structures for the aggregate and distributional data, we estimate the time series behavior of the functional, i.e., distributional, data using Bayesian techniques and link aggregate and distributional data without imposing a structural macroeconomic model.

The state-space representation lends itself naturally to the use of the Kalman filter for Bayesian estimation of the state-space model. This has several important advantages. It al-

allows us to use and merge numerous micro-datasets that refer to the same economic variable but with different operationalized measures, e.g., differences in the sources of income covered. These different operationalizations, along with sampling uncertainty, are captured by the measurement error in the observation equation of the state-space model. Having an observation equation also allows for the combination of micro-datasets with different sampling frequencies and also allows us to exploit the information on the evolution of distributions even from microdata that contain only a subset of the variables of interest.

Finally, we overcome the limited availability of high-frequency distributional information and construct estimates of business cycle fluctuations for joint distributions at any point in time, including periods where microdata on distributions are unavailable. The estimated state-space model allows us to construct synthetic high-frequency distributional data by means of the Kalman smoother. The synthetic data itself, while originally functional data, can be expressed approximately in the form of repeated cross sections of microdata containing consumption, income, and wealth observations of synthetically constructed households. These households represent groups of granularity that the researcher can flexibly specify.

The method is general and can be applied to any kind of distributional dynamics. To demonstrate the power of the novel method, we study the dynamics of the joint distribution of consumption, income, and wealth. We apply the new estimation technique to a rich set of U.S. household microdata from the *Panel Study of Income Dynamics* (PSID), the *Survey of Consumer Finances* (SCF), the *Consumption Expenditure Survey* (CEX), the *Survey of Income and Programme Participation* (SIPP), and the *Current Population Survey* (CPS). We complement the microdata with a comprehensive set of macroeconomic time series. In using the various microdata jointly, our method overcomes three existing challenges. First, only the PSID contains all three variables of interest: consumption, income, and wealth. In the other micro-datasets, at least one of the variables is absent. At the same time, all of the micro-datasets contain information on the joint distribution of consumption, income, and wealth. Second, all of these datasets are available at different frequencies. Third, they differ in sampling approaches and details of their measurement concepts. Our method deals with all three challenges.

From the estimation of the state-space model on these data, we then construct high-frequency synthetic distributional data represented by groups of households. Each group is defined by a particular combination of quantiles of consumption, income, and wealth. Over time, the conditional expectations for each quantile changes, and so do the consumption, income, and wealth of each group. The population weight reflects how likely it is to observe combinations of quantiles, and therefore also the weight changes over time. Thus, the dynamics of the population weights induce the dynamics of the cross-sectional correlations in the three variables. We construct the detrended business cycle variations of the joint distribution in consumption,

income, and wealth from 1962 to 2021.

We carefully validate each step of the estimation procedure. First, we show that the factor representation of the distributional data imposes almost no loss of information compared to the information provided by the microdata when observed. Second, we validate the choice of priors in Bayesian estimation, particularly with respect to measurement error. We show that the state-space model is consistent with the sampling uncertainty of the observed distributional data at the sampling points. Furthermore, we show that the model closely predicts the distributional data, even when unobserved, through significant comovement with aggregates. Specifically, we show this for the consumption distributions of the CEX and the wealth distributions of the SCF. Third, we compare the prediction for the dynamics of income and wealth distribution with that implied by the distributional flow of funds methodology (DFA, see Batty, Bricker, Briggs, Friedman, Nemschoff, Nielsen, Sommer, & Volz, 2020) and that estimated by the World Inequality Database (Alvaredo, Atkinson, Chancel, Piketty, Saez, & Zucman, 2016; Piketty, Saez, & Zucman, 2018). We conclude that our method is capable of producing reliable estimates of distributional data at the business cycle frequency.

Finally, we demonstrate an application of the new method and show the dynamics of consumption along the joint distribution of income and wealth. This application highlights two directions in which the method is able to expand on existing research: First, it provides novel empirical estimates that remain unobserved in existing data. Second, it provides new empirical model targets for guiding model building and parameter choices for macroeconomic models with rich heterogeneity. The observed dynamics are also economically interesting. Relative to a representative average household, we show that consumption of the middle class hardly moves in any recession, whereas consumption of the poor and the rich show significant cyclical movement relative to the average. Comparing the Dotcom-recession, the Great Recession and the Covid Recession, we document that the recessions differ qualitatively in terms of consumption dynamics of the rich and the poor relative to the average household. The Great Recession brought consumption losses of the wealth rich, who had even gained in terms of consumption during the Dotcom Recession. The Covid recession primarily brought about consumption losses of the income-rich, more so than of the wealth-rich. Wealth exerted a mitigating effect on the consumption response and appeared to have offered a buffer for consumption dynamics during the most recent recession, especially for the income-rich.

The remainder of this paper is organized as follows: Section 2 provides an overview of the relevant literature. Section 3 develops our estimation method and Section 4 evaluates its quality. Section 5 provides an application of the novel method. Finally, Section 6 concludes the paper. An appendix follows.

2 Literature

The paper most closely related to ours is Chang, Chen, and Schorfheide (2024), which develops a Bayesian state-space approach to estimate the coevolution of aggregate variables and the marginal distribution of earnings. We differ from this paper in three important ways. First, we develop a method that is suitable for dealing with the evolution of *joint* distributions over time (distributions of consumption, income, and wealth). Second, we follow Kneip and Utikal (2001) and Tsay (2016) and Ramsay and Silverman (2005, Chapter 8) in not approximating the distribution functions by a fixed set of basis functions, but rather determine the basis functions based on a principal component analysis (see also Meeks & Monti, 2023, for a macroeconomic application of this method). Third, we focus on the production of high-frequency synthetic distribution data, dealing with missing observations and the mixed frequencies of aggregate and microdata.

The latter focus on generating new microdata relates our work to the large body of empirical literature on trends and fluctuations in inequality that took off after the seminal paper by Piketty and Saez (2003): Blanchet, Saez, and Zucman (2022) propose a methodology for producing high-frequency (monthly), timely income and wealth distribution statistics for the United States from 1976 to the present. The paper matches the CPS and SCF microdata with individual tax data collected from Piketty, Saez, and Zucman (2018) to produce a harmonized set of monthly microfiles representing synthetic adults, whose income/wealth data are consistent with national accounts and whose distribution reflects only publicly observed data. The paper emphasizes the timeliness of its data, which facilitates policymaking and public discourse on social inequality (e.g., age, race, gender, etc.) as well as income and wealth inequality. Most of this work focuses primarily on income or wealth separately, and thus often concentrates on marginal distributions, emphasizing specific moments of these distributions, such as top wealth shares. The latter is a widespread feature of the literature. For example, Smith, Zidar, and Zwick (2021), which uses administrative data and the capitalization method of Saez and Zucman (2016), also provides high-frequency estimates, but focuses more on the top of the wealth distribution, albeit at greater wealth granularity.

Similarly to what we do, Batty et al. (2020) also construct an extensive synthetic dataset of quarterly estimates since 1989 on the balance sheet of US households; they rely on the granularity of the wealth module of the SCF and aggregate information from the Financial Accounts. They use Chow-Lin/Fernández type models (see Chow & Lin, 1971; Fernandez, 1981; Litterman, 1983) in this endeavor. The advantage of our state-space approach is that it can explicitly deal with sampling uncertainty by treating the underlying microdata as samples of a time series of distribution functions. This means that we can explicitly deal with sampling uncertainty in a dynamic setting and combine different microdata sources that contain the same economic

objects with slightly different operationalizations of measurement. Moreover, by combining many microdata sources, we can go a step further and obtain the business cycle fluctuations in the joint distributions of consumption, income, and wealth going back to the 1960s.

With the estimated joint distribution, we can speak to the large macroeconomic literature that has established the importance of heterogeneity for modeling macroeconomic dynamics (Bayer, Born, & Luetticke, 2024; Bayer et al., 2019; Kaplan, Moll, & Violante, 2018) and complement it with the still missing descriptions of the short-run dynamics of the consumption, income, and wealth distribution. Our new method aims to fill this gap in the literature. Data on the joint distribution allow for an analysis of the dynamics of the joint distribution that provides important information for model building and thus extends the rapidly growing literature that examines the impact of policy shocks on the marginal distributions of consumption, income, or wealth and their aggregate feedbacks (see for example Bartscher, Schularick, Kuhn, & Wachtel, 2022; Berger, Bocola, & Dovis, 2023; Chang & Schorfheide, 2024; Cloyne, Ferreira, & Surico, 2020; Coibion, Gorodnichenko, Kueng, & Silvia, 2017; Holm, Paul, & Tischbirek, 2021). McKay and Wolf (2023) surveys the empirical literature on the effects of monetary policy on inequality. Indirectly, we also confirm the theoretical conjecture empirically that the factor structure in the time series of the distribution of consumption, income, and wealth is relatively low-dimensional, driven by a small number of aggregate factors.

The paper addresses the large methodological and global literature that focuses on estimating high-frequency time series using related series and lower-frequency measures. Common methods for estimating such balanced time series include interpolation (Denton, 1971; Friedman, 1962), regression-based methods with autocorrelated errors (Chow & Lin, 1971; Fernandez, 1981; Litterman, 1983), dynamic Chow-Lin models, and structural multivariate time series models that allow for endogeneity of related series (Di Fonzo, 2003; Gregoir, 2003; Salazar & Weale, 1999; Silva & Cardoso, 2001).

The paper takes a resurfaced approach; specifically, formulating the estimation in a state space framework (Harvey & Chung, 2000; Harvey, 1990; Harvey & Pierse, 1984; Moauro & Savio, 2005; Mönch, Uhlig, et al., 2005; Proietti, 2006) of functional data (see e.g. Chang, Kim, & Park, 2016; Diebold & Li, 2006; Inoue & Rossi, 2021; Kneip & Utikal, 2001; Otto & Salish, 2022, for economic applications). The main advantages of using a functional state-space model lie in its (1) flexibility: with appropriate manipulation, it can encompass the other models, (2) its ability to use the well-studied Kalman filter, its results and intuitive diagnostics, and (3) its dynamic nature, which is not the case with widely used models such as the Chow-Lin/Fernandez models.

3 Method

This section describes a general method for generating high-frequency estimates of joint distributions of economic variables of interest over a large number of micro-units. This method uses microeconomic and aggregate data as inputs. It requires only the joint observation of the microeconomic variables in at least one dataset over several, but potentially infrequent, time periods. The developed method treats the distributional data as functional data in a time-series state-space framework with unobserved states. In the following, we describe the method using the example of the joint distribution of consumption, income, and wealth, an important macroeconomic application.

3.1 Distributions as time series of functional data

We consider a sequence $\Xi_t(w)$ of multidimensional distribution functions, indexed by $t \in \mathbb{T} := \{1 \dots T\}$, defined over a d -dimensional *vector* $w \in \mathbb{R}^d$. In the case of our application, we have $d = 3$, where w is a vector of consumption, income, and wealth at the household level. In addition to this sequence of distribution *functions*, there is a sequence of real-valued vectors Y_t of stationary aggregate data. In the following exposition, we assume that Y_t is observed at all times $t \in \mathbb{T}$. The extension to missing observations in Y_t is standard.

From the distributions $\Xi_t(w)$, we observe only randomly drawn samples. We allow these samples to come from different sampling procedures or to have different operationalizations of the underlying theoretical variables. For example, the Panel Study of Income Dynamics (PSID) and the Survey of Consumer Finances (SCF) use different sampling procedures and slightly different concepts of wealth and income. We index each of the sampling procedures/datasets by $j = 1 \dots J$. All of these different datasets are typically not observed in all time periods. Instead, dataset j is only observed in a particular subset $\mathcal{T}_j \subset \mathbb{T}$. In addition, not all samples contain all variables of interest, but may contain only a subset $\mathcal{D}_j \subseteq \mathbb{D} := \{1 \dots d\}$ of variables. For example, the Current Population Survey (CPS) provides only income information but neither wealth nor consumption. However, at least one dataset, j , must contain all the variables of interest for $\mathcal{D}_j = \mathbb{D}$. In our application, such a dataset is the PSID, which contains information on consumption, income, and wealth (at least for some years).

Our goal is to obtain estimates of the joint distribution functions, $\hat{\Xi}_t(w)$, $\forall t \in \mathbb{T}$, by efficiently combining information from the various related microdata sources and aggregate information, Y_t . We assume that there is a time series structure such that the density $d\Xi_t$ evolves according to the functional difference equation

$$d\Xi_{t+1} = G(d\Xi_t, Y_t) + \epsilon_t, \quad (1)$$

where Y_t are observed aggregate data (including lags); G determines the dynamics of the system, and ϵ_t are the corresponding shocks to the functional equation.¹ This structure arises naturally in so-called HANK models (see e.g. Bayer, Born, & Luetticke, 2024; Kaplan, Moll, & Violante, 2018). For our purposes, it is sufficient to understand the reduced form relationship between Y_t and Ξ_{t+1} and thus treat Y_t as exogenous.²

Viewing the J sampling procedures as capturing the same fundamental object Ξ_t but with some measurement error, ν_t , allows us to combine the data in a systematic way. This means that a dataset gives us an estimate

$$d\tilde{\Xi}_t^j = \int_{\mathbb{D} \setminus \mathcal{D}_j} d\Xi_t + \nu_t \quad \text{for } t \in \mathcal{T}_j \quad (2)$$

The measurement error then captures sampling uncertainty, differences in sampling procedures, and differences in operationalization of economic concepts. The integral $\int_{\mathbb{D} \setminus \mathcal{D}_j}$ reflects that those variables unobserved in dataset j have been integrated out.

3.2 Implementing the Estimation

Estimating Equation (1) directly is not feasible because it is an infinite-dimensional nonlinear functional difference equation and, of course, we only observe samples of the distribution functions, not the functional data itself. Our innovation is to overcome this challenge by making it possible to estimate (1) using traditional Bayesian techniques and a Kalman filter (Section 3.2.4). This requires transforming (1) into a linearized (infinite-dimensional) state-space model (Section 3.2.3), which is estimable once we reduce its dimensionality by finding an appropriate factor representation (Section 3.2.2). However, first, we need to operationalize the measurement of the distribution functions as they appear in Equation (2) by transforming the microdata samples into estimates of the distribution functions themselves (Section 3.2.1). In doing so, we have to account for changes in the effective support of the distributions and deal with the unobservability of some of the micro-variables in some datasets.

3.2.1 Transforming the microdata

Handling changes in scale One challenge in working with distributional data is that the magnitude of the variables of interest in w , and thus the support of Ξ , changes over time. We deal with this in two ways. First, to deal with level changes, we rescale the vector w observed for individual i in the micro-dataset j at time t , w_{ijt} , by its dataset- and time-specific mean

¹For Ξ_{t+1} to be a distribution, we assume that $\int G(dF, \cdot)(w)dw = 1$, $\int \epsilon_t(w)dw = 0$, and $G(dF, Y_t)(w) \geq -\epsilon_t(w)$.

²This requires Y_t to be observed every period and without measurement error.

\bar{w}_{jt} .³ Second, to deal with changes in the width of the support, we decompose the distributional data into its marginals and a copula. Copulas, by definition, have a constant support (hypercubes of $[0, 1]$). By representing the marginals by their quantile functions (i.e., the inverse of the marginal cumulative distribution function), we again achieve constant support by construction. This quantile and copula representation contains the same information as Ξ , but makes the support of all functions time-invariant.

This is based on the fact that any multivariate cumulative distribution function $\Xi_t(w)$ can be written in terms of its marginal distributions $\Xi_{mt}(w)$, along the dimension $m \in \mathbb{D}$, and a copula, $C_t : [0, 1]^d \rightarrow [0, 1]$ with uniform marginals.⁴ The copula captures the dependence structure between the random variables in w and is invariant to monotone transformations in w . For our application, the copula is

$$\begin{aligned} C_t(u_1, \dots, u_d) &= P_t(U_1 \leq u_1, \dots, U_d \leq u_d) \\ &= P_t(w_1 \leq \Xi_{1t}^{-1}(u_1), \dots, w_d \leq \Xi_{dt}^{-1}(u_d)) \quad \forall t \in \mathbb{T} \end{aligned} \quad (3)$$

where $U_m \sim U[0, 1]$ for $m \in \mathbb{D}$ are the uniform marginals generated by taking the probability integral transform of each component m of w s.t. $U_m = \Xi_m(w_m) \sim U[0, 1]$.⁵ The second line highlights the quantile functions or the inverse transform of the univariate CDFs, $\Xi_{mt}^{-1}(w_m)$, where:

$$\Xi_{mt}^{-1}(u_m) = \inf\{w_m \in \mathbb{R} : \Xi_m(w_m) \geq u_m\} \quad \forall t \in \mathbb{T}, m \in \mathbb{D}. \quad (4)$$

Finally, for C_t to be a copula, the constraint must hold for all $k \in \{1 \dots d\}$ that when integrating out all but one dimension (the marginal distribution) k , the copula is identical to the value of the marginal distribution:

$$C_t(1, \dots, u_k, \dots, 1) = u_k. \quad (5)$$

For the actual estimation of the copula and the quantile functions, we will rely on a series estimator. Series estimators project the functions of interest onto some infinite-dimensional space of polynomials, namely the space of squared-integrable functions \mathcal{L}^2 . In the following, we will treat the copula densities dC_t and quantile functions Ξ_{mt}^{-1} as elements of \mathcal{L}^2 , and project them onto a space of shifted orthonormal Legendre polynomials $\{Q_m : m \in \mathbb{N}\}$ for some order m , satisfying

$$\int_{[0,1]} Q_m(x) Q_k(x) dx = \delta_{mk}, \quad (6)$$

³Estimating the model relative to the dataset-specific time means also allows us to flexibly match per capita/household aggregate targets to the synthetic data our estimation produces. This simply requires that the constructed synthetic high-frequency distribution data be scaled back not by the dataset-specific time average, but by the appropriate aggregate target. We can also produce a consensus estimate of the business cycle component across all datasets by using average fixed effects when scaling back and not adding back any trend.

⁴The theory behind the decomposition is laid down by Sklar (1959, 1973).

⁵Any distributional transform that produces a uniform cdf will do. See Rüschendorf (2009) for details.

with $\delta_{mk} = 1$ if $m = k$ (reflecting the normalization condition) and 0 otherwise by orthogonality. Classical Legendre polynomials are defined on $[-1, 1]$, so the shift is to the same domain as the copula densities and quantile functions—in the support $[0, 1]$.⁶ The orthonormal shifted Legendre polynomials Q_m themselves are defined on $[0, 1]$ by

$$Q_m(x) = \sqrt{2m+1} L_m(2x-1), \quad (7)$$

where L_m are the well-known classical Legendre polynomials defined by

$$L_0 = 1, \quad L_1(x) = x, \quad (8)$$

$$(m+1)L_{m+1}(x) = (2m+1)xL_m(x) - mL_{m-1}(x). \quad (9)$$

For the quantile function for variable m and copula density, this means the following representation:

$$\Xi_{mt}^{-1}(u_m) = \sum_{o \in \mathbb{N}} \xi_{ot}^m Q_o(u) \quad (10)$$

$$dC_t(u_1, u_2, \dots, u_d) = \sum_{(o_1, \dots, o_d) \in \mathbb{N}^d} \kappa_{(o_1, \dots, o_d), t} \prod_{m=1}^d Q_{o_m}(u_m) \quad (11)$$

where both functions are represented as infinite sums of polynomials over the order of the polynomials $o \in \mathbb{N}$. Because of orthonormality (6) and using Equation (10), we obtain:

$$\int_0^1 \Xi_{mt}^{-1} Q_o(u) du = \int_0^1 \sum_{j \in \mathbb{N}} \xi_{jt}^m Q_j(u) Q_o(u) du = \sum_{j \in \mathbb{N}} \delta_{jo} \xi_{jt}^m = \xi_{ot}^m. \quad (12)$$

This means that the coefficients can be obtained as the inner products of the functions and the Legendre polynomials of some order o . The same holds true for the copula density:

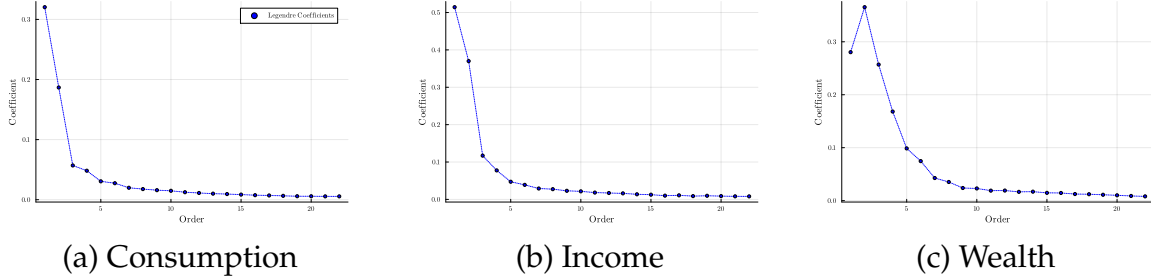
$$\int_{[0,1]^d} \prod_{m=1}^d Q_{o_m}(u_m) dC_t du_1, \dots, du_d = \kappa_{o_1, \dots, o_d, t} \quad (13)$$

For the estimation of these coefficients in practice, we rely on the uniformity of ranks and

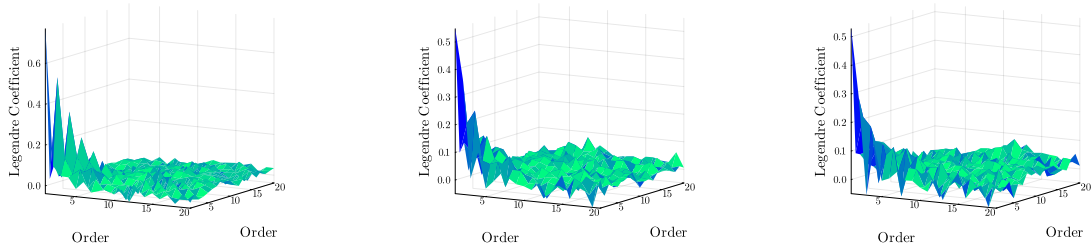
⁶Further details on our series estimator can be found in Bakam and Pommeret (2023).

Figure 1: Legendre Coefficients Across the Distributional Data

Legendre Coefficients across Quantile Functions



Legendre Coefficients across Copulas



(a) Consumption & Income (b) Consumption & Wealth (c) Income & Wealth

Notes: Figure presents two panels. The first panel presents the coefficients (in dots) on the Legendre polynomials from estimating the quantile function, in increasing order. The second panel presents the coefficients (as a surface) on the Legendre polynomials from estimating the copula density in lexicographic order. Data are from the 2019 PSID.

that ranks are within $[0, 1]$ and replace the inner product by sample averages:

$$\hat{\xi}_{ot}^m := N^{-1} \sum_i w_{mit} Q_o(u_{mit}) \quad (14)$$

$$\hat{\kappa}_{o_1, \dots, o_d, t} := N^{-1} \sum_i \left(\prod_{m=1}^d Q_{o_m}(u_{mit}) \right) \quad (15)$$

where u_{mit} are the data ranks of w_{mit} , the sample analogue of Ξ_{mt}^{-1} for observation i .

In the estimation, we truncate the sums in Equations (10) and (11) at a given maximal order \mathbb{O} and by orthonormality of the polynomials, the kept coefficients are not affected by the truncation. The coefficients in our case indeed decrease rapidly with each polynomial as we will show in Figure 1. Coefficients beyond order 10 are negligibly small.⁷

⁷A small coefficient in absolute value does imply the contribution of the corresponding polynomial to be small in an R^2 sense.

Dealing with partial unobservability of the microdata Another difficulty is that the microdata, at certain points in time or always, may not contain the entire vector w as an observable; but instead, a subset. Since w is not completely observed, this means that we cannot generate the microdata estimate of the full d -dimensional copula. However, we can still estimate copulas with the unobserved dimensions integrated out — lower-dimensional copulas.

The representation in the form of Legendre polynomials is very useful in this respect. First, we need to show that the density of the higher-dimensional copula must be equal to the lower-dimensional one when we integrate out the "missing" dimension d :

$$\int_0^1 dC(u_1, \dots, u_d) du_d \stackrel{!}{=} dC(u_1, \dots, u_{d-1}) \quad (16)$$

In this effort, we write out the integrand using Equation (11) and make use that the first (shifted) Legendre polynomial integrates to one while all others integrate to zero to obtain:

$$\begin{aligned} \int_0^1 dC(u_1, \dots, u_d) du_d &= \int_0^1 \sum_{o_1} \cdots \sum_{o_d} \kappa_{(o_1, \dots, o_d)} \prod_{m=1}^d Q_{o_m}(u_m) du_d \\ &= \sum_{o_1} \cdots \sum_{o_{d-1}} \sum_{\mathbf{o}_d} \kappa_{(o_1, \dots, o_{d-1}, \mathbf{o}_d)} \left(\prod_{m=1}^{d-1} Q_{o_m}(u_m) \right) \int_0^1 Q_{\mathbf{o}_d}(u_d) du_d \\ &= \sum_{o_1} \cdots \sum_{o_{d-1}} \kappa_{(o_1, \dots, o_{d-1}, 1)} \left(\prod_{m=1}^{d-1} Q_{o_m}(u_m) \right) \end{aligned} \quad (17)$$

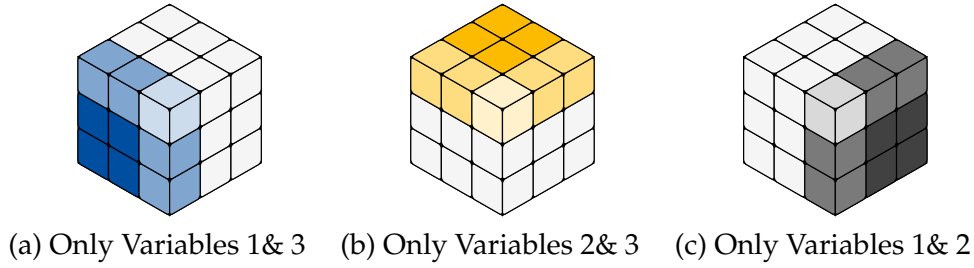
In words, the polynomial coefficients of the lower-dimensional copula density are identical to the leading "slice" of the higher-dimensional copula. This means that when a dataset does only obtain two out of the three variables of interest, we still obtain a measurement of a subset of the coefficients from this data, see Figure 2.⁸

3.2.2 Dealing with the Curse of Dimensionality

Vectorizing the coefficients for each time period, t , leaves us with sequences of coefficients, $\theta_t^j = (\xi_{o_1, t}^{j, m}, \dots, \kappa_{(o_1, \dots, o_d), t}^j)$, for each cross-sectional dataset, j . For example, with $d = 3$ dimensions in our application—consumption, income, and wealth and using polynomials up to order ten ($O = 10$) to organize the data—the copula density would be represented by a vector with $O^d - d \times (O - 1) + 1 = 972$ variable entries for each time point and $d \times (O - 1) + 1 = 28$ invariable. These 28 invariable entries are redundant due to the constraints imposed by C being a copula. In addition, there would be $d \times O = 30$ coefficients of the polynomials representing the quantile functions, which we collect in θ_t^j as well.

⁸By the same line of argument, a copula requires $\kappa_{(1, \dots, 1)} = 1$ and $\kappa_{(1, \dots, j, \dots, 1)} = 0$ (i.e., only a single order is not one).

Figure 2: Geometric Representation of Partially Observed Copula Density Coefficients



Notes: Figure shows three cubes. A cube can be interpreted as an array of copula coefficients $\kappa_{(o_1, \dots, o_d), t}^j$ for some dataset j at time t . Each cube corresponds to a scenario where one variable is missing in the estimation of the copula density. The light edge denotes the (1,1,1) coordinate. For each scenario, the white boxes are coefficients we cannot estimate. The slightly colored boxes correspond to the immutable coefficients, which have fixed values independent of data. The darker colored boxes are scenario specific and correspond to (time-varying) coefficients that need to be estimated.

From this example, it is clear that the dimensionality of $\theta_t^j \in \mathbb{R}^N$, $N = O^d + (d - 1)$ for a dataset with d variables and polynomial order O is too large to formulate and estimate a time series model directly in terms of θ_t^j itself. For this purpose, we postulate (and then estimate) a dynamic factor model for θ . This is where another advantage of the polynomial representation comes in handy: The variance (over time) of a coefficient is proportional to its contribution to fluctuations of the function (in the L^2 sense). Put simply: The polynomial coefficient provides a useful form of standardization that provides a natural metric and allows us to uncover the factor structure behind the time series changes in the distributions. This factor structure finally allows us to overcome the curse of dimensionality in the distributional data.

For this purpose, all (free) coefficients of the polynomial representation of dC_t^j (and separately of the quantiles) are horizontally concatenated:

$$\theta^j = \begin{bmatrix} \theta_{1,1}^j & & \theta_{1,T}^j \\ & \ddots & \\ \theta_{N,1}^j & & \theta_{N,T}^j \end{bmatrix}$$

and perform principal component analysis (PCA) (a singular value decomposition), which nonparametrically reduces the dimensionality of the data (Breitung & Eickmeier, 2006).⁹

Before performing this model reduction, we detrend and standardize the distribution data θ^j separately by data source j and distribution objects $o \in \{1, \dots, d + 1\}$ (d quantile functions and a copula) and obtain standardized measures $\tilde{\theta}^j$.¹⁰ This takes care of dataset-specific ef-

⁹Performing principal component analysis on the polynomial coefficient domain or the observed data is equivalent to standardizing the data (Chen, Er, & Wu, 2005).

¹⁰ The normalization of the coefficient n in the object $o(n)$ in the dataset j is given by $\tilde{\theta}_{nt}^j = \left(\frac{\theta_{nt}^j - \mu_n^j}{\sigma_{o(n)}^j} \right)$ where μ_n^j are the specific means and $\sigma_{o(n)}^j$ are the standard deviations of all coefficients of the object $o(n)$ (copula, quantile functions) to which the coefficient n refers. This removes dataset- and coefficient-specific fixed effects, μ_n^j . This can

fects. We store the information needed to transform $\tilde{\theta}^j$ back to the originally observed objects to obtain source-specific predictions. For example, the income quantiles (that would be one object o) in the SCF and the PSID (two sources j) may be permanently different due to differences in sample design and operationalization. We standardize by object (and not by coefficient) because the polynomial coefficients represent standardized contributions to the object's structure (e.g., the quantile function or copula). As a result, additional standardization within the same object is unnecessary, since their relative magnitudes are naturally balanced by the properties of the polynomial basis. This leaves us with the standardized observation $\tilde{\theta}_{nt}^j$ of coefficient n in data source j at time t .

Let $\tilde{\theta}$ denote the horizontal concatenation of the $\tilde{\theta}^j$. The PCA of $\tilde{\theta}$ provides us with a projection matrix $\Gamma \in \mathbb{R}^{N \times R}$, with full column rank R , that projects $R \ll N$ factors into the $N = O^d + (d-1)$ dimensional distributional data. More specifically, we decompose $\tilde{\theta}$ into latent orthogonal factors $\begin{bmatrix} F & f \end{bmatrix}'$ (ordered by importance) and their time constant loadings $\begin{bmatrix} \Gamma & \gamma \end{bmatrix}$. This decomposition is unique up to the scale of each factor, which allows us to normalize the loadings so that all factors have unit variance. The factors obtained are then divided into "important" and "unimportant" factors according to their contribution to the total variance (measured by their singular value):

$$\tilde{\theta} = \begin{bmatrix} \Gamma & \gamma \end{bmatrix} \begin{bmatrix} F & f \end{bmatrix} \quad (18)$$

where F represents the R important factors, which capture almost all of the variation in the data, and f the $N - R$ less important factors, which can be interpreted as some measurement noise. This step, in a sense, identifies an ideal functional basis (see Kneip & Utikal, 2001; Tsay, 2016) (the columns of Γ) to approximate the changes in the distribution over time and reduces the dimensionality of the data entering the state-space model.

In practice, some coefficients in $\tilde{\theta}^j$ may be impossible to construct in a data source j because that data source does not contain information on the corresponding variable. Consequently, these coefficients remain unobserved. For example, the SCF does not contain information on consumption. In Appendix A, we discuss alternative estimators to handle such unobservability. We implement these estimators and follow the statistical literature in choosing the estimator with the highest marginal data density.¹¹ In Appendix B, we describe the PCA on the

capture, e.g., permanent differences in sampling procedure and operationalization.

¹¹An EM algorithm (see e.g., Bańbura & Modugno, 2014; Barigozzi & Luciani, 2019; Doz, Giannone, & Reichlin, 2012) would be another alternative to our two-step approach to first estimate Γ outside the model and then, estimate the actual state-space model in Bayesian fashion. The high dimensionality of our setting, where N is very large and the setup with multiple measurements (i.e., SCF, PSID, etc.) for, in principle, the same object presents a challenge for an EM algorithm, in that the algorithm does not converge well. Furthermore, it has been shown to have a sensitivity to initialization, which is further exacerbated by the high dimensionality (McLachlan & Krishnan, 2008; Wang et al., 2015; Wu, 1983).

aggregate data to further reduce the dimensionality of the controls. The important aggregate factors retained are denoted by Y .

3.2.3 Factor State Space Model and Measurement

With this preprocessing of the data, we can turn to estimating the state-space model that captures the evolution of the distributional factors. Specifically, we postulate the following state space model

$$F_t = AF_{t-1} + BY_t + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \Omega). \quad (19)$$

Since the factors are orthogonal by construction, we restrict the innovations of the dynamic model ϵ_t to be uncorrelated, both serially and between factors, so that Ω is a diagonal matrix with diagonal entries $\Omega_{11}, \dots, \Omega_{RR}$. The loading matrix B on the aggregate controls Y_t , as well as the law of motion matrix A , are not constrained.

Since the factors are not directly observable, we complement the factor model with an observation equation for each dataset j :

$$\tilde{\theta}_t^j = H_t^j(\Gamma F_t + v_t^j), \quad v_t^j \sim \mathcal{N}(0, \Sigma_j^{1/2} S^j \Sigma_j^{1/2'}) \quad (20)$$

where S^j is a diagonal scaling matrix with the n -th diagonal entry $s_{o(n)}^j$ and Σ_j is a positive semidefinite (covariance) matrix. Stacking the datasets j then yields the complete observation equation.

The observation equation (20) translates the factors into observations of the distribution $\tilde{\theta}_t^j$ for dataset j via our estimated projection matrix Γ and the selector matrix H_t^j , which indicates whether (parts of) the distribution are observed in data source j at time t . This logical matrix $H_t^j \in \{0, 1\}^{N \times N}$ indicates whether a given variable is observed in a given dataset (following Durbin & Koopman, 2012).

The measurement error, v_t^j , is composed of sampling uncertainty and other errors that reflect the fact that a given dataset has its specific operationalizations of the common economic variables being measured. Differences in operationalizations can not only shift the level of a particular measurement (which we capture through fixed effects, see Footnote 10), but can also become differentially important over time.¹² In principle, we would need to estimate a

¹²For example, the PSID and SCF differ in the way they ask respondents about their business wealth (Pfeffer, Schoeni, Kennickell, & Andreski, 2016). PSID and CPS differ in the sampling unit, making household/family income sensitive to labor supply patterns (Gouskova, Andreski, & Schoeni, 2010). Similarly, differences in the propensity to sample business owners between the different datasets make income sensitive to relative changes in business and labor income (Kim & Stafford, 2000). Finally, the CEX and the PSID differ in the consumption categories covered in the survey, with the PSID being much coarser (Insolera, Simmert, & Johnson, 2021)

measure error variance for each coefficient and dataset; however, this would be too large a number of parameters to estimate within the time series model. For this reason, we do two things. First, we assume that the correlation structure of all measurement errors is the same as the correlation structure for sampling uncertainty. Under this assumption, the matrix Σ_j can be estimated outside the time series model using bootstraps or the supplied replication weights to estimate the covariance from sampling uncertainty by data source j .¹³

Second, we estimate the N elements of the diagonal matrix S^j within the time series model with the restriction that its entries vary only by dataset j and object $o(n)$. The factors s_o^j in matrix S scale the standard deviation of the measurement error in Equation (20) compared to the sampling uncertainty for each object o (quantile functions and copula) in each dataset j .

Since we have an external estimate $\hat{\Sigma}_j$, we rewrite the observation equations as

$$\begin{aligned}\hat{\Sigma}_j^{-1/2}\tilde{\theta}_t^j &= \hat{\Sigma}_j^{-1/2}H_t^j\hat{\Sigma}_j^{1/2}(\hat{\Sigma}_j^{-1/2}\Gamma F_t + \tilde{v}_t^j) \\ \hat{\Sigma}_j^{-1/2}\tilde{\theta}_t^j &= \hat{\Sigma}_j^{-1/2}H_t^j\hat{\Sigma}_j^{1/2}(\hat{\Sigma}_j^{-1/2}\Gamma F_t + \tilde{v}_t^j) \quad \tilde{v}_t^j \sim \mathcal{N}(0, S^j) \\ \tilde{\theta}_t^j &= \tilde{H}_t^j(\tilde{\Gamma}_j F_t + \tilde{v}_t^j), \quad \tilde{H}_t^j := \hat{\Sigma}_j^{-1/2}H_t^j\hat{\Sigma}_j^{1/2}, \quad \tilde{\Gamma}_j := \hat{\Sigma}_j^{-1/2}\Gamma.\end{aligned}\tag{21}$$

Equations (19) and (21) form a linear system of equations that can be estimated using standard Bayesian techniques and the Kalman filter, which we explain next.

3.2.4 Bayesian Estimation

We need to estimate the (vector) autocorrelation A of the factors, the loading matrix B on the aggregate controls, the variance-covariance matrix of the shocks to the factors Ω , and the variance-covariance matrix $\Sigma_j^{1/2}S^j\Sigma_j^{1/2'}$ of the measurement errors. The covariance structure $\Sigma^{1/2}$ is estimated outside the time series model, as noted above, while the scaling matrices S^j are estimated within the model. Given the size of the A, B, Ω , and S matrices, we use a Bayesian approach to estimate the system. We do this by shrinking all entries of B and the non-diagonal entries of A to zero if they are not needed to explain the data. For the diagonal entries of S , we apply the restrictions described in the last subsection.

The estimation is then one of a standard Bayesian VAR with mixed frequency data. We collect all parameters in the parameter vector ψ and formulate prior probability $p_{\text{prior}}(\psi)$. For

¹³To do this, we draw bootstrap samples (or equivalently use the supplied replication weights) for each dataset j , $\{\tilde{\theta}_{t,b}^j\}_{b=1}^B$, for each period t . Then, we demean the bootstrap samples b for each j and t and compute the average within-time variance-covariance matrix $\hat{\Sigma}_j$ pooling the demeaned bootstrap samples of the dataset j . If an object o is unobserved in dataset j , we set the covariance terms to zero and the diagonal elements to one to still be able to compute $\Sigma_j^{-1/2}$. In our application, for example, this means that in the PSID, where we observe consumption, income, and wealth, we estimate a full $(N \times N)$ variance covariance matrix Σ_{PSID} . In the CEX, where we only observe consumption and income, we bootstrap the variance covariance matrix for the objects related to these two variables. We set the off-diagonal entries for wealth-related objects to zero and the diagonal elements to one.

the state equation, we use a Minnesota prior on matrices A and B :

$$\begin{pmatrix} \text{vec}(A) \\ \text{vec}(B) \end{pmatrix} \sim \mathcal{MN}(\mu_{Minn}, V_{Minn}). \quad (22)$$

We specify the parameters of the prior distributions, the hyperparameters, as follows: We set the vector of expected values μ_{Minn} so that all but the autocorrelation terms in A have an expected value of zero. The expected values for the autocorrelations (main diagonal of A) are governed by a single hyperparameter κ_3 , which is optimized over using the methodology of Giannone, Lenza, and Primiceri (2015). The hyperprior on κ_3 reflects the quarterly nature of our data and the typically high persistence in aggregate economic time series, but also maintains the stationarity of the model. The choice of the variance-covariance hyperparameters in V_{Minn} and further details on their hyperpriors are discussed in detail in the Appendix C using a variant of the original Minnesota prior (Doan, Litterman, & Sims, 1984; Litterman, 1980).

For Ω , the variances of shocks to the factors, we specify the prior for each of the diagonal elements as a normal distribution with mean μ_Ω , such that the a priori long-run variance of each factor is $1.0 = \frac{\mu_\Omega}{1-\kappa_3^2}$, consistent with our prior for autocorrelation (in the matrix A) and factor normalization to unit variance; see Equation (19). Further information on the hyperparameters governing Ω is relegated to the Appendix C.

Turning to the state equation, for S , we assume that $\log(\exp(s_o^j) - 1)$ follows a normal distribution. We set the mode of this distribution to $\log(\exp(1) - 1) \approx 0.54$, which would imply only sampling uncertainty and no additional measurement error reflecting conceptual differences. With this prior, three-fourths of the distribution falls between 0.0 and 2.0, with values below 1.0 allowing for the possibility that our estimate of $\hat{\Sigma}_j$ are too large, which is important to allow for, since they are estimates themselves.¹⁴

Likelihood and sampling With this prior on ψ , we obtain the model-likelihood $p(\tilde{\theta}|\psi)$ using a Kalman filter. The posterior log-likelihood is then calculated as the sum of the prior log-probability and the data log-likelihood. To sample from the potentially complex, multi-modal, high-dimensional posterior distribution, we employ the Differential-Independence Mixture Ensemble (DIME) sampler from Boehl (2024). Details and convergence results are in Appendix D.

¹⁴Identical priors across datasets mean that we do not a priori prioritize a conceptual measure for object o in one dataset over another. Setting different priors is generally possible if a particular measurement concept should be prioritized on theoretical grounds.

3.3 Estimating the High-Frequency Fluctuations in the Distributional Data

Given the estimated parameters (the posterior mean), we use the Kalman smoother to estimate the sequence of unobserved factors \hat{F}_t . With these generated factors \hat{F}_t , we obtain a consensus estimate of the standardized polynomial coefficients of the functional distribution data $\hat{\theta}_t$ by premultiplying the projection matrix Γ . This gives us, for each data source j , a predicted high-frequency sequence of quantile functions, $\hat{\Xi}_{jmt}^{-1}$, and copula densities, $d\hat{C}_{jt}$.

Together, they describe the sequence of joint distributions, $\hat{\Xi}_{jt}$, as functional data. In our application, we approximate all functions by polynomials of up to order ten. With the estimated sequences of coefficients at hand, we then generate arbitrary groups of households formed by a range of ranks and obtain their weight by integrating over the copula densities. Similarly, we then obtain average realizations of variables for these groups by integrating over the quantile functions (i.e., by forming conditional expectations). This implementation implies that, without the need for sampling, we obtain an output that can be immediately interpreted as synthetic microdata. For each cell given by a combination of a consumption quantile, an income quantile, and a wealth quantile, we interpret the vector

$$X_{it} = \begin{pmatrix} c_{it} \\ y_{it} \\ w_{it} \\ \omega_{it} \end{pmatrix} = \begin{pmatrix} \int_{u \in U_i^c} \hat{\Xi}_{jct}^{-1}(u) du \\ \int_{u \in U_i^y} \hat{\Xi}_{jyt}^{-1}(u) du \\ \int_{u \in U_i^w} \hat{\Xi}_{jw}^{-1}(u) du \\ \iiint_{(u_c, u_y, u_w) \in U_i^c \times U_i^y \times U_i^w} d\hat{C}_{jt}(u^c, u^y, u^w) \end{pmatrix} \quad (23)$$

as data for a synthetic (representative) household i , where $(U_i^c \times U_i^y \times U_i^w)$ is the quantile combination that defines household i , e.g. the first decile in consumption c , the third decile in income y , and the seventh decile in wealth w . The mass, ω , of the households in that cell defines a weight for that synthetic household.¹⁵ We obtain a consensus estimate across datasets by simply averaging the $d\hat{C}_{jt}$ and $\hat{\Xi}_{jmt}^{-1}$ over datasets j .¹⁶

4 Application

We apply our method to estimate the joint distribution of consumption, income, and wealth at the household level for the United States (U.S.) from 1962 to 2024. Commonly used microdata for the U.S. are the *Consumer Expenditure Survey* (CEX), *Current Population Survey* (CPS), *Panel Study of Income Dynamics* (PSID), *Survey of Consumer Finances* (SCF); including the histor-

¹⁵The integrals can be calculated very efficiently as time varying linear combinations of the time invariant integrals of the basis functions.

¹⁶Alternatively, given the estimated measurement error variance for each data source and object type, one could use the inverse measurement error standard deviations as weights in averaging across datasets.

Table 1: Micro Data Sources and their Sample Periods

Object	CEX	CPS	SCF	PSID	SIPP
Consumption quantiles	1984Q4 - 2021Q4	-	-	1999Q2-2021Q2	-
Income quantiles	1984Q4 - 2021Q4	1967Q4-2022Q4	1962Q3-2022Q3	1968Q2-2021Q2	1983Q3-2022Q4
Wealth quantiles	-	-	1962Q3-2022Q3	1983Q2-2021Q2	1983Q3-2022Q4
Copula densities	1984Q4 - 2021Q4	-	1962Q3-2022Q3	1983Q2-2021Q2	1983Q3-2022Q4

Notes: The table reports the sample periods we use for the different micro datasets across the different objects.

ical backfiles (SCF+), and the *Survey of Income and Programme Participation* (SIPP). These will be the ones we use for the application.¹⁷

We abstain from any sample selection in all of these datasets. For the CEX, we pool all data for a given year to remove seasonality. We date the CEX and CPS to quarter 4; the PSID is assumed to reflect quarter 2; the SCF is dated to quarter 3, and the SIPP data are aggregated to quarterly level and then naturally assigned to the respective quarter.¹⁸ Table 1 lists the distributional objects from each dataset, together with the sampling period. Note again that we require at least one dataset j that includes all objects, which in our case is the PSID between 1999 to 2021.

In terms of aggregate data, we use a wide range of standard business cycle data (GDP, consumption, employment, etc.) as well as data on household balance sheets, expectations, asset prices, and interest rates from McCracken and Ng (2021). We include data from 1962Q3 to 2024Q1. The starting point of the aggregate data determines the earliest date for the sample periods of the microdata used. From these time series, we extract the 21 most important factors. Details on factor selection can be found in the Appendix B and Appendix E. Again, we choose the number of aggregate factors based on marginal data densities. The estimated parameter vector of the state space model in the posterior mean can be found in Appendix F.

4.1 Reliability Analysis

In a first pass, we check the reliability of our estimation procedure at each step. First, we show that no relevant information is lost by using the factor model for the distribution (Sections 3.2.1 and 3.2.2). Second, we show that our state-space model typically implies estimates for the distributions that are within the confidence bounds of the microdata with approximately the probability corresponding to the confidence bounds. In other words, the recon-

¹⁷The methodology can handle permanent and time-varying differences in operationalizations across concepts, however, notable structural breaks in the CPS (1992) and SIPP (2013) were dealt with conservatively by treating the subsamples (before and after the break) as separate measurements of the distribution. This potentially avoids the introduction of artificial dynamics.

¹⁸Documentation on the timing of the CPS can be found [here](#), pg. 6. For the timing of the SCF, see [here](#), pg. 33. For the PSID, interview dates are provided as variables, with the vast majority of interviews falling in quarter 2.

structed time series support our choice of priors for the measurement errors (Section 3.2.4). Third, we rerun the estimation omitting some of the microdata samples at selected points in time. We then show that the reconstructed, synthetic microdata agree well with the omitted microdata. In other words, we show that the estimated state-space model is informative about the time-series fluctuations in the distributional data (i.e., this verifies the steps in Sections 3.2.3 and 3.2.4). Finally, we also show that our reconstructed distributional data agree well with the cyclical fluctuations in the distribution of wealth in the World Inequality Database (WID) and the Distributional Financial Accounts (DFAs). Note that all trends will be fitted by construction, see Section 3.2.1.

4.1.1 Precision of the Factor Model

The first step in our procedure is to estimate Legendre polynomial coefficients for the functional representation of the distribution. Then, we estimate the factor structure in these data. Since we only retain “important” factors, we potentially introduce an approximation error resulting from forcing “unimportant” factors f_t to take time-averaged values. The size of the approximation error can be controlled by choosing how many factors to keep. We choose to retain the eight most important factors, which explain 99% of the (business cycle frequency) variation of the distribution (i.e., of $\tilde{\theta}$ to be precise).

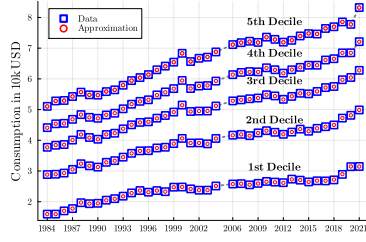
The different panels in Figure 3 visualize the approximation error in our application by showing the implied deciles. The figure compares the observed conditional decile means for consumption, income, and wealth (squares) with their approximated counterparts (circles).

We find that the factor model with its eight main factors is very close to the distributional dynamics over time: The circles are typically entered around the midpoint of the squares. Figure 4 compares the copula over time between the approximation and the raw data. We do this in terms of the Kullback-Leibler divergence. The dashed black line shows how distant the actual distribution is from its long-term average (how much variation is there to capture), and the solid line shows the difference between the actual distribution and the approximation based on the important factors only (how much the factors do not capture). The Kullback-Leibler divergence of the actual copula from its long-run average is between 0.075 and 0.12 (between 1999 and 2019),¹⁹ while the divergence between the approximation and the actual distribution is almost two orders of magnitude smaller. To put this simple: There are significant fluctuations in the copulas over time, but the factors are able to capture them well.

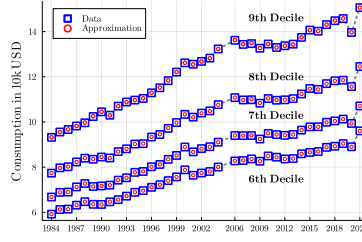
¹⁹The Kullback-Leibler divergence for the pandemic year 2021 is even 0.2.

Figure 3:
Comparison of Quantile Functions in Raw and Approximated Data (Important Factors)

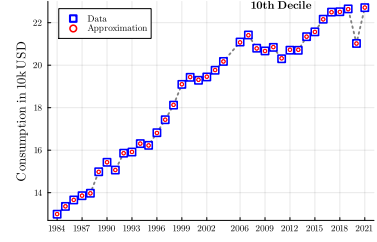
Mean Consumption



(a) 1st to 5th Decile

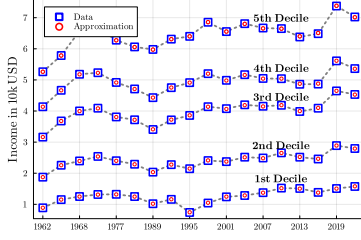


(b) 6th to 9th Decile

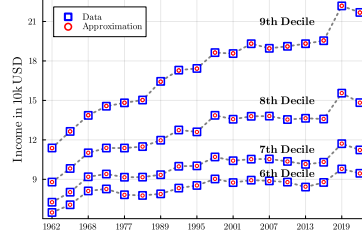


(c) Top Decile

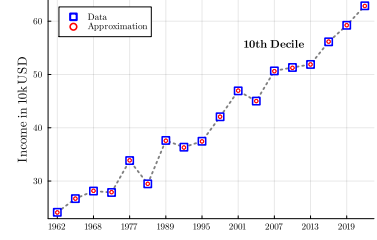
Mean Income



(d) 1st to 5th Decile

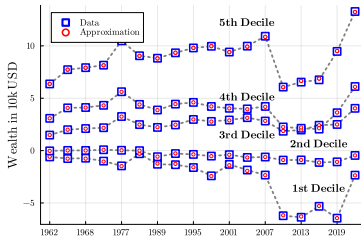


(e) 6th to 9th Decile

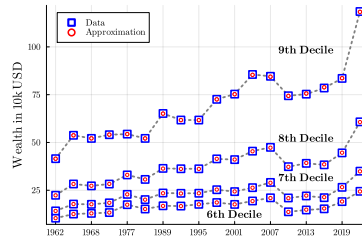


(f) Top Decile

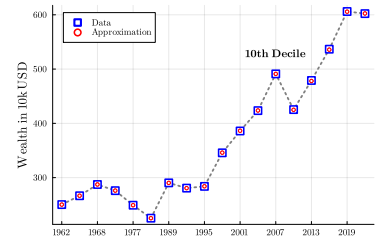
Mean Wealth



(g) 1st to 5th Decile



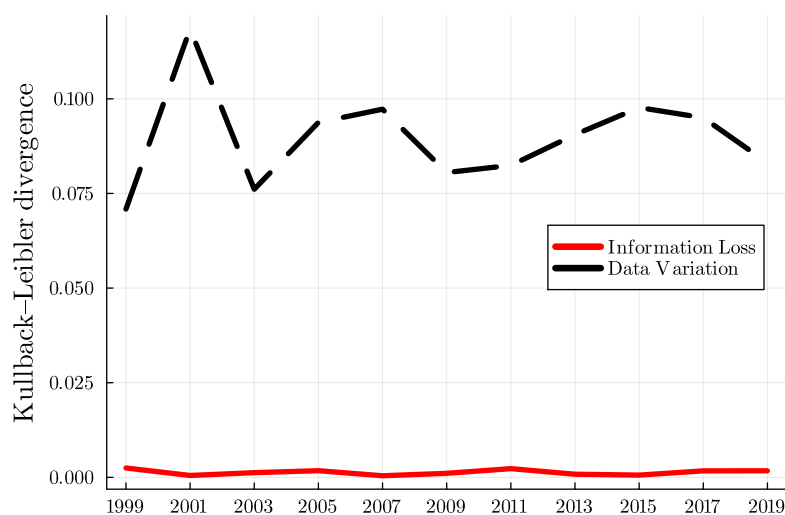
(h) 6th to 9th Decile



(i) Top Decile

Notes: Figure shows the quantile functions (mean within decile) for consumption, income, and wealth deciles from the survey data (squares) and approximation (dots) using only the fluctuations in the most important factors in (18). Top row shows quantile functions for CEX consumption. Middle and bottom row show quantile functions for SCF income and wealth. Dotted lines show linear interpolation between survey waves.

Figure 4: Comparison of the Raw data and Approximated copula (Important Factors)



Notes: Figure shows the Kullback-Leibler divergence for two copulas relative to the raw data copula. The black dashed line represents the divergence between the time-averaged copula of the PSID and the raw data copula for each survey year. The red solid line represents the divergence between the copula obtained by allowing only the most important factors in Equation (18) to fluctuate and the raw data copula for each survey year.

4.1.2 Validation of Hyperparameter Choices

To evaluate our hyperparameter choices, the Bayesian estimation priors for the measurement error variances, we compare the series resulting from the Kalman smoother after estimation with the actual point estimates and their confidence bounds from the survey data.

Intuitively, if the prior mean for the measurement error variance is too low, it will force the estimator to exactly match each survey estimate of the distribution, despite the fact that each survey estimate is itself subject to measurement error. Thus, we should expect the smoother estimate to fall within the confidence bounds of each sample estimate at most with the corresponding confidence level of the bounds. The fact that the confidence level is an upper bound reflects that the estimated measurement error captures not only the sampling uncertainty that the confidence bounds capture, but also conceptual differences.

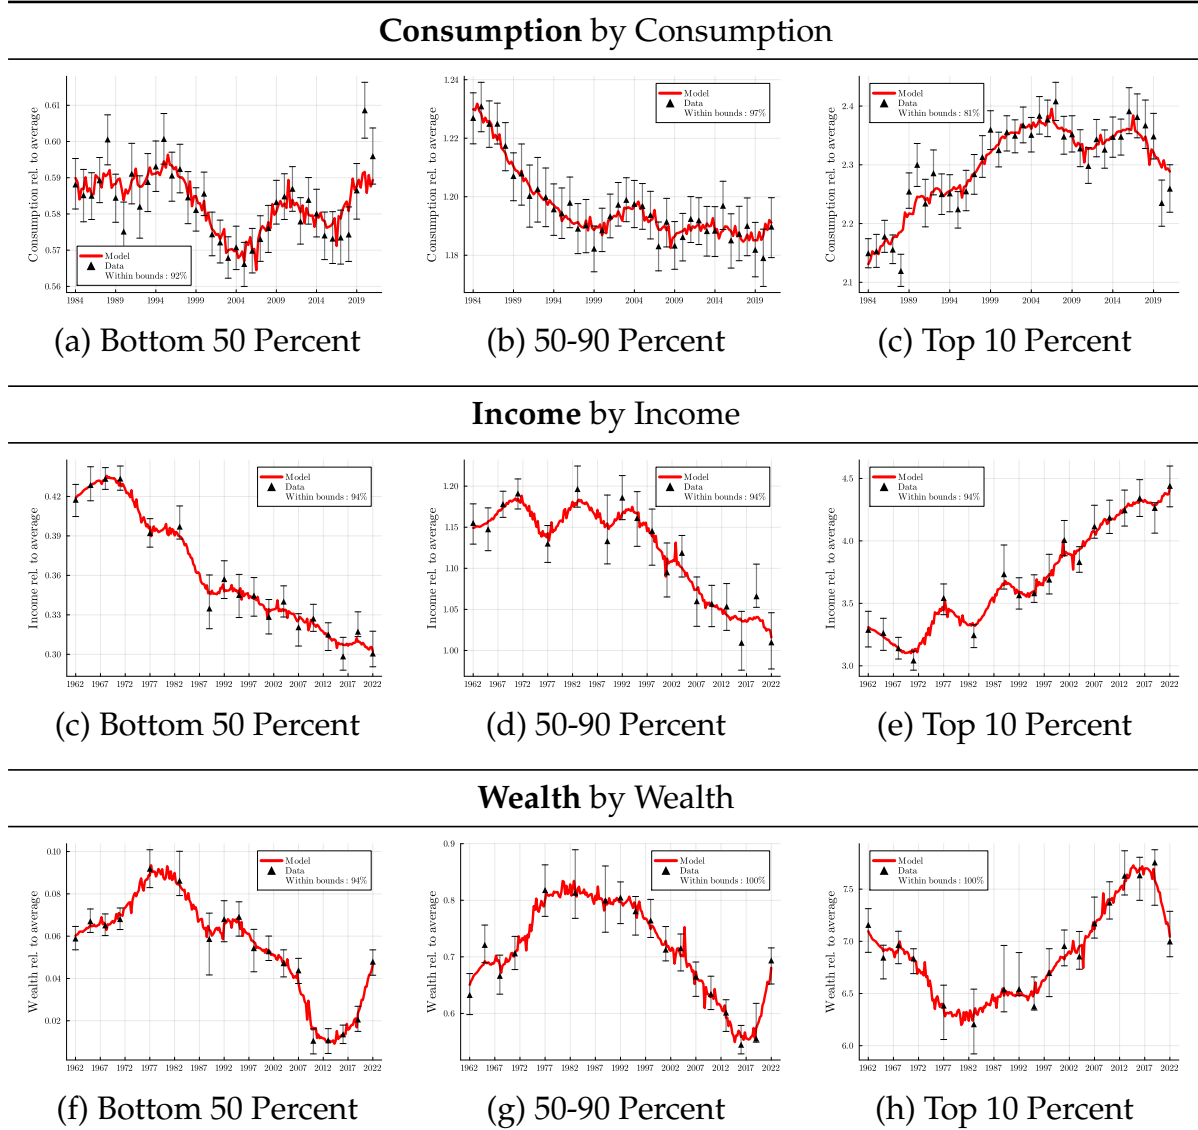
Choosing narrow measurement errors would overstate precision and potentially limit comovement with aggregates. As a result, it would drive parameter estimates for B , which captures this comovement, toward zero. Another reason not to be conservative with the measurement errors is that allowing for measurement error also accounts for the fact that we combine data from different sources to produce a consensus estimate. These different data sources, despite their individual detrending, may produce some temporarily divergent estimates of the distributions. Without sufficient measurement error, the consensus estimate is then forced to oscillate between these different distribution estimates over short time intervals, rather than capturing their co-movements.

On the other hand, if the prior mean for the measurement error variance is too high, the estimator will treat the data as uninformative, and the smoother will miss the survey estimates more often and to a much greater extent than implied by its confidence bounds. We validate the choice of hyperparameters graphically for income and wealth in the SCF data and provide comprehensive summary statistics across all datasets and estimates.

Figure 5 shows average consumption (top row), income (middle row) and wealth (bottom row) for the bottom 50 percent (first column), the next 40 percent (second column), and the top 10 percent (last column) of the respective distributions. It shows the point estimates from the surveys (black triangles), along with their 95% confidence limits, and the results from the Kalman smoother based on our estimates of the parameters of Equation(19). Overall, the smoothed estimates fall outside their respective confidence bounds in only four out of the 108 observations ((c) to (h)). This is a probability slightly smaller than the confidence level.

Table 2 provides a comprehensive summary of this validation approach. For all quantile functions and copulas, we report for each dataset and survey year how often the respective smoother estimate is within the confidence limits. Again, we use a confidence level of 95%. For the quantile functions, we find overall a modest difference (one to three percent) between

Figure 5: Comparison of Smoothed Distributional Data and Direct Survey Estimates



Notes: Figure shows the average consumption, income, and wealth for the bottom 50 percent, 50-90 percent, and top 10 percent of households of the respective distribution. Dots show the estimates from the individual survey waves together with 95% confidence bounds. The solid red line shows the baseline estimate from the Kalman smoother at the posterior mode. Consumption shows CEX data and reconstruction. Income and wealth show SCF data and reconstruction. The legend reports for each panel the share of smoothed estimates within the confidence bounds of the survey waves.

Table 2: Deviations of Smoothed Estimates and Microdata: Fraction within Confidence Bounds

Measure	CEX	CPS	SCF	SIPP	PSID	Overall
Consumption quantiles	87%	—%	—%	—%	100%	90%
Income quantiles	92%	75%	93%	52%	98%	70%
Wealth quantiles	—%	—%	92%	55%	100%	67%
Copula densities	100%	—%	97%	94%	97%	97%

Notes: The table reports, by microdata and object, the fraction of estimates from the Kalman smoother at the posterior mode that fall within the 95% bootstrapped confidence intervals for the respective microdata. Quantile and copula estimates are defined on a decile grid.

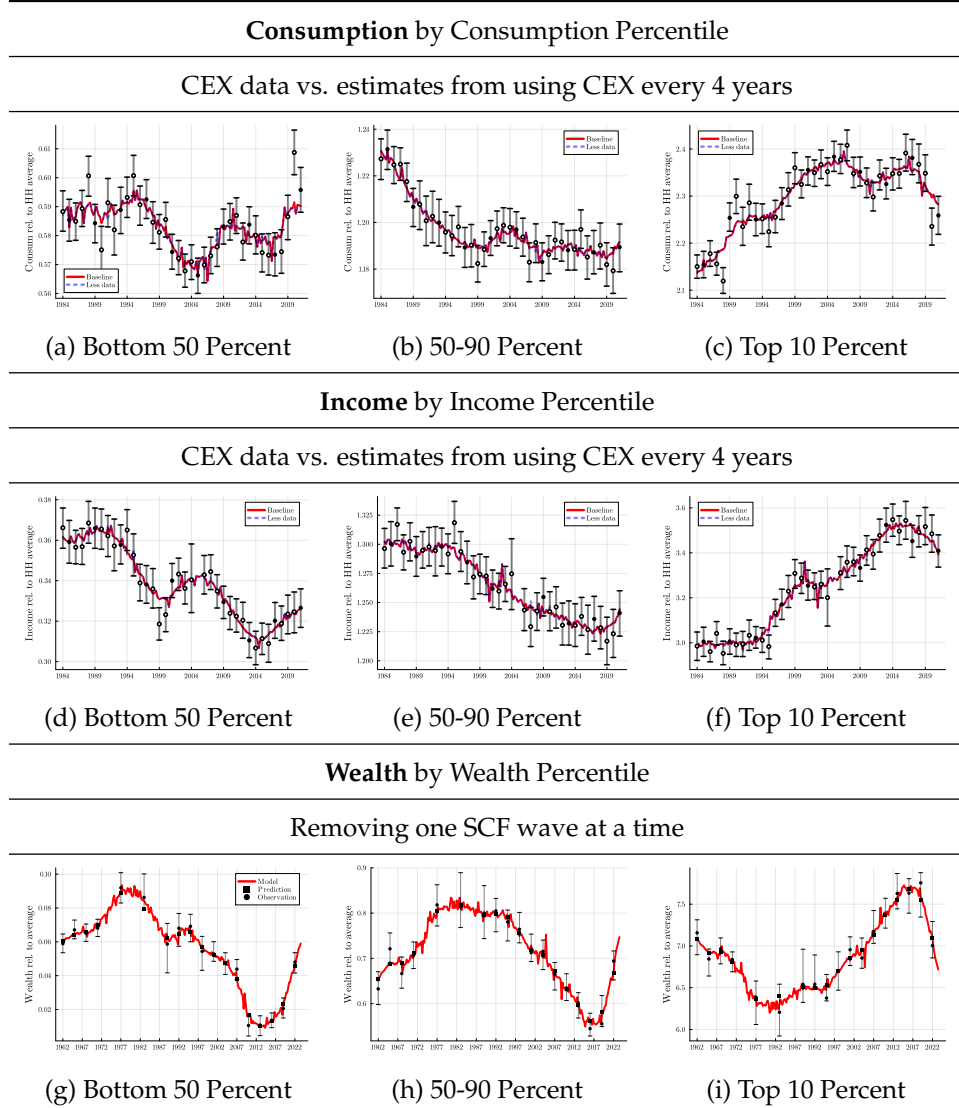
the confidence level and the fraction of smoothed estimates that fall outside the confidence bounds. Only for the SIPP and to some extent the CPS, our estimator suggests a significant measurement error beyond sampling uncertainty reflecting differences in sample design and income/wealth measurement.

4.1.3 Predictability of Distributional Data

To validate how well the method can predict distributional dynamics, we compare how well the model predicts unused microdata. We run two conceptually different, but related, experiments. The first experiment fixes the model parameters to that of the baseline and removes some waves of microdata as inputs when running the Kalman smoother. This experiment answers the question of how difficult it is to predict the microdata. The second experiment re-estimates the model entirely with the scarcer data. This experiment takes into account that changes in the data also imply changes in the model parameter estimates. We show the results of these experiments in Figures 6 and 7.

For the first experiment (Figure 6), we first include only every fourth CEX survey year in the estimation, reducing the number of CEX survey years included in the Kalman smoother from 38 to 10. We consider both consumption (Panels (a) - (c)) and income (Panels (d) - (f)), showing the average consumption/income of the top 10%, next 40% and bottom 50% (in terms of consumption/income). There is a large sampling uncertainty around the CEX data, whose 95% confidence intervals are displayed as error bars. For this reason, even in the data-rich specification (with all annual CEX data), the smoothed estimate regularly deviates from the raw distributional data, with correlations of the two around 95%. The correlation of the smoothed data using only every fourth survey year with the data using every survey year is very high, meaning that the model can predict the consumption distribution well. Given that we have less wealth data, we run 17 experiments dropping just one SCF wave at a time. This gives us 17 smoothed distributional data series alongside the one using all SCF waves. The bottom row

Figure 6: Predictability of Distributional Data (given the Baseline Parameters)



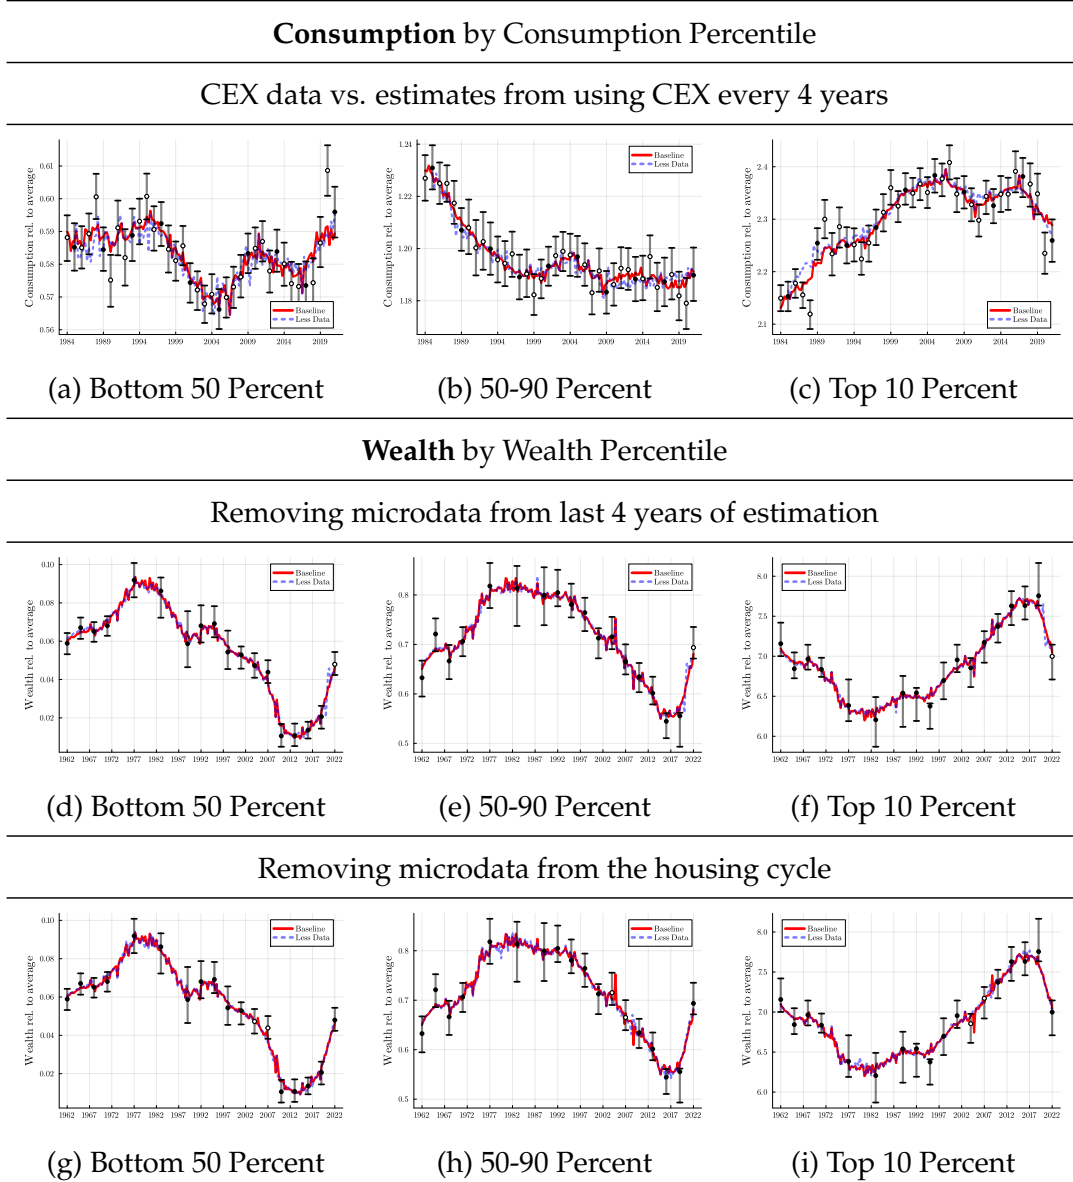
Notes: Figure shows baseline model estimates for consumption (panels (a) to (c)), income (panels (d) to (f)), and wealth (panels (g) to (i)) for different samples. Baseline estimates using all data is always shown as solid red line. Panel (a) to (f): *less data* (dashed blue line) shows smoothed estimates when CEX microdata only enters the smoother every fourth year (black solid dots). Empty dots denote the observations removed from the Kalman smoother. Panel (g) to (i): shows smoothed estimates when only a single SCF wave has been dropped in the Kalman smoother. The black squares show the prediction of the dropped data at the survey wave and the dots show the estimate from the survey data of this wave. Error bars in all figures indicate 95% confidence bounds for each individual survey sample.

of Figure 6 illustrates the experiment. The solid lines are estimates from using all the SCF data. The squares represent the smoothed predictions corresponding to the timing of each omitted survey wave. For example, the square for 1992 is the prediction for the wealth distribution where the 1992 SCF survey *did not* enter the smoother. The circle shows the direct estimate of the corresponding survey (with its confidence limits). The fact that the squares are virtually on top of the solid line implies that, conditional on the model, a single observation of the distributional data has little effect on the smoothed series. In other words, aggregate factors must be important.

The results of the second exercise can be found in Figure 7. Again, we start by dropping three out of four years in the CEX. One can view this experiment as motivated by the fact that countries in Europe only survey consumption every four years, and so we want to find out how well our model would perform in such environment. The first row of Figure 7 shows the corresponding results. The second and third row focus on the distributional dynamics of wealth, since wealth is notoriously the least frequently observed data. Here, we remove a series of microdata observations. In the second row, this is the last four years of the estimation period, 2020Q1 to 2024Q1. The purpose of this exercise is to assess the predictability of the distributional data using our method in terms of a nowcast of the wealth distribution. The third experiment drops all microdata over the housing cycle of the first decade of the 21st century between 2004Q4 and 2009Q4—arguably a period characterized by large swings in the wealth distribution (Kuhn, Schularick, & Steins, 2020).²⁰ In both experiments, we let only the aggregate data inform the estimated distributional dynamics for the period in which we remove the microdata. For both experiments, we find that the prediction that uses all microdata and the prediction that omits four years of microdata are very close. Even in the period where we drop the microdata information, the distributional dynamics of wealth are well captured by aggregate factors, in line with previous research (Bayer, Born, & Luetticke, 2024; Kuhn, Schularick, & Steins, 2020).

²⁰We also perform a second housing cycle experiment from 2007Q4 to 2011Q4 to test the predictability of the recovery phase. Appendix G, Table 7 provides correlations of baseline estimates and all missing-data model estimates, including this second housing experiment.

Figure 7: Predictability of Distributional Data (Re-estimating the Model)



Notes: Figure shows baseline model estimate for consumption (panels (a) to (c)) and wealth (panels (d) to (i)) for different samples. Baseline estimates using all data is always shown as solid red line. Panel (a) to (c): *less data* (dashed blue line) shows the smoothed estimate that results from a re-estimation of the model (and Kalman smoother) when CEX microdata enters only every fourth year (black solid dots). Panel (d) to (f): *less data* (dashed blue line) shows smoothed estimates when the last 4 years of all microdata have been dropped in model estimation (2020Q1 to 2024Q1) and Kalman smoother. Empty dots denote the observations removed from the re-estimation and Kalman smoother. Panel (g) to (i): Same exercise as (d) to (f) but dropping the observations over the house price cycle (2004Q4 and 2009Q4). Error bars in all figures indicate 95% confidence bounds for each individual survey sample.

4.2 Comparison with External Estimates from Other Sources/Methods

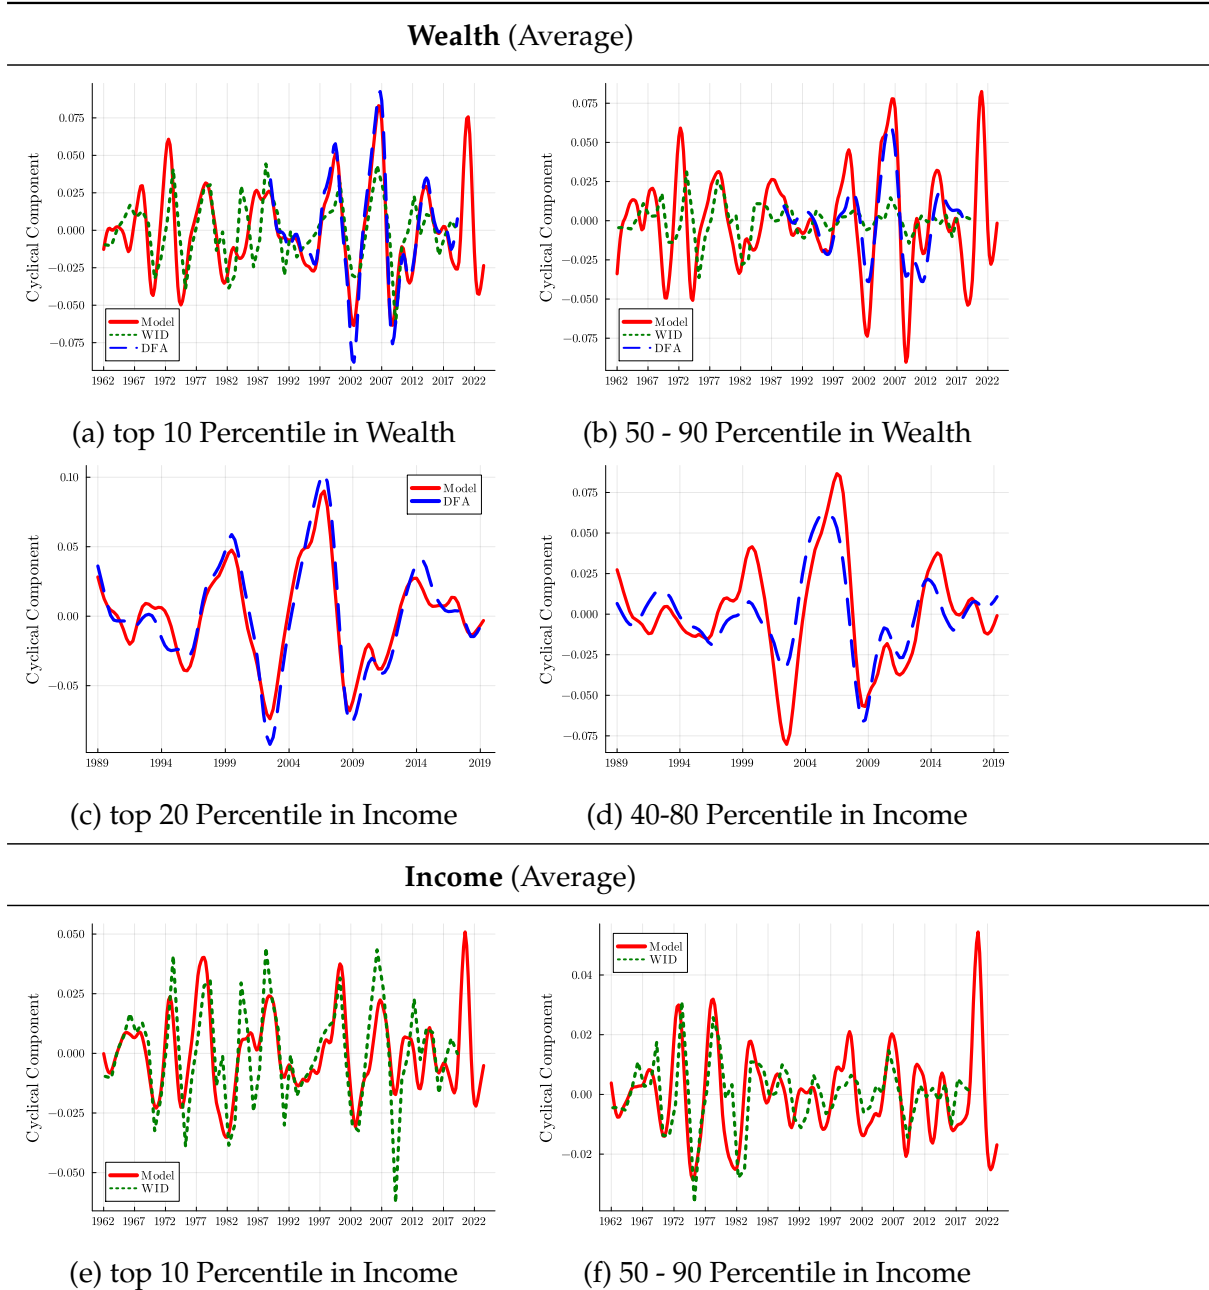
The preceding sections established that our method predicts well out of sample and that estimates lie within the sampling variability used to estimate the distributional data. We now turn to evaluating whether our estimates align with established benchmarks and prevailing understandings of income and wealth. Since the true distributional dynamics are inherently unobservable, we compare our high-frequency distributional estimates with estimates of the cyclical component of income and wealth distribution from the *Distributional Financial Accounts* (DFA) (Batty et al., 2020) and the *World Inequality Database* (WID) (Piketty, Saez, & Zucman, 2018). The former is based exclusively on the SCF as microdata and uses a different estimation technique to produce high-frequency distributional estimates related to household wealth. The latter is based on annual tax data on tax units, but does not use a time series framework to generate higher frequency data. To compare the correlational structure of our estimates, we rely on the DFA estimates of average wealth by income, which allow us to compare the estimates of the joint distribution along these two dimensions.

Figure 8 shows the results of this comparison for the cyclical fluctuations (in logs).²¹ The first two panels compare the wealth by wealth group data from the DFA and the WID. We focus on the wealth of the wealth-richest 10 percent and the wealth of the next 40 percent because the poorest half of the population has wealth very close to zero (Kuhn & Ríos-Rull, 2016).

We find that our method produces smoothed estimates that are close to both alternative estimates and well within the range of the DFA and WID (see Table 6 for correlations). DFA also estimates a time series of wealth by income group, shown in subfigures (c) and (d). Special consideration for these panels needs to be acknowledged, as these model estimates additionally depend on the coefficients related to the correlational structure, represented by the copula density. Under the baseline, the number of coefficients needed to estimate the copula at each point in time is $O^d - d \times (O - 1) + 1 = 972$. Despite perhaps the greater margin for error, we again see model estimates in close correlation to its DFA counterpart. Finally, subfigures (e) and (f) look at the income of the richest 10 percent and the next 40 percent, which is only available in the WID. Again, we find a strong comovement between our estimates and the WID estimates.

²¹We define the cyclical component as the difference between the log of the raw series and its HP-filtered counterpart with the smoothness parameter λ set to 1600 for quarterly series and 6 for annual series.

Figure 8: Comparison of Cyclical Component of Distributional Data to External Sources



Notes: Figure shows the cyclical component of (log) average wealth of (a) the wealthiest 10 percent, (b) the next wealthiest 40 percent, (c) the 20 percent income-richest households, and (d) the next 40 percent income-richest households. Bottom row shows the cyclical component of (log) average income of (e) the income-richest 10 percent and (f) the next income-richest 40 percent. Red lines show cyclical components from baseline model at quarterly frequency. Dotted green line show annual data from the *World Inequality Database* (WID). Dashed blue lines show quarterly data from the *Distributional Financial Accounts* (DFA). Cyclical components are obtained by an HP-filter with smoothing parameter $\lambda = 6$ for annual data and $\lambda = 1,600$ for quarterly data.

5 Consumption Dynamics along the Income and Wealth Distribution over the Business Cycle

The preceding exercises, be it the out-of-sample performance under various settings and the model predictions aligning with external estimates, demonstrated the model’s capacity in deriving reliable estimates. Building on these results, we now turn to the application of estimating consumption dynamics. Consumption dynamics are a crucial component of many theories of aggregate fluctuations and their heterogeneity by income and wealth play a key role in shaping these theories. A key challenge for these theories is therefore the lack of available evidence on high-frequency consumption dynamics along the income and wealth distribution. In this section, the current application therefore demonstrates how our method can be used to fill such gaps.

Given its macroeconomic importance, there is already an existing literature estimating consumption dynamics based on two main approaches: the direct method, which draws on household-level consumption data (e.g., Cloyne, Ferreira, & Surico, 2020; Coibion et al., 2017), and the indirect method, which imputes consumption through a budget identity (e.g., Eika, Mogstad, & Vestad, 2020; Fagereng, Holm, & Natvik, 2021; Holm, Paul, & Tischbirek, 2021). Both approaches come with challenges: the direct approach requires rich sampling variation, while the indirect approach is susceptible to errors due to its reliance on assumptions and multiple measurements.

The strength of our method is that it circumvents both approaches by adopting a functional approach and with some additional structure, addresses measurement error concerns.²² On top, the estimated distributional data are of high frequency and contain consumption along the income and wealth distribution, thus providing insights into business cycle dynamics that are absent from existing U.S. data. With these high-dimensional data, we trace the key distributional dynamics underlying macroeconomic dynamics that play the important role in heterogeneous agent business-cycle models (e.g., Bayer, Born, & Luetticke, 2024; Kaplan, Moll, & Violante, 2018).

5.1 Consumption dynamics over three Recessions

The existing theoretical literature has demonstrated that a key driver of macroeconomic dynamics and a determinant of macroeconomic stabilization policies is the consumption dy-

²²To clarify, our method estimates functions directly, leveraging microdata, macrodata, and the model to identify coefficients (weights) free from measurement error. This approach determines the shape of the marginals and copula density, enabling integration over sparsely sampled intervals without the limitations of the direct method. Additionally, by not imposing an identity when estimating consumption, we avoid the potential errors inherent in the indirect approach.

namics of households during recessions. In addition, the heterogeneous-agent literature has emphasized the importance of heterogeneity of consumption dynamics along the income and wealth distribution. In this effort, we trace out the consumption dynamics along the income and wealth distribution for the three most recent recessions in the United States: the dot-com recession of the early 2000s, the Financial crisis near the end of the first decade of the 21st century, and most recently, the Covid recession. With this, we contribute to discussions concerning different heterogeneous-agent-model mechanisms, whose microfoundation still remains constrained by data limitations as high-frequency data on consumption dynamics along the income and wealth distribution had remained unavailable. In the main part, we focus on the three most recent recessions but report results for previous recessions in Appendix H.

For the three recessions, we consider the consumption dynamics of the bottom 50%, the 50% to 90%, and the top 10% of the income and wealth distribution. To compare the business cycle dynamics of consumption for the different income and wealth groups, we express income relative to the economy-wide average for each period and then index these relative consumption dynamics to the quarter preceding the recession.^{23,24} Given that income is the primary source of consumption financing for most households, Figure 9 first analyzes how consumption evolves across income groups during the last three U.S. recessions.

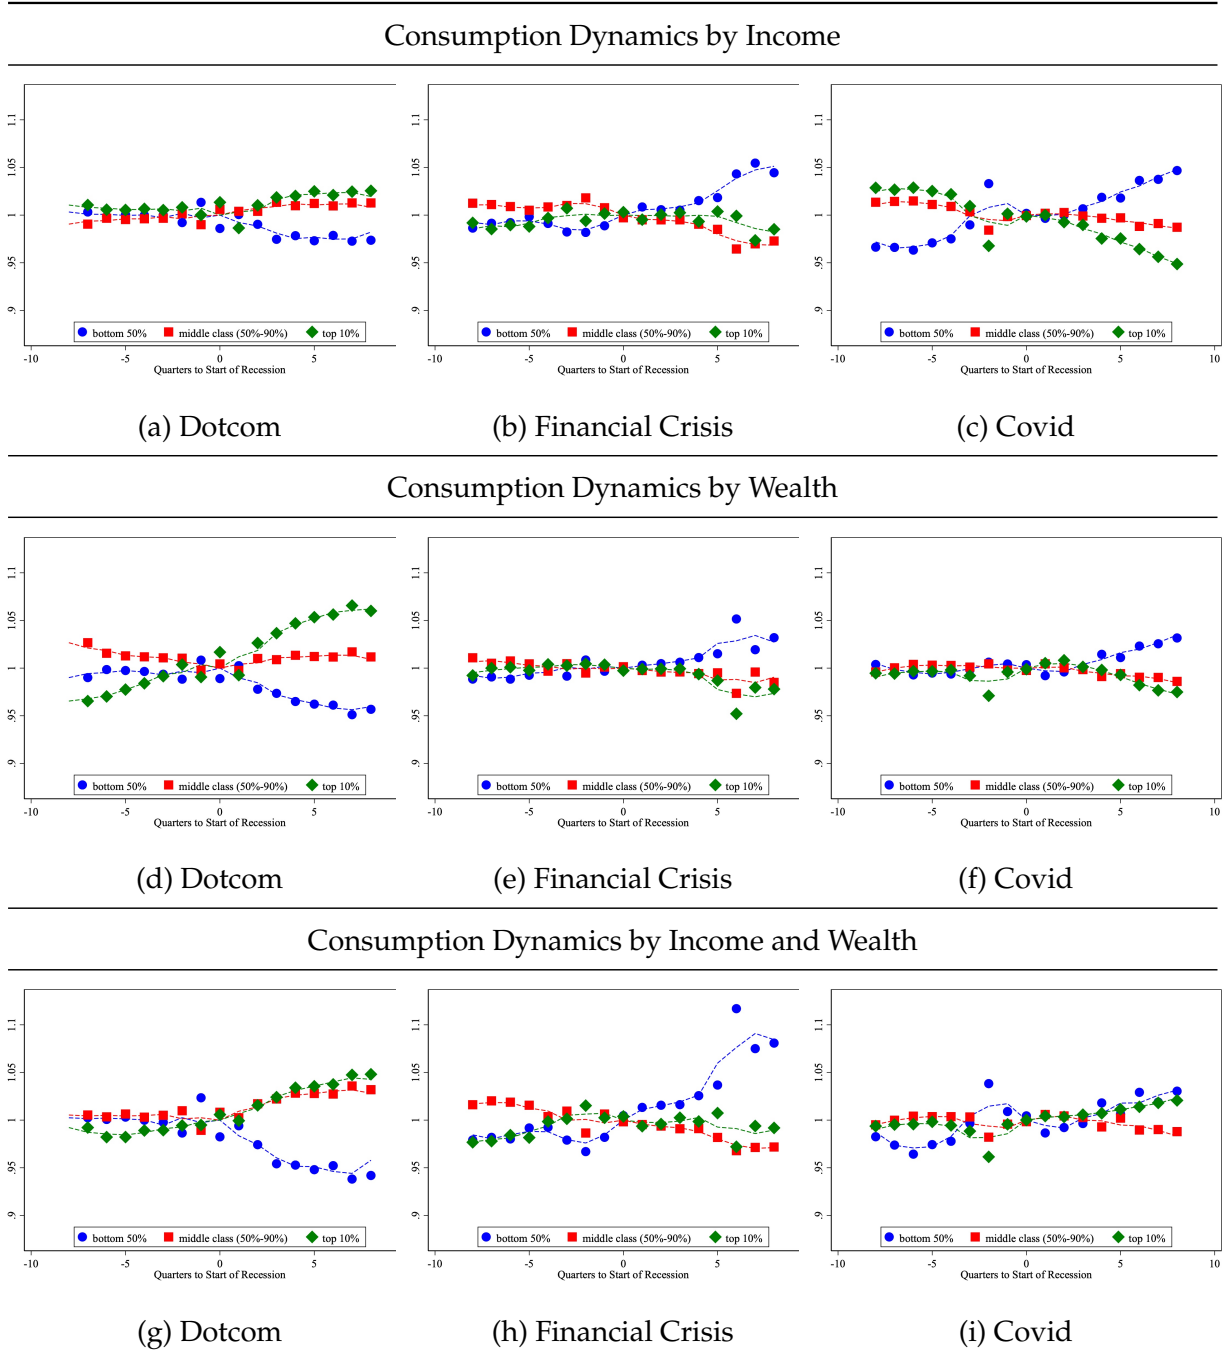
In Panels (a) to (c), we make several interesting observations. First, outside the Covid recession, the top half of the income distribution shows very similar consumption dynamics over recessions (green and red squares). Consumption cyclicalities of the top and middle class follow that of the average household before each recession, but the onset of each recession brought different responses: relative increases for the Dotcom crisis, minute declines during the financial crises, and pronounced declines in the top 10% during Covid. In fact, the top 10% (relative to average) lost 5% relative to peak post Covid—a striking pattern for recessions in the 21st century—though in line with some findings from the recent literature (e.g., Chetty, Friedman, and Stepner, 2024).

Second, the results support the idea that the bottom 50% are poorly self-insured against aggregate risk (blue circles). This conclusion is supported by comparing consumption responses from the Dotcom era when no fiscal intervention took place to when the government implemented supportive fiscal policy during the Great Recession and Covid. The data support that the fiscal policy during the two recent recessions served as insurance for low-income households and sustained their consumption. Strikingly, we find that the bottom 50% of the income distribution saw even an over 5% increase in their consumption relative to average consump-

²³We construct symmetric 3-quarter moving averages and use this moving average for the normalization to the pre-recession quarter.

²⁴We rely on the NBER business cycle dates with the Dotcom recession starting in 2001Q1, Financial Crisis in 2007Q4, and Covid in 2019Q4.

Figure 9: Comparison of Consumption Dynamics during Recessions



Notes: Relative consumption dynamics during recessions along the income and wealth distribution. Consumption dynamics of each income/wealth group are shown relative to average household consumption. These relative consumption time series for each group are indexed to the beginning of the recession. The horizontal axes shows changes of consumption over time relative to the change of average consumption over time. The horizontal axis shows the time relative to the start of the recession. The recessions are the Dotcom recession in 2001Q1, Financial Crisis in 2007Q4, and Covid in 2019Q4.

tion in the aftermath of the Covid recession.

Figure 9 also reports consumption dynamics by wealth group in panels (d) to (f). Looking along the wealth distribution, we observe a very different picture. The Financial Crisis and the

Covid recession show similar and rather weak consumption responses across wealth groups. The aftermath of the Dotcom recession, however, reveals strong divergent dynamics, plausibly driven by large swings in asset prices (Kuhn, Schularick, & Steins, 2020). In the aftermath of the Dotcom recession, we find a flipped picture with the top 10% of households showing much higher consumption growth than the bottom 50% of the wealth distribution. One potential explanation are realized gains from short-selling (Lamont & Stein, 2004; Ofek & Richardson, 2003), and rising house prices during the early 2000s.

Finally, the last panel leverages the particular strength of our novel synthetic data as it looks at the quarterly consumption dynamics in both the income and wealth distribution. Bottom 50% here means bottom 50% in both income and wealth. Middle and top are defined analogously. Hence, we are moving along the main diagonal of the joint distribution. We find the patterns of the joint distribution to be qualitatively and quantitatively different to the responses along the marginal distributions.

Starting with the Dotcom episode in panel (g), we find that when bottom-income households are also at the bottom of the wealth distribution, their consumption response is more sharp but still persistent. For middle-class income households who are also middle-class wealth, we find the same consumption response as the top 10 of the wealth distribution in panel (d). The pattern changes dramatically in the Financial Crisis where consumption for households in the bottom 50% of both income and wealth now saw a nearly 10% increase in consumption six quarters into the recession. By contrast, consumption of the bottom 50% strongly declined during the Dot-Com recession. This response of being jointly in the income and wealth bottom 50% is notably larger than when households were classified as bottom in only income or wealth. This is different to the consumption responses of middle- and top households that offer no evidence that variations in income and wealth were complementary or substitutable in influencing their consumption behavior. Finally, during the COVID-19 recession, having both lower income and wealth appeared to dampen the consumption response (relative to panel (c)). Households at the top of both the income and wealth distributions appears not only smoothed, but enjoyed a modest positive consumption response—similar to that of the bottom 50%. The middle class appears not to respond to more or less income or wealth. Hence, the results for the Covid recession suggest a muting effect of wealth on the consumption response.

To conclude, two key insights emerge. First, consumption dynamics differ significantly between the income and wealth rich and poor and when we consider the joint distribution. Whereas some recessions reduce consumption inequality such as the Covid recession, others such as the Dotcom recession increase consumption inequality. Second, asset prices are a likely important driver of the differential consumption dynamics by income and wealth.

Comparing the Dotcom recession with large swings in asset prices before and after and the Covid crisis with a strong fiscal response and income support programs, we find very different consumption patterns by income and wealth that a one-dimensional analysis of only consumption by income or by wealth would not have detected. Our novel data allow us to identify such underlying differences in the potential drivers of recession dynamics by jointly studying consumption dynamics by income and wealth. While a more detailed analysis and contrasting dynamics to theoretical mechanisms appears fascinating and pressing, these steps on the research agenda go beyond the scope of the current paper that lays the foundation for future research in this direction by providing the tools for such analyses. Future work will have to explore in more detail which of the proposed economic mechanisms in the rich class of heterogeneous agent models accounts for these new facts that our synthetic distributional data have uncovered.

6 Conclusion

In this paper, we presented a new method to derive synthetic distributional consumption, income, and wealth data. The method contributes to the modern theory of macroeconomic dynamics that has the joint distribution of consumption, income, and wealth as a key determinant of aggregate dynamics. Our method closes a gap as it provides a method to study the empirical distributional dynamics as counterpart to the existing theory at business-cycle frequency over time. We have shown that the method is able to incorporate information from various microdata sources independent of their frequency and coverage of variables. By forecasting out of sample, we show that our method can generate joint distributional information at high frequency with a good precision. We show that the derived data can shed new light on the question of how business cycle fluctuations and the distribution of consumption, income, and wealth interact.

References

- Ahn, S. C., & Horenstein, A. R. (2013). Eigenvalue ratio test for the number of factors. *Econometrica*, 81(3), 1203–1227.
- Alvaredo, F., Atkinson, A., Chancel, L., Piketty, T., Saez, E., & Zucman, G. (2016). Distributional national accounts (dina) guidelines: Concepts and methods used in wid. world.
- Andersen, A. L., Johannesen, N., Jørgensen, M., & Peydró, J.-L. (2021). Monetary policy and inequality.
- Attanasio, O., Hurst, E., & Pistaferri, L. (2014). The evolution of income, consumption, and leisure inequality in the united states, 1980–2010. *Improving the measurement of consumer expenditures* (pp. 100–140). University of Chicago Press.
- Attanasio, O., & Pistaferri, L. (2014). Consumption inequality over the last half century: Some evidence using the new psid consumption measure. *American Economic Review*, 104(5), 122–126.
- Auclert, A., Bardóczy, B., Rognlie, M., & Straub, L. (2021). Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5), 2375–2408.
- Bai, J., & Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, 70(1), 191–221.
- Bai, J., & Ng, S. (2019). Rank regularized estimation of approximate factor models. *Journal of Econometrics*, 212(1), 78–96.
- Bai, J., & Ng, S. (2021). Matrix completion, counterfactuals, and factor analysis of missing data. *Journal of the American Statistical Association*, 116(536), 1746–1763.
- Bakam, Y. I. N., & Pommeret, D. (2023). Nonparametric estimation of copulas and copula densities by orthogonal projections. *Econometrics and Statistics*.
- Bañbura, M., & Modugno, M. (2014). Maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data. *Journal of applied econometrics*, 29(1), 133–160.
- Barigozzi, M., & Luciani, M. (2019). Quasi maximum likelihood estimation and inference of large approximate dynamic factor models via the em algorithm. *arXiv preprint arXiv:1910.03821*.
- Bartscher, A. K., Schularick, M., Kuhn, M., & Wachtel, P. (2022). Monetary policy and racial inequality. *Brookings Papers on Economic Activity*, 2022(1), 1–63.

- Batty, M., Bricker, J., Briggs, J., Friedman, S., Nemschoff, D., Nielsen, E., Sommer, K., & Volz, A. H. (2020). The distributional financial accounts of the united states. *Measuring and understanding the distribution and intra/inter-generational mobility of income and wealth*. University of Chicago Press.
- Bayer, C., Born, B., & Luetticke, R. (2024). Shocks, frictions, and inequality in us business cycles. *American Economic Review*, 114(5), 1211–1247.
- Bayer, C., Luetticke, R., Pham-Dao, L., & Tjaden, V. (2019). Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. *Econometrica*, 87(1), 255–290.
- Berger, D., Bocola, L., & Dovis, A. (2023). Imperfect risk sharing and the business cycle. *The Quarterly Journal of Economics*, 138(3), 1765–1815.
- Bhandari, A., Evans, D., Golosov, M., & Sargent, T. J. (2021). Inequality, business cycles, and monetary-fiscal policy. *Econometrica*, 89(6), 2559–2599.
- Blanchet, T., Saez, E., & Zucman, G. (2022). *Real-time inequality* (tech. rep.). National Bureau of Economic Research.
- Boehl, G. (2024). Dime mcmc: A swiss army knife for bayesian inference. *Journal of Econometrics*.
- Breitung, J., & Eickmeier, S. (2006). Dynamic factor models. *Allgemeines Statistisches Archiv*, 90(1), 27–42.
- Chang, M., Chen, X., & Schorfheide, F. (2024). Heterogeneity and aggregate fluctuations. *Journal of Political Economy*, forthcoming.
- Chang, M., & Schorfheide, F. (2024). *On the effects of monetary policy shocks on income and consumption heterogeneity* (Working Paper No. 32166). National Bureau of Economic Research.
- Chang, Y., Kim, C. S., & Park, J. Y. (2016). Nonstationarity in time series of state densities. *Journal of Econometrics*, 192(1), 152–167.
- Chen, W., Er, M. J., & Wu, S. (2005). Pca and lda in dct domain. *Pattern Recognition Letters*, 26(15), 2474–2482.
- Chetty, R., Friedman, J. N., & Stepner, M. (2024). The economic impacts of covid-19: Evidence from a new public database built using private sector data. *The Quarterly Journal of Economics*, 139(2), 829–889.
- Chodorow-Reich, G., Nenov, P. T., & Simsek, A. (2021). Stock market wealth and the real economy: A local labor market approach. *American Economic Review*, 111(5), 1613–57.

- Chow, G. C., & Lin, A.-I. (1971). Best linear unbiased interpolation, distribution, and extrapolation of time series by related series. *The review of Economics and Statistics*, 372–375.
- Cloyne, J., Ferreira, C., & Surico, P. (2020). Monetary policy when households have debt: New evidence on the transmission mechanism. *The Review of Economic Studies*, 87(1), 102–129.
- Coibion, O., Gorodnichenko, Y., Kueng, L., & Silvia, J. (2017). Innocent bystanders? monetary policy and inequality. *Journal of Monetary Economics*, 88, 70–89.
- Curtin, R. T., Juster, T., & Morgan, J. N. (1989). Survey estimates of wealth: An assessment of quality. *The measurement of saving, investment, and wealth* (pp. 473–552). University of Chicago Press, 1989.
- Cutler, D. M., Katz, L. F., Card, D., & Hall, R. E. (1991). Macroeconomic performance and the disadvantaged. *Brookings papers on economic activity*, 1991(2), 1–74.
- Czajka, J. L., Jacobson, J. E., & Cody, S. (2003). Survey estimates of wealth: A comparative analysis and review of the survey of income and program participation. *Soc. Sec. Bull.*, 65, 63.
- Denton, F. T. (1971). Adjustment of monthly or quarterly series to annual totals: An approach based on quadratic minimization. *Journal of the american statistical association*, 66(333), 99–102.
- Di Fonzo, T. (2003). Constrained retropolation of high-frequency data using related series: A simple dynamic model approach. *Statistical Methods and Applications*, 12(1), 109–119.
- Di Maggio, M., Kermani, A., & Majlesi, K. (2020). Stock market returns and consumption. *The Journal of Finance*, 75(6), 3175–3219.
- Diebold, F. X., & Li, C. (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics*, 130(2), 337–364.
- Doan, T., Litterman, R., & Sims, C. (1984). Forecasting and conditional projection using realistic prior distributions. *Econometric reviews*, 3(1), 1–100.
- Doz, C., Giannone, D., & Reichlin, L. (2012). A quasi–maximum likelihood approach for large, approximate dynamic factor models. *Review of economics and statistics*, 94(4), 1014–1024.
- Durbin, J., & Koopman, S. J. (2012). *Time series analysis by state space methods* (Vol. 38). OUP Oxford.
- Eika, L., Mogstad, M., & Vestad, O. L. (2020). What can we learn about household consumption expenditure from data on income and assets? *Journal of Public Economics*, 189, 104163.

- Fagereng, A., Holm, M. B., & Natvik, G. J. (2021). Mpc heterogeneity and household balance sheets. *American Economic Journal: Macroeconomics*, 13(4), 1–54.
- Fernandez, R. B. (1981). A methodological note on the estimation of time series. *The Review of Economics and Statistics*, 63(3), 471–476.
- Flavin, M., & Yamashita, T. (2002). Owner-occupied housing and the composition of the household portfolio. *American Economic Review*, 92(1), 345–362.
- Freyaldenhoven, S. (2022). Factor models with local factors—determining the number of relevant factors. *Journal of Econometrics*, 229(1), 80–102.
- Friedman, M. (1962). The interpolation of time series by related series. *Journal of the American Statistical Association*, 57(300), 729–757.
- Gagliardini, P., Ossola, E., & Scaillet, O. (2019). A diagnostic criterion for approximate factor structure. *Journal of Econometrics*, 212(2), 503–521.
- Giannone, D., Lenza, M., & Primiceri, G. E. (2015). Prior selection for vector autoregressions. *Review of Economics and Statistics*, 97(2), 436–451.
- Gouskova, E., Andreski, P., & Schoeni, R. F. (2010). *Comparing estimates of family income in the panel study of income dynamics and the march current population survey, 1968-2007*. Survey Research Center, Institute for Social Research, University of ...
- Gregoir, S. (2003). Propositions pour une désagrégation temporelle basée sur des modèles dynamiques simples. *Paris-Bercy*, 141.
- Gronau, Q. F., Sarafoglou, A., Matzke, D., Ly, A., Boehm, U., Marsman, M., Leslie, D. S., Forster, J. J., Wagenmakers, E.-J., & Steingroever, H. (2017). A tutorial on bridge sampling. *Journal of mathematical psychology*, 81, 80–97.
- Hamilton, J. D., & Xi, J. (2022). Principal component analysis for nonstationary series.
- Harvey, A., & Chung, C.-H. (2000). Estimating the underlying change in unemployment in the uk. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 163(3), 303–309.
- Harvey, A. C. (1990). Forecasting, structural time series models and the kalman filter.
- Harvey, A. C., & Pierse, R. G. (1984). Estimating missing observations in economic time series. *Journal of the American statistical Association*, 79(385), 125–131.
- Holm, M. B., Paul, P., & Tischbirek, A. (2021). The transmission of monetary policy under the microscope. *Journal of Political Economy*, 129(10), 2861–2904.
- Inoue, A., & Rossi, B. (2021). The effects of conventional and unconventional monetary policy: A new approach. *Quantitative Economics*, 12(4), 1085–1138.

- Insolera, N. E., Simmert, B. A., & Johnson, D. S. (2021). An overview of data comparisons between psid and other us household surveys. *Technical Series Paper*, 21–02.
- Kaplan, G., Moll, B., & Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, 108(3), 697–743.
- Kim, Y.-S., & Stafford, F. P. (2000). The quality of psid income data in the 1990s and beyond. *Technical Series Paper*.
- Kneip, A., & Utikal, K. J. (2001). Inference for density families using functional principal component analysis. *Journal of the American Statistical Association*, 96(454), 519–542.
- Kuhn, M., & Ríos-Rull, J.-V. (2016). 2013 update on the us earnings, income, and wealth distributional facts: A view from macroeconomics. *Federal Reserve Bank of Minneapolis Quarterly Review*, 37(1), 2–73.
- Kuhn, M., Schularick, M., & Steins, U. I. (2020). Income and wealth inequality in america, 1949–2016. *Journal of Political Economy*, 128(9), 3469–3519.
- Lamont, O. A., & Stein, J. C. (2004). Aggregate short interest and market valuations. *American Economic Review*, 94(2), 29–32.
- Litterman, R. B. (1980). *Techniques for forecasting with vector autoregressions* (Doctoral dissertation). Ph. D. thesis, University of Minnesota.
- Litterman, R. B. (1983). A random walk, markov model for the distribution of time series. *Journal of Business & Economic Statistics*, 1(2), 169–173.
- McCracken, M. W., & Ng, S. (2021). FRED-QD: A Quarterly Database for Macroeconomic Research. *Review*, 103(1), 1–44.
- McKay, A., & Wolf, C. K. (2023). Monetary policy and inequality. *Journal of Economic Perspectives*, 37(1), 121–144.
- McLachlan, G. J., & Krishnan, T. (2008). *The em algorithm and extensions*. John Wiley & Sons.
- Meeks, R., & Monti, F. (2023). Heterogeneous beliefs and the phillips curve. *Journal of Monetary Economics*, 139, 41–54.
- Meng, X.-L., & Wong, W. H. (1996). Simulating ratios of normalizing constants via a simple identity: A theoretical exploration. *Statistica Sinica*, 831–860.
- Mian, A. R., Straub, L., & Sufi, A. (2020). *The saving glut of the rich* (tech. rep.). National Bureau of Economic Research.
- Moauero, F., & Savio, G. (2005). Temporal disaggregation using multivariate structural time series models. *The Econometrics Journal*, 8(2), 214–234.

- Mönch, E., Uhlig, H. et al. (2005). Towards a monthly business cycle chronology for the euro area. *Journal of Business Cycle Measurement and Analysis*, 2005(1), 43–69.
- Ofek, E., & Richardson, M. (2003). Dotcom mania: The rise and fall of internet stock prices. *the Journal of Finance*, 58(3), 1113–1137.
- Onatski, A., & Wang, C. (2021). Spurious factor analysis. *Econometrica*, 89(2), 591–614.
- Otto, S., & Salish, N. (2022). Approximate factor models for functional time series. *arXiv preprint arXiv:2201.02532*.
- Pfeffer, F. T., Schoeni, R. F., Kennickell, A., & Andreski, P. (2016). Measuring wealth and wealth inequality: Comparing two us surveys. *Journal of economic and social measurement*, 41(2), 103–120.
- Piketty, T., & Saez, E. (2003). Income inequality in the united states, 1913–1998. *The Quarterly journal of economics*, 118(1), 1–41.
- Piketty, T., Saez, E., & Zucman, G. (2018). Distributional national accounts: Methods and estimates for the united states. *The Quarterly Journal of Economics*, 133(2), 553–609.
- Proietti, T. (2006). Temporal disaggregation by state space methods: Dynamic regression methods revisited. *The Econometrics Journal*, 9(3), 357–372.
- Ramsay, J. O., & Silverman, B. W. (2005). *Functional data analysis*. Springer New York.
- Rüschendorf, L. (2009). On the distributional transform, sklar’s theorem, and the empirical copula process. *Journal of statistical planning and inference*, 139(11), 3921–3927.
- Saez, E., & Zucman, G. (2016). Wealth inequality in the united states since 1913: Evidence from capitalized income tax data. *The Quarterly Journal of Economics*, 131(2), 519–578.
- Salazar, E., & Weale, M. (1999). Monthly data and short-term forecasting: An assessment of monthly data in a var model. *Journal of Forecasting*, 18(7), 447–462.
- Silva, J. S., & Cardoso, F. (2001). The chow-lin method using dynamic models. *Economic modelling*, 18(2), 269–280.
- Skinner, J. (1987). A superior measure of consumption from the panel study of income dynamics. *Economics Letters*, 23(2), 213–216.
- Sklar, A. (1959). Vol. 8 of fonctions de répartition à n dimensions et leurs marges, 229–231. *Paris: Publications de l’Institut de statistique de l’Université de Paris*.
- Sklar, A. (1973). Random variables, joint distribution functions, and copulas. *Kybernetika*, 9(6), 449–460.

- Smith, M., Zidar, O. M., & Zwick, E. (2021). *Top wealth in america: New estimates and implications for taxing the rich* (tech. rep.). National Bureau of Economic Research.
- Tsay, R. S. (2016). Some methods for analyzing big dependent data. *Journal of Business & Economic Statistics*, 34(4), 673–688.
- Wang, Z., Gu, Q., Ning, Y., & Liu, H. (2015). High dimensional em algorithm: Statistical optimization and asymptotic normality. *Advances in neural information processing systems*, 28.
- Wu, C. J. (1983). On the convergence properties of the em algorithm. *The Annals of statistics*, 95–103.

A Estimation of Factor Structures and Marginal Data Densities

A.1 Alternative Methods to estimating the Projection Matrix Γ

A key input to our procedure is the mapping from factors to observables, defined by the projection matrix Γ . To estimate Γ , one method would be based on complete data alone. For our case, this means using exclusively the PSID data to estimate Γ based on the PCA of $\tilde{\theta}^{PSID}$. We consider different scenarios where we retain a different numbers of factors each time. Specifically, for this approach, we estimate the model with 3, 6, 7, and 8 distributional factors, obtained from the 12 PSID waves.

However, this approach runs the risk of not fully representing the entire subspace that characterizes the evolution of the joint distributions because of the scarcity of PSID data. The literature has suggested alternatives for taking also incomplete data into account in factor models. In particular, we explore one alternative estimator for Γ : we consider the Tall-Wide algorithm of Bai and Ng (2021). We compare the quality of the model estimates under the alternative methods using marginal data densities for model comparison.

A.1.1 Tall-Wide Algorithm of Bai and Ng (2021)

In alternative approach to incorporating an additional block of data to the complete data is provided by Bai and Ng (2021). The paper tackles this problem by identifying two blocks of data within some larger T by N matrix—a tall block (data observed for all periods) and a wide block (time periods for which entire distribution is observed). For our setting, this means we add a tall block to our estimation of Γ , while the wide block would be what we call the complete data.²⁵ The tall block we add incorporates data on lower dimensional copulas of (income, consumption) and (income, wealth), as well as additional data on the marginals. This would be variation coming from both the CEX and SCF.

This alternative estimator for Γ , since it relies on more data, may improve model fit. For our application, we extract the factors that explain 80% of the summable variation, as well as 85%. This corresponds to 10 and 12 factors; however, as mentioned before, it is unclear the degree of variation shared among this set of factors relative to the original set of factors. Furthermore, keeping more factors from this estimator runs the risk of over-parameterizing the model. Table 3 presents the marginal likelihoods of these TW projection matrices. Results suggest these models have marginally inferior fit than models that only rely on complete data.

²⁵It is important to note that, given some data, such an estimator can consistently estimate the common component without making any assumptions on the nature of missingness. We refer the reader to the paper for more details on how the exact procedure is performed.

Table 3: Model Comparison

Model	Harmonic Mean	Bridge
<i>TW Projection Matrix</i>		
12 distributional factors, 21 agg. factors	−27,622.8	−28,246.1
10 distributional factors, 21 agg. factors	−28,135.3	−28,871.0
<i>Ordinary Projection Matrix</i>		
8 distributional factors, 15 agg. factors	−25,906.1	−24,207.3
8 distributional factors, 21 agg. factors	−25,061.3	−23,307.4
8 distributional factors, 25 agg. factors	−27,189.2	−24,518.6
7 distributional factors, 21 agg. factors	−31,033.1	−29,621.4
6 distributional factors, 21 agg. factors	−36,913.6	−35,359.3
3 distributional factors, 10 agg. factors	−44,714.9	−45,465.6

Notes: The table reports the log of the marginal data densities (MDD) across different model specifications. The MDD is estimated using two estimators: (1) Harmonic Mean and (2) a Bridge sampler. The baseline model is represented with bold numbers. Higher values means most efficient.

A.2 Marginal Data Densities and Optimal Factor Structures

We estimate various versions of the state space model for different factor loadings Γ , that differ both in the number of factors and the estimation of Γ itself, and different numbers of aggregate factors. For each estimated model, we calculate the marginal data density (MDD) in order to discriminate between these alternatives. The set of models considered rely on the same data, but differ in the size of the parameter space. The MDD will effectively internalize these two features, selecting the model for which we can expect the best forecasting performance, while penalizing models through a lower density due to larger parameter spaces.

The marginal data densities $p(\tilde{\theta})$ are estimated using Geweke’s modified harmonic mean estimator. Albeit standard, there may be concerns, however, that given the dimensionality of the model, that such an estimator may not be appropriately approximated by a Gaussian distribution. Secondly, the estimator may not be numerically stable, as it requires the inversion of matrices – in this case, large. In this sense, we also estimate the density using a bridge sampler (Gronau et al., 2017; Meng & Wong, 1996), which is numerically stable (does not require any inversions) and may better internalize the shape of the posterior.

Table 3 covers the span of proposed models that were ultimately assessed, as well as the marginal data densities over the different estimators. We find that for the model with the projection matrix from the complete PSID data, using the maximum number of 8 factors achieves the highest MDD with 21 aggregate factors. Increasing the number of aggregate factors or decreasing it reduces the MDD. Using the factor loadings Γ from the Bai and Ng (2021) TW-

algorithm also achieves a lower MDD.

B Data

The construction of these estimates relies on a great deal of data. An advantage with our method, however, is that it can incorporate these different microdata and their various differences in generating consensus estimates of the distributional data. Below, we describe the data, all expressed in 2019 dollars, and explain the mappings across data to ensure measures are at some base comparability (Curtin, Juster, & Morgan, 1989; Czajka, Jacobson, & Cody, 2003; Pfeffer et al., 2016). See Table 1 for information on their availability.

B.1 SIPP

The SIPP panel is a nationally representative, individual-level survey known for providing high-frequency dynamics on employment, earnings, wealth, household composition and program participation. For the data cleaning, the data is aggregated to the household-level, at quarterly frequency.

Income. For the 2014 releases and onward, we use the `THTOTINC` variable for income. For data releases prior, we sum over (1) earnings (`ws1_am`, `ws1_am`) (2) property/investment income (`tpprpinc`) (3) unemployment (`tuc1amt`, `tuc2amt`, `tuc3amt`) and (4) transfers (`tptrninc`, `tpscininc`, `twicamt`, `tfs_am`, `tssi_amt`) to construct household income.

Wealth. For the 2014 releases and onward, we use the `THNETWORTH` variable for wealth. For data releases prior, wealth is defined as total assets (`hhtwlth`) net total liabilities (`hhusdbt`, `hhscdbt`).

B.2 SCF+

The Survey of Consumer Finances (SCF), since its inception in 1983, is seen as the data gold mine for household information on income and wealth; however, due to the research excavations of Kuhn, Schularick, and Steins (2020), we are able to combine these triennial cross-sections with historical waves of the SCF; hence the name SCF+. Kuhn, Schularick, and Steins (2020) mention "... the SCF+ is the first dataset that makes it possible to study the joint distributions of income and wealth over the long run". Thus, it goes without saying how requisite this is for our study. Below we describe the concepts in turn.

Income. Our definition of income follows Kuhn, Schularick, and Steins (2020), which consists of the following components: (1) labor income (i.e., earnings) (2) income from public transfers (3) income from professional practice and self-employment (4) income from rents (5) dividend income and (6) business/farm income. A different taxonomy that illustrates these components are taxable and transfer income.

Assets. Total assets include (1) liquid assets such as a household's checking and savings account, CDs, call/money market accounts, short-term government bonds, and mutual funds (2) illiquid assets such as housing and other real estate minus debt on that properties respectively, automobiles (3) defined-contribution retirement plans (4) the cash value of life insurance (5) stocks and (6) business equity.

Debt. We define debt of a household as the sum of personal (mostly unsecured) debt and housing (mortgage) debt. Housing debt includes debt from all properties and any loans made against the housing e.g., through HELOCs. Personal debt includes car loans, education loans, any loans from relatives, credit card debt, medical debt and legal debt.

Wealth. Wealth is total assets net total debt of a household.

B.3 PSID

The Panel Study of Income Dynamics complements the SCF+ extraordinarily well, as they take our estimations beyond more than half a century. In comparison to the post-1983 SCF, a deeper analysis of their similarity can be found in Pfeffer et al. (2016).

Income. The PSID has collected family income annually from 1968 to 1996 and then biennially from 1997 to 2021. Its measure of income is the sum of taxable income, transfers, and social security for the reference person, the spouse/partner (if any) and other members of the family.²⁶

Assets. Data collection on household wealth took place in 1984, 1989, 1994, and then every wave beginning in 1999. The data on assets is split into liquid and illiquid assets. Albeit minor, the definition of liquid assets will vary between datasets, so careful attention here. Liquid assets for the PSID includes checking and savings accounts, short-term instruments such as money-market accounts, certificates of deposit, and treasury bills. Illiquid assets include business equity, financial assets held in mutual funds, stocks, bond funds, investment funds; real assets held in real estate, vehicles like motor homes, boats, trailers, and cars; and retirement wealth in private annuities or IRAs.

²⁶In the PSID, a family is a group of people living together who are economically interdependent.

Debt. For the PSID, we achieve the same debt split: personal and mortgage debt. This includes all kinds of real-estate debt, and unsecured debt such as credit card debt, student loans, medical debt, legal debt, and loans from relatives.

Wealth. Wealth is total assets net total debt of a household.

Consumption. Studying papers such as Attanasio, Hurst, and Pistaferri (2014), Attanasio and Pistaferri (2014), Cutler et al. (1991), Flavin and Yamashita (2002), and Skinner (1987), we define consumption as the sum of these expenditures: food, rent (for renters), housing rental equivalence (for home-owners), utilities, health, public transport, education, and childcare. We set the housing rental equivalence to be 6% of the home market value reported by households in the PSID. Consumption data is only available from 1999 in a biennial interval.

B.4 CPS

We use the Community Population Survey (CPS) Annual Social and Economic Supplement (ASEC). The sample is designed primarily to produce estimates of the labor force characteristics and runs from 1962 to 2022.

Income. Income data are collected as part of the ASEC for the months of February, March and April as a supplement to the regular CPS monthly labor force interviews. The ASEC asks each person in the sample who is 15 years old and over about the amount of income received from a list of sources in the previous calendar year. We treat these observations as being observed in quarter four of the previous calendar year. For details on top-coding, see cps.ipums.org/cps/topcodes_tables.shtml.

B.5 CEX

The Consumption Expenditure Survey (CEX) is the most comprehensive household survey in the U.S. for recording the consumption habits of households. The CEX has two components: the interview survey (IS) and the diary survey (DS). The interview survey has sufficiently rich data on what we need, so we only use data from this component. Within this component, there are several files, each of which pertain to a topic, from which we can extract information. The following table breaks down each category of consumption, defining which UCCs belong to which category and which file it can be found in. All of these categories will combine to make the consumption variable. The table will also define wealth concepts of the CEX we use in our study. Since each household consumption record is with respect to a UCC, we find this presentation most apropos.

Item	UCCs / FMLI label	File
	<i>Consumption</i>	
Food	190904, 790220, 190901, 190902, 190903, 790410, 790430, 200900, 790330, 790420, 800700, 790230, 790240	MTBI
Rent	210110, 800710	MTBI
Utilities	250111, 250112, 250113, 250114, 250211, 250212, 250213, 250214, 250221, 250222, 250223, 250224, 250901, 250902, 250903, 250904, 250911, 250912, 250913, 250914, 260111, 260112, 260113, 260114, 260211, 260212, 260213, 260214, 270211, 270212, 270213, 270214, 270310, 270411, 270412, 270413, 270414, 270101, 270102, 270104, 270105, 270310, 270311, 690116, 270901, 270902, 270903, 270904	MTBI
Health	570110, 570111, 570210, 570220, 570230, 560110, 560210, 560310, 560330, 560400, 340906, 540000, 550110, 550320, 550330, 550340, 570901, 570903, 570240, 580111, 580112, 580113, 580114, 580311, 580312, 580901, 580903, 580904, 580905, 580906, 580400, 580907	MTBI
Public Transport	520531, 520532, 530311, 530312, 530501, 530902, 530210, 530411, 530412, 520511, 520512, 520521, 520522, 520542, 520902, 520903, 520904, 520905, 520906, 520907, 530110, 530901, 520110, 520310	MTBI

Education	210310, 370903, 390901, 660110, 660210, 660310, 660900, 670110, 670210, 670901, 670902, 800802, 800804, 690111, 690112, 660410, 660902, 670410, 670903, 690114, 690310	MTBI
Child care	340210, 340211, 340212, 670310, 660901	MTBI
Rental Equiva- lence	910050, 800721 (market value of home), SIMHOUSX, RENTEQVX	FMLI, MTBI
Gas & Vehicle Re- pairs	470111, 470112, 470113, 470220, 470211, 470212, 480110, 480212, 480213, 480214, 490110, 490211, 490212, 490221, 490231, 490232, 490311, 490312, 490313, 490314, 490318, 490319, 490411, 490412, 490413, 490501, 490502, 490900, 520410, 480215, 620113	MTBI
Other Concepts		
Housing Debt	QBLNCM1X, QBLNCM2X, QBLNCM3X, QBLNCM1G, PRINAMTX	MOR
Personal Debt	6001, 6002 (1990–2013), 5400, 5500, 5600, CREDITX, STUDNTX, OTHLONX, CREDITX1, CREDITX5, QBALNM1X	MTBI, ITBI, FMLI, FN2
Liquid Assets	SAVACCTX, CKBKACTX, USBNDX, 920010, 920020, 920030, 5100, LIQUIDX	FMLI, ITBI
Financial Assets	5800, 920040, STOCKX, SECESTX, OTHASTX	FMLI, ITBI
Income	FINCBTAX	FMLI

Notes: Table shows, by item, the identifiers necessary to construct each component of consumption, income and wealth for the CEX. The location of these identifiers can be found under the *File* column.

B.6 Aggregates

Together with the microdata, we specify a model component that captures the various aggregate shocks that buffer the joint distribution of consumption, income, and wealth. This is represented in the state equation of the state-space model. The aggregate data we rely on to extract this information comes from the FRED-QD. This has various macro-data on industrial production, employment, housing, inventories, prices, earnings, productivity, household expectations, household balance sheets, interest rates, credit, etc. You can find more information on research.stlouisfed.org/econ/mccracken/fred-databases/.

Before performing the PCA on the aggregates, we are careful to check each series for non-stationarity. Recent literature has placed emphasis on the identifiability of orthogonal factors in high-dimensional settings, in particular for macroeconomic aggregates, and finds non-stationarity to be the culprit of spurious variation (Hamilton & Xi, 2022; Onatski & Wang, 2021). Running the PCA on the non-stationary data will erroneously find that a large set of aggregates is confined to just a few factors. Taking note, we first remove any variation due to seasonality using the X13-ARIMA and closely follow the transformations (to induce stationarity) proposed by McCracken and Ng (2021). The resulting series satisfy an Augmented-Dickey-Fuller Test with a significance level of $\alpha = 0.05$ and are visually inspected for abnormalities.

The set of now stationary aggregates are concatenated with four of its lags to form a data matrix of quintuple the size and then column-wise standardized. A PCA on this block of data is performed and 21 orthogonal factors are kept. The number of factors chosen is based on Freyaldenhoven (2022). The baseline model estimation includes these 21 factors as inputs Y_t . More on the selection of factors can be found in Appendix E.

C Minnesota Prior

Given the size of the model, we use a Bayesian approach. This means we choose to regularize the model using priors defined on A , B , and Ω , which are all the parameter matrices for the state equation and S for the measurement equation. This section will only cover the prior on the state equation, which is defined in Block (22). To represent the uncertainty surrounding these parameters, the following Minnesota prior is proposed:

$$\begin{pmatrix} \text{vec}(A) \\ \text{vec}(B) \end{pmatrix} \sim \mathcal{MN}(\mu_{Minn}, V_{Minn}), \quad (24)$$

$$A_{ij} = \begin{cases} \kappa_3 & \text{for the first lag of the state variable, } i = j, \\ 0 & \text{for the exogenous terms, } i \neq j, \end{cases}$$

$$B = \mathbf{0},$$

$$V_{Minn, ii} = \begin{cases} \frac{\kappa_0}{l^2} & \text{for own lags of the respective state variable } i, \\ \frac{\kappa_0 \kappa_1}{l^2} \times \frac{\hat{\sigma}_{ii}^2}{\hat{\sigma}_{jj}^2} & \text{for lags of the other state variables } j, \\ \kappa_0 \kappa_2 \times \sigma_{ii}^2 & \text{for aggregate factors,} \end{cases} \quad (25)$$

$$\kappa_1 \sim \mathcal{U}[0.2, 0.99] \quad \kappa_2 \sim \mathcal{U}[0.2, 0.99] \quad \kappa_3 \sim \mathcal{U}[-.99, .99] \quad \kappa_4 \sim \mathcal{N}[0, 1] \quad , \quad (26)$$

where the prior distribution on A and B is a multivariate-normal with μ_{Minn} the mean of the distribution and V_{Minn} the diagonal variance-covariance matrix. All elements in μ_{Minn} are shrunk to zero except for the diagonal elements of A , whose values are set to κ_3 . κ_3 controls the level of persistence of the state law of motion. Since the model is stationary, it means the eigenvalues of A must be in absolute value less than 1. Since apriori A is diagonal, this reduces to simply ensuring its diagonal elements are in absolute value less than 1. The prior on κ_3 is defined to represent precisely this.

Governing the variance V_{Minn} are a set of hyperparameters $\{\kappa_i\}_{i=0}^2$, which ultimately govern the tightness on the prior. As shown in (25), the main hyperparameter here is κ_0 , as it appears as a scaling term in each case in V_{Minn} . For κ_0 , we set it to 0.05 following Figure 1 of Giannone, Lenza, and Primiceri (2015).²⁷ κ_1 and κ_2 govern the importance of cross-lags and

²⁷Given the size of the model, including κ_0 in the hyperparameter MCMC sampling led to extreme shrinkage. In some comparative static exercise, Giannone, Lenza, and Primiceri (2015) showed this is the expected behavior for a BVAR, but no comment on state-space models. Keeping this extreme shrinkage would mute all cyclical movements coming from the aggregate factors and cross-lag factors and fix the state equation at the prior mean. Also, with extremely tight priors, the sampling procedure would operate in an extremely bounded space, which invites larger rejection rates due to draws falling outside the prior support and, thereby, decreasing sampling

aggregates. The priors on κ_1 and κ_2 are set to say that we are not informed about the value of these parameters, but are certain the variation from cross-lags and aggregates are important, but proportionally less than that of a factor’s own-lags. Finally, $\hat{\sigma}_{ii}^2$ is the estimated variance of the residuals from a least squares estimation of estimated factor i on 1 lag, where factor i is estimated from the PCA.

Also part of the state equation is Ω , which is the variance-covariance matrix of the one step ahead forecast errors. A prior is defined over the diagonal elements of Ω . As stated in the text, the persistence of the factors, κ_3 , determines the first moment of the distribution of these diagonal elements. An additional hyperparameter, κ_4 , determines the variance of these diagonal elements and is set to be a weakly informative prior on the variances.

We estimate each hyperparameter $\{\kappa_i\}_{i=1}^4$ using the methodology of Giannone, Lenza, and Primiceri (2015). Table 5 summarizes the role of each hyperparameter.

Table 5: Hyperparameters on Minnesota Prior

Hyperparameter	Description
κ_0	controls overall tightness of prior variances.
κ_1	the size of the prior variance of factors, not corresponding to own lags.
κ_2	the size of the prior variance of aggregate factors
κ_3	persistence of state LOM
κ_4	variance on diagonal elements of Ω

D Details on MCMC

To estimate the posterior distribution of the parameters and subsequently sample, we employ the DIME sampler from Boehl (2024). The sampler is particularly advantageous for dealing with potentially complex, high-dimensional, multi-modal posterior distributions, especially when these distributions have ex-ante unknown properties. Traditional MCMC methods often struggle with such distributions due to their reliance on gradient-based optimization or difficulties in converging efficiently. DIME addresses these issues by combining the strengths of global multi-start optimizers with the robustness of Monte Carlo methods, allowing it to explore the typical set of the posterior distribution more quickly and effectively.

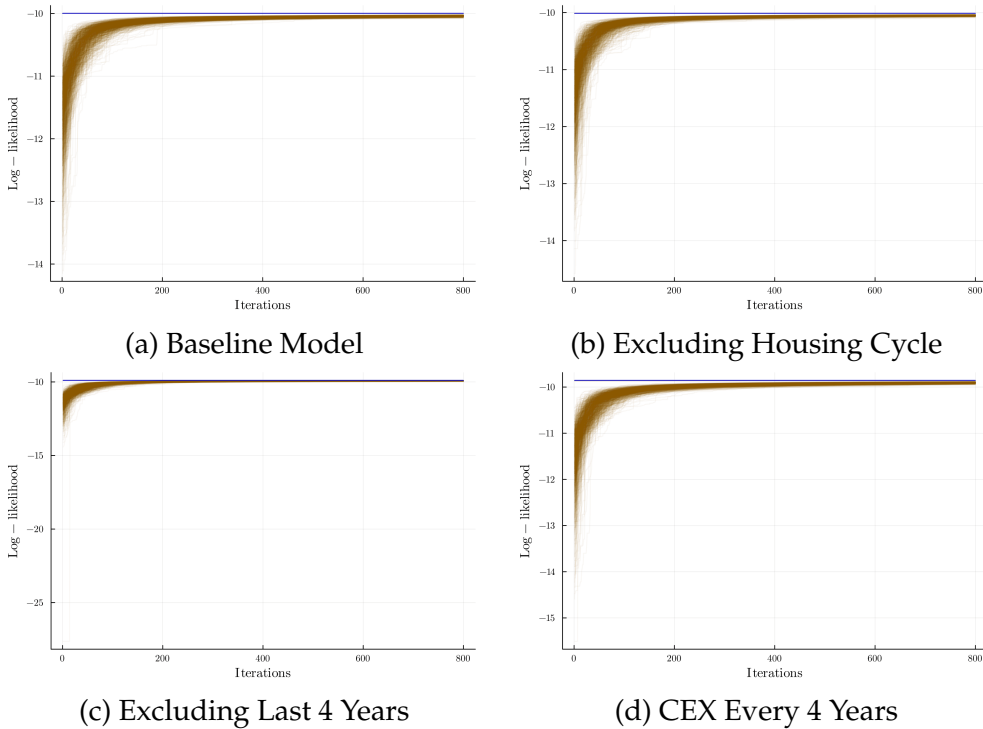
To initialize the sampler, we first run a tentative mode finder (not necessary, but makes process 4-5 times faster) and then let an ensemble of $5n$ chains run for 800 – 1200 iterations, for n the size of the parameter vector. The last 25% of draws are kept as the posterior distribution.

efficiency. Nonetheless, by using a global optimizer with a less constrained, but still informed prior centered on 0.05 should alleviate these concerns.

There is a single tuning parameter χ that dictates (for each iteration and for each chain) the probability of mixture between the local and global transition kernel. We set $\chi = 0.1$, which means with 10% probability, we draw the global transition kernel.

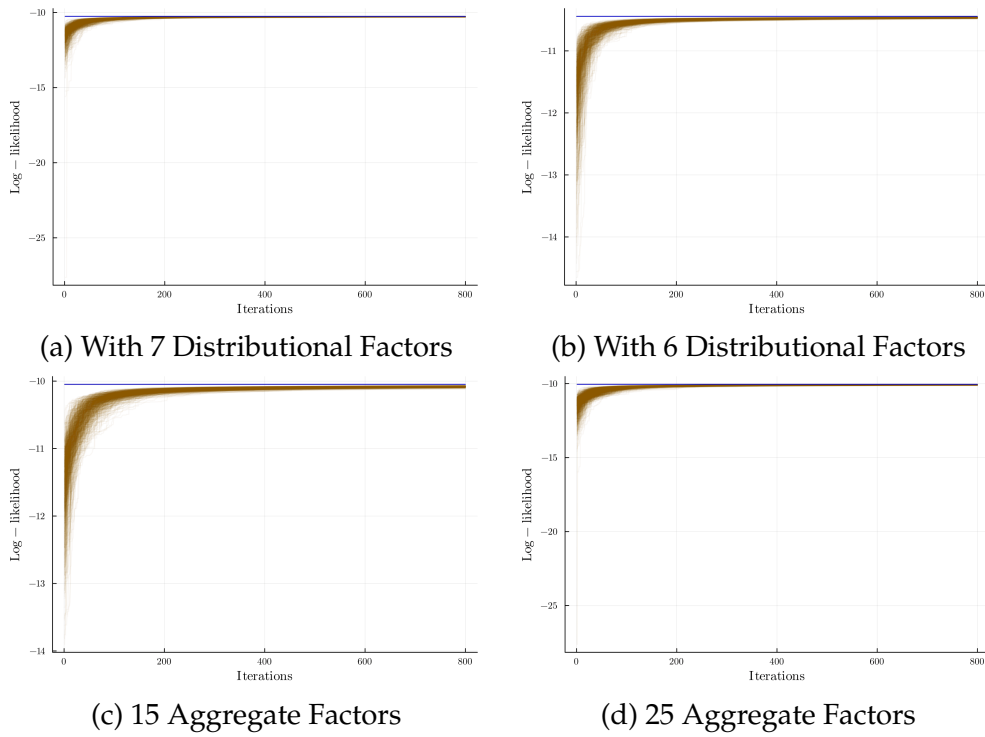
Figure 10 and 11 show traces of the (scaled) log-likelihood of all chains over the iterations. The plot clearly shows signs of convergence. The sampler implementation also returns the current log-weight on the history of the proposal distribution and the standard deviation of likelihoods. The log-weight measures how much the current ensemble of MCMC samples influences the proposal distribution. Early in the sampling process, the log-weight should be greater than zero, indicating that the current samples strongly influence the proposal distribution, allowing it to adapt to the target distribution. As the sampling progresses and the chains begin to converge, the log-weight should be very close to zero, indicating that the influence of the current samples decreases and the proposal distribution stabilizes. After the sampling, the log-weight is always very close to 0 (around $1e-7$ to be exact). Similar intuition can be applied for the interpretation of the standard deviation. Standard deviations are around 1% the size of the mode-likelihood. The acceptance rates are within acceptable range, suggesting that the sampler is effectively exploring the parameter space.

Figure 10: Converging Chains: Missing Data



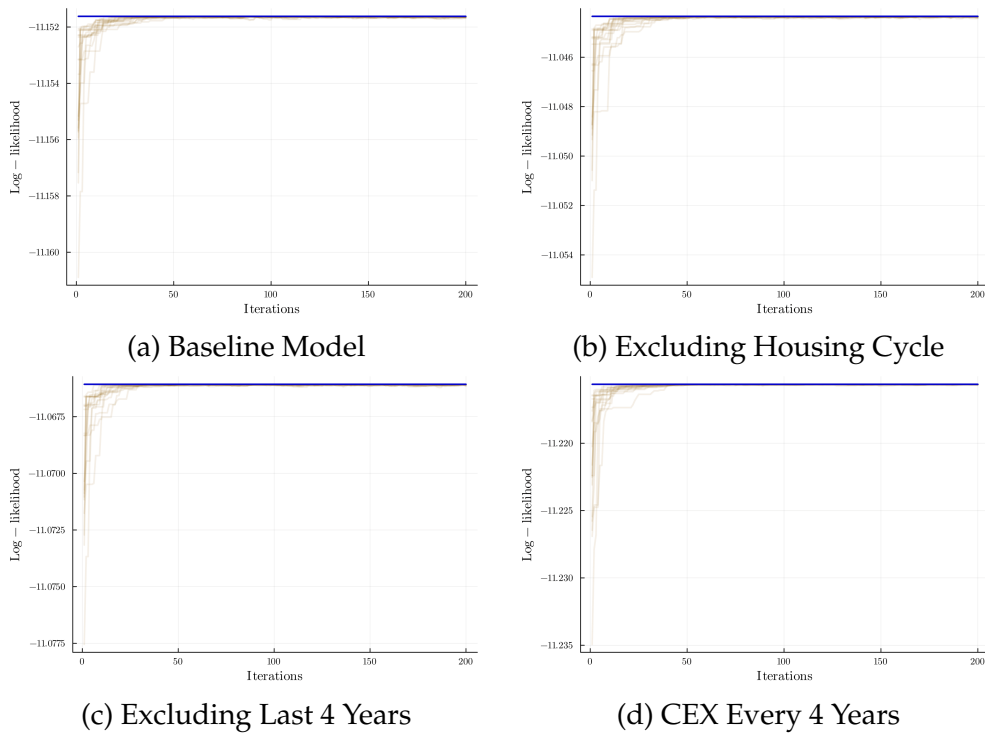
Notes: Figure shows, for each model with missing data, the evolution of the ensemble of chains in terms of log-likelihood. Log-likelihood values are scaled to ensure visualization of the chains' evolution. Last 25% of draws (of each chain) are kept as samples of the posterior. Refer to the text for the different model specifications.

Figure 11: Converging Chains: Different Factor Structures



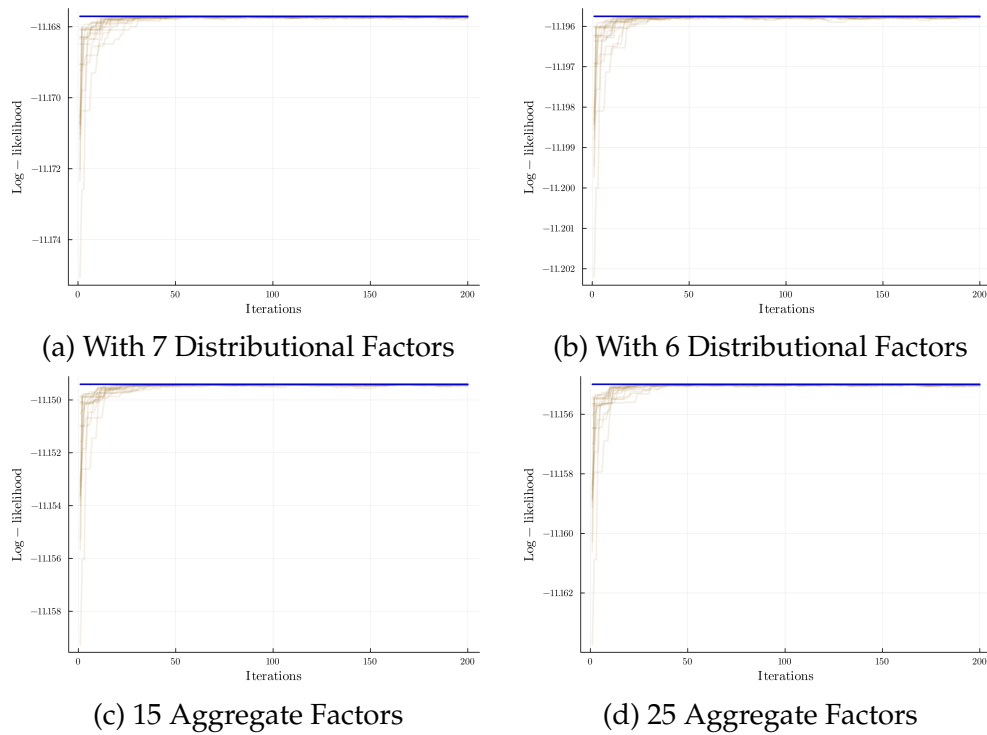
Notes: Figure shows, for each model of varying factors, the evolution of the ensemble of chains in terms of log-likelihood. Log-likelihood values are scaled to ensure visualization of the chains' evolution. Last 25% of draws (of each chain) are kept as samples of the posterior. Refer to the text for the different model specifications.

Figure 12: Hyperparameter Convergence: Missing Data



Notes: Figure shows, for each model with missing data, the evolution of the ensemble of chains in terms of log-likelihood. Log-likelihood values are scaled to ensure visualization of the chains' evolution. Last 25% of draws (of each chain) are kept as samples of the posterior. Refer to the text for the different model specifications.

Figure 13: Hyperparameter Convergence: Different Factor Structures



Notes: Figure shows, for each model of varying factors, the evolution of the ensemble of chains in terms of log-likelihood. Log-likelihood values are scaled to ensure visualization of the chains' evolution. Last 25% of draws (of each chain) are kept as samples of the posterior. Refer to the text for the different model specifications.

E Factor Selection

The estimation of the joint distribution of consumption, income, and wealth necessitates samples of this distribution and a comprehensive set of macroeconomic data to inform its dynamics. Macroeconomic theory and empirical validation suggest that the distribution and business cycle fluctuations are driven by a smaller set of underlying factors. Consequently, a significant area of macro-econometric research has produced several estimators to identify these factors and give recommendations on how many factors to retain to explain the data variation, under different settings.²⁸ It is additionally crucial to motivate the specific macroeconomic data we use, since, alongside the microdata, they will determine the cyclicity of the distributional data. We elaborate on this interdependency in the following section.

E.1 Choice of Factor Representation

The goal of the factor decomposition of distributional data is to inform the size of the state-space model and ensure that these factors can accurately reconstruct the cyclical movements observed in the data. As illustrated in Figure 3, the factor decomposition effectively reconstructs the data with minimal to no information loss. Consequently, we remain agnostic about the specific factors retained or their interpretation, opting to retain enough factors to replicate the data on average. Factors may consist of components that explain general or local movements within the distribution (Freyaldenhoven, 2022), possess eigenvalues greater than or less than 1, or induce weak to no cross-correlation in the unretained factors.²⁹ This observation underscores that solely retaining factors that explain common movements, have eigenvalues greater than 1, or focus on specific subsets of the factor space may inadequately capture the heterogeneity-rich cyclicity of consumption, income, and wealth, which is paramount in this study.

For the estimation of business cycle fluctuations, the selection of macroeconomic data must account for the rich heterogeneity present in the distributional data. The conventional approach to estimating business cycle fluctuations relies on the FRED-QD dataset—a common starting point of 200+ time series for exploratory factor analysis in macroeconomics. Many studies will then estimate the common component of these macroeconomic time series, consisting of a set of factors and their respective loadings, and define it as the most relevant movements in the macroeconomy. For a given estimator, the number of factors (in the common component) from projecting the FRED-QD dataset will vary and explain around 40 – 50% of

²⁸For references, see Ahn and Horenstein (2013), Bai and Ng (2002, 2019), Freyaldenhoven (2022), and Gagliardini, Ossola, and Scaillet (2019).

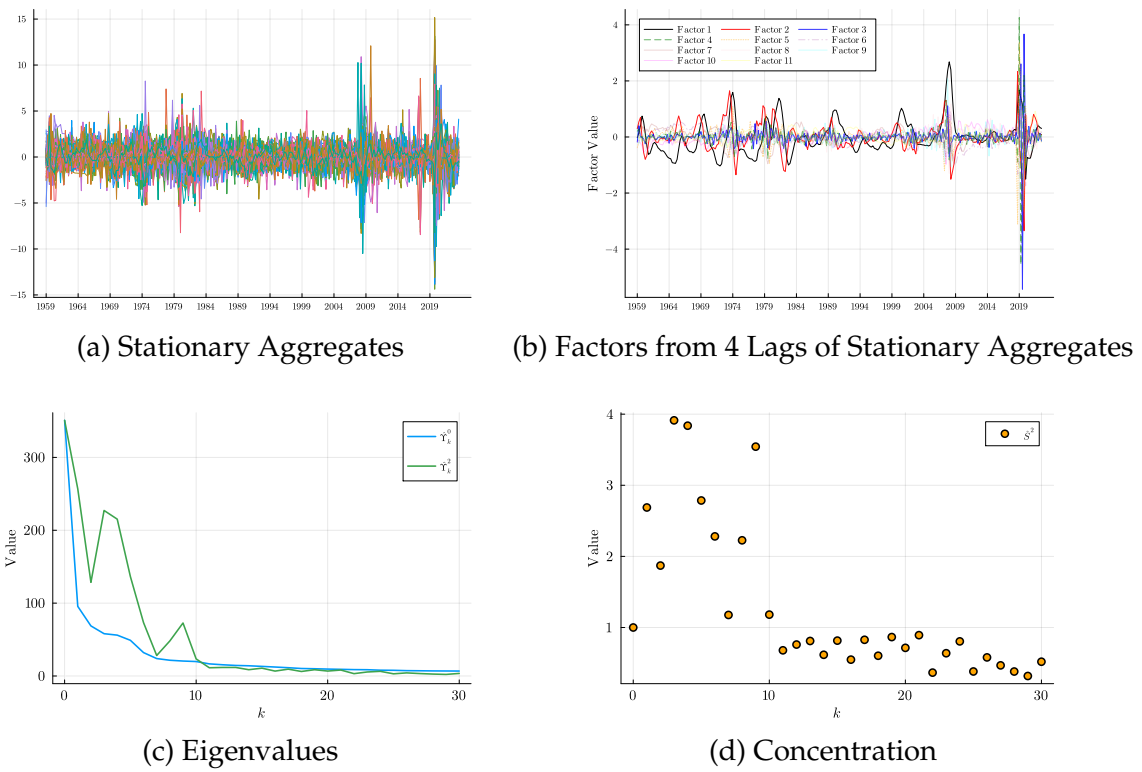
²⁹Our decomposition of the joint distribution into its correlational structure and marginal distributions implies the potential existence of local factors. Any local change in a bona fide distribution inherently represents a global change.

the (summable) data variation.

We adopt this approach, but use a more conservative estimator, which augments the standard estimated set of factors with *local* factors (Freyaldenhoven, 2022). These local factors only explain a subset of the data, so in this sense, they are *not* common; but carry nonetheless relatively large loadings on this subset. This approach would capture the most pervasive business cycle fluctuations as well as the granular movements, both potentially necessary to explain the distributional dynamics. Using this estimator, we find that six factors are sufficient to explain the cyclical movement in the aggregate data (Panel (a)).

However, to inform the dynamics of the distributional factors, we specify, on top, four lags of aggregate information. Figure 14, Panel (b) presents the factors from performing the PCA on these five quarters of data (over 1000 time series). Figure 14, Panel (c) plots the unweighted eigenvalues $\hat{\Upsilon}_k^0$ and the eigenvalues accounting for the contribution of the loadings $\hat{\Upsilon}_k^2$. Panel (d) plots \hat{S}^2 , which measures how concentrated the corresponding eigenvector is on its z largest entries. Taking into account these two plots, we find that around 10 factors are sufficient to explain the aggregate data, but around 20 to sufficiently capture everything. It is around 10, 21 and 25 factors that $\hat{\Upsilon}_k^2 \approx \hat{\Upsilon}_k^0$. We ultimately settle on testing 10, 15, 21, and 25 factors, since it is around this region that the concentration hovers around 1 before fading.

Figure 14: Eigenvalue Analysis of Aggregates



Notes: Figure shows an overview of the aggregate data. Panel (a) shows each aggregate series reduced to its stationary component. Panel (b) are the factors from stationary aggregate data at time t to $t - 4$. Panel (c) are the resulting eigenvalues weighted by the respective eigenvector loadings. Panel (d) are the weight contributions to the eigenvalues based on eigenvector loadings. Data used are from FRED-QD 1959Q1 to 2024Q1.

F Estimated Parameters

The parameters of the model are defined by the law of motion of states A , control variables B , the variance covariance matrix of the state process noise Ω and the scaling matrix S . The resulting parameter vector ψ from the estimated baseline model of consumption, income and wealth is summarized below. Parameter values for Ω and Δ are transformed (in the code) using a soft-plus transformation.

$$A = \begin{bmatrix} -0.0336949 & -0.0137318 & -0.0138072 & -0.00272868 & -0.00255301 & 0.00233189 & -0.00132588 & -0.000223374 \\ -0.133908 & 0.158993 & -0.00228447 & -0.00525042 & 0.003576 & -0.000414065 & 0.000928783 & 0.000110617 \\ -0.0187872 & -0.0189666 & -0.0422979 & -0.00229514 & 0.00107028 & 0.00289869 & -0.00196483 & 0.000492708 \\ 0.215591 & -0.101937 & -0.0197191 & 0.0568578 & 0.00639829 & -0.0127118 & 0.00510572 & 0.00419775 \\ 0.0535823 & -0.00423982 & 0.0488145 & -0.0229334 & 0.0969417 & -0.00900602 & -0.00343 & -0.00373173 \\ 0.116814 & 0.110363 & 0.000466101 & -0.016026 & 0.137886 & 0.186213 & -0.00189404 & 0.00182888 \\ -0.256287 & -0.0433632 & 0.126973 & -0.292554 & 0.205178 & -0.00238534 & 0.0746469 & -0.00958551 \\ 0.456505 & 0.0741462 & 0.29013 & 0.0636951 & -0.328791 & 0.0168262 & -0.00285531 & 0.0739768 \end{bmatrix}$$

$$B' = \begin{bmatrix} 0.012406 & 0.00815823 & -0.00518083 & -0.0532562 & -0.00356801 & 0.0153967 & 0.00828925 & -0.0197842 \\ 0.00481314 & 0.00608261 & -0.0292164 & 0.00320076 & 0.0213526 & -0.0360846 & 0.0268246 & -0.0813989 \\ -0.00905672 & 0.0257137 & -0.0147278 & -0.0179019 & -0.0313208 & 0.0876856 & 0.0314477 & -0.0611468 \\ -0.00482005 & 0.0198363 & 0.00905425 & 0.01084 & -0.0184017 & 0.0893193 & -0.0741054 & -0.156888 \\ -0.00738746 & 0.0151689 & -0.00695897 & 0.0309455 & -0.0366866 & -0.00781706 & -0.00527809 & -0.12059 \\ -0.00219059 & -0.0136928 & 0.00609896 & -0.00196025 & 0.0376343 & 0.0995637 & 0.0233585 & 0.103249 \\ -0.00831088 & 0.013039 & -0.0119729 & -0.0009803 & 0.0377599 & 0.0374526 & 0.0339219 & -0.0795743 \\ -0.00756725 & 0.00598843 & 0.00741306 & -0.0371864 & 0.0201494 & 0.106674 & 0.201247 & 0.024687 \\ 0.00318964 & -0.0177609 & 0.0134544 & 0.0000320882 & -0.0216779 & 0.0257637 & -0.00855473 & 0.193853 \\ -0.00469136 & -0.0131196 & 0.00733017 & -0.0223295 & 0.0333039 & -0.0676264 & 0.0293511 & -0.304486 \\ 0.00500845 & 0.000918726 & 0.00256625 & -0.00839652 & -0.0140791 & -0.116656 & 0.113768 & 0.00938289 \\ -0.00851789 & 0.00902402 & 0.0110001 & -0.0227373 & -0.000787335 & 0.0806602 & -0.00473101 & 0.00703768 \\ -0.0103545 & -0.0124117 & 0.00288321 & -0.0190253 & 0.00688511 & -0.0436931 & 0.054844 & -0.0379028 \\ 0.00456194 & 0.00362161 & 0.00241241 & -0.00715349 & -0.0367603 & -0.0445555 & -0.0182123 & 0.334412 \\ -0.00307826 & -0.00947954 & 0.00437902 & -0.0902256 & 0.0351224 & -0.0787602 & 0.131506 & -0.00415562 \\ 0.00412828 & 0.0145638 & 0.0156836 & 0.106691 & -0.00401169 & 0.040731 & 0.0301459 & 0.0431524 \\ 0.00693329 & 0.00313172 & 0.0171261 & 0.0599548 & 0.0195106 & 0.047194 & -0.154379 & -0.0387055 \\ -0.00984089 & -0.010498 & 0.00443422 & 0.0302647 & -0.00654655 & -0.0327508 & 0.0401239 & -0.048429 \\ 0.00270008 & 0.00336171 & 0.0201289 & 0.0218016 & -0.0304612 & 0.0182255 & 0.038053 & 0.0216766 \\ 0.00457309 & -0.00159936 & -0.0106156 & 0.0287034 & 0.00231657 & -0.0135371 & 0.0154398 & 0.192843 \\ -0.00894874 & -0.00588382 & 0.0119702 & 0.0426665 & 0.0000728098 & 0.14174 & -0.0777418 & 0.188022 \end{bmatrix}$$

$$\begin{aligned}
diag(\Omega) = & \begin{pmatrix} 0.476533 \\ 0.476878 \\ 0.476613 \\ 0.476863 \\ 0.477093 \\ 0.476164 \\ 0.476253 \\ 0.476779 \end{pmatrix} \\
diag(\Delta) = & \begin{pmatrix} \text{CEX copula} - 1.66422 \\ \text{CEX consumption} - 3.31417 \\ \text{CEX income} - 0.0685886 \\ \text{CPS1 income} - 18.0599 \\ \text{CPS2 income} - 13.253 \\ \text{PSID copula} - -3.90535 \\ \text{PSID consumption} - -4.69801 \\ \text{PSID income} - 0.327083 \\ \text{PSID wealth} - -3.33474 \\ \text{SCF copula} - -0.747217 \\ \text{SCF income} - 3.04712 \\ \text{SCF wealth} - -1.16332 \\ \text{SIPP1 copula} - 1.4332 \\ \text{SIPP1 income} - 21.5057 \\ \text{SIPP1 wealth} - 1.89892 \\ \text{SIPP2 copula} - 0.670052 \\ \text{SIPP2 income} - 5.0387 \\ \text{SIPP2 wealth} - 12.4339 \end{pmatrix}
\end{aligned}$$

G Correlations

The following tables provide the correlations of the synthetic data with external sources and the actual data when we leave out some datasets in the construction of the synthetic data.

Table 6: Correlations with External Estimators

	Baseline-WID	Baseline-DFA	WID-DFA
Bottom			
Income	0.77	-	-
Wealth	0.45	0.86	0.28
Middle			
Income	0.4	-	-
Wealth	0.23	0.62	0.51
Top			
Income	0.72	-	-
Wealth	0.69	0.98	0.76

Notes: The table reports correlations between the **SCF** model-implied cycle estimates of the baseline model, the WID and the DFA. Cycle estimates are the difference between the raw series (in logs) and a series obtained by an HP-filter with smoothing parameter $\lambda = 6$ for annual data and $\lambda = 1,600$ for quarterly data. Bottom, Middle, and Top correspond to the Bottom 50, Next 40, and Top 10 of the income/wealth distribution.

Table 7: Out of Sample Performance

Excluding Housing Cycle									
Condition	Bottom			Middle			Top		
	C	I	W	C	I	W	C	I	W
Entire Series	-	0.97	0.97	-	0.96	0.89	-	0.98	0.91
Specific Timeframe	-	0.59	0.83	-	0.76	0.51	-	0.76	0.64

Excluding Housing Cycle #2									
Condition	Bottom			Middle			Top		
	C	I	W	C	I	W	C	I	W
Entire Series	-	0.98	0.96	-	0.96	0.93	-	0.98	0.95
Specific Timeframe	-	0.97	0.94	-	0.9	0.84	-	0.97	0.92

Excluding Last 4 Years									
Condition	Bottom			Middle			Top		
	C	I	W	C	I	W	C	I	W
Entire Series	-	0.97	0.9	-	0.94	0.88	-	0.97	0.86
Specific Timeframe	-	0.93	0.69	-	0.63	0.88	-	0.88	0.83

Every 4 Years									
Condition	Bottom			Middle			Top		
	C	I	W	C	I	W	C	I	W
Entire Series	0.7	0.7	-	0.73	0.66	-	0.7	0.71	-
Specific Timeframe	0.75	0.7	-	0.75	0.65	-	0.72	0.7	-

Notes: Table reports correlations between the baseline cyclical estimates and the missing-data models, split by panel. Entire Series reports correlations for all estimates in the estimation timeframe. Specific Timeframe reports correlations of estimates in periods where data was intentionally left out of the estimation of the specific missing-data model. **Bottom**, **Middle**, and **Top** are the bottom 50, next 40, and top 10 of the respective distribution, denoted by **I**, **W**, and **C**. **I** is for income, **W** is for wealth, and **C** is for consumption. For the first three panels, correlations are made between SCF model estimates (no consumption). The final panel presents correlations from CEX model estimates (no wealth). Specific models (in header) are discussed in Section 4.

H Consumption Dynamics: pre-2000s Recessions

In Section 5, we presented consumption dynamics along the income and wealth distribution for the three most recent recessions. In this appendix, we provide consumption dynamics for all the earlier recessions covered by our microdata. This demonstrates the power of the approach — we can also generate estimates *outside* (versus only in-between) the sampling periods of the survey data. For example, PSID consumption begins in 1999, but with aggregate and other microdata data, we can generate consumption estimates back to 1962Q3.³⁰ Here, we consider the four recessions pre-2000s: (1) the Gulf War Recession starting in 1990Q3, (2) the Double-Dip Recession with its second dip in 1981Q3, (3) the Oil Crises in 1973Q4, and the (4) Unemployment Recession in 1969Q4. The consumption dynamics are again by different income, wealth and joint income and wealth groups and are shown in Figure 15. Since these estimates do not include a trend component, the plots only show the consumption of the respective group relative to consumption at the beginning of the recession.

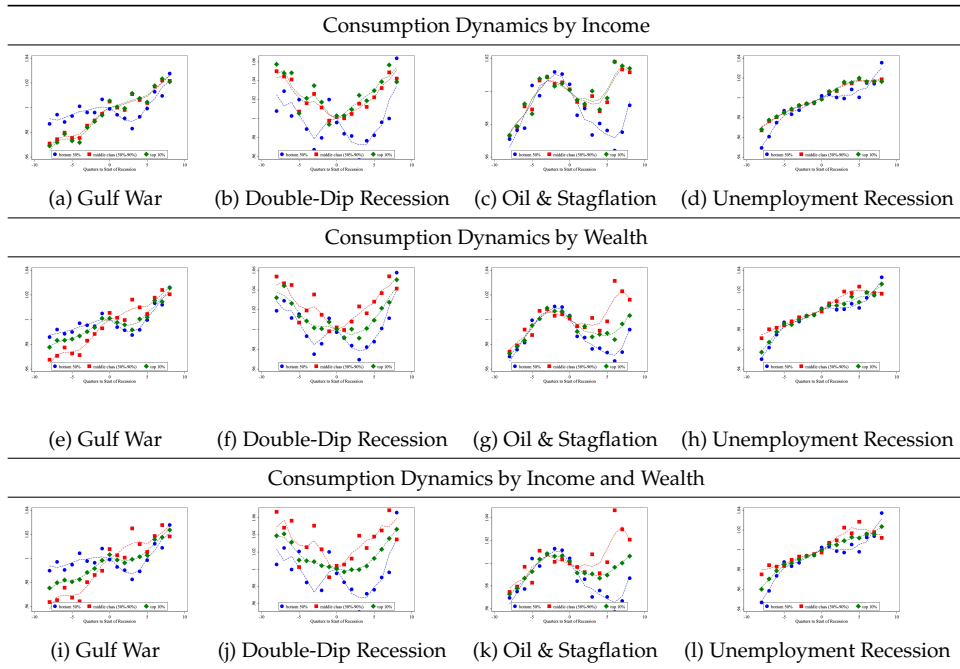
The results in Figure 15 show distinct patterns in consumption dynamics during four pre-2000 recessions. During the first year of the Gulf War Recession, panel (a) shows higher-income households were largely unaffected. They were able to smooth their consumption over general price increases arising from oil price hikes; however, the supply chain disruption propagated into financial markets, impacting the consumption of higher-wealth households through asset price fluctuations, as shown in panel (e), though having higher income mutes the consumption dip (panel (i)).

In the Double-Dip Recession, the U.S. experienced a short decline in economic output followed by a short recovery, only again to be followed by a steeper, longer period of economic contraction. Economic historians point to double-digit inflation, contraction monetary policy (20% FFR) as a result, and bank deregulation as an interaction, all contributing to historically high unemployment. As shown in panels (b), (f), and (j), the bottom half of each distribution were unable to smooth through these effects, experiencing both dips in consumption. The top half of each distribution, relative to peak, saw declines in consumption around the first dip (before 0), but quickly recovered around the second dip.

The Oil Crises in the early 1970s pushed the U.S. economy into an unprecedented situation. Disruptions in the oil supply chain caused the price of crude oil to quintuple, an event unprecedented in the past century. Global stock indices plunged, losing roughly 40% of their value. Economic growth stalled, and unemployment almost doubled to around 9%. As soaring oil prices fueled inflation across the economy and unemployment surged, the standard approach

³⁰The state-space model generates estimates for the cyclical component of the distribution from 1962Q3 to 2024Q1. Time variation is truncated if one wishes to add the trend component, which is estimated outside the model and specific to the dataset of choice.

Figure 15: Consumption Dynamics during Recessions (pre-2000s)



Notes: Figure shows plots of relative consumption dynamics during recessions along the income and wealth distribution. Consumption dynamics of each income (wealth) group are indexed to the beginning of the recession. **These series are not relative to the average household as in Figure 9.** The horizontal axis shows the time relative to the start of the recession. The recessions are the Gulf War Recession beginning in 1990Q3, the second dip of the Double-Dip Recession beginning in 1981Q3, the Oil Crises / Stagflation episode beginning in 1973Q4, and the Unemployment Recession beginning in 1969Q4.

of tightening monetary policy to combat inflation risked worsening the downturn. This rare combination of high inflation and stagnant growth became known as stagflation. How did households cope? No group could fully insulate itself from these shocks, with the bottom half of income and wealth distributions suffering the most severe impacts. However, middle-class households recovered relatively quickly, while higher-wealth households appeared to experience a slower rebound (panel (g)).

The last recession we examine is the Unemployment Recession, characterized by an increase in unemployment from 3.5% to 6%. Among the recessions covered, it was the mildest, with a relatively small rise in unemployment and a brief downturn in financial markets. Recovery was rapid, with consumption showing a slight stagnation compared to pre-recession levels but rebounding quickly, particularly for the bottom half of each income and wealth distribution.