

Nonlinear Micro Income Processes with Macro Shocks^{*}

MARTÍN ALMUZARA[†]

MANUEL ARELLANO[§]

RICHARD BLUNDELL[‡]

STÉPHANE BONHOMME[¶]

June 2026

[\[Link to supplemental appendix\]](#)

Abstract

We propose a nonlinear framework that exploits both macro and micro data to study the dynamic transmission of aggregate and idiosyncratic shocks to household income. We develop identification arguments and estimation methods that combine macro time series and time series of panels. Using U.S. aggregate data and the PSID, we find that business-cycle fluctuations modulate the persistence of heterogeneous individual histories and the risk faced by households. We document how aggregate and idiosyncratic shocks propagate over time for households in different macro and micro states. Lastly, we quantify the welfare cost of aggregate and idiosyncratic risk.

Keywords: Income process, business cycle, persistence, exposure to aggregate shocks.

^{*}We have greatly benefited from discussions with Anmol Bhandari, Florin Bilbiie, Mariacristina De Nardi, Gianluca Violante, Jonathan Heathcote, Luigi Pistaferri, and Frank Schorfheide, and with audience members at various venues. Arellano acknowledges financial support from grant PID2024-158202NB-I00 funded by MICIU/AEI/10.13039/501100011033 and ERDF/EU. Blundell acknowledges support from the ESRC Centre for the Microeconomic Analysis of Public Policy (CPP) at IFS, grant number ES/M010147/1. The views expressed in this paper do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

[†]FRBNY: martin.almuzara@ny.frb.org

[‡]UCL and IFS: r.blundell@ucl.ac.uk

[§]CEMFI: arellano@cemfi.es

[¶]University of Chicago: sbonhomme@chicago.edu

1 Introduction

In this paper, we propose a nonlinear framework to study the dynamic transmission of aggregate and idiosyncratic shocks to income by leveraging both macro and micro data. Our approach makes it possible to empirically examine how business-cycle fluctuations modulate the persistence of heterogeneous individual histories and the risk faced by households. We also consider questions such as how aggregate and idiosyncratic shocks propagate over time for units in different macro and micro states, and how these shocks contribute to the cost of business-cycle risk. Answering these questions is important. They are essential to documenting the dynamics of income inequality over the business cycle. Furthermore, how the incomes of heterogeneous agents respond to macro and micro shocks is key for consumer and firm behavior, and for the design of optimal monetary and fiscal policies (Bhandari, Evans, Golosov, and Sargent, 2021, Acharya, Challe, and Dogra, 2023).

The literature on income risk has uncovered significant nonlinearities in the dynamics of individual incomes (Arellano, Blundell, and Bonhomme, 2017; Guvenen, Karahan, Özkan, and Song, 2021) and in their variation over the business cycle (Guvenen, Ozkan, and Song, 2014). Moreover, a growing recent literature investigates the heterogeneous effects of monetary policy shocks on individual-level outcomes (Holm, Paul, and Tischbirek, 2021; Andersen, Johannesen, Jørgensen, and Peydró, 2023; Amberg, Jansson, Klein, and Rogantini Picco, 2022). Yet, a methodology for modeling the interaction between micro and macro shocks capable of integrating nonlinearities in the life-cycle and business-cycle dynamics of income is still lacking. This is our main contribution.

We consider a nonlinear Markovian micro process for (log) persistent income η_{it} with a macro state variable Z_t of the following form:

$$\begin{aligned}\eta_{it} &= Q_\eta(\eta_{i,t-1}, Z_t, Z_{t-1}, u_{it}), \\ Z_t &= Q_Z(Z_{t-1}, V_t),\end{aligned}$$

where u_{it} and V_t are micro and macro shocks, and η_{it} and Z_t are potentially unobserved. A measurement system connects these two latent variables to observed micro and macro data, specifically, a flexible persistent-transitory model for the micro states and a dynamic factor model for the macro states. Our triangular formulation has the potential to allow for feedback from the micro to the macro level, as Z_t can incorporate distributional characteristics of the micro data. However, we assume that unit-level shocks u_{it} are atomistic, in the sense that

they are independent of aggregate shocks V_t at all leads and lags.

Based on our income process we will highlight two key quantities. The first one is income persistence:

$$\rho_{it} = \frac{\partial \eta_{it}}{\partial \eta_{i,t-1}} \equiv \frac{\partial Q_\eta(\eta_{i,t-1}, Z_t, Z_{t-1}, u_{it})}{\partial \eta_{i,t-1}}.$$

Here, ρ_{it} is a measure of nonlinear persistence (Arellano et al., 2017, ABB) that may vary depending on the position in the income distribution and the idiosyncratic shocks hitting the household. Moreover, unlike in ABB, our model allows for the aggregate state to affect the whole income process and, thus, for the shape of persistence to be different in good or bad times.

The second quantity is the elasticity of individual persistent income to the aggregate business-cycle state Z_t :

$$\beta_{it} = \frac{\partial \eta_{it}}{\partial Z_t} \equiv \frac{\partial Q_\eta(\eta_{i,t-1}, Z_t, Z_{t-1}, u_{it})}{\partial Z_t},$$

where the derivative holds Z_{t-1} fixed because we are interested in the “shock” component of Z_t . In our setup, β_{it} is a measure of a household’s exposure to shocks to the aggregate state that is heterogeneous along both the income distribution and business-cycle conditions. In addition, β_{it} may vary with the idiosyncratic shock u_{it} , and so the impact of an aggregate shock may differ depending on idiosyncratic events such as a job loss or a promotion.¹

Documenting how ρ_{it} s and β_{it} s vary across income histories and over the business cycle allows us to paint a rich picture of the interaction between individual and aggregate income dynamics. Our model also allows us to flexibly measure how features of the income process such as (conditional) dispersion and skewness depend on business-cycle conditions. Since the model completely specifies the law of motion of individual income η_{it} and the aggregate state Z_t , it can be used for impulse response analysis and to quantify the cost of both micro and macro sources of income risk, as we illustrate empirically.

We study the nonparametric identification of this model by bringing together macroeconometric and microeconomic techniques. The micro side builds on ABB. On the macro

¹Our measure is different from, but related to, others in the growing empirical literature on the heterogeneous effects of recessions. For example, Guvenen, Schulhofer-Wohl, Song, and Yogo (2017) document variation in individual income exposure to aggregate income—referred to as “worker betas”—across age and income groups, while Patterson (2023) examines how these heterogeneous exposures relate to differences in marginal propensities to consume across demographic groups.

side we use techniques for factor models estimated from time series aggregates to recover a measure of the aggregate state Z_t (Stock and Watson, 2016). To combine the two parts, we create a time series of short panels, such that each panel is representative of the economy in that period. Thus, identification relies on three dimensions of our combined dataset: large cross-sections, short individual panels, and long time series spanning sufficient cyclical fluctuations.

We propose an approach to estimation and inference that uses a flexible parametric version of the model and can be implemented with stable simulation-based algorithms. This approach was first introduced in Arellano and Bonhomme (2016) and was adapted to a setup with time-varying latent variables in ABB. Here it is further extended to a long time series of short panels involving both micro and macro latent variables. Our stochastic EM algorithm iterates between draws of latent variables from their posterior distributions evaluated at current parameter values, and updates of parameters from regressions based on those draws.

We also develop a methodology for impulse response function analysis in our nonlinear context where we want to measure the importance of macro and micro shocks and their interactions. We start by considering an experiment in which we directly perturb a state variable at some point in time. We then compare the trajectory of the system following the perturbation with a baseline trajectory in the absence of perturbation. Since we wish to obtain comparability of impulse responses across households with different income processes, we need to consider ways of introducing comparable perturbations. We do so via a set of *rules* that map perturbations to a common system of units, and we show that different perturbation experiments can be associated with different formulations of local shocks.

We take our model to quarterly macro time series data for the U.S. and a time series of panels that we construct from the Panel Study of Income Dynamics (PSID) spanning the period 1970-2019, thus covering seven recessions. PSID waves are annual through 1997 and biennial afterwards. For consistency of the microdata, we then form sequences of biennial subpanels of four waves each covering all available years, in the spirit of Storesletten, Telmer, and Yaron (2004). Each subpanel has its own heterogeneous initial conditions. Compared to standard long panel approaches, this *time series of panels* approach has the advantage of mitigating concerns over the representativeness of the data, and of facilitating a transparent analysis of identification based on macro and micro data. In this paper, our primary focus is on disposable household income net of taxes and transfers. However, for comparability with other studies, we also present estimates based on male earnings and household earnings

before taxes and transfers.²

Empirical results. Our analysis yields a number of novel empirical insights. To begin, our results illustrate the asymmetric impact of the business cycle on income persistence. Specifically, income persistence ρ_{it} increases for low-income households and decreases for high-income households during recessions. That is, during a downturn, it is harder for low-income earners to leave the low-income state, whereas for high-income earners remaining high-income becomes more difficult. The cyclical variation we find in the persistence of past income histories coexists with the ample variability along income and micro-shock distributions uncovered by ABB.

Our results also highlight the presence of heterogeneous exposures to aggregate shocks. The coefficients β_{it} tend to be higher when associated with bad idiosyncratic shocks. Moreover, we find that β_{it} s are *countercyclical*: they are higher in recessions and lower in expansions. This is important because the cyclical behavior of income elasticities to the macro state—particularly the self-amplifying nature of negative aggregate shocks—has major implications for the cost of aggregate income risk, as we argue below.

In addition, we document two main facts about income skewness. First, we find that left skewness is countercyclical, consistent with the findings in [Guvenen et al. \(2014\)](#). Second, we find that skewness decreases with income at any point in time, consistent with the findings in ABB, but differentially so depending on the phase of the business cycle. This *tale of two skewnesses* is a clear reflection of the nonlinear transmission of micro and macro shocks. In this and other dimensions, business-cycle variability is most pronounced for male earnings and, to a lesser extent, for household earnings and disposable income, in that order—that is, from the income measure with the least insurance to the one with the most—but the patterns are qualitatively similar for all three measures.

Concerning impulse response functions, we find a large direct impact of macro shocks but with generally short-lived effects. These impacts, however, are highly heterogeneous across income measures (larger for male earnings, smaller for disposable income) and along the persistent income distribution (larger at the bottom, smaller in the middle), and they interact with idiosyncratic uncertainty by compounding the negative consequences of bad micro shocks. In contrast, micro impulse responses decay slowly, with different degrees of

²Earnings-based results are limited to units with positive earnings. The share of households with zero earnings stays below 2% for most of the period, rising slightly during the Great Recession. For men, zero earners remain under 6% until the Great Recession, when the rate nears 8%. See [Figure 2](#) and its discussion.

persistence depending on the initial level of income. In fact, as we show, there is a tight link between impulse responses and β_{it} s and ρ_{it} s.

We explore heterogeneity of our estimates along two dimensions: age and education. Our model allows for flexible dependence on the former, and we show that the cyclical patterns of persistence and exposure to aggregate shocks are present at all age groups. However, while our baseline model allows for rich, subpanel-specific heterogeneity in initial conditions, it assumes that the transition function of persistent income is homogeneous across households. We extend our specification and separately estimate our model for two education groups (high-school and college education), and find that the patterns uncovered in our main specification are reproduced within these education groups.

Finally, we use our framework to quantify the cost of both macro and micro sources of risk. For macro risk, we compute the (percent) compensating variation that equalizes the expected lifetime utility of income with and without macro shocks (as in [Lucas, 1987, 2003](#)), where (rational) expectations are based on our estimated income process. In this exercise, the cyclical behavior of the nonlinear income exposure to macro shocks β_{it} is a key determinant of the cost of business cycles. In the presence of empirically plausible amplification effects, the yearly cost of macro risk can be as high as 4.5 percent of income, whereas it is negligible under a linear transmission of macro shocks. This is a novel channel through which macro uncertainty can lead to welfare losses at the household level, distinct from curvature in the utility function. Although, as expected, the cost of micro risk is higher, for many units (mainly young and low-income) macro shocks explain a nontrivial fraction of the total cost of risk. It is worth emphasizing that an income process linear in aggregate shocks would miss all of the rich cyclical patterns documented here. Thus, a core lesson from our paper is the importance of accounting for nonlinearities in β_{it} when the goal is to study the consequences of aggregate fluctuations.

Selected literature. Our paper contributes to several strands of the literature. First, we build on the vast literature on income dynamics, both with and without business cycles, e.g., [Gottschalk and Moffitt \(1994, 2009\)](#), [Meghir and Pistaferri \(2004\)](#), [Blundell, Pistaferri, and Preston \(2008\)](#), [Browning, Ejrnæs, and Álvarez \(2010\)](#), [Altonji, Smith, and Vidangos \(2013\)](#); see [Blundell, Bollinger, Hokayem, and Ziliak \(2024\)](#) for a comprehensive review. Within that body of work, our paper is most closely related to [Storesletten et al. \(2004\)](#), [Guvenen et al. \(2014\)](#), [Arellano et al. \(2017\)](#), [Guvenen, McKay, and Ryan \(2023\)](#), [Halvorsen, Holter, Ozkan, and Storesletten \(2024\)](#), [Guvenen, Pistaferri, and Violante \(2022\)](#), and the

multi-country GRID project. Relative to this work, we are the first to develop a framework capable of integrating aggregate shocks and rich nonlinear dynamics at the micro level.

Second, we contribute to the literature on estimating heterogeneous agents models using micro data (Arellano and Bonhomme, 2017; Liu and Plagborg-Møller, 2023; Fernández-Villaverde, Hurtado, and Nuño, 2023). Compared to them, we offer a principled approach to building nonlinear reduced forms for structural models when agents face potentially latent macro-level uncertainty. More generally, we add to an early literature on the combination of household survey data with time series data (Tobin, 1950; Chetty, 1968; Maddala, 1971) and to recent work on the econometrics of models with aggregate shocks (Hahn, Kuersteiner, and Mazzocco, 2020; Almuzara and Sancibrián, 2024) by developing novel tools for identification analysis and estimation in a time series of panels framework.

An important reference is Chang, Chen, and Schorfheide (2024), who propose functional vector autoregression methods to combine macro aggregates with repeated cross-sections. In a setup with both macro and micro data, the effects of macro shocks at the unit level and their impact on cross-sectional distributions are distinct but related empirical objects. While we focus on the former, Chang et al. (2024) aim at the latter. From this perspective, our papers are complementary. Another recent related reference is Sargent and Selvakumar (2025), who propose a dynamic mode decomposition method to study distributional dynamics of income and consumption. Compared to both papers, a distinctive feature of our work is the use of panel data to analyze household dynamics.

Lastly, our empirical results speak to a vast macro literature on inequality and aggregate fluctuations (e.g., Krusell and Smith, 1998; Krueger, Mitman, and Perri, 2016; Ahn, Kaplan, Moll, Winberry, and Wolf, 2018; Bhandari et al., 2021), and to work that seeks to quantify the welfare cost of business cycles (Lucas, 1987, 2003; Storesletten, Telmer, and Yaron, 2001; Otrok, 2001; Barlevy, 2004; Galí, Gertler, and López-Salido, 2007; Krebs, 2007). In particular, we document new empirical patterns about the interaction between cyclical variation and nonlinearities in the income processes that have the potential to amplify the welfare consequences of recessions and expansions.

Outline. The paper is organized as follows. Section 2 outlines our framework of analysis and the quantities of empirical interest. Section 3 introduces the statistical population by means of a time series of panels and studies identification. Section 4 details our estimation strategy. In Section 5, we present empirical results on nonlinear persistence, exposure to aggregate shocks, and skewness over the business cycle. Section 6 develops methodology for

nonlinear impulse response functions of macro and micro shocks, while Section 7 focuses on quantifying the idiosyncratic and aggregate components of income risk. Section 8 discusses estimates by age and education groups. Finally, Section 9 concludes. Additional material can be found in the Supplemental Appendix.

2 Framework

In this section, we describe our framework of analysis that combines time series aggregates with longitudinal micro-level survey data to analyze the impact of macro shocks on nonlinear income processes.

2.1 Model

We model log income y_{it} of household i at time t as the sum of a persistent component η_{it} and a transitory component ε_{it} ,

$$y_{it} = \eta_{it} + \varepsilon_{it}.$$

While the researcher observes log income y_{it} , the two components η_{it} and ε_{it} are latent.³

Following Arellano et al. (2017, ABB), the persistent component is a flexible first-order Markov process and the transitory component is serially independent (at biennial frequency). However, unlike ABB, our focus is on understanding how aggregate conditions affect income trajectories. For this purpose, we introduce a time series aggregate Z_t , a macro state variable that affects η_{it} and ε_{it} .

In our empirical analysis we allow both processes to depend on observable heterogeneity x_{it} (for example, we include age to capture life-cycle patterns and, in some cases, education) but we abstract from them here for ease of exposition. Section 8 examines age and education heterogeneity.

Persistent component. We specify the persistent component as

$$\eta_{it} = Q_{\eta}(\eta_{i,t-1}, Z_t, Z_{t-1}, u_{it}). \tag{1}$$

³We build on a long tradition that uses panel data to decompose income into persistent and transitory components (Hause, 1977; Lillard and Willis, 1978; MaCurdy, 1982); see Arellano (2014) for a survey.

In (1), Q_η is strictly increasing in its last argument, and the idiosyncratic shock u_{it} is i.i.d. uniform on $(0, 1)$ and independent of $(\eta_{i,t-1}, Z_t, Z_{t-1})$. Therefore, for all $\tau \in (0, 1)$, $Q_\eta(\eta_{i,t-1}, Z_t, Z_{t-1}, \tau)$ is the conditional τ -quantile of η_{it} . This model allows for a general nonlinear relationship between the persistent income component η_{it} and its various determinants: lagged income $\eta_{i,t-1}$, the aggregate factor Z_t , and the idiosyncratic shock u_{it} .

To aid interpretation, it is useful to consider the following special case of Equation (1):

$$\eta_{it} = \rho\eta_{i,t-1} + \gamma\eta_{i,t-1}Z_t + \delta Z_t + \underbrace{g(Z_t, u_{it})}_{=\zeta_{it}}, \quad (1')$$

where we have omitted the dependence on Z_{t-1} , which can be thought of as fixed in the discussion. The quantity $\zeta_{it} = g(Z_t, u_{it})$ is a *composite income shock*, measured in log income units, which exhibits aggregate variation due to the presence of Z_t and idiosyncratic variation driven by u_{it} .

In model (1'), the impact of a marginal change in the aggregate state Z_t on the persistent component of income (the household's aggregate exposure β_{it}) can be decomposed as

$$\frac{\partial \eta_{it}}{\partial Z_t} = \underbrace{\gamma\eta_{i,t-1} + \delta}_{\text{income heterogeneity}} + \underbrace{\frac{\partial g(Z_t, u_{it})}{\partial Z_t}}_{\text{shock distribution}},$$

where, implicitly, Z_{t-1} is held fixed so the derivative reflects variation in the aggregate shock V_t . The first term shows that the impact of an aggregate shock can vary along the income distribution, a point highlighted in an extensive literature in macroeconomics; the second shows that aggregate shocks may alter the distribution of income shocks, e.g., by affecting their variance (as in [Storesletten et al., 2004](#)) or skewness (as in [Guisen et al., 2014](#)).⁴

Compared to (1'), which permits interactions between Z_t and past income $\eta_{i,t-1}$ and between Z_t and the shock u_{it} , model (1) allows for a third type of interaction between past income $\eta_{i,t-1}$ and the shock u_{it} . Model (1) is therefore better able to capture nonlinearities in income persistence, including the observation from ABB that different shocks to the persistent component may be associated with different degrees of persistence. In addition, model (1) can generate a significantly richer relationship between macro and micro responses, which is key to empirically fit the shape of individual impulse responses to aggregate shocks.

⁴For a special case consider $g(Z_t, u_{it}) = \sigma(Z_t)q(u_{it})$, where $\frac{\partial g(Z_t, u_{it})}{\partial Z_t} = \frac{\partial \sigma(Z_t)}{\partial Z_t}q(u_{it})$ reflects the impact of Z_t on the variance of the composite shock distribution $\sigma^2(Z_t)$.

Transitory component and initial condition. We specify the transitory component as

$$\varepsilon_{it} = Q_{\varepsilon,t}(v_{it}), \quad (2)$$

with $Q_{\varepsilon,t}$ strictly increasing, and the shock v_{it} i.i.d. uniform on $(0, 1)$ and independent of u_{it} . The function $Q_{\varepsilon,t}$ may vary over time in unrestricted ways reflecting general aggregate effects on transitory income shocks (coming from Z_t or other factors). Our specification can also accommodate non-Gaussianity in the density of transitory shocks. In practice, ε_{it} will likely be a mix of substantive transitory shocks and measurement error, and without further assumptions our approach will not allow us to distinguish between the two. For this reason, in our empirical analysis we mostly focus on the properties of the persistent component.

Lastly, we specify the individual initial condition of the persistent income process as

$$\eta_{i,t_0} = Q_{\text{init},t_0}(\nu_{i,t_0}),$$

with Q_{init,t_0} strictly increasing, and ν_{i,t_0} uniformly distributed on $(0, 1)$ and independent of u_{it} and v_{it} . With a similar rationale as for $Q_{\varepsilon,t}$, we let Q_{init,t_0} depend flexibly on the initial time period t_0 (which may differ across individuals).

Since we permit general time-variation in both the transitory component and the initial condition, other factors beyond the business cycle Z_t can influence the evolution of income in our model. Moreover, we will develop an identification framework (in Section 3) where we follow individual units only for short periods, and the flexible specification of initial conditions Q_{init,t_0} in each subpanel will be important to account for *ex-ante* forms of unobservable heterogeneity across households. Thus, as $Q_{\varepsilon,t}$ and Q_{init,t_0} reflect a mix of business-cycle and other factors, we interpret them as rich, possibly nonstationary time-varying “controls”, with our primary goal being instead to document the dynamics of the persistent component captured by Q_η .

Macro state. We model the aggregate state variable Z_t as a stationary first-order Markov process of the form

$$Z_t = Q_Z(Z_{t-1}, V_t), \quad (3)$$

where V_t is a vector of aggregate shocks, i.i.d. over t . In this formulation, the aggregate state Z_t will typically display persistence, reflecting at each point in time the dynamic impact of

current and past shocks $\{V_\tau\}_{\tau \leq t}$.⁵ In what follows, we assume that the researcher observes Z_t directly, although we later discuss extensions of our framework where Z_t is unobservable and has to be inferred from macro data, and we will account for Z_t being latent and estimated in our empirical analysis.

Since our interest is in how the business cycle interacts with individual dynamics, we will infer a one-dimensional state of the business cycle from a collection of macroeconomic indicators (including GDP, investment and unemployment).⁶ However, our methodology allows for a multivariate Z_t that includes additional macro states (e.g., monetary and financial factors) and even summaries of micro-level distributions (e.g., wealth as in heterogeneous agents models). The limit is statistical rather than conceptual, as a higher-dimensional Z_t can quickly exhaust the degrees of freedom offered by a short time series. Our focus on nonlinear effects, which we will show are key to quantifying aggregate risk, puts additional demands on the data and further motivates considering a scalar Z_t .

We can view Equations (1) and (3) as a nonlinear first-order vector autoregressive (VAR) model for $\{\eta_{it}\}_i$ and Z_t , where i indexes the relevant time- t cross-sectional population. What remains to specify is the link between the macro and micro sides, which we do in Assumption 3 below where we impose that idiosyncratic income shocks and aggregate shocks are independent of each other. Note that this assumption does not rule out feedback from the micro to the macro side. Instead, what we require is that no single unit in the sample has aggregate effects, that is, that units are *atomistic*. To be precise, atomicity restricts how the unit-level shocks u_{it} and aggregate shocks V_t relate to each other, but not the dependence between aggregate summaries of the $\{\eta_{it}\}_i$ -distribution and Z_t , so it allows for feedback from micro to macro to operate through the latter.⁷

2.2 Objects of empirical interest

A primary goal of the empirical analysis is to quantify the effect of aggregate shocks on nonlinear income processes. Our framework allows us to recover several quantities of interest.

⁵Since Z_t can be a vector, the restriction to a first-order process in (3) is without loss of generality—any p -th order Markov process can be cast as a first-order Markov process in companion form. In fact, our empirical analysis uses an AR(2) process for Z_t .

⁶The idea of extracting a low-dimensional summary of the covariation of multiple economic time series is at the core of the concept of business cycle (Burns and Mitchell, 1946). The use of linear factor models for that purpose also has a long history in economics; see Stock and Watson (2016) for a survey.

⁷Our framework can also accommodate sectoral and regional variation by including sector-level or region-level aggregate factors in Z_t .

A first quantity is the following measure of *nonlinear persistence* which extends the ABB persistence to a setup with aggregate shocks:

$$\rho(u_{it}, \eta_{i,t-1}, Z_t, Z_{t-1}) \equiv \left. \frac{\partial Q_\eta(\eta, Z_t, Z_{t-1}, u_{it})}{\partial \eta} \right|_{\eta=\eta_{i,t-1}}. \quad (4)$$

That is, persistence is measured by how a change in the micro state $\eta_{i,t-1}$ affects its next-period value. In a linear autoregressive model, $\rho(u_{it}, \eta_{i,t-1}, Z_t, Z_{t-1}) = \rho$ is constant and equal to the autoregressive root of the process. In contrast, in our nonlinear process, this measure is state-dependent: persistence may vary with the past position in the income distribution $\eta_{i,t-1}$ and with current and past aggregate conditions Z_t and Z_{t-1} .

A key feature of ρ is that it captures how persistence changes with the micro shock u_{it} , and how u_{it} interacts with the remaining determinants of persistent income. In micro panels, a robust finding documented in ABB is that persistence decreases for good-shocks/low- η and bad-shocks/high- η combinations. This reflects the fact that a good shock arriving in a low-income state has sometimes the power to erase a bad income history; the reverse holds for high- η households reached by a bad u . This feature is absent from linear income processes. But the framework of this paper adds an extra layer. In (4), the entire shape of ρ as a function of u_{it} and $\eta_{i,t-1}$ —the intensity with which past income histories are wiped out by big shocks—can change with the aggregate state of the economy Z_t . This is important because we find evidence of business-cycle variation in nonlinear persistence.

A second quantity is the persistent income *nonlinear exposure to aggregate shocks*:

$$\beta(u_{it}, \eta_{i,t-1}, Z_t, Z_{t-1}) \equiv \left. \frac{\partial Q_\eta(\eta_{i,t-1}, Z, Z_{t-1}, u_{it})}{\partial Z} \right|_{Z=Z_t}. \quad (5)$$

This measure captures how the exposure of a household to the aggregate state varies with the determinants of persistent income η_{it} . The aggregate exposure may vary along the persistent income distribution $\eta_{i,t-1}$ or as a function of the idiosyncratic shock u_{it} the household is hit with. It may also display cyclical patterns as it depends on the aggregate states Z_t and Z_{t-1} . It is worth noting that this measure is not available in models of income risk that exclude macro shocks, and in most specifications that do include aggregate uncertainty it is typically restricted to be constant, thus ruling out cyclical exposures and interactions with micro-level states and shocks.

Conceptually, β is a partial derivative that holds other determinants of income fixed. In particular, since we are holding Z_{t-1} fixed, β can be interpreted as the impact of a change in

the shock to Z_t which is V_t . There is therefore a natural connection between β and impulse responses with respect to V_t that we explore in Section 6. This illustrates why it is important to include the lagged aggregate state Z_t as a conditioning variable in Q_η .

A third quantity of interest is a measure of *conditional skewness*, $\text{sk}(\eta_{i,t-1}, Z_t, Z_{t-1})$, that we define as

$$\frac{Q_\eta(\eta_{i,t-1}, Z_t, Z_{t-1}, 0.9) + Q_\eta(\eta_{i,t-1}, Z_t, Z_{t-1}, 0.1) - 2Q_\eta(\eta_{i,t-1}, Z_t, Z_{t-1}, 0.5)}{Q_\eta(\eta_{i,t-1}, Z_t, Z_{t-1}, 0.9) - Q_\eta(\eta_{i,t-1}, Z_t, Z_{t-1}, 0.1)}, \quad (6)$$

which is the Kelley skewness of the predictive distribution of η_{it} . In micro panels, skewness tends to be decreasing in past $\eta_{i,t-1}$; that is, income risk is tilted to the upside for the low- η and to the downside for the high- η . As for persistence, the novelty of our paper is to allow us to trace how skewness changes with the aggregate state of the economy Z_t (and Z_{t-1}). How skewness varies with Z_t is a central empirical question. [Güvenen et al. \(2014\)](#), for example, find that recessions have a sizable negative impact on the skewness of one-year income growth. Our framework enhances previous analyses in various directions. It allows us to quantify the cyclical patterns of persistent income risk directly, without the need to proxy it by income changes. It also allows us to examine the role of the household's position in the income distribution. Hence, our approach provides a unified treatment of the two types of income skewness that have been highlighted in the literature: across the income distribution (indexed by $\eta_{i,t-1}$) and over the business cycle (indexed by Z_t).⁸

In addition to exploring the cyclical behavior of nonlinear income processes, our setup allows us to recover impulse response functions (IRFs) to both macro shocks (i.e., a shock to Z_t such as a recession or boom) and micro shocks (i.e., a shock u_{it} such as a promotion or demotion), as we discuss in Section 6. Finally, our framework allows us to quantify the contribution of macro and micro shocks to the cost of income risk, the topic of Section 7.

3 Identification based on time series of panels

Here, we describe the statistical population concept for our framework which is based on what we refer to as a *time series of panels*. This section focuses on identification and we defer estimation and inference to Section 4.

⁸ Similarly, we can define measures of conditional dispersion and kurtosis. We report estimates of these quantities in our empirical analysis.

3.1 Time series of panels

We begin by introducing the notion of a time series of panels, a data structure whose infinite-sample counterpart is the statistical population of our setting. The researcher observes a time series of aggregate indicators

$$Z_t \in \mathcal{Z}, \quad t = 0, \dots, T + S,$$

and a sequence of panel datasets

$$Y_t \in \mathcal{Y}_t, \quad t = 1, \dots, T,$$

where each $Y_t = \{y_{i,t+s} : i \in \mathcal{I}_t, 0 \leq s \leq S - 1\}$ contains individual observations on units in a sample $\mathcal{I}_t$ of size N_t , representative of the population of interest at time t . The length of each panel, S , is chosen to preserve the representativeness of the samples $\mathcal{I}_t$ while ensuring identification of the parameters of the model. Our approach relies on large cross-sections (large N_t), short subpanels (small and fixed S), and a long time series spanning sufficient aggregate variation coexisting with an evolving population of micro units (large T).

An advantage of our time series of panels approach, compared to standard long panel approaches, is that it mitigates concerns over the representativeness of the data when attrition depends on income shocks or covariates such as age. For example, a long panel of households spanning the years 1970 to 2019 (as in our empirical analysis) will overrepresent more stable units with higher incomes compared to a representative sample in any given year. In particular, our time series of panels approach allows for heterogeneous initial conditions in each subpanel. A precursor to our time series of panels approach is the analysis of cyclical income risk in the U.S. by [Storesletten et al. \(2004\)](#).⁹

We are now ready to describe the statistical population. For conciseness, we use $\{A_t\}$ as shorthand for the time series process $\{A_t : -\infty < t < \infty\}$.

Assumption 1 (Macro states). *There is a vector-valued macro stochastic process $\{Z_t\}$ that is stationary and satisfies (3) with V_t i.i.d. over t .*

Assumption 2 (Micro processes). *There is a cross-section of scalar micro stochastic processes $\{(\eta_{it}, \varepsilon_{it})\}$, where $\eta_{it} \in \mathcal{A}$ and $\varepsilon_{it} \in \mathcal{E}$, that satisfies (1) and (2) with u_{it} and v_{it}*

⁹Another advantage of this approach is that it allows us to handle covariates that would otherwise be collinear with time, e.g., age. Moreover, it may be expensive (sometimes impossible) to keep track of the same units for very long periods.

mutually independent, i.i.d. over t , and uniformly distributed on $(0, 1)$.

Assumption 3 (Atomicity). *For each unit i , the persistent micro shock u_{it} is independent of the macro shocks V_τ at all lags and leads.*

In Assumption 1 we assume that Z_t is a stationary process, so $Q_\eta(\cdot, Z_t, Z_{t-1}, \cdot)$ (which is our main empirical focus) is stationary as well. In contrast, the initial condition and transitory shock parts of the micro processes defined in Assumption 2 are unrestricted; in fact, we treat Q_{init, t_0} and $Q_{\varepsilon, t}$ as non-stochastic time-varying parameters. Lastly, Assumption 3 requires each individual to be atomistic, in the sense that no single idiosyncratic shock has influence on the aggregate shock. Nonetheless, as discussed in Section 2, our framework permits feedback from the micro to the macro side of the model, as summaries of micro distributions may be included in Z_t .

3.2 Identification

The identification question is to determine the unknown functions Q_η and Q_Z from the probability distribution of observable quantities. As anticipated, our argument applies to $N_t, T \rightarrow \infty$ with S fixed. With large T , identification of Q_Z is immediate since the latter is the transition kernel of a stationary Markov process and here we assume that Z_t is observed by the researcher. In our implementation, we will infer Z_t from a collection of aggregate indicators using a dynamic factor model, in which case identification of Q_Z will follow from standard arguments in the time series literature.

We establish identification of Q_η in two steps. In the first step (the “micro” step), we consider the cumulative distribution function (CDF) F_t of the variables $y_{i,t}, \dots, y_{i,t+S-1}$ for a representative household i in subpanel t . This CDF is conditional on the particular realizations of the macro state in the subpanel, say $z_t, \dots, z_{t+S-1}$. Based on F_t , we rely on the arguments in ABB to obtain sufficient conditions for the identification of the persistent component process

$$Q_\eta(\eta, z_{t+s}, z_{t+s-1}, u), \quad \text{for all } \eta \in \mathcal{A}, \quad u \in (0, 1), \quad s = 1, \dots, S-1.$$

In other words, the function Q_η is identified, but only at the values of Z_t that occurred during the time horizon of the subpanel. This step, which requires $S \geq 4$ together with regularity conditions that we detail in Supplemental Appendix A, highlights a key advantage of our time series of panels framework compared to approaches that rely on repeated cross-

sections: observing the same unit over a number of periods makes it possible to separate permanent from transitory components by exploiting information on individual transitions.

Next, in the second step (the “macro” step), we use the fact that Z_t is a stationary Markov process that visits all possible states with positive probability. Hence, now considering the time series dimension $t = 1, 2, \dots$, we obtain the identification of

$$Q_\eta(\eta, z, \tilde{z}, u), \quad \text{for all } \eta \in \mathcal{A}, \quad u \in (0, 1), \quad z \in \mathcal{Z}, \quad \tilde{z} \in \mathcal{Z}.$$

That is, the persistent income process is identified for all possible realizations of the aggregate state. This second step (which we formalize in Supplemental Appendix A) is essential, since it is the one that allows us to identify how the time-variation in Q_η relates to the business cycle Z_t , the main goal of our analysis.

Our identification analysis clarifies the role played by each element in our framework: large cross-sections, together with the subpanel length S , help recover the time- t short-panel distribution F_t ; the subpanel length S balances the ability to identify latent variable distributions with the credibility of sample representativeness; finally, long time series of both external macro data and micro data allow the researcher to link the time- t cross-sections to the aggregate factor of interest.

4 Empirical specification and estimation

In this section, we specify a flexible class of parametric models for our framework and outline a suitable estimation technique. Our empirical strategy relies on parametric specifications of the functions Q_η , $Q_{\varepsilon,t}$, $Q_{\text{init},t}$ to create nonlinear representations of income processes. We build on the stochastic EM method of [Arellano and Bonhomme \(2016\)](#) and its extension to time-varying latent variables in ABB, further extending it to a long time series of short panels involving both macro and micro latent variables.

In our empirical specification, we also account for the age x_{it} of the household head, and we specify the quantile function of the persistent component η_{it} parametrically as

$$Q_\eta(\eta, z, \tilde{z}, x, u) = \sum_{k=1}^K \Theta_k(u; \theta) \psi_k(\eta, x, z, \tilde{z}). \quad (7)$$

In this specification, ψ_k are known polynomial functions; in practice, tensor products of Hermite polynomials. In turn, $\Theta_k(\cdot; \theta)$ is a piecewise-linear spline function with nodes $(\bar{u}_1, \dots, \bar{u}_L)$

(as in [Wei and Carroll, 2009](#)), augmented with an exponential modeling in the tail intervals $(0, \bar{u}_1)$ and $(\bar{u}_L, 1)$ (as in [Arellano and Bonhomme, 2016](#)). The vector θ , which collects all intercept, slope, and tail parameters, shapes the measures of persistence, aggregate exposure, dispersion, and skewness we introduced in Section 2, as well as the IRFs and risk decompositions we will introduce in Sections 6 and 7.

To complete the empirical specification, we model the quantile function $Q_{\text{init},t}$ of subpanel- t initial persistent component, η_{it} , using a functional form similar to (7), and we model the quantile function $Q_{\varepsilon,t}$ of the transitory components, ε_{it} , using finite mixtures of Gaussians.¹⁰ We denote the vector containing all parameters in $Q_{\text{init},t}$ and $Q_{\varepsilon,t}$ as δ_t . Specifying the quantile functions of transitory shocks and initial conditions as flexible functions of time allows us to absorb long-term trends and other unmodeled sources of nonstationarity in the data. We provide details about the exact specification in Supplemental Appendix C.

Our model of income dynamics is parametric, indexed by the parameters θ and $\delta_1, \dots, \delta_T$. Although one could obtain an arbitrarily accurate approximation to any quantile function Q_η by increasing the degree of the polynomial specification K in (7), as in nonparametric methods, we adopt a flexible parametric perspective in which θ is finite-dimensional (and the δ_t 's are too). While the identification argument of Section 3 covers nonparametric models, the precision with which the objects of interest can be estimated is limited by the time series length T , or more concretely, by the scarcity of recessions and booms in the sample. In that context, our flexible parametric approach balances flexibility with statistical power.

Our model implies moment restrictions that involve the latent variables η_{it} . For example, given (7), we have for all $u \in \{\bar{u}_1, \dots, \bar{u}_L\}$,

$$E \left[\sum_{\tau=t+1}^{t+S-1} \left(\sum_{k=1}^K \frac{\partial \Theta_k(u; \theta)}{\partial \theta} \psi_k(\eta_{i,\tau-1}, x_{i\tau}, Z_\tau, Z_{\tau-1}) \right) \times \left(u - \mathbf{1} \left\{ \eta_{i\tau} \leq \sum_{k=1}^K \Theta_k(u; \theta) \psi_k(\eta_{i,\tau-1}, x_{i\tau}, Z_\tau, Z_{\tau-1}) \right\} \right) \right] = 0. \quad (8)$$

By adding the moment restrictions corresponding to tail parameters in θ and the δ_t parameters appearing in $Q_{\varepsilon,t}$ and $Q_{\text{init},t}$, we obtain a system of moment conditions on all the model parameters $\theta, \delta_1, \dots, \delta_T$.

¹⁰An advantage of this specification of $Q_{\varepsilon,t}$ compared to a piecewise-linear quantile specification is that it ensures that the (integrated) moments we use in estimation are smooth functions of the parameters. We also estimated the nonlinear model using quantile specifications for transitory shocks and found similar empirical results; see the working paper version ([Almuzara, Arellano, Blundell, and Bonhomme, 2025](#)).

To transform these moments involving latent variables into moments involving observed variables, we integrate out the latent persistent components η_{it} with respect to their posterior density under the model. For example, this gives for all $u \in \{\bar{u}_1, \dots, \bar{u}_L\}$,

$$E \left[\int \dots \int \sum_{\tau=t+1}^{t+S-1} \left(\sum_{k=1}^K \frac{\partial \Theta_k(u; \theta)}{\partial \theta} \psi_k(\eta_{i,\tau-1}, x_{i\tau}, Z_\tau, Z_{\tau-1}) \right) \right. \\ \times \left(u - \mathbf{1} \left\{ \eta_{i\tau} \leq \sum_{k=1}^K \Theta_k(u; \theta) \psi_k(\eta_{i,\tau-1}, x_{i\tau}, Z_\tau, Z_{\tau-1}) \right\} \right) \\ \left. \times \underbrace{f(\eta_{it}, \dots, \eta_{i,t+S-1} \mid y_{it}, x_{it}, \dots, y_{i,t+S-1}, x_{i,t+S-1}; \theta, \delta_t)}_{\text{posterior density}} d\eta_{it} \dots d\eta_{i,t+S-1} \right] = 0, \quad (9)$$

where the posterior density is fully determined via Bayes rule through our parametric model. Unlike (8), (9) only involves observed variables and can thus be used for generalized method of moments (GMM) estimation. We provide a complete list of the moments we rely on in Supplemental Appendix C.

The system of integrated moment conditions such as (9) is a standard one, except for the fact that the number of parameters $\theta, \delta_1, \dots, \delta_T$ grows with the sample size. As our identification analysis makes clear, a growing number T of periods is needed to recover the impact of the business cycle on household income. In Supplemental Appendix B, we study the asymptotic behavior of the GMM estimator based on integrated moments in an approximation where the number T of subpanels grows, the length of each subpanel S is kept fixed, and the number of cross-sectional units N_t is large relative to T . This asymptotic regime is intended to capture the dimensions of our empirical setting. Under assumptions that we detail, we show that the GMM estimator $\hat{\theta}$ is consistent and asymptotically normal, and that the presence of the growing parameter sequence $\delta_1, \dots, \delta_T$ does not cause an incidental parameter problem.

For GMM estimation, we need to find the parameter values that set the empirical moments to zero. This requires approximating the integrated moments. For this purpose, we rely on a stochastic Expectation Maximization (EM) algorithm as in Arellano and Bonhomme (2016) and ABB. The basic structure is as follows.

Algorithm 1. Initialize parameters $\hat{\theta}^{(0)}$ and $\{\hat{\delta}_t^{(0)}\}_{t=1}^T$. For $j = 1, \dots, J$, iterate between

- 1) *Stochastic E-step:* independently over units i and subpanels t , draw $\eta_{it}^{(j)}, \dots, \eta_{i,t+S-1}^{(j)}$ from the posterior $f(\eta_{it}, \dots, \eta_{i,t+S-1} \mid y_{it}, x_{it}, \dots, y_{i,t+S-1}, x_{i,t+S-1}; \hat{\theta}^{(j-1)}, \hat{\delta}_t^{(j-1)})$; and

2) *M step*: given the draws, update the parameters to $\hat{\theta}^{(j)}, \hat{\delta}_1^{(j)}, \dots, \hat{\delta}_T^{(j)}$ by proceeding as if $\eta_{i,t+s}$ were observed, equal to $\eta_{i,t+s}^{(j)}$.

For some $\mu \in (0, 1)$, set $\hat{\theta} = (\mu J)^{-1} \sum_{j=(1-\mu)J}^J \hat{\theta}^{(j)}$ and $\hat{\delta}_t = (\mu J)^{-1} \sum_{j=(1-\mu)J}^J \hat{\delta}_t^{(j)}$.

The stochastic E-step amounts to drawing latent variables from a joint posterior distribution. For this purpose, we use a Sequential Monte Carlo approach (Creal, 2012, Arellano, Blundell, Bonhomme, and Light, 2023). In turn, in the M step we treat the latent component draws as if they were the true η_{it} 's. For example, pooling (8) across subpanels t , we update ℓ -by- ℓ the parameters $\Theta_k(\bar{u}_\ell; \theta) = \theta_{k\ell}$, for $k = 1, \dots, K$, as

$$\begin{aligned} \{\hat{\theta}_{k\ell}^{(j)}\} = \underset{\{\theta_{k\ell}\}}{\operatorname{argmin}} \sum_{t=1}^T \sum_{i \in \mathcal{I}_t} \sum_{\tau=t+1}^{t+S-1} \left(\eta_{i\tau}^{(j)} - \sum_{k=1}^K \theta_{k\ell} \psi_k \left(\eta_{i,\tau-1}^{(j)}, x_{i\tau}, Z_\tau, Z_{\tau-1} \right) \right) \\ \times \left(\bar{u}_\ell - \mathbf{1} \left\{ \eta_{i\tau}^{(j)} \leq \sum_{k=1}^K \theta_{k\ell} \psi_k \left(\eta_{i,\tau-1}^{(j)}, x_{i\tau}, Z_\tau, Z_{\tau-1} \right) \right\} \right), \end{aligned}$$

which is a standard quantile regression (Koenker and Bassett, 1978). Likewise, we update the tail parameters in θ using exponential regressions, and proceed similarly to update the δ_t parameters. We provide full implementation details in Supplemental Appendix C.

For a large enough number of draws J , our asymptotic analysis justifies the use of conventional GMM inference. Taking advantage of the fact that the model is parametric, we rely on a parametric bootstrap method, for three main reasons. First, while we use a large number of draws in our implementation, the bootstrap can capture the additional variability due to simulation error. Second, the bootstrap seamlessly accounts for features of our setting such as the fact that the subpanels partially overlap. Third, our asymptotic theory in Supplemental Appendix B shows that the convergence rate of the GMM estimator is affected by the degree of dependence between cross-sectional units. The bootstrap gives us a simple way of checking the impact of allowing for cross-sectional dependence by specifying a parametric model for it and simulating bootstrap samples under dependence. In a robustness check, we also report confidence bands that allow for parametric cross-sectional dependence.

Extension: a measurement system for Z_t . Thus far we have treated the aggregate state Z_t as observed to the researcher. In our application, we follow a common approach in macroeconomics and estimate Z_t based on a set of aggregate indicators W_t using a linear dynamic factor structure:

$$W_t = \Lambda Z_t + e_t, \quad Z_t = \Phi Z_{t-1} + \sigma_V V_t = Q_Z(Z_{t-1}, V_t), \quad (10)$$

with V_t an i.i.d. standard normal vector of shocks independent of e_t at all lags and leads, and with the elements of e_t specified as mutually independent Gaussian autoregressive processes. In our empirical analysis, W_t consists of variables that are informative about aggregate fluctuations (including GDP and the unemployment rate) while Z_t is a scalar summary indicator of the state of the business cycle. One interpretation of Z_t draws on [Angeletos, Collard, and Dellas \(2020\)](#), who find that a single serially uncorrelated process, what they call the main business cycle (MBC) shock, accounts for the largest share of the unpredictable variation in each of the variables W_t at business cycle frequencies.

To account for the fact that Z_t is estimated, we modify our estimation and inference approach in three ways. First, we estimate the Z_t process based on a Bayesian approach using Gibbs sampling. Second, we modify the Stochastic E-step of Algorithm 1 by adding a prior step where we draw $Z_t^{(j)}$, for $t = 1, \dots, T+S-1$, from the posterior distribution of $\{Z_t\}$ given $\{W_t\}$. Lastly, we incorporate these additional steps into our parametric bootstrap procedure. Hence, bootstrap variability also reflects uncertainty in the realization of the aggregate factor Z_t —even though in our application we find this additional source of parameter uncertainty to be small. Our description of the estimation and inference procedure in Supplemental Appendix C accounts for these extensions.

Finally, even though our focus in this paper is on income processes, it is important to highlight that this method offers a flexible, general-purpose way to estimate rich models of micro-level latent variables subject to macro shocks, and it can be applied more generally.

5 Macro shocks and nonlinear micro income processes

We now present our empirical analysis of the effects of macro shocks on the nonlinear income process. We discuss the macro data and business-cycle state in Section 5.1, followed by a description of our main source of micro data, the PSID, in Section 5.2. We then turn to the quantification of aggregate effects on measures of nonlinear persistence (Section 5.3), exposure to aggregate shocks (Section 5.4), and conditional skewness (Section 5.5).

5.1 Macro data and the aggregate state

We extract the aggregate state Z_t from a vector of macro observables W_t informative about the business cycle: GDP, consumption, investment, the unemployment rate, and hours worked. The data are quarterly and cover 1960Q1–2019Q4. GDP, consumption, investment and hours are measured in log per capita terms while the unemployment rate is multiplied

by -1 to make it procyclical as the other entries of W_t . All series are detrended by removing a two-sided 40-quarter low-pass filter from them.¹¹

For estimation, we normalize the loading on GDP to one. Therefore, the business cycle process Z_t is measured as a log deviation from its per-capita growth trend. We specify Z_t and each entry of e_t in (10) as independent AR(2) processes and estimate the macro parameters λ by Gibbs sampling using a flat prior. Supplemental Appendix C.4 discusses the details. Figure 1 shows the time series evolution of W_t and estimates of Z_t .




FIGURE 1. Business cycle indicators and estimates of Z_t

Note: We plot real GDP, consumption and investment per capita, the unemployment rate and hours worked alongside our estimate of the business cycle factor Z_t . We standardize them to mean zero and unit variance. Red shaded areas are 90% pointwise probability bands and gray areas indicate NBER-dated recessions.

Three takeaways from the figure are as follows. First, and not surprisingly, the variables included in W_t show strong comovements. As a result, filtering Z_t is statistically easy, as evidenced by the tight 90% probability bands drawn around the estimate. Second, there is substantial cyclical variation in Z_t , and this remains true when we consider temporally aggregated versions of it. For example, Figure 1 marks with a diamond the estimates of Z_t for the fourth quarter of each year. They clearly retain a meaningful share of the aggregate variation we see at the quarterly frequency. This is important because income data at the micro level refer to annual income. Third, there is evidence of some asymmetry between recessions and expansions, particularly in the last 40 years: Recessions are deep and occur suddenly, while expansions are mild and unfold gradually. It would be possible to enrich model (3) to accommodate that feature, although this is unlikely to affect our estimates of Z_t given their precision. Nonetheless, when reporting summaries that condition on the

¹¹The data come from Federal Reserve Economic Data (FRED). To construct W_t we took from Angeletos et al. (2020) the subset of variables most informative about the MBC shock, although our results are robust to alternative specifications that expand W_t with other macro indicators.

aggregate state in this section, we will adopt the notion of a typical recession, $\tilde{Z}_r = -2\sigma_Z$, a steady state, $\tilde{Z}_{ss} = 0$, and a typical expansion, $\tilde{Z}_e = \sigma_Z$, where $\sigma_Z^2 = \text{Var}(Z_t)$. This choice does not affect the results on impulse responses and risk calculations in Sections 6 and 7.

In this context, a central object is the IRF of the macro state variable Z_t to the shock V_t . We will return to it when discussing IRFs in Section 6.2.¹²

5.2 Micro data and cyclical patterns in the PSID

As we outlined in Section 3, the micro data for our study take the form of a time series of panels. To this end, we draw from the long history of panel data on income and earnings available in the Panel Study of Income Dynamics (PSID).

Established in 1968, the PSID initially surveyed a nationally representative set of U.S. households. Ever since, it has followed those families and, as their children come of age and form independent households, incorporated those new units into the panel. With the periodic addition of refresher and immigrant samples, the PSID design aims to capture the process of household formation and dissolution in the U.S. economy over time. It also accords with our notion of time series of panels in Section 3 as it allows us to form a long sequence of short subpanels, each reflecting representativeness at a point in time, which taken together span a rich history of aggregate fluctuations.

Interviews were conducted annually between 1968 and 1997, and biennially starting in 1999. Whether annual or biennial, the year k interview (which typically occurs between March and November) asks the household to report annual income for year $k - 1$. We make use of all available waves beginning in 1970 and ending in 2019: we exclude 1968 and 1969 as some income and demographic variables were not collected in those waves, and we stop at 2019 before the COVID-19 pandemic. We then form a time series of panels where each subpanel has $S = 4$ biennial waves, although, crucially, we exploit all years. We use biennial panels because we are constrained by the change in data collection frequency after 1999. This also has the advantage of making the assumption of serially uncorrelated ε_{it} more plausible. Nonetheless, prior to 1999, the panels we construct have base years that cover both even and odd years, so we make use of all available annual observations. Given the biennial nature of the subpanels, $t - 1$ below is understood as two years before t , and with a slight abuse of notation we denote by Z_t the value of the macro state in the fourth quarter of t . This

¹²We also discuss IRFs of W_t to V_t and compare them to the IRFs estimated in [Angeletos et al. \(2020\)](#) for the MBC shock in Supplemental Appendix G.3. Empirically, IRFs of W_t to V_t are quantitatively similar to IRFs to the MBC shock, albeit slightly less persistent.

convention is relevant for interpreting the income process and summaries.

In our data selection, we keep households whose representative person is male, married and aged 25 to 60, and whose income is positive.¹³ We consider three different measures of income: male earnings, household earnings, and disposable income. Male earnings are the labor income of the household head. Household earnings include the labor income of the spouse. Disposable income adds transfers and subtracts federal taxes. We take transfer income as reported in the PSID, which includes welfare payments, social security income, and unemployment benefits. We compute taxes with the tax functions estimated by [Borella, De Nardi, Pak, Russo, and Yang \(2023\)](#). Lastly, we deflate income to 2016 dollars using the consumer price index (CPI-U).¹⁴

We construct y_{it} by residualizing log income against the following covariates interacted with a quadratic time trend: education, race, family size, number of children, an indicator for the presence of dependents, and state of residence. Except for family size and number of children, the other variables are treated as categorical. We then add back to all units the same constant representing a household with typical covariates.¹⁵ We also exclude age from this operation, which, as anticipated, enters Q_η and Q_{init,t_0} in the form of x_{it} .¹⁶

Netting out the effect of covariates interacted with time trends in the construction of y_{it} allows us to capture part of the low-frequency changes during the period. In addition, the processes of initial conditions, Q_{init,t_0} , and transitory shocks, $Q_{\varepsilon,t}$, both depend on parameters that vary flexibly over time, thus capturing additional temporal variation. A central assumption in our approach is that, net of these factors, the dynamic transition process of persistent income, Q_η , is stationary.

Descriptives. Figure 2 offers a first look at the micro data using repeated cross-sections ($S = 1$). The period covers seven NBER recessions (the two downturns of the 1980s merged into a single bar). Panel (a) shows the fraction of zeros that would be dropped to compute y_{it} . It shows that most of the time less than 2% of the sample reports zero annual household earnings (blue). This applies to recessions too, except for the Great Recession when

¹³In addition to the main sample representative of the U.S. population, the PSID maintains a low-income sample known as the Survey of Economic Opportunities. Our analysis makes use only of the former.

¹⁴Beyond partialling out state-of-residence dummies, we do not explicitly account for differences in state taxes, but one possibility is to use the tax functions in [Fleck, Heathcote, Storesletten, and Violante \(2025\)](#).

¹⁵This is a white college graduate with family size four, two children, no dependents, who resides in the state of New York. This level shift has no impact on the dynamics of η_{it} or any of the summaries of interest.

¹⁶We estimated a version of the model with x_{it} as an argument of $Q_{\varepsilon,t}$, and found little effect of x_{it} . Given this, we report results based on the more parsimonious specification.

the fraction briefly exceeded 2%. The story is different for male earnings, for which zeros constitute a larger fraction and there seems to be a steeper upward trend over time.




FIGURE 2. Distributional characteristics of income measures over time

Note: We show the fraction of zeros and the 10th, 50th and 90th percentiles (conditional on positive income) for our income measures in thousands of 2016 dollars. Shaded areas are NBER recessions.

Given the demographic criteria in our data selection, practically no unit reports zero disposable income for a whole year. Since this is also the relevant measure for studying the effects of the business cycle on consumption and welfare, disposable income will be our primary focus below. We report results for male and household earnings for comparability, with the caveat that doing so abstracts from the extensive margin. To the extent that this margin is cyclical and recessions imply larger drops in labor income than what we observe, our results for male and household earnings should be taken as a lower bound on the impact of negative macro shocks along those margins.

Panels (b) to (d) in Figure 2 illustrate the evolution of low, middle and high incomes over time. One highlight is that there are divergent secular trends in the income distribution, with fast growth at the top and stagnation (or even decline) at the bottom. These trends have been extensively studied in the literature, usually linking them to changes in female labor force participation, structural transformation, reforms to the tax and social security systems, and other phenomena. But importantly for us, these trends coexist with sizable cyclical variation. A quick glance reveals that incomes fall in recessions in ways that are heterogeneous over the income distribution (more at the bottom than in the middle), across income measures (more for male earnings than for disposable income) and depending on the severity of the downturn (the double-dip recession of the 80s and the Great Recession being the worst). Figure 2 is also indicative of the redistributive role of taxes and transfers—as going from household earnings to disposable income pushes low incomes up and higher incomes down—and of the insurance role of both spousal income and taxation—which tend to mitigate income losses during recessions.

Model fit. We fit our model to the macro and micro data with the empirical specification detailed in Supplemental Appendix C, where we also explain how estimation and inference are implemented. We have assessed the performance of the estimator in terms of fit to the micro data and by means of Monte Carlo simulations. In addition, in Supplemental Appendix D we show that our model matches the time series evolution of quantiles and moments of observed income levels and growth rates well.

5.3 Nonlinear persistence: ρ tilts with the aggregate state

We now turn to quantifying empirically the effect of aggregate shocks on different aspects of the nonlinear income process. We begin by discussing the persistence of past income histories and how it changes over the business cycle. To that end, we use the measure of nonlinear persistence $\rho(u, \eta, Z_t, Z_{t-1}, x)$ defined in Equation (4) in Section 2.

The top panels of Figure 3 show our estimates of $\rho(u, \eta, Z_t, Z_{t-1}, x)$ for the three measures of income we consider. We plot persistence by quantiles of the current shock $u = u_{it}$ and past persistent income $\eta = \eta_{i,t-1}$. Age $x = x_{it}$ is averaged out, $Z_{t-1} = \tilde{Z}_{ss}$ is held at its steady-state value, and Z_t takes on different values representing a typical recession $\tilde{Z}_r$ (red), the steady state $\tilde{Z}_{ss}$ (yellow), or a typical mild expansion $\tilde{Z}_e$ (blue). On the bottom panels, we report the recession minus the expansion persistence. If the recession persistence is above (below) its expansion counterpart according to a one-sided pointwise test at the 5% significance level, we indicate this in magenta (cyan).¹⁷

We highlight two takeaways from Figure 3. First, it confirms the nonlinear persistence pattern uncovered in ABB over a much longer history and across income measures. For example, in the steady state surface for disposable income, the average persistence is 0.92. It falls to 0.90 for a unit in the 90th η -percentile hit with a bad shock $u = 0.1$ and to 0.62 for a unit in the 10th η -percentile hit with a good shock $u = 0.9$. As already discussed, since ρ is a measure of how closely related current and past incomes are, this captures the fact that a big shock of a given direction may sometimes erase an entire income history.

The second takeaway is that the ABB persistence pattern is itself macro state-dependent, as aggregate shocks *tilt* the estimated persistence surface. Comparing recessions and expansions, ρ increases for low- η and decreases for high- η units (particularly those affected by bad shocks) during a downturn. In words, a low- η household finds it more difficult to leave the

¹⁷We present confidence bands accounting for cross-sectional dependence in Supplemental Appendix E, which are very similar to the ones relying on independence.




FIGURE 3. Nonlinear persistence.

Note: The top panels report the persistence measure $\rho(u, \eta, Z_t, Z_{t-1}, x)$ by quantile of the shock $u = u_{it}$ and past persistent income $\eta = \eta_{i,t-1}$. Here, age $x = x_{it}$ is averaged out, $Z_{t-1} = \tilde{Z}_{ss}$ and Z_t is a recession $\tilde{Z}_r$, the steady state $\tilde{Z}_{ss}$ or an expansion $\tilde{Z}_e$ (see Section 5.1). The bottom panels show the gap in persistence between recession and expansion, $\rho(u, \eta, \tilde{Z}_r, \tilde{Z}_{ss}, x) - \rho(u, \eta, \tilde{Z}_e, \tilde{Z}_{ss}, x)$. The difference is painted magenta (cyan) if it is statistically positive (negative) at the 5% significance level (one-sided test).

low-income state in the midst of a contraction, while a high- η household finds it harder to remain high-income. These effects are statistically non-negligible. For disposable income, ρ decreases by 0.03 for the 90th η -percentile with $u = 0.1$ and increases by 0.05 for the 10th η -percentile with $u = 0.9$ as we move from $Z_t = \tilde{Z}_e$ to $Z_t = \tilde{Z}_r$.

These patterns are quite similar across income measures. The increase in persistence at low incomes during recessions tends to be slightly larger for our earnings measures than for disposable income, possibly suggestive of the insurance role of taxes and transfers in attenuating the impact of negative aggregate shocks for the left tail of the income distribution—although the comparison abstracts from extensive margin effects.

5.4 Exposure to aggregate shocks: β is countercyclical

A second key aspect of the nonlinear income process is the persistent income exposure to macro shocks, measured by the nonlinear aggregate exposure coefficient $\beta(u, \eta, Z_t, Z_{t-1}, x)$ given in Equation (5) of Section 2. Figure 4 presents our estimates.




FIGURE 4. Nonlinear exposure to aggregate shocks.

Note: We report the exposure coefficient $\beta(u, \eta, Z_t, Z_{t-1}, x)$ by quantile of the shock $u = u_{it}$ averaged over persistent income $\eta = \eta_{i,t-1}$. Age $x = x_{it}$ is averaged out, $Z_{t-1} = \tilde{Z}_{ss}$ and Z_t is a recession $\tilde{Z}_r$, the steady state $\tilde{Z}_{ss}$ or an expansion $\tilde{Z}_e$ (see Section 5.1). Shaded areas represent 90% confidence bands.

In Figure 4, we show $\beta(u, \eta, Z_t, Z_{t-1}, x)$ as a function of the micro-level shock $u = u_{it}$ that occurs at the time of the macro shock. We average past income $\eta = \eta_{i,t-1}$ and age $x = x_{it}$, and fix $Z_{t-1} = \tilde{Z}_{ss}$, allowing Z_t to take values compatible with a recession $\tilde{Z}_r$, steady state $\tilde{Z}_{ss}$ or expansion $\tilde{Z}_e$ (discussed in Section 5.1).¹⁸

Three patterns stand out. First, the exposure to aggregate shocks depends on Z_t and is countercyclical. Averaged over u and η , in our disposable income calculations, β is 1.4 in a recession, 0.8 in steady state and 0.5 in an expansion. Given the normalization of Z_t in Section 5.1, these numbers can be interpreted as the elasticity of persistent income to an aggregate shock that implies a one percentage point change in GDP per capita relative to its trend. Accordingly, a negative macro shock leads to an aggregate reduction in the persistent component of disposable income that is more or less than one-for-one the fall in GDP depending on whether the economy is already in a recession when the shock hits. This is a major form of aggregate state-dependence and one that is ruled out by models in which aggregate shocks enter additively. What is more, we argue in Section 7 that this nonlinearity plays a key role in the cost of business-cycle risk.

¹⁸Confidence bands that account for cross-sectional dependence in Supplemental Appendix E show additional variability compared to the ones that use independence, although the qualitative conclusions remain.

Second, the aggregate exposure coefficient varies across income measures, with male earnings the most and disposable income the least sensitive to Z_t . For example, the recession average β is 2.4 for male earnings, 1.9 for household earnings and 1.4 for disposable income. The same applies to the steady-state and expansion β , and at different quantiles of $u = u_{it}$ and $\eta = \eta_{i,t-1}$, as seen in additional results in Supplemental Appendix F.

The last pattern to analyze is the interaction between macro and micro shocks. Units affected by bad micro shocks $u = u_{it}$ at the time when the macro shock occurs are relatively more exposed during recessions (and generally less during expansions) than units subject to neutral or good shocks. In other words, the consequences of recessions are not evenly distributed but are borne mostly by those hit by adverse idiosyncratic shocks as the downturn unfolds. This is another feature often ruled out by linear models with aggregate shocks.

5.5 A tale of two skewnesses

We conclude by discussing the conditional skewness measure introduced in Equation (6) of Section 2. Measures of dispersion and kurtosis complete the picture of aggregate effects on the income process. They appear to be less cyclical than the skewness (in line with Guvenen et al., 2014), and we report them in Supplemental Appendix F.

The upper panels of Figure 5 show $\text{sk}(\eta, Z_t, Z_{t-1}, x)$ for different quantiles of $\eta = \eta_{i,t-1}$ with $Z_{t-1} = \tilde{Z}_{ss}$ when Z_t ranges from recession $\tilde{Z}_r$, through steady state $\tilde{Z}_{ss}$ to expansion $\tilde{Z}_e$. Age $x = x_{it}$ is averaged out. The lower panels display the difference in skewness between recession and expansion including 90% pointwise confidence bands.¹⁹

Figure 5 provides empirical support for a tale of two skewnesses suggested by our descriptive analysis. The ABB skewness pattern, by which $\text{sk}(\eta, Z_t, Z_{t-1}, x)$ decreases for higher η , is accompanied here by a cyclical skewness pattern: Recessions tend to shift $\text{sk}(\eta, Z_t, Z_{t-1}, x)$ downward, often into negative values, at all levels of η , often by a large margin. Mirroring nonlinear persistence and exposure to aggregate shocks, cyclical shifts in skewness are strongest for male earnings and weakest for disposable income, that is, after spousal earnings and the tax-transfer system are taken into account.

Since $\text{sk}(\eta, Z_t, Z_{t-1}, x)$ is a measure of the relative strength of upside and downside income shocks, our results characterize recessions as periods when downside shocks become more prevalent throughout the income distribution, but heterogeneously so.

¹⁹The confidence bands accounting for cross-sectional dependence in Supplemental Appendix E are similar to the ones relying on independence in this case.




FIGURE 5. Conditional skewness.

Note: The upper panels report skewness $sk(\eta, Z_t, Z_{t-1}, x)$ by past persistent income $\eta = \eta_{i,t-1}$ where age $x = x_{it}$ is averaged out, $Z_{t-1} = \tilde{Z}_{ss}$ and Z_t is a recession $\tilde{Z}_r$, the steady state $\tilde{Z}_{ss}$ or an expansion $\tilde{Z}_e$ (see Section 5.1). The lower panels show the gap in skewness between recession and expansion, $sk(\eta, \tilde{Z}_r, \tilde{Z}_{ss}, x) - sk(\eta, \tilde{Z}_e, \tilde{Z}_{ss}, x)$. Shaded areas represent 90% confidence bands.

6 Macro and micro impulse responses

In this section, we develop new methodology for measuring impulse responses to macro and micro shocks. We first present our approach and then discuss empirical estimates.

6.1 Measuring the propagation of macro and micro shocks

To develop the idea, we focus on the persistent component η and omit covariates x . From Equations (1) and (3) we get by recursive substitution the representation

$$\eta_{i,t+h} = q_{\eta,h}(\mathbf{u}_{it}^h, \mathbf{V}_{t+1}^{h-1}, \eta_{i,t-1}, Z_t, Z_{t-1}), \quad h = 0, 1, \dots, \quad (11)$$

where $\mathbf{u}_{it}^\ell = (u_{it}, \dots, u_{i,t+\ell})$ and $\mathbf{V}_t^\ell = (V_t, \dots, V_{t+\ell})$. With the distribution of micro and macro shocks, Equation (11) determines the predictive distribution of $\eta_{i,t+h}$ given initial states $(\eta_{i,t-1}, Z_t, Z_{t-1})$. In what follows we assume Z_t is scalar, although it is straightforward to generalize the derivation to the multivariate case.

The macroeconometric tradition often defines macro and micro impulse responses as

$$E[\eta_{i,t+h} \mid V_t = 1] - E[\eta_{i,t+h} \mid V_t = 0] \text{ and } E[\eta_{i,t+h} \mid u_{it} = 1] - E[\eta_{i,t+h} \mid u_{it} = 0],$$

respectively. In linear models, they coincide with the derivatives of $q_{\eta,h}$ with respect to shocks V_t and u_{it} . In our setup, however, this approach suffers from various shortcomings. First, our model features significant nonlinearities in the persistence and interactions between macro and micro states and shocks; impulse responses should account for those. Second, changes in one unit of V_t or u_{it} need not be directly comparable. Third, in our panel data setup, a change of given size in u_{it} may have different impacts for different individuals.

To address these issues, we define impulse responses as the impact on the predictive distribution of $\eta_{i,t+h}$ of perturbations to *past states* (as opposed to shocks), extending Gallant, Rossi, and Tauchen (1993) beyond the time series setup. For this purpose, we introduce a *rule* (denoted g below) that maps perturbations onto a consistent system of measurement. As a result, impulse responses are indexed by past states—allowing us to document nonlinearities—and depend on the normalization rule—allowing us to achieve comparability across shocks and units.

IRFs via perturbations. Let us define a benchmark value for one of the state variables: Z^b for Z_t for the macro impulse response; η^b for η_{it} for the micro impulse response. Let us then define the new state value Z^p (resp., η^p) by means of the perturbation $\Delta = Z^p - Z^b$ (resp., $\Delta = \eta^p - \eta^b$). Considering experiments that perturb a single state at a time, we define impulse responses as suitably scaled differences in the expected trajectory of $\eta_{i,t+h}$ for marginal perturbations Δ , holding every other past and current state constant.

To select the perturbation Δ , we introduce a rule that maps it to a system of comparable units of change π . The rule is given by a function g such that

$$g(Z^b + \Delta(\pi)) - g(Z^b) = \pi \quad \text{or} \quad g(\eta^b + \Delta(\pi)) - g(\eta^b) = \pi.$$

The rule g may depend on the benchmark value and the reference value for the remaining

states but we omit the dependence from the notation. We will focus on the *unit rule*:

$$g(z) = z, \quad \text{leading to } \Delta(\pi) = \pi.$$

The unit rule is natural if the perturbed state is measured in money terms, or if it is in logs in which case the impulse consists of the same approximate $100 \times \pi\%$ change applied to all individuals. Other choices are possible, depending on the empirical context. For example, under a *rank rule*, $g(z)$ is the CDF of the perturbed state given the other state and past states $\{\eta_{i,t-1-\ell}, Z_{t-\ell}\}_{\ell>1}$. The rank rule is appropriate when the state is a concept with no natural unit of measurement, such as welfare. For the micro IRF, g can be thought of as an income transfer program implemented by a social planner or policymaker.

Given a rule g , the macro and micro impulse responses are, respectively,

$$\begin{aligned} \text{IRF}_{\eta Z}(h, \pi) &= \frac{E[\eta_{i,t+h} \mid \eta_{i,t-1}, Z_t = Z^b + \Delta(\pi), Z_{t-1}] - E[\eta_{i,t+h} \mid \eta_{i,t-1}, Z_t = Z^b, Z_{t-1}]}{\pi}, \\ \text{IRF}_{\eta\eta}(h, \pi) &= \frac{E[\eta_{i,t+h} \mid \eta_{it} = \eta^b + \Delta(\pi), Z_{t+1}, Z_t] - E[\eta_{i,t+h} \mid \eta_{it} = \eta^b, Z_{t+1}, Z_t]}{\pi}, \end{aligned}$$

where we hold $(\eta_{i,t-1}, Z_t, Z_{t-1})$, for $\text{IRF}_{\eta Z}$, and $(\eta_{it}, Z_{t+1}, Z_t)$, for $\text{IRF}_{\eta\eta}$, fixed.

Impulse responses are functions of the parameters that govern Q_η and Q_Z , the estimation of which we described in Section 4. In addition, it is possible to link impulse responses to the ρ and β measures of Section 5, and to derivatives with respect to locally-defined shocks. We explore those links in Supplemental Appendix G.

6.2 Empirical estimates of impulse responses

Macro impulse responses. We begin with estimates of $\text{IRF}_{\eta Z}$ for a negative perturbation to Z_t around the steady state benchmark $Z^b = \tilde{Z}_{ss}$. This emulates an aggregate shock that tips the economy into a recession. We calibrate $\pi = -2\sigma_V$ with $\sigma_V^2 = \text{Var}(Z_t \mid Z_{t-1})$ to match a shock comparable to the Great Recession and, to facilitate interpretation, we multiply impulse responses by -1 . Figure 6 shows the results.

It is instructive to consider first the trajectory of Z_t in panel (d), which is annualized by averaging the quarterly responses and scaled by the standard deviation of log GDP. The normalization translates responses into log deviations from the GDP per capita trend. By this account, our underlying experiment implies a GDP roughly 2.3 and 1.3 percentage points (pp) below trend in years zero and one after the shock, respectively, returning to trend




FIGURE 6. Macro impulse responses

Note: Panels (a), (b) and (c) show IRFs of η_{it} to a negative macro shock for different income measures with $Z_t^b = Z_{t-1} = \tilde{Z}_{ss}$ and $\eta_{i,t-1}$ set to the 10th (low), 50th (middle) and 90th (high) percentiles of the persistent income distribution. Panel (d) shows the IRF of Z_t annualized and scaled to detrended log GDP per capita. IRFs are multiplied by -1 . Shaded areas are 90% pointwise confidence bands.

afterwards with a slight overshoot in year three.

Next, panels (a), (b) and (c) report macro impulse responses of η_{it} for our three income measures and three initial levels of income: the 10th, 50th and 90th percentiles of the $\eta_{i,t-1}$ -distribution. We highlight the following takeaways. First, the responses are quantitatively consistent with the dynamics of GDP described above. For example, averaging over the distribution of $\eta_{i,t-1}$, male earnings fall by 2.8 pp on impact following the macro shock, while household earnings and disposable income fall by 2.3 and 1.4 pp. Furthermore, both disposable income and household earnings are near their pre-shock trends after two years, but the effects on male earnings are slightly more persistent.²⁰

Second, household earnings appear less cyclically sensitive than male earnings but more than disposable income. This mimics the discussion of Section 5 and is suggestive of the role of spousal income and the tax-transfer system as potential sources of insurance against

²⁰Supplemental Appendix G.4 presents local projection estimates of income responses to V_t that uncover similar patterns. They also point to a significant response at $h = 1$ that our biennial setup cannot measure.

aggregate shocks. Third, the macro IRF is U-shaped in $\eta_{i,t-1}$, with lower (but still significant) responses in the middle of the income distribution.

An advantage of our framework is the possibility to measure the interplay between macro and micro shocks. This is illustrated in Figure 7 where we compute a modified version of $\text{IRF}_{\eta Z}(h, \pi)$ that further conditions on the micro shock u_{it} contemporaneous to the macro perturbation, averaged over the distribution of $\eta_{i,t-1}$. This allows us to quantify the micro quantile treatment effects associated with aggregate shocks.




FIGURE 7. Interaction between macro impulse responses and micro shocks

Note: Panels (a), (b) and (c) show macro IRFs conditional on the micro shock u_{it} averaged over the distribution of $\eta_{i,t-1}$ for disposable income, household earnings, and male earnings with $Z_t^b = Z_{t-1} = \tilde{Z}_{ss}$. A magenta (cyan) area indicates the response is statistically negative (positive) at the 5% significance level.

Figure 7 shows sizable heterogeneity in the impact of macro shocks along the micro-rank distribution. For disposable income, an individual subject to a bad shock $u_{it} = 0.1$ suffers an expected income loss of 1.9 pp, well above the average of 1.4 pp. The number is 1.3 and 1.4 pp for $u_{it} = 0.5$ and $u_{it} = 0.9$. We find similar (but more pronounced) U-shaped patterns in our household and male earnings estimates; see panels (b) and (c).

Micro impulse responses. We conclude this section with estimates of the micro impulse responses $\text{IRF}_{\eta\eta}$ for a negative perturbation π that implies a 10% reduction in η_{it} . We hold Z_{t+1} and Z_t at their steady state value $\tilde{Z}_{ss}$ and multiply responses by -0.1 for ease of interpretation. See Figure 8.

The main takeaway, consistent across income measures, is that micro responses decay slowly and at different rates that depend on the initial level of income, with less (more) persistence for low- η (high- η) units. This reflects an intrinsic connection between nonlinear persistence ρ and $\text{IRF}_{\eta\eta}$ that we discuss in Supplemental Appendix G.

Positive shocks. Supplemental Appendix G.5 reports IRFs to positive perturbations.




FIGURE 8. Micro impulse responses

Note: Panels (a), (b) and (c) display IRFs of η_{it} to a negative micro shock for different income measures with $Z_{t+1} = Z_t = \tilde{Z}_{ss}$ and η_{it}^b set to the 10th (low), 50th (middle) and 90th (high) percentiles of the distribution of η_{it} . Shaded areas are 90% pointwise confidence bands.

7 Risk quantification

In this section, we use our framework to quantify aggregate and idiosyncratic contributions to income risk. We conclude by discussing empirical estimates.

7.1 Measuring the contribution to risk of macro and micro shocks

Our approach to risk quantification relies on the indirect utility of persistent income. Let $U(e^{\eta_{it}})$ be the period utility of a household whose persistent log income component at time t is η_{it} . We focus on indirect utility as we do not have consistent data on consumption, and abstract away from the transitory component ε_{it} as it may contain measurement error. In addition, we rely on disposable income but we do not consider additional sources of insurance beyond taxes and transfers. Following [Lucas \(1987, 2003\)](#), we measure the risk contribution of shocks by compensating variation: the value CV such that

$$E^* \left[\sum_{h=1}^H \delta^h U \left((1 - CV) e^{\eta_{i,t+h}} \right) \middle| \eta_{it}, Z_t \right] = E \left[\sum_{h=1}^H \delta^h U \left(e^{\eta_{i,t+h}} \right) \middle| \eta_{it}, Z_t \right]. \quad (12)$$

In (12), E denotes expectations over the actual persistent income process, while E^* denotes expectations under a counterfactual income process where the macro or the micro shocks have been removed. Here CV may depend on (η_{it}, Z_t) but we leave it implicit.

The quantity CV measures the fraction of income the agent would be willing to forgo in every period in order to eliminate a source (macro or micro) of income risk. Much of the literature emphasizes the role of curvature in preferences for the cost of risk. A common

finding is that log-utility with an exponential income process implies little aggregate risk, and that high risk aversion is needed to obtain even moderate costs of business cycles. Here we highlight a different channel that we find matters greatly for inferring the magnitude of aggregate income risk: the nonlinear relationship between the income process and the aggregate factor. Indeed, our model allows for a general nonlinear relationship between η_{it} and (Z_t, Z_{t-1}) ; see (1). These nonlinearities, and more specifically a countercyclical β (as documented in Section 5.4), can generate larger welfare costs of risk than usually found.

To gain intuition, we compute a small-noise expansion of CV around a baseline that sets shocks to their median value. We present below our calculations setting $H = 1$ for the measure of macro risk CV_{macro} where the experiment eliminates the macro shocks only.²¹ Let $\eta_{i,t+1}(Z_{t+1}) = Q_\eta(\eta_{it}, Z_{t+1}, Z_t, x_{i,t+1}, 0.5)$ and $\beta_{i,t+1}(Z_{t+1}) = \beta(0.5, \eta_{it}, Z_{t+1}, Z_t, x_{i,t+1})$ be the quantile function (1) and the exposure to aggregate shocks (5) at a median micro shock $u_{i,t+1} = 0.5$, leaving the dependence on $(\eta_{it}, Z_t, x_{i,t+1})$ implicit. Let $\eta_{i,t+1}$ and $\beta_{i,t+1}$ denote their values in the absence of a macro shock. Letting $\gamma(\eta) = -e^\eta U''(e^\eta)/U'(e^\eta)$ be the relative risk aversion, we obtain

$$CV_{\text{macro}} \approx \underbrace{\frac{\sigma_V^2}{2} (\gamma(\eta_{i,t+1}) - 1) \beta_{i,t+1}^2}_{\text{risk aversion channel}} + \underbrace{\frac{\sigma_V^2}{2} \left(-\frac{\partial \beta_{i,t+1}}{\partial Z_{t+1}} \right)}_{\text{cyclical exposure channel}}. \quad (13)$$

Equation (13) shows that CV_{macro} captures the combination of two effects: the curvature of the utility function and nonlinearities in the exposure to aggregate shocks. When risk aversion is moderate, the cost of business-cycle risk depends crucially on how nonlinear the income process is with respect to Z_{t+1} . For log utility (i.e., $\gamma(\eta) = 1$) and a linear income process, CV_{macro} is approximately zero, but it can be large when aggregate shock exposures are countercyclical, that is, when $\partial \beta_{i,t+1} / \partial Z_{t+1} < 0$. This reflects the self-amplifying nature of recessions and aligns with the empirical findings we discussed in Section 5.4.

7.2 Empirical estimates of macro and micro costs of risk

We present below our compensating variation estimates for eliminating macro and micro shocks for disposable income.²² To investigate the role of nonlinearities in the aggregate

²¹Details of the expansion and the formula for arbitrary H can be found in Supplemental Appendix H. An analogous approximation applies to the measure of micro risk CV_{micro} .

²²In a model without endogenous labor supply, this is the relevant income measure to study consumption and welfare. In that context, the indirect utility of persistent income $U(e^{\eta_{it}})$ can be rationalized by a

shocks, in addition to computing risk measures from the full model estimated in Section 5, we also consider a version in which $\beta(u, \eta, Z_t, Z_{t-1}, x)$ is constant in (Z_t, Z_{t-1}) , even though it can depend on (u, η, x) .²³ In the restricted model, β is acyclical in contrast to the countercyclical β we find in our full model.




FIGURE 9. Risk quantification (disposable income).

Note: We show compensating variation for aggregate (upper panels) and idiosyncratic risk (lower panels) for various ages $x = x_{it}$ and initial incomes η_{it} . Our full model (where β is countercyclical) is on the left and a restricted acyclical- β model on the right. Lifetime utility is $\sum_{h=1}^{(65-x_{it})/2} \delta^h \frac{e^{(1-\gamma)\eta_{i,t+h}}}{(1-\gamma)}$ with $\delta = (0.96)^2$ and $\gamma = 3$. Gray shaded areas are 90% confidence bands.

variety of consumption functions: e.g., hand-to-mouth households if ε_{it} is mostly measurement error or unconstrained permanent-income households if $E_t[e^{y_{i,t+h}} | \eta_{it}, \varepsilon_{it}, Z_t]$ is proportional to $e^{\eta_{it}}$.

²³We achieve this by restricting ψ to a first-order polynomial in $(\tilde{Z}, Z)$ in the specification described at the end of Section 4. We fit this model to the data using the same approach as for our full model.

Figure 9 illustrates the results. Left panels show risk calculations from our full model while right panels show the restricted model, with macro risk above and micro risk below. Period utility is $U(c) = c^{1-\gamma}/(1-\gamma)$ with risk aversion coefficient $\gamma = 3$. We set the biennial discount factor to $\delta = (0.96)^2$ and we calibrate the horizon for lifetime utility to the number of biennial periods until a notional retirement age of 65, i.e., $H = (65 - x_{it})/2$.

The main highlight from Figure 9 is the striking difference in the cost of business-cycle risk between linear and nonlinear income exposures to aggregate shocks. In line with our analytical derivation, with an acyclical β (right), CV_{macro} is virtually zero, supporting the conclusion in Lucas (2003) that eliminating business-cycle fluctuations is second-order from a welfare point of view. In stark contrast, a countercyclical β (left) produces CV_{macro} orders of magnitude higher, with households willing to give up between 1.9% and 4.5% of their income each period in order to avoid cyclical fluctuations. Since our full model nests the restricted one and CV_{macro} is statistically non-zero in the former, the evidence favors the view that aggregate shocks have non-negligible welfare costs.²⁴ It is worth highlighting that the statistical uncertainty for the aggregate risk measure is affected by the presence of cross-sectional dependence; see Supplemental Appendix E for confidence bands that account for such dependence.

Turning to the cost of idiosyncratic risk in Figure 9, perhaps unsurprisingly, CV_{micro} is higher than CV_{macro} even in the countercyclical- β case, but there is ample variation over age and across the income distribution. To give a sense, the cost of macro risk is around 10% of the cost of micro risk for older and richer households; it is more than 40% for young low-income households. This aligns with the tale of two skewnesses documented in Section 5.5: For low- η units, micro risk is primarily upside risk while aggregate shocks (given their self-amplifying effects) carry downside risk. As η increases (or as households age), conditional skewness becomes more negative and the cost of micro risk increases. Estimates of micro risk are precisely estimated, even when accounting for the presence of cross-sectional dependence (see Supplemental Appendix E).

Interestingly, the restricted model and our full model have similar implications for CV_{micro} . The nonlinearities that underlie the aggregate and idiosyncratic components of income risk are separate features, targeted by different parts of our framework. A natural question is whether agents have different amounts of information about the macro- and micro-relevant

²⁴This conclusion changes very little for larger values of the risk aversion coefficient γ . Instead, the cost of risk under countercyclical β is higher for larger persistence and variance of aggregate shocks, confirming that the interaction between $\partial\beta_{it}/\partial Z_t$ and marginal utility is, empirically, the relevant dimension of the income process for risk quantification.

features, and how much this matters for risk quantification. A second issue is the impact of the assumption that Z_t follows a linear process for welfare calculations. A third question is how households having access to various forms of insurance, including self-insurance, would modify the quantification of aggregate and idiosyncratic risk. We leave these important questions for future research.

8 Exploring heterogeneity by age and education

In this section, we examine the age and education dimensions more closely. Turning first to age, our nonlinear model allows for flexible patterns over the life cycle as age already enters Q_η . Moreover, the risk calculations reported in Figure 9 indicate an important role for age, with young low-income households suffering a larger cost of macro risk. A full set of estimates for our main target quantities split by age, for young (25 years old), prime-age (35 to 45) and older workers (55), is provided in Supplemental Appendix I. Here we highlight some findings for persistence and skewness; see Figure 10.

Panel (a) of Figure 10 shows that nonlinear persistence for low income households hit by good shocks is lower for the young, reflecting a higher probability of upward income mobility. However, the cyclical patterns of persistence do not vary strongly with age. The results in Supplemental Appendix I show that these cyclical patterns also hold for exposure coefficients, although for male earnings the young face higher exposure to aggregate shocks in recessions especially compared to older workers. The results for skewness, presented in panel (b) of Figure 10, show that older workers see the strongest negative skewness, especially those with higher past persistent income $\eta_{i,t-1}$. For male earnings, the larger negative skewness of older workers displays a steeper gradient with past income in recessions compared to expansions. However, for disposable income—that is, including spousal earnings, taxes and transfers—the pattern of skewness shows little difference between recessions and expansions, even though age differences remain. Adding in these income sources appears to mitigate the impact of recessions on skewness.

In addition to age, our model captures various forms of household heterogeneity through a flexible conditioning on past persistent income $\eta_{i,t-1}$ and a rich, subpanel-specific modeling of initial conditions η_{i,t_0} . However, our baseline model assumes that, given age and the aggregate state, Q_η is homogeneous across households. To relax this assumption, we have estimated the full model separately for two education groups: households whose head has at most a high school degree, and those with some college education. In Supplemental Appendix




FIGURE 10. Persistence and skewness by age

Note: Panel (a) reports the persistence measure $\rho(u, \eta, Z_t, Z_{t-1}, x)$ by quantile of the shock $u = u_{it}$ and past persistent income $\eta = \eta_{i,t-1}$ for disposable income and male earnings with age $x = x_{it}$ set to either 25 or 55, $Z_{t-1} = \tilde{Z}_{ss}$ and Z_t is a recession $\tilde{Z}_r$, the steady state $\tilde{Z}_{ss}$ or an expansion $\tilde{Z}_e$ (see Section 5.1). Panel (b) reports the skewness measure $sk(\eta, Z_t, Z_{t-1}, x)$ by quantile of past persistent income $\eta = \eta_{i,t-1}$ with the same format. Shaded areas in panel (b) represent 90% confidence bands.

I we present results based on this specification where parameters may differ across education groups. We find that neither nonlinear persistence nor exposure to aggregate shocks differ markedly by education. In addition, our compensating variation estimates of the cost of aggregate and idiosyncratic risk are similar for the two groups. That being said, statistical uncertainty is larger within groups, as reflected by the size of confidence bands.²⁵

²⁵Moreover, variation along the income distribution (for example, in persistence) is now conditional on education, and hence not directly comparable to the income variation in our main estimates (in Figure 3).

9 Conclusion

In this paper, we propose a nonlinear framework for income processes with aggregate shocks. While our interest is in the impact of macro shocks on nonlinear income risk, our approach provides a basis for building rich nonlinear reduced forms for heterogeneous-agent models with both macro and micro uncertainty. The nonlinear reduced-form relationships our approach can uncover are useful to assess the fit of such models and their implications. Allowing for nonlinear income dynamics that permit rich interactions between aggregate and idiosyncratic shocks, we study identification, estimation, and inference tools that leverage both macro and micro data.

In our empirical analysis of U.S. panel data on income, we find that the nonlinear persistence and conditional skewness patterns documented in previous work are affected by business-cycle fluctuations. The persistence surface tilts in recessions, decreasing for high-income households hit by a bad shock and increasing for low-income households, while skewness declines throughout the income distribution. We also find evidence of nonlinear exposures to aggregate shocks, with higher sensitivity of income to macro shocks during recessions and nontrivial interactions between macro and micro shocks. Measures of persistence, skewness and exposure also display systematic differences by age. Our results further suggest that nonlinearities with respect to aggregate shocks matter for risk quantification. One important avenue to explore in future research is how to account for model uncertainty in our risk measurement approach.

Three natural extensions, which we leave to future work, are particularly promising. The first is to study the pass-through of aggregate and idiosyncratic income shocks to consumption. This would allow us to examine the degree of insurance against persistent and short-lived income shocks across the business cycle. We would follow [Arellano et al. \(2017\)](#) and [Arellano et al. \(2023\)](#), where consumption data from the PSID were combined with a nonlinear income process to quantify the impact of income shocks on consumption. Those studies found that consumption responses differ by initial wealth and household heterogeneity. The framework developed in this paper would enable estimation of the transmission of aggregate shocks, reflecting not only the business cycle but also variation in monetary and fiscal policy, to consumption, as well as the differential impact of idiosyncratic income shocks by aggregate state.

The second extension relates to the primary components of household income: male and female earnings. These two earnings measures embody the labor supply responses of the

spouses. In a linear partial insurance framework, [Blundell, Pistaferri, and Saporta-Eksten \(2016\)](#) investigate the responses of spousal hours of work and consumption to wage shocks in the PSID and uncover an important role for family labor supply in insuring permanent and transitory wage shocks. However, they do not consider nonlinear wage processes and ignore the impact of aggregate factors. The additional insights that would come from extending the family labor supply insurance framework to incorporate nonlinear wage dynamics and the business cycle make this an exciting area for future research.

A third extension, related to the second one, concerns the zeros in the income data. Figure 2 showed a small but growing and cyclical proportion of zeros for male earnings in our couples sample, a smaller and more stable proportion of zeros for household earnings, and no zeros for disposable income. This is one reason why we have focused on disposable income. Nonetheless, we have also drawn conclusions in comparison with male and household earnings. Modeling the extensive margin requires some assumptions on the process of selection. A natural framework is to assume random missingness conditional on the history of the income process. But it is likely that aggregate shocks also play a role here. [Braxton, Herkenhoff, Rothbaum, and Schmidt \(2021\)](#) develop a Kalman filter and EM algorithm approach to incorporate observations with zero (or missing) earnings. By specifying a law of motion for persistent earnings, they show how assuming a model for earnings upon re-entry can deliver identification under selection. In another recent paper, [Gobillon, Magnac, and Roux \(2022\)](#) study earnings dynamics in French administrative data assuming zeros are missing at random conditional on a set of factors drawn from earnings history. The nonlinear framework with aggregate factors developed in our paper provides a promising setting for incorporating the extensive margin into the study of income dynamics over the cycle.

References

- ACHARYA, S., E. CHALLE, AND K. DOGRA (2023): “Optimal monetary policy according to HANK,” *American Economic Review*, 113, 1741–1782.
- AHN, S., G. KAPLAN, B. MOLL, T. WINBERRY, AND C. WOLF (2018): “When Inequality Matters for Macro and Macro Matters for Inequality.” *NBER Macroeconomics Annual*, 32, 1–75.
- ALMUZARA, M., M. ARELLANO, R. BLUNDELL, AND S. BONHOMME (2025): “Nonlinear Micro Income Processes with Macro Shocks.” Working paper.

- ALMUZARA, M. AND V. SANCIBRIÁN (2024): “Micro Responses to Macro Shocks.” Federal Reserve Bank of New York Staff Report 1090.
- ALTONJI, J. G., A. A. SMITH, AND I. VIDANGOS (2013): “Modeling earnings dynamics.” *Econometrica*, 81, 1395–1454.
- AMBERG, N., T. JANSSON, M. KLEIN, AND A. ROGANTINI PICCO (2022): “Five Facts about the Distributional Income Effects of Monetary Policy Shocks.” *AER: Insights*, 4, 289–304.
- ANDERSEN, A. L., N. JOHANNESSEN, M. JØRGENSEN, AND J.-L. PEYDRÓ (2023): “Monetary Policy and Inequality.” *Journal of Finance*, 78, 2945–2989.
- ANGELETOS, G.-M., F. COLLARD, AND H. DELLAS (2020): “Business-Cycle Anatomy.” *American Economic Review*, 110, 3030–3070.
- ARELLANO, M. (2014): “Uncertainty, persistence and heterogeneity: A panel data perspective.” *Journal of the European Economic Association*, 12, 1127–1153.
- ARELLANO, M., R. BLUNDELL, AND S. BONHOMME (2017): “Earnings and consumption dynamics: A nonlinear panel data framework.” *Econometrica*, 85, 693–734.
- ARELLANO, M., R. BLUNDELL, S. BONHOMME, AND J. LIGHT (2023): “Heterogeneity of Consumption Responses to Income Shocks in the Presence of Nonlinear Persistence,” *Journal of Econometrics*, 1–45.
- ARELLANO, M. AND S. BONHOMME (2016): “Nonlinear panel data estimation via quantile regressions.” *Econometrics Journal*, 19, C61–C94.
- (2017): “Nonlinear panel data methods for dynamic heterogeneous agent models.” *Annual Review of Economics*, 9, 471–496.
- BARLEVY, G. (2004): “The Cost of Business Cycles Under Endogenous Growth.” *American Economic Review*, 94, 964–990.
- BHANDARI, A., D. EVANS, M. GOLOSOV, AND T. J. SARGENT (2021): “Inequality, Business Cycles, and Monetary-Fiscal Policy.” *Econometrica*, 89, 2559–2599.
- BLUNDELL, R., C. R. BOLLINGER, C. HOKAYEM, AND J. P. ZILIAK (2024): “Interpreting Cohort Profiles of Lifecycle Earnings Volatility.” Working Paper CES-24-21.
- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): “Consumption inequality and partial insurance.” *American Economic Review*, 98, 1887–1921.

- BLUNDELL, R., L. PISTAFERRI, AND I. SAPORTA-EKSTEN (2016): “Consumption inequality and family labor supply.” *American Economic Review*, 106, 387–435.
- BORELLA, M., M. DE NARDI, M. PAK, N. RUSSO, AND F. YANG (2023): “The Importance of Modeling Income Taxes Over Time.” *Journal of the European Economic Association*, 21, 2237–2286.
- BRAXTON, J. C., K. F. HERKENHOFF, J. L. ROTHBAUM, AND L. D. W. SCHMIDT (2021): “Changing Income Risk across the US Skill Distribution: Evidence from a Generalized Kalman Filter.” NBER Working Paper 29567.
- BROWNING, M., M. EJRNE, AND J. ÁLVAREZ (2010): “Modelling income processes with lots of heterogeneity.” *Review of Economic Studies*, 77, 1353–1381.
- BURNS, A. F. AND W. C. MITCHELL (1946): *Measuring Business Cycles*, NBER.
- CHANG, M., X. CHEN, AND F. SCHORFHEIDE (2024): “Heterogeneity and Aggregate Fluctuations.” *Journal of Political Economy*, 132.
- CHETTY, V. K. (1968): “Pooling of Time Series and Cross Section Data.” *Econometrica*, 36, 279–290.
- CREAL, D. (2012): “A Survey of Sequential Monte Carlo Methods for Economics and Finance.” *Econometric Reviews*, 31, 245–296.
- FERNÁNDEZ-VILLAYERDE, J., S. HURTADO, AND G. NUÑO (2023): “Financial Frictions and the Wealth Distribution.” *Econometrica*, 91, 869–901.
- FLECK, J., J. HEATHCOTE, K. STORESLETTEN, AND G. L. VIOLANTE (2025): “Fiscal Progressivity of the U.S. Federal and State Governments.” NBER Working Paper 33385.
- GALLANT, A., P. ROSSI, AND G. TAUCHEN (1993): “Nonlinear dynamic structures,” *Econometrica*, 61, 871–907.
- GALÍ, J., M. GERTLER, AND J. D. LÓPEZ-SALIDO (2007): “Markups, Gaps, and the Welfare Costs of Business Fluctuations.” *Review of Economics and Statistics*, 89, 44–59.
- GOBILLON, L., T. MAGNAC, AND S. ROUX (2022): “Lifecycle Wages and Human Capital Investments: Selection and Missing Data.” Working Paper TSE Working Paper 22-1299.
- GOTTSCHALK, P. AND R. MOFFITT (1994): “The Growth of Earnings Instability in the U.S. Labor Market.” *Brookings Papers on Economic Activity*, 1994, 217–272.

- GOTTSCHALK, P. AND R. A. MOFFITT (2009): “The rising instability of U.S. earnings.” *Journal of Economic Perspectives*, 23, 3–24.
- GUVENEN, F., F. KARAHAN, S. ÖZKAN, AND J. SONG (2021): “What Do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Dynamics?” *Econometrica*, 89, 2303–2339.
- GUVENEN, F., A. MCKAY, AND C. B. RYAN (2023): “A Tractable Income Process for Business Cycle Analysis.” NBER Working Paper 30959.
- GUVENEN, F., S. OZKAN, AND J. SONG (2014): “The nature of countercyclical income risk.” *Journal of Political Economy*, 122, 621–660.
- GUVENEN, F., L. PISTAFERRI, AND G. L. VIOLANTE (2022): “Global trends in income inequality and income dynamics: New insights from GRID.” *Quantitative Economics*, 13, 1321–1360.
- GUVENEN, F., S. SCHULHOFER-WOHL, J. SONG, AND M. YOGO (2017): “Worker Betas: Five Facts about Systematic Earnings Risk.” *American Economic Review*, 107, 398–403.
- HAHN, J., G. KUERSTEINER, AND M. MAZZOCCO (2020): “Estimation with Aggregate Shocks.” *Review of Economic Studies*, 87, 1365–1398.
- HALVORSEN, E., H. A. HOLTER, S. OZKAN, AND K. STORESLETTEN (2024): “Dissecting Idiosyncratic Earnings Risk.” *Journal of the European Economic Association*, 22, 617–668.
- HAUSE, J. C. (1977): “The Covariance Structure of Earnings and the On-The-Job Training Hypothesis.” in *Annals of Economic and Social Measurement, Volume 6, number 4*, NBER, 335–365.
- HOLM, M. B., P. PAUL, AND A. TISCHBIREK (2021): “The Transmission of Monetary Policy under the Microscope.” *Journal of Political Economy*, 129, 2861–2904.
- KOENKER, R. AND G. BASSETT (1978): “Regression quantiles.” *Econometrica*, 46, 33–50.
- KREBS, T. (2007): “Job Displacement Risk and the Cost of Business Cycles.” *American Economic Review*, 97, 664–686.
- KRUEGER, D., K. MITMAN, AND F. PERRI (2016): “Macroeconomics and Household Heterogeneity.” in *Handbook of Macroeconomics*, ed. by J. Taylor and H. Uhlig, North-Holland, vol. 2A.
- KRUSELL, P. AND A. A. SMITH (1998): “Income and wealth heterogeneity in the macroeconomy.” *Journal of Political Economy*, 106, 867–896.
- LILLARD, L. A. AND R. J. WILLIS (1978): “Dynamics aspects of earning mobility.” *Econometrica*, 46, 985–1012.

- LIU, L. AND M. PLAGBORG-MØLLER (2023): “Full-information estimation of heterogeneous agent models using macro and micro data.” *Quantitative Economics*, 14, 1–35.
- LUCAS, ROBERT E., J. (1987): *Models of Business Cycles*, Basil Blackwell.
- (2003): “Macroeconomic Priorities.” *American Economic Review*, 93, 1–14.
- MACCURDY, T. E. (1982): “The use of time series processes to model the error structure of earnings in longitudinal data analysis.” *Journal of Econometrics*, 18, 83–114.
- MADDALA, G. S. (1971): “The Likelihood Approach to Pooling Cross-Section and Time-Series Data.” *Econometrica*, 39, 939–953.
- MEGHIR, C. AND L. PISTAFERRI (2004): “Income variance dynamics and heterogeneity.” *Econometrica*, 72, 1–32.
- OTROK, C. (2001): “On Measuring the Welfare Cost of Business Cycles.” *Journal of Monetary Economics*, 47, 61–92.
- PATTERSON, C. (2023): “The Matching Multiplier and the Amplification of Recessions.” *American Economic Review*, 113, 982–1012.
- SARGENT, T. J. AND Y. J. SELVAKUMAR (2025): “Aggregate Shocks and Cross-Section Dynamics: Quantifying Redistribution and Insurance in U.S. Household Data.” Working paper.
- STOCK, J. H. AND M. W. WATSON (2016): “Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics.” in *Handbook of Macroeconomics*, ed. by J. B. Taylor and H. Uhlig, Elsevier, vol. 2, chap. 8.
- STORESLETTEN, K., C. I. TELMER, AND A. YARON (2001): “The Welfare Cost of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk.” *European Economic Review*, 45, 1311–1339.
- (2004): “Cyclical dynamics in idiosyncratic labor market risk.” *Journal of Political Economy*, 112, 695–717.
- TOBIN, J. (1950): “A Statistical Demand Function for Food in the U.S.A.” *Journal of the Royal Statistical Society. Series A (General)*, 113, 113–141.
- WEI, Y. AND R. J. CARROLL (2009): “Quantile regression with measurement error,” *Journal of the American Statistical Association*, 104, 1129–1143.