

Inequality and Business Cycles*

Florin Bilbiie[†]

Giorgio E. Primiceri[‡]

Andrea Tambalotti[§]

June 2024. First version: May 2022

Abstract

We quantify the connection between inequality and business cycles in a medium-scale New Keynesian model with tractable household heterogeneity, estimated with aggregate and cross-sectional data. We find that inequality substantially amplifies cyclical fluctuations. The primary source of this amplification is cyclical precautionary saving behavior. Savers reduce their consumption to insure themselves against the idiosyncratic risk of large income drops, which rises in recessions.

1 Introduction

Inequality is one of the defining issues of our time. Economic research on this topic has traditionally focused on its secular evolution, since most measures of inequality tend to move slowly over time. For instance, figure 1 shows the evolution of the distribution of labor earnings in the U.S. since 1967, from [Heathcote et al. \(2020\)](#). The dramatic and steady rise of income inequality over the past five decades is evident. But its cyclicity is at least as striking. This paper focuses on the latter.

*We thank Manuel Amador, Marco Bassetto, Christian Bayer, Javier Bianchi, Luigi Bocola, Benjamin Born, Tobias Broer, Davide Debortoli, Marco Del Negro, Ruben Dominguez-Diaz, Axelle Ferrière, Jordi Galí, Jose Elias Gallegos, Jonathan Heathcote, Diego Känzig, Fabrizio Perri, Carolin Pflueger, Christian Pröbsting, Matt Rognlie, Frank Schorfheide, Ludwig Straub, Paolo Surico, Gianluca Violante and many seminar and conference participants for comments and suggestions. Giovanni Sciacovelli has provided invaluable research assistance. Bilbiie is grateful to the Janeway Institute and the Keynes Fund for financial support. Tambalotti's contribution was part of a continued collaboration based on work done prior to joining Amazon. This paper and its contents are not related to Amazon and do not reflect the position of the company and its subsidiaries.

[†]University of Cambridge and CEPR

[‡]Northwestern University, CEPR and NBER

[§]Amazon




Figure 1: The evolution of the U.S. labor earnings distribution. Source: [Heathcote et al. \(2020\)](#).

The question that we address is if inequality matters for business cycles, and through which channels. A glance at the figure above suggests that the answer must be yes, given that the dispersion in incomes that it depicts varies between booms and recessions by similar magnitudes as the trend does over several decades. But do these fluctuations in inequality simply reflect the underlying business cycle dynamics of aggregate variables, or is the counter-cyclicality of inequality itself one of the mechanisms that shape how the U.S. economy experiences expansions and recessions?

To answer these questions, we estimate a dynamic stochastic general equilibrium (DSGE) model with heterogeneous households, which includes all the ingredients that the literature has shown to be necessary to fit the aggregate data (e.g. [Christiano et al., 2005](#), [Smets and Wouters, 2007](#) and [Justiniano et al., 2010](#)). To keep the framework as close as possible to the representative agent New Keynesian (RANK) models that have so far dominated business cycle studies in macroeconomics, we adopt the stylized view of household heterogeneity proposed by [Bilbiie \(2018, 2020\)](#), which is based on a long tradition of two-agent models, as in [Campbell and Mankiw \(1989\)](#), [Galí et al. \(2007\)](#), [Bilbiie \(2008\)](#) and [Eggertsson and Krugman \(2012\)](#). The two types of households in our model are savers (S) and hand-to-mouth (H). The former choose consumption and saving by maximizing intertemporal utility as standard representative con-

sumers. The latter consume all their post-tax income in every period.

The model features two main mechanisms through which heterogeneity might alter the nature of cyclical fluctuations, relative to an economy with a representative household. First, the presence of agents with high marginal propensity to consume (MPC) amplifies business cycles if their income is relatively more exposed to aggregate fluctuations, i.e. if inequality is counter-cyclical. Second, S agents in the model have higher income than H agents, but they face the risk of becoming the poorer H agents with a given probability in every period. This risk may rise in recessions, giving them a precautionary reason to save, which further propagates business cycles.

To evaluate the strengths of these propagation mechanisms, we estimate the model using the same aggregate data of [Smets and Wouters \(2007\)](#) and [Justiniano et al. \(2010\)](#), as well as information on the evolution of the cross-sectional standard deviation of labor income and post-tax income from the Current Population Survey (the data underlying figure 1). We also incorporate a priori information based on microeconomic studies on aspects of the model that affect the interaction between its cross-sectional and time-series behavior. For example, we calibrate the share of H agents to be in line with the average marginal propensity to consume out of tax rebates and other transitory income shocks estimated by [Souleles et al. \(2006\)](#), [Parker et al. \(2013\)](#) and others. And we choose the probability with which savers might become hand-to-mouth to match the standard deviation and kurtosis of the annual change in (log) individual incomes measured by [Guvenen et al. \(2021\)](#).

With the estimated model in hand, we can quantify the role of several aspects of inequality in the amplification and propagation of business cycles. We do so through counterfactual simulations that eliminate features of the model with heterogeneity that the literature has identified as potentially relevant for the transmission of shocks. We initially consider a counterfactual in which the marginal utility of consumption is equalized between the two classes of agents, as it would be in an economy with complete markets. This is our “no-inequality” benchmark. In this economy, business cycle fluctuations are less pronounced than in the data. The standard deviation of GDP growth is 27 percent less than in the data, while that of de-trended hours worked is 36 percent less. In addition, we document that this version of the model with perfect risk sharing fits the data worse than the baseline. The good fit of our model with tractable heterogeneity stands out in light of the poor forecasting accuracy of prominent HANK models, relative to their representative-agent counterpart, as demonstrated by [Acharya et al. \(2023\)](#).

We then show that most of the reduction in volatility experienced in the more equal econ-

omy is due to the elimination of individual risk and its cyclical pattern, rather than the cyclical-ity of inequality interacting with MPC heterogeneity. In fact, an economy in which inequality in marginal utilities is constant, but S agents are still subject to cyclical risk of becoming H , behaves very similarly to the baseline economy with cyclical inequality. The intuition of why cyclical inequality plays a quantitatively limited role in amplifying business cycles is not specific to our two-agent model, but generalizes to economies with more complex dimensions of heterogeneity. Matching both average MPCs from micro estimates and the high level of inequality in the data requires that high-MPC agents consume only a small fraction of total consumption. Therefore, their behavior has only a modest impact on aggregate fluctuations, despite the fact that their income is indeed more exposed to business cycles. This result is in line with the empirical findings of [Chang et al. \(2024\)](#), which are based on the estimation of a functional vector autoregression with aggregate and cross-sectional variables. Observe that the introduction of wealthy hand-to-mouth agents, as in [Kaplan et al. \(2014\)](#), is unlikely to change the nature of this result. These middle-class agents, in fact, do consume a larger fraction of total consumption, but their income is not very exposed to business cycles ([Heathcote et al., 2020](#)), limiting the bite of their high MPC. Finally, we stress that the fact that MPC heterogeneity plays a quantitatively small role in amplifying business cycle fluctuations does not at all imply that it is an unimportant feature of the model more generally. For example, the presence of high-MPC agents in our model is key to study questions related to the effectiveness of fiscal transfers, which is consistent with a large body of literature ([Kaplan and Violante, 2022](#)).

In sum, according to our estimated model, the primary channel through which inequality shapes business cycles is not its cyclicity, but its long-run level and the resulting effect of risk on the behavior of savers. Quantitatively, this risk is mostly driven by the decline in consumption associated with falling to the bottom of the income distribution. In reality, this occurrence is often associated with unemployment, as discussed by [Heathcote et al. \(2020\)](#), and modeled explicitly by [Challe et al. \(2017\)](#), [Ravn and Sterk \(2017\)](#), [Haan et al. \(2018\)](#), [Ravn and Sterk \(2020\)](#), [Broer et al. \(2022a\)](#), and [Graves \(2023\)](#).

Contribution and related literature

We quantify the role of household heterogeneity in amplifying business cycles, and assess what fraction of this amplification is due to the cyclical fluctuations in inequality versus idiosyncratic risk. Our paper is the first to do it in an estimated general equilibrium model. As such, our contribution is complementary to several recent articles that study related aspects of the same

problem.

For example, [Bayer et al. \(2019\)](#) use data from the Survey of Consumer Finances and the Survey of Income and Program Participation to document the positive impact of an increase in idiosyncratic uncertainty on the share of liquid assets in household portfolios. They reproduce these empirical patterns within a calibrated heterogeneous-agent new-Keynesian (HANK) model, and employ it to infer the effect of higher idiosyncratic uncertainty on economic activity. Instead, we start from a general equilibrium model estimated using macro and micro data, and use it to quantify the propagation role of countercyclical inequality and risk. Remarkably, the estimated time-varying probability of a large income fall in our model (our measure of income risk) has a 0.3 correlation with the standard deviation of persistent income shocks extracted by [Bayer et al. \(2019\)](#) from the micro data for the period from 1979 to 2013.

The channel of precautionary saving in response to idiosyncratic risk is of course central to the entire HANK literature, but very few papers in addition to [Bayer et al. \(2019\)](#) quantify its role in the propagation of fluctuations. Two important exceptions are [Challe et al. \(2017\)](#) and [Ravn and Sterk \(2020\)](#). The latter endogenize unemployment risk within a stylized HANK model, and our estimation results provide empirical support to their modeling effort. [Challe et al. \(2017\)](#) use an estimated tractable HANK model to isolate the role of unemployment risk in the recessions of 1990-91, 2001 and 2008-09, concluding that it played a large role only in the last one. Our notion of income risk is broader than the one of [Ravn and Sterk \(2020\)](#) and [Challe et al. \(2017\)](#), and its measurement informed by data on earnings and post-tax income inequality, resulting into its more systematic contribution to business cycles. In addition, and differently from [Challe et al. \(2017\)](#) and [Ravn and Sterk \(2020\)](#), our approach disentangles the amplification of aggregate fluctuations due to countercyclical risk from that due to countercyclical inequality interacting with heterogeneous MPCs ([Bilbiie, 2008](#), [Auclert, 2019](#), [Patterson, 2019](#)).

Importantly, several studies provide external or indirect evidence consistent with our finding about the importance of countercyclical risk for the amplification of business cycles. For example, [Carroll et al. \(2019\)](#) show that the U.S. saving rate goes up in recessions due to higher labor income/unemployment risk. Along similar lines, [Dominguez-Diaz \(2022\)](#) documents a rise in liquid savings in the form of bank deposits in response to an exogenous increase in income risk. In finance, a large literature argues that countercyclical labor income risk is key to match certain properties of asset prices, following the seminal work of [Constantinides and Duffie \(1996\)](#). More recently, [Pflueger et al. \(2020\)](#) show that a measure of perceived risk extracted from the cross section of stocks is strongly correlated with the real interest rate on safe assets, and that it

forecasts real investment and unemployment. This finding is consistent with ours—indeed, our estimated measure of income risk is tightly connected with the real rate and has a correlation of 0.42 with the measure of risk perception of [Pflueger et al. \(2020\)](#).

Our work is also very related to papers that estimate HANK models with time series data. We are of course not the first to estimate a HANK model, but our paper differs from its predecessors both methodologically and in terms of its substantive conclusions. For example, [Auclert et al. \(2020\)](#) concentrate on the role of inattention in replicating both the high average MPCs at the micro level and the hump-shaped impulse responses to monetary policy shocks at the aggregate level. They also stress the role of investment for the general propagation of exogenous disturbances. But these authors do not study the quantitative role of countercyclical income risk and precautionary saving motives in amplifying business cycles, which is the main result of our paper. Their model also abstracts from cyclical fluctuations in earnings inequality.

[Bayer et al. \(2020\)](#) illustrate how to conduct inference in HANK models with Bayesian methods using the state-space representation. They apply their methodology to the estimation of a medium-scale DSGE, augmented with heterogeneous households and shocks to taxes and income risk. These authors show that the model can explain the medium-frequency movements in the 10th percentile of the income and wealth distribution, even when it is estimated only using aggregate data. However, they are not concerned with evaluating the importance of market incompleteness for the propagation and amplification of aggregate fluctuations, which is instead the primary focus of our paper.

In this respect, our work is more connected to that of [Berger et al. \(2019\)](#), who assess the role of imperfect risk sharing for business cycles by performing a “wedge-accounting” exercise. Their methodology has the advantage of being robust to some possible model misspecifications. However, it is not suitable for counterfactual and policy analysis, since the wedges are endogenous objects.

At a more general level, our paper contributes to the literature exploring the macroeconomic implications of household heterogeneity in models with nominal rigidities. Existing quantitative HANK models with rich heterogeneity have predominantly been used to study the transmission of specific policies and shocks. Examples of this line of work include [Oh and Reis \(2012\)](#), [McKay and Reis \(2016\)](#), [McKay et al. \(2016\)](#), [Gornemann et al. \(2016\)](#), [Guerrieri and Lorenzoni \(2017\)](#), [Kaplan et al. \(2018\)](#), [Auclert \(2019\)](#), [Bayer et al. \(2019\)](#), [Auclert et al. \(2018\)](#) and [Hagedorn et al. \(2019\)](#).

The solution of these models is typically computationally demanding. Therefore, for simplicity, we build on their simpler two-agent version, which has been progressively enriched to capture quantitatively relevant propagation channels (e.g. [Galí et al., 2007](#), [Bilbiie, 2008, 2018, 2020](#), [Nistico 2016](#)). [Debortoli and Galí \(2018\)](#) conduct a formal comparison of a two-agent models and those with a richer structure of heterogeneity. Alternative approaches to simplifying heterogeneity and gain analytical insights into the transmission mechanism of these models include [Werning \(2015\)](#), [Acharya and Dogra \(2020\)](#), [Broer et al. \(2020\)](#), [Ravn and Sterk \(2020\)](#), and [Debortoli and Galí \(2023\)](#). Notice that our findings on the amplification of business cycles through heterogeneity differ from those of [Debortoli and Galí \(2023\)](#), because income risk is acyclical in their setting and consumption risk is thus concentrated among poor households who have little impact on the aggregate. On the contrary, our economy delivers amplification because it features countercyclical income risk faced by mid- and high-consumption agents, who save to self-insure against potentially large income drops.

This difference highlights a key advantage of the simple two-agent structure of our model, relative to other quantitative HANK DSGEs, that is the ability to perform laser-focused counterfactual experiments to shed light on various mechanisms through which heterogeneity, inequality and market incompleteness shape aggregate fluctuations. On the contrary, in HANK models with many agents, completing the markets to compute the counterfactual allocation under perfect risk sharing is challenging, and it can only be done by rewriting the model from scratch with a representative consumer. An additional benefit is that our approach entails only minimal extensions of the benchmark DSGE model of [Christiano et al. \(2005\)](#) and [Smets and Wouters \(2007\)](#) that is routinely used for policy analysis. Finally, our model fits the data well, which is not a common feature of HANK economies, as documented by [Acharya et al. \(2023\)](#). Naturally, however, this two-agent modeling strategy is only an approximation of a more complex reality with much broader dimensions of household heterogeneity. Therefore, it is a complementary tool to the models with rich heterogeneity reviewed above.

2 The Model

The model builds on [Justiniano et al. \(2010\)](#), JPT hereafter, and [Bilbiie \(2018\)](#). Its business cycle structure comes from the representative consumer economy in JPT. It includes nominal rigidities, and a host of real frictions and shocks necessary to fit the aggregate data, as first demonstrated by [Christiano et al. \(2005\)](#) and [Smets and Wouters \(2007\)](#). We add two important features

to this familiar framework, to address the role of income and consumption inequality in business cycles. First, households differ in their ability to access consumption smoothing opportunities and in the quality of their labor endowments. Second, firms combine the different kinds of labor supplied by households with capital through a constant elasticity of substitution (CES) production function inspired by [Krusell et al. \(2000\)](#). The degree of heterogeneity in the model economy is kept to a minimum to maximize its tractability. Yet, this limited degree of heterogeneity is sufficient to capture some of the salient cyclical features of inequality and several key channels through which inequality and idiosyncratic income risk shape business cycles.

The economy is populated by two classes of consumers: H , for hand-to-mouth, and S , for savers. Hand-to-mouth agents consume their entire disposable income, and their MPC out of any change in after-tax income is fixed at one. The savers are standard permanent-income maximizers, along the lines of the representative household in typical medium-scale DSGEs. Similar to those agents, they consume, supply labor, and accumulate capital that they rent to firms. Unlike typical representative agents, though, the savers face the risk of becoming hand-to-mouth. The realization of this idiosyncratic risk is unpleasant for two reasons. First, the H cannot smooth consumption. Second, their income and consumption are lower in steady state. Therefore, H agents tend to be income poor and their marginal utility is higher than that of S agents. As a result, the latter engage in precautionary saving to mitigate the adverse consequences of idiosyncratic shocks.

S and H agents are also endowed with two different types of labor, which are combined with capital in a CES production function. According to our estimates of this function, S hours are complementary to capital in production, while H hours are substitutes, like skilled and unskilled labor in [Krusell et al. \(2000\)](#). As a result of this difference in their substitution patterns, H workers face a more cyclical labor demand than S agents. The demand for S hours lags behind that for H work in response to positive business cycle shocks because capital accumulates slowly, restraining the increase in demand for its complementary labor input. In contrast, H labor is more readily available for production when demand surges, since it is less reliant on slowly accumulating capital to be productive. As a result, H wages, hours, and labor income are highly responsive to the business cycle. They surge in booms, but plunge in recessions. S workers instead face more stable labor demand conditions, which, together with a relatively inelastic labor supply, produce a smoother labor income. On net, these differences imply that H workers are more exposed to the business cycle and labor income inequality is strongly counter-cyclical, consistent with a growing body of empirical evidence (e.g. [Heathcote et al., 2010, 2020, 2023](#),

Guvenen et al., 2017, Patterson, 2019).

In summary, our framework enriches a standard medium-scale DSGE model of business cycles with a stylized theory of earnings, income and consumption inequality. This theory nevertheless captures many of the channels through which inequality affects business cycles that have been identified in the literature reviewed above. One group of households has higher and more stable labor income, access to consumption smoothing opportunities and therefore a low MPC out of transitory income shocks. Another group of households cannot smooth consumption and has lower income that is also more exposed to business cycle fluctuations. In the data, individuals with these characteristics tend to have low skills and are more likely to experience unemployment in recessions. This is an important reason why their income is low and cyclical, their wealth minimal, and thus their MPC is higher. Large literatures have developed models that capture the underlying mechanisms responsible for the cross-sectional correlation among two or more of these characteristics. Here, we abstract entirely from those mechanisms, such as human capital and skill accumulation, labor market frictions leading to unemployment, and financial market imperfections limiting the opportunity to smooth consumption. Instead, we take as given an extreme view of the observed correlation: our H agents have low and very cyclical incomes and a high MPC, while the opposite is true for S households. Even with this extreme assumption, our model can match several stylized facts about the interaction between income and consumption inequality and aggregate fluctuations. Therefore, we consider it a useful laboratory to inspect and quantify some of the channels through which inequality affects the transmission of business cycle shocks, as well as the extent to which business cycles shape observed inequality.

2.1 Households

There is a continuum of households of two types, H and S , with constant measure θ and $1 - \theta$ respectively. Households of type S stay S with probability s_t per period. They face the risk of becoming H with probability $1 - s_t$. This probability can vary over the cycle according to the function $s_t = s_0 (X_t/X_t^*)^{s_1}$, whose argument is a measure of the GDP gap defined below. Households of type H face a probability of switching back to S such that the measure of the two groups remains constant over time.

2.1.1 Consumption and saving decisions

Each household group $h = H, S$ includes a continuum of members indexed by $j \in [0, 1]$, each endowed with differentiated labor $L_{h,t}(j)$ earning a nominal wage $W_{h,t}(j)$, and consuming $C_{h,t}(j)$. The instantaneous utility of the j^{th} household of type $h = S, H$ is

$$u_{h,t}(j) = b_t \left[\log(C_{h,t}(j) - \eta_h C_{h,t-1}) - \varphi_h \frac{L_{h,t}(j)^{1+\nu_h}}{1+\nu_h} \right],$$

where the habit $C_{h,t-1}$ is external and specific to the household type, and b_t is an intertemporal preference shock.

To maximize tractability, we assume that all saving and consumption decisions are made by a single utilitarian family head who seeks to maximize the average intertemporal welfare of all members of both household groups,

$$U_t = (1 - \theta) u_{S,t} + \theta u_{H,t} + \beta E_t U_{t+1}. \quad (1)$$

This family head, however, faces limits to risk sharing: she can freely transfer resources across members of the same household group (after idiosyncratic uncertainty is realized), but not across groups, as in [Bilbiie \(2018\)](#).¹

Given these assumptions, S households are essentially the same as the representative permanent-income household in JPT, except for the risk that they face of becoming H . In particular, they all consume the same amount $C_{S,t}$, even though they are subject to idiosyncratic risk and their members are endowed with heterogeneous labor. They save in one-period nominal bonds and accumulate capital, subject to the budget constraint

$$P_t C_{S,t} + P_t I_{S,t} + B_{S,t} \leq s_t s_{t-1} R_{t-1} B_{S,t-1} + \Pi_{S,t} + \int_0^1 W_{S,t}(j) L_{S,t}(j) dj + r_t^k u_t \bar{K}_{S,t-1} - P_t a(u_t) \bar{K}_{S,t-1} + T_{S,t}, \quad (2)$$

where $I_{S,t}$ is investment, $B_{S,t}$ is end-of-period- t holdings of government bonds, $\Pi_{S,t}$ is the profit accruing to households from ownership of the firms, $T_{S,t}$ is nominal net lump-sum transfers (or taxes if negative) and ς_t is a shock to the spread between the nominal interest rate set by the central bank, R_{t-1} , and the return on capital r_t^k , as in [Smets and Wouters \(2007\)](#). The different workers in household S pool their labor income, which explains the integral on the right-hand-side of the budget constraint. All households of type S also share the interest income from the

¹See also [Heathcote and Perri \(2018\)](#) and [Bilbiie and Ragot \(2017\)](#) for similar “family head” assumptions yielding tractability.

fraction s_t of bonds accumulated in the previous period. The remaining $1 - s_t$ share of $t - 1$ bonds is owned by today's H households who were of type S in $t - 1$. Therefore, it is part of H income, as shown below.

S households own capital and choose the capital utilization rate, u_t , which transforms physical capital $\bar{K}_{S,t-1}$ into effective capital according to

$$K_{S,t} = u_t \bar{K}_{S,t-1}.$$

Effective capital is then rented to firms at the rate r_t^k . The cost of capital utilization is $a(u_t)$ per unit of physical capital. In steady state, $u = 1$, $a(1) = 0$ and $\chi \equiv \frac{a''(1)}{a'(1)}$. In the log-linear approximation of the model solution this curvature is the only parameter that matters for the dynamics. The physical capital accumulation equation is

$$\bar{K}_{S,t} = (1 - \delta)\bar{K}_{S,t-1} + \mu_t \left(1 - \Psi \left(\frac{I_{S,t}}{I_{S,t-1}} \right) \right) I_{S,t},$$

where δ is the depreciation rate and μ_t is an investment shock. The function Ψ captures the presence of adjustment costs in investment. In steady state, $\Psi = \Psi' = 0$ and $\Psi'' > 0$. Capital is assumed to be illiquid, so that S households who become H cannot access this asset.

Before describing the intertemporal choices of S for holding the liquid asset (bonds), we need to specify the budget constraint for households of the other type H , that S households face the risk of becoming, which they self-insure against.

Households of type H are hand-to-mouth. They consume all their income, which consists of labor income, transfers and their per-capita share of the return on the bonds accumulated in the previous period, when a fraction $1 - s_t$ of them was of type S . Like the S , also the H households pool all of their income, including the interest income contributed by those that used to be S in the previous period. As a result, their consumption is

$$P_t C_{H,t} = \int_0^1 W_{H,t}(j) L_{H,t}(j) dj + T_{H,t} + (1 - s_t) \frac{1 - \theta}{\theta} \varsigma_{t-1} R_{t-1} B_{S,t-1}. \quad (3)$$

The solution of the problem of the family head—who maximizes intertemporal aggregate welfare [1](#), subject to the budget constraints [2](#) and [3](#)—yields the Euler equation

$$\Lambda_{S,t} = \beta \varsigma_t R_t E_t [s_{t+1} \Lambda_{S,t+1} + (1 - s_{t+1}) \Lambda_{H,t+1}],$$

where $\Lambda_{S,t}$ is the marginal utility of nominal income for the representative household of type S , which is equal to

$$P_t \Lambda_{S,t} = \frac{b_t}{C_{S,t} - \eta_S C_{S,t-1}}.$$

Marginal utility is the same for all households of type S due to the assumptions on income sharing described above.

This Euler equation illustrates how savers engage in precautionary saving due to the idiosyncratic risk of becoming H . Assume, as will be the case in the steady state's calibration, that the marginal utility of income is higher for H than for S . Then, a positive probability of becoming H next period increases the expected marginal utility on the right-hand-side of the equation compared to the case with no risk, in which that probability is 0. As a result, today's marginal utility must be higher with than without risk. Consumption must then be lower and saving higher, given income. In summary, a higher risk of turning from S to H increases desired saving if the H state is unpleasant, as measured by a higher marginal utility. If this risk rises in recessions, the resulting precautionary saving motive will be counter-cyclical. This effect is further amplified if state H is worse in recessions relative to state S , or, in other words, if inequality in marginal utilities is counter-cyclical. Together, the counter-cyclicalities of the switching probability and of inequality result in counter-cyclical risk.

Section 2.5 shows how the consumption behavior of the S and H households described above can be integrated into an aggregate Euler equation that disciplines the evolution of average marginal utility.

2.1.2 Wage setting

Households of type h face a demand for their work j of the form

$$L_{h,t}(j) = \left(\frac{W_{h,t}(j)}{W_{h,t}} \right)^{-\frac{1+\lambda_{w,t}}{\lambda_{w,t}}} L_{h,t},$$

where $\lambda_{w,t}$ is a wage markup or labor supply shock. $L_{h,t}$ is the usual CES aggregate of the underlying j varieties

$$L_{h,t} = \left[\int_0^1 L_{h,t}(j)^{\frac{1}{1+\lambda_{w,t}}} \mathbf{d}j \right]^{1+\lambda_{w,t}},$$

and $W_{h,t}$ is its minimum price

$$W_{h,t} = \left[\int_0^1 W_{h,t}(j)^{-\frac{1}{\lambda_{w,t}}} \mathbf{d}j \right]^{-\lambda_{w,t}}.$$

The aggregates $L_{S,t}$ and $L_{H,t}$ are the labor inputs used by firms in their production, as explained below.

We assume that the wage for all j workers in households of type h is set by a “representative” h union that chooses it optimally, given the demand it faces and a Calvo friction. It solves the problem

$$\begin{aligned} \max_{W_{h,t}(j)} E_t & \left\{ \sum_{s=0}^{\infty} \xi_w^s \beta^s \left[\Lambda_{h,t+s} W_{h,t}(j) \Pi_{t,t+s}^w L_{h,t,t+s}(j) + \right. \right. \\ & \left. \left. - b_t \left(\log(C_{h,t} - \eta_h C_{h,t-1}) + \varphi_h \frac{L_{h,t}(j)^{1+\nu_h}}{1+\nu_h} \right) \right] \right\} \\ \text{s.t. } L_{h,t,t+s}(j) & = \left(\frac{W_{h,t}(j) \Pi_{t,t+s}^w}{W_{h,t+s}} \right)^{-\frac{1+\lambda_{w,t+s}}{\lambda_{w,t+s}}} L_{h,t+s}, \end{aligned}$$

where

$$\Pi_{t,t+s}^w = \prod_{k=1}^s \left[(\pi e^\gamma)^{1-\iota_w} (\pi_{t+k-1} e^{z_{t+k-1}})^{\iota_w} \right]$$

is an indexation factor. When making this choice, which is intertemporal due to nominal wage rigidity, the union weighs the utilities of j workers that are members of households of type h today and in the future, even if their identities change over time due to idiosyncratic risk. It follows from this assumption that the income generated by supplying labor of type j in period $t + s$, namely $W_{h,t}(j) L_{h,t,t+s}(j)$, is evaluated at the marginal utility of income of households of type h at that time, $\Lambda_{h,t+s}$. This setup is the same as in [Ascari et al. \(2017\)](#). As a result, the model features two New Keynesian wage Phillips curves, which govern the supply of the two labor composites $L_{S,t}$ and $L_{H,t}$. The marginal rate of substitution between consumption and leisure that drives wage setting in these equations is of the familiar form

$$\text{MRS}_{h,t} = \varphi_h b_t \frac{L_{h,t}^{\nu_h}}{\Lambda_{h,t}}.$$

2.2 Firms

Firms are standard New Keynesian monopolistic competitors indexed by i , subject to a Calvo price-setting friction. Their production function is

$$Y_t(i) = \left\{ \omega A_t^\sigma \mathcal{L}_{H,t}^\sigma(i) + (1 - \omega) \left[\alpha a_{k,t} K_t^\varrho(i) + (1 - \alpha) A_t^\varrho \mathcal{L}_{S,t}^\varrho(i) \right]^{\frac{\sigma}{\varrho}} \right\}^{\frac{1}{\sigma}} - A_t F,$$

where the parameters σ and ϱ control the elasticity of substitution between inputs, and F is a fixed cost. The stochastic processes $a_{k,t}$ and A_t represent two forms of technological progress, capital and labor augmenting. The first is stationary. A_t is not and it is responsible for trend growth in the economy along a balanced growth path. $K_t(i)$ is the amount of effective capital demanded by firm i , so that capital market clearing gives

$$\int_0^1 K_t(i) di = (1 - \theta) K_{S,t}.$$

Similarly, $\mathcal{L}_{h,t}(i)$ represents the amount of composite labor of type h demanded by firm i , which is connected to total labor supply by the market clearing conditions

$$\theta \left[\int_0^1 L_{H,t}(j)^{\frac{1}{1+\lambda_{w,t}}} dj \right]^{1+\lambda_{w,t}} = \theta L_{H,t} = \mathcal{L}_{H,t} \equiv \int_0^1 \mathcal{L}_{H,t}(i) di$$

and

$$(1 - \theta) \left[\int_0^1 L_{S,t}(j)^{\frac{1}{1+\lambda_{w,t}}} dj \right]^{1+\lambda_{w,t}} = (1 - \theta) L_{S,t} = \mathcal{L}_{S,t} \equiv \int_0^1 \mathcal{L}_{S,t}(i) di.$$

Total hours worked are then

$$\theta L_{H,t} + (1 - \theta) L_{S,t} \equiv L_t = \mathcal{L}_t \equiv \mathcal{L}_{H,t} + \mathcal{L}_{S,t}$$

and the aggregate hourly wage is

$$W_t = \frac{W_{H,t} \theta L_{H,t} + W_{S,t} (1 - \theta) L_{S,t}}{L_t}.$$

Total hours and the aggregate wage are the model objects with a direct counterpart in the data used in estimation, as described in more detail below.

Firms set their prices optimally, given a Calvo friction with indexation and the demand for

their output

$$Y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\frac{1+\lambda_{p,t}}{\lambda_{p,t}}} Y_t,$$

where $P_t(i)$ is the price of differentiated good i and

$$P_t = \left[\int_0^1 P_t(i)^{-\frac{1}{\lambda_{p,t}}} di \right]^{-\lambda_{p,t}}$$

is the price of the final good that is assembled according to

$$Y_t = \left[\int_0^1 Y_t(i)^{\frac{1}{1+\lambda_{p,t}}} di \right]^{1+\lambda_{p,t}}.$$

Variation in the elasticity of substitution in this aggregator, $\lambda_{p,t}$, is a source of markup shocks.

This structure results in a standard New Keynesian Phillips curve with marginal cost

$$MC_t = \frac{1}{\omega} A_t^{-\sigma} \left\{ \omega A_t^\sigma + (1 - \omega) \left[\alpha a_{k,t} \left(\frac{K_t(i)}{\mathcal{L}_{H,t}(i)} \right)^\varrho + (1 - \alpha) A_t^\varrho \left(\frac{\mathcal{L}_{S,t}(i)}{\mathcal{L}_{H,t}(i)} \right)^\varrho \right]^{\frac{\sigma}{\varrho}} \right\}^{1-\frac{1}{\sigma}} W_{H,t}.$$

2.3 Fiscal and monetary policy

The model's fiscal authority chooses government spending, taxes and transfers to balance its budget every period

$$\theta T_{H,t} + (1 - \theta) T_{S,t} + P_t G_t = 0.$$

Transfers are determined implicitly by the rule

$$\frac{\mathcal{I}_{H,t} + T_{H,t}}{\mathcal{I}_{S,t} + T_{S,t}} = \tau_0 \left(\frac{\mathcal{I}_{H,t}}{\mathcal{I}_{S,t}} \Big/ \frac{\mathcal{I}_H}{\mathcal{I}_S} \right)^{\tau_1},$$

where $\mathcal{I}_{h,t}$ is the pre-tax income of the representative household of type h . This transfer rule is clearly not a realistic description of actual tax arrangements. However, it provides a flexible parametrization of the pass-through from pre- to post-tax income inequality, similar to the “progressivity function” in [Heathcote et al. \(2017\)](#), [Ferriere and Navarro \(2020\)](#), and [Auclert et al. \(2018\)](#). This pass-through is one of the features of the model economy that we will vary in the counterfactual exercises described below. This representation of the transfer function makes the extent of that variation more transparent, even if it does not allow to map it back into specific features of an existing fiscal system. For example, $\tau_1 = 0$ corresponds to the case in which transfers eliminate any existing movement in the ratio of pre-tax incomes, resulting in a perfectly stable

degree of disposable income inequality. In this case, τ_0 parametrizes the resulting constant level of disposable income inequality between the two types of consumers in the model. Conversely, $\tau_1 = 1$ implies that post-tax income ratios are as volatile as income ratios before transfers, even if $\tau_0 \neq 1$ still allows for the fiscal system to reduce the level of inequality in steady state. Finally, the assumption of a balanced budget implies that government bonds are in zero net supply. Since they are the only liquid assets in the economy, we are for now focusing on an equilibrium with no liquidity.

Monetary policy sets the policy rate according to the interest rate rule

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R} \right)^{\rho_R} \left[\left(\frac{\prod_{j=0}^3 \pi_{t-j}}{\pi} \right)^{\phi_\pi} \left(\frac{X_t}{X_t^*} \right)^{\phi_X} \right]^{1-\rho_R} \left[\frac{X_t/X_{t-1}}{X_t^*/X_{t-1}^*} \right]^{\phi_{dX}} \eta_{mp,t},$$

where X_t is GDP, defined as

$$X_t \equiv C_t + (1 - \theta) I_{S,t} + G_t,$$

and X_t^* is a statistical measure of potential GDP, obtained using the exponential filter of [Curdia et al. \(2015\)](#).

2.4 Resource constraint

GDP differs from final output because the latter must also cover the capital adjustment cost, as shown by the resource constraint

$$C_t + (1 - \theta) I_{S,t} + G_t + (1 - \theta) a(u_t) \bar{K}_{S,t-1} = Y_t.$$

2.5 Inequality and risk in the aggregate Euler equation

This section summarizes some of the key features of our model economy through its aggregate Euler equation. This equation intuitively highlights the mechanisms through which inequality and risk affect aggregate consumption and saving behavior. Therefore, it also provides a useful organizing framework to interpret the exercises conducted in section 5, where we will isolate the quantitative role of inequality and risk on aggregate fluctuations through counterfactual simulations.

Our model's aggregate Euler equation can be written as

$$\Lambda_t = r_t E_t \Lambda_{t+1} \tag{4}$$

$$+ \theta (\Lambda_{H,t} - r_t E_t \Lambda_{H,t+1}) \tag{5}$$

$$+ (1 - \theta) r_t E_t [(1 - s_{t+1}) (\Lambda_{H,t+1} - \Lambda_{S,t+1})], \tag{6}$$

where $\Lambda_t \equiv \theta \Lambda_{H,t} + (1 - \theta) \Lambda_{S,t}$ denotes “aggregate” marginal utility of nominal income and $r_t \equiv \beta \varsigma_t R_t$. The first line of the expression represents the standard intertemporal Euler equation of a representative agent model, whose marginal utility is expected to decline when interest rates rise. The two terms on the remaining lines constitute a wedge that is related to the presence of heterogeneity—our model's version of the Euler equation wedge also derived in [Berger et al. \(2019\)](#) and [Debortoli and Galí \(2023\)](#). This wedge would collapse to zero if the marginal utilities of the two types of agents were the same, i.e. $\Lambda_{H,t} = \Lambda_{S,t}$. The resulting allocation would be the one observed under perfect consumption insurance. We refer to this as the “no-inequality” case, even if heterogeneity in incomes, wages and hours worked remain. But the key mechanisms through which inequality affects consumption and saving behavior, and therefore the transmission of business cycles, are absent in this “no-inequality” benchmark.

The term on the second line of the equation captures the effect of the presence of constrained agents. If H households could optimize, their marginal utility would satisfy an Euler equation like the one in the first line of the expression, and this term would be equal to zero. But the H agents cannot smooth consumption. Therefore, their Euler equation wedge is not zero and varies over time, distorting the aggregate Euler equation. Intuitively, a higher share θ of constrained agents in the economy increases the quantitative relevance of this distortion.

The term on the last line of the equation is the wedge component that is due to the presence of idiosyncratic risk. It is positive—thus raising average marginal utility and depressing expenditures—if S households run the risk of becoming H (i.e. if $s_t < 1$) and if this switch leads to an increase in marginal utility (i.e. if $\Lambda_{H,t} > \Lambda_{S,t}$). In our quantitative model, this term is large, since matching U.S. inequality data requires $\Lambda_{H,t}$ and $\Lambda_{S,t}$ to be far from each other. In addition, the term (6) is pro-cyclical because both inequality in marginal utilities *and* the chances of becoming H rise in recessions, pushing current marginal utility and aggregate saving higher, and consumption lower. As we will see in section 5, the time variation in the probability of switching plays a crucial quantitative role in the amplification of business cycles. This is because its effect

is larger the larger is long-run inequality in marginal utilities, as evident from expression (6).² Intuitively, the stakes involved in falling down the cross-sectional income ladder, as represented in our model by the marginal utility gap between S and H agents, are much larger than those associated with the exposure to aggregate fluctuations.

3 Data and Inference

We estimate the model parameters using nine time series that provide information on both the aggregate economy and the evolution of inequality. The aggregate variables are: the growth rate of per capita (i) real GDP, (ii) real consumption, and (iii) real investment; (iv) the logarithm of hours worked, also per capita; (v) the growth rate of real hourly wages; (vi) the inflation rate; and (vii) the federal funds rate. The inequality variables are the cross-sectional standard deviations of the logarithm of (viii) (pre-tax) labor earnings and (ix) disposable income, both de-trended. Appendix A provides details on the data sources and construction of the series used in estimation. Series (i) to (vii) are quarterly, from 1954:III to 2019:IV. They provide information primarily on the coefficients of the medium-scale DSGE structure at the core of our model, such as the degree of nominal rigidities, habit formation, investment adjustment costs, the monetary policy rule coefficients, as well as the persistence and variance of the shocks.

Series (viii) is only observed at an annual frequency, starting in 1967. It helps to pin down the elasticities of substitution between the three inputs in the production function, since those are related to the relative movements of demand for labor of type H and S , given the dynamics of capital accumulation. Figure 2 plots the level of labor earnings inequality together with its cyclical pattern, which is calculated using a band-pass filter that extracts fluctuations with periodicities lower than 30 years. The bottom panel of the figure also plots the cyclical behavior of aggregate hours worked, obtained through the same filter, as an indicator of the phases of the business cycle. The comparison between these two series highlights the strong counter-cyclicity of labor earnings inequality: it surges when hours fall in recessions and it declines when hours rise in expansions. Heathcote et al. (2020) document that this counter-cyclicity in income inequality is mostly driven by the higher exposure of *both* the wages and hours of workers with low earnings to business cycle fluctuations, compared to those of workers in the rest of the earnings distribution. During recessions, workers with low earnings work fewer hours and

²To see this, notice that the dynamics of (6) can be approximated up to first order by the sum of two terms: (i) the cyclical variation of the switching probability $1 - s_{t+1}$ around a steady state with inequality $\Lambda_H - \Lambda_S$; and (ii) the cyclical variation of inequality $\Lambda_{H,t+1} - \Lambda_{S,t+1}$ around a steady state with constant risk $1 - s$.




Figure 2: Labor income inequality: level and cyclical fluctuations. Details on data sources and construction are provided in appendix A. The de-trending of inequality and hours in the bottom panel uses a band-pass filter that extracts fluctuations with periodicities lower than 30 years.

earn lower wages, while the wages and especially the hours of all other workers are quite stable. This correlation between the relative hours and wages of low income workers over the business cycle suggests that the demand for their labor services is more cyclical than that of other workers. Our model is consistent with this pattern when labor of type H —the one supplied by workers with lower income—is more substitutable with capital than that of type S .

Series (ix) is available since 1978, also at an annual frequency. Figure 3 plots the level of this variable, which has been rising steadily over the sample, together with its cyclical fluctuations in the bottom panel. Disposable income inequality is also counter-cyclical, but substantially less so than pre-tax labor earnings inequality, pointing to the important role of fiscal redistribution. Together with the information on pre-tax income inequality, therefore, series (ix) is crucial to identify the parameters that characterize the model's tax and transfer system.

We characterize the posterior distribution of the coefficients by combining the likelihood




Figure 3: Disposable income inequality: level and cyclical fluctuations. Details on data sources and construction are provided in appendix A. The de-trending of inequality and hours in the bottom panel uses a band-pass filter that extracts fluctuations with periodicities lower than 30 years.

function of the log-linearized model with prior information. The likelihood function is evaluated using the Kalman filter, which can easily handle the mixed frequency (quarterly and annual) of the aggregate and inequality data used in estimation. Conditional on the information contained in the observable variables, the Kalman filter and smoother also return historical estimated paths for all the model unobservable variables, including the shocks, which are the key input in the counterfactual experiments discussed below.

We fix two important parameters using micro information not contained in our dataset. First, the share of H households (θ) is set to 0.2, consistent with the empirical consensus on the average marginal propensity to consume out of transitory income shocks (e.g. Souleles et al., 2006, Parker et al., 2013 and Kaplan et al., 2014). Second, we choose the steady-state probability with which S agents become H , s_0 , to be consistent with some key cross-sectional properties of the annual change in (log) individual incomes. According to Guvenen et al. (2021), the cross-sectional standard deviation and kurtosis of this variable are 0.51 and 15, respectively. We can closely match both statistics by setting $s_0 = 0.987$. In addition, we assume that the standard deviations of labor income and disposable income are observed with some noise. The variance of this measurement error is set to 1/16 of the variance of the respective series.

The priors on the other coefficients are reported in table 1. They are in line with those adopted in previous studies, such as Smets and Wouters (2007) and Justiniano et al. (2010, 2011, 2013). For the parameters related to inequality that do not appear in typical medium-scale DSGE models, we have chosen relatively dispersed priors. For example, the prior on the elasticities of substitution between inputs in the production function, σ and ϱ , is a Gamma distribution with mode equal to 1, corresponding to the Cobb-Douglas case, and standard deviation equal to 2. The prior on the pass-through from pre-tax to disposable income inequality, τ_1 , is a Uniform on $[0, 2]$. The prior on s_1 , the elasticity of the function determining the probability with which households switch type with respect to economic conditions, is a (shifted and scaled) Beta distribution, with mean 0.5 and standard deviation equal to 0.25.³

Finally, we introduce prior information on α , ω and τ_0 through priors on three objects. The first object is the steady-state capital share, whose prior distribution is a Beta with mean 0.33 and standard deviation 0.05. The second object is the cross-sectional standard deviation of log labor earnings in steady state, whose prior is a Gamma distribution with mean 0.9 and standard deviation 0.1. The mean of this distribution reflects the average level of inequality in labor earnings

³The objective of shifting and rescaling a standard Beta distribution is to define a density for s_0 over the support $[-0.4, 0.4]$, which guarantees that the probability of switching is non-negative.

from the top panel of figure 2. This inequality trends higher over our sample, as it is well-known, but the model abstracts from this trend to focus instead on cyclical fluctuations, as captured by the de-trended standard deviation in the lower panel of the figure. Finally, the third object on which we incorporate prior information is the ratio between the cross-sectional standard deviations of the logarithms of disposable income and labor earnings, as reported in figures 3 and 2. The two series both trend up, but their trend is similar. As a result, their ratio is around 2/3 across the available sample. Therefore, we set the prior distribution on this object to a Gamma distribution with mean 0.66 and standard deviation 0.025.

Table 1: Prior distributions and posterior parameter estimates.

		prior			posterior			
		distribution	mean	SD	mode	median	5%	95%
production								
α	KORV share parameter	B	0.5	0.25	0.97	0.96	0.93	0.99
ω	KORV share parameter	B	0.5	0.25	0.48	0.46	0.39	0.55
σ	KORV elasticity	G	2.56	2	9.38	9.18	7.17	12.97
ϱ	KORV elasticity	G	2.56	2	0.38	0.40	0.34	0.43
λ_p	SS price markup	N	0.15	0.025	0.24	0.24	0.20	0.28
ι_p	price indexation	B	0.5	0.15	0.11	0.12	0.06	0.22
ξ_p	price stickiness	B	0.66	0.15	0.80	0.79	0.73	0.83
households and unions								
$100(\beta^{-1} - 1)$	discount factor	G	0.25	0.1	0.24	0.25	0.13	0.41
η_S	habit formation (S)	B	0.5	0.1	0.96	0.96	0.94	0.97
η_H	habit formation (H)	B	0.5	0.1	0.94	0.94	0.91	0.96
$100 \log L_{ss}$	SS log hours	N	0	0.5	-0.03	-0.05	-0.82	0.76
ν_S	inverse Frisch (S)	G	3	2	0.82	1.01	0.60	1.70
ν_H	inverse Frisch (H)	G	3	2	0.08	0.15	0.04	0.30
χ	util. cost elasticity	G	5	1	3.23	3.30	1.94	4.87
Ψ''	investment adj. cost	G	4	1	1.60	1.67	1.25	2.28
$\frac{1}{2} + \frac{5}{4}s_1$	risk elasticity	B	0.5	0.25	0.96	0.94	0.86	0.99
λ_w	SS wage markup	N	0.15	0.025	0.15	0.15	0.11	0.19

Table 1: Prior distributions and posterior parameter estimates.

		prior			posterior				
ι_w	wage indexation	B	0.5	0.15	0.10	0.12	0.06	0.21	
$\xi_{\omega,S}$	wage stickiness (S)	B	0.66	0.15	0.70	0.68	0.59	0.74	
$\xi_{\omega,H}$	wage stickiness (H)	B	0.66	0.15	0.86	0.85	0.80	0.88	
policy									
100 ($\pi - 1$)	SS inflation	N	0.5	0.1	0.39	0.42	0.28	0.57	
ρ_R	interest rate inertia	B	0.6	0.2	0.77	0.76	0.68	0.81	
ϕ_π	reaction inflation	N	1.7	0.3	1.97	1.96	1.77	2.18	
ϕ_X	reaction gap	N	0.125	0.05	0.23	0.21	0.17	0.26	
ϕ_{dX}	reaction gap growth	N	0.125	0.05	0.08	0.08	0.05	0.11	
τ_0	SS redistribution	U	1	0.58	0.35	0.37	0.31	0.44	
τ_1	cyclical redistribution	U	1	0.58	0.21	0.21	0.19	0.23	
exogenous processes									
100 γ	SS tech. growth	N	0.5	0.025	0.48	0.48	0.45	0.52	
ρ_z	autocorr. neutral tech.	B	0.4	0.2	0.22	0.22	0.09	0.37	
ρ_k	autocorr. capital tech.	B	0.6	0.2	0.98	0.97	0.95	0.99	
ρ_{mp}	autocorr. mp	B	0.4	0.2	0.21	0.25	0.13	0.42	
ρ_g	autocorr. gov. spending	B	0.6	0.2	0.99	0.99	0.99	0.99	
ρ_μ	autocorr. investment	B	0.6	0.2	0.46	0.56	0.27	0.87	
ρ_ζ	autocorr. risk premium	B	0.6	0.2	0.96	0.95	0.93	0.98	
ρ_b	autocorr. disc. factor	B	0.6	0.2	0.14	0.14	0.07	0.24	
ρ_p	autocorr. price mk	B	0.4	0.2	0.98	0.98	0.96	0.99	
ρ_w	autocorr. wage mk	B	0.4	0.2	0.99	0.99	0.99	0.99	
ρ_p^{ma}	MA price mk	B	0.5	0.2	0.86	0.83	0.73	0.90	
ρ_w^{ma}	MA wage mk	B	0.5	0.2	0.94	0.94	0.90	0.96	
100 σ_z	std neutral tech.	IG1	1	1	0.84	0.81	0.71	0.92	
100 σ_k	std capital tech.	IG1	0.5	1	0.77	0.83	0.68	1.00	
100 σ_{mp}	std mp	IG1	0.15	1	0.20	0.21	0.19	0.22	
100 σ_g	std gov. spending	IG1	1	1	1.51	1.52	1.41	1.64	

Table 1: Prior distributions and posterior parameter estimates.

		prior			posterior			
100 σ_μ	std investment	IG1	1	1	2.43	2.55	1.01	3.99
100 σ_ζ	std risk premium	IG1	0.5	1	0.23	0.23	0.15	0.33
100 σ_b	std disc. factor	IG1	0.15	1	0.20	0.20	0.17	0.22
100 σ_p	std price mk	IG1	0.15	1	0.16	0.17	0.15	0.19
100 σ_w	std wage mk	IG1	0.15	1	0.31	0.31	0.27	0.37

4 Estimation Results

This section discusses some key features of the estimated model. First, we look at parameter estimates, focusing on coefficients that are not present in representative agent models. Second, we analyze the cyclical properties of income risk, and compare them to the available micro evidence. Third, we present impulse responses to monetary policy shocks to illustrate how inequality evolves in response to a typical demand shock, and how its presence in the model changes the response of other variables. We concentrate on monetary shocks because their propagation has received a lot of attention in the HANK literature, following the seminal contribution of [Kaplan et al. \(2018\)](#). Fourth, we examine the sources of fluctuations in the model through a variance decomposition.

4.1 Parameter estimates

Columns 6 to 9 of table 1 report some properties of the posterior of the model parameters. Most of these estimates are in line with those in previous studies, such as [Smets and Wouters \(2007\)](#) and [Justiniano et al. \(2010, 2011, 2013\)](#). The main exception is the degree of habit formation, which is higher.

Among the parameters related to the presence of inequality, the most interesting estimates are those of σ and ϱ , which determine the elasticities of substitution between inputs in the production function. Their posteriors imply that labor of type H is a substitute for capital, while that of type S is complementary with it. These estimates allow the model to match the pronounced

counter-cyclicality of inequality in labor earnings observed in the data. Relative to the estimates of the equivalent parameters in [Krusell et al. \(2000\)](#), our posterior implies a stronger degree of substitutability between H labor and capital, which corresponds to a high value for σ . Another important parameter is s_1 , which is estimated to be positive. This implies that the probability of remaining type S increases in good times, making income risk counter-cyclical in the estimated model.

4.2 Counter-cyclical income risk

The counter-cyclicality of income risk plays a central role in our quantitative assessment of the effect of inequality on business cycles, as discussed in section 5. To support the empirical plausibility of that assessment, this subsection shows that the cyclical properties of income risk implied by our estimated model are consistent with the available micro evidence.

[Guvenen et al. \(2014\)](#) provide a wealth of empirical evidence on the cyclical behavior of labor income risk in the U.S., based on a large panel of earning histories from the U.S. Social Security Administration records. Using data from 1978 to 2011, they show that the cross-sectional left skewness of individual labor income growth is strongly counter-cyclical. Namely, the left tail of the distribution becomes longer and/or fatter in recessions, as illustrated in their figure 6. Intuitively, this means that large drops in labor income are more likely in recessions than in expansions. Qualitatively, this pattern is also present in our model, since H income falls relative to that of S types in recessions, at the same time as the probability of becoming H increases.⁴

To check more precisely how our estimated model compares with this empirical evidence, we simulate the labor income dynamics of a large panel of households in our economy. We already know that the cross-sectional distribution of the *level* of their income has two points with mass θ and $1 - \theta$ at the income of S and H types respectively. But if we follow individual households over time as they switch between the two types, and as aggregate fluctuations buffet their respective incomes, the resulting histories produce a richer panel of income *growth rates*, whose moments can in principle match those in the data. Figure 4 presents the time series of the skewness of the 1- and 5-year changes in log annual labor earnings from this simulated panel, as in figure 6 of [Guvenen et al. \(2014\)](#). The aggregate shocks that underlie this simulation are those that the model estimates to have driven the observed evolution of the observable variables since 1979. The model does not reproduce the fact that the changes in labor earnings have negative

⁴See also [Storesletten et al. \(2004\)](#) for earlier evidence on the countercyclicality of individual labor income risk.




Figure 4: Model implied skewness of the 1- and 5-year labor income growth of individual earners. The shaded vertical bars denote NBER recessions.

skewness on average. The mechanical reason is that every switch from S to H must correspond to a switch from H to S in our model, so as to keep the measure of each type unchanged. And the negative and positive income changes associated with those switches have the same average size, resulting in a symmetric distribution of income changes. However, more importantly for our analysis of business cycles, the left skewness of the 1- and 5-year changes is strongly countercyclical: it increases during recessions and tends to recover during booms.⁵ On average over the recessions in the sample, the skewness is 0.18 and 0.09 points lower than in the expansions at the 1- and 5-year horizon, in line with the quantitative findings of Guvenen et al. (2014).

This is a remarkable success, given the simplicity of the income process in our model. Intuitively, the reason why it works is that it approximates an environment with two main kinds of labor income changes, as in the parametric process estimated by Guvenen et al. (2014) in their section 7. In the data, workers who stay in their jobs draw relatively small income changes that are mostly determined by aggregate conditions. In our model, these are the agents whose type does not change. In contrast, workers who change jobs draw their new earnings from a much wider distribution. These are the agents who switch type in our model. The mixing of these two distributions, together with the especially large negative income changes due to unemployment

⁵Recall that an increase in the left skewness of a distribution (i.e. an increase in the length and/or fatness of its left tail) means that its skewness falls since it becomes more negative.

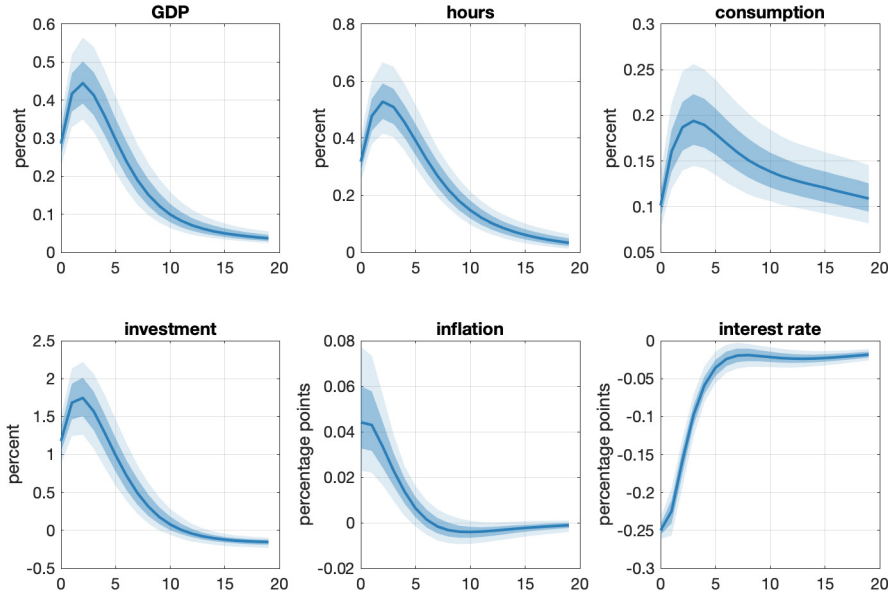


Figure 5: Impulse responses to a monetary policy shock: aggregate variables. The solid lines are posterior medians, while the shaded areas correspond to 68- and 95-percent posterior credible regions.

spells, give the empirical distribution of individual income changes its negative skew. And the fact that income declines, including those generated by job loss, are more frequent in recessions makes that skew countercyclical. In our model, these dynamics are replicated by the fact that, in recessions, the probability of becoming H increases at the same time as H earnings fall relative to those of S agents.

4.3 Impulse responses to a monetary policy shock

Figure 5 plots the response of the key macroeconomic variables used in the model's estimation to a 25-basis-point surprise cut in the federal funds rate. As in RANK models, and in Structural VARs, GDP, consumption, investment, hours and inflation all increase. How does heterogeneity affect this propagation?

Figure 6 provides an important element of the answer to this question. It reports the response of labor income, hours worked, consumption and disposable income for H and S agents separately. Before taxes, the labor income of H households responds almost three times as much as that of the S agents to the boom generated by lower interest rates, as a percentage of their respective steady-state incomes. The same is true for hours worked, whose relative movements




Figure 6: Impulse responses to a monetary policy shock: income, hours and consumption of H and S households. The solid lines are posterior medians, while the shaded areas correspond to 68- and 95-percent posterior credible regions.

therefore are primarily responsible for the reduction in relative labor income inequality associated with expansionary monetary policy. In contrast, the responses of disposable income are much closer to each other, due to the inequality-reducing effect of fiscal redistribution. Nevertheless, H consumption rises about four times as much as that of S agents because the latter save much of the transitory increase in their income, while the former consume it all. These results are in line with the empirical findings on the response of inequality to monetary policy in [Coibion et al. \(2017\)](#) for the U.S., and [Slacalek et al. \(2020\)](#), [Lenza and Slacalek \(2021\)](#), and [Broer et al. \(2022b\)](#) for the Euro Area and Germany.

4.4 Dynamics of consumption inequality

Overall, the expansion in aggregate demand associated with lower interest rates, which is representative of how other key demand shocks transmit in the model, makes inequality of incomes, hours and consumption counter-cyclical. These patterns of inequality are consistent with the data, with the possible exception of consumption. Micro data on consumption in the U.S. are far from ideal to make precise statements about the strength of the counter-cyclicity of its standard deviation. For example, [Attanasio and Pistaferri \(2016\)](#) survey several measures of con-

sumption inequality, documenting their similarities and differences. Their figure 1 highlights that some of these measures, such as the one proposed by [Aguiar and Bils \(2015\)](#), are substantially more cyclical than others. Along similar lines, [Coibion et al. \(2017\)](#) find that consumption inequality responds more than income inequality to a monetary policy shock. Nonetheless, on balance, our reading of the literature is that consumption inequality tends to be less counter-cyclical than that of disposable income, since most households have access to some consumption smoothing opportunity. In our stylized model, in contrast, H agents do not have any smoothing opportunity and consume all their disposable income in every period. But as shown in section 5, the behavior of H agents is not central to the business cycle properties of our estimated model. The saving behavior of S types worried to become H is the key. Nevertheless, modifying the model to give H agents some access to saving opportunities so as to reduce the sensitivity of consumption inequality to shocks might be a worthwhile extension of the current framework.

4.5 Sources of fluctuations

Figure 7 provides information about the drivers of economic fluctuations in the estimated model. We group those drivers into three categories. Markup shocks include price and wage markup disturbances; technology shocks comprise labor- and capital-augmenting productivity innovations; demand shocks are investment, risk premium, intertemporal preference, government spending and monetary policy shocks. Consistent with the typical findings of medium-scale DSGE studies, demand shocks are the main source of fluctuations for most variables, including GDP growth and inflation. The main exceptions are hours worked and labor income inequality, for which markup shocks are more prominent. In the case of hours, this result is due to their pronounced low-frequency trend, which is mainly accounted for by wage markup shocks, as illustrated by [Justiniano et al. \(2010\)](#). In the model with heterogeneity, wage markup shocks also contribute a low frequency component to labor income inequality, which mostly accounts for the large role that those disturbances play in the unconditional variance of this variable.

5 Inequality and Risk in Business Cycles

In this section, we use our estimated model to quantify the impact of inequality and risk on business cycle fluctuations. We do so through a set of counterfactual experiments constructed




Figure 7: Posterior distribution of the decomposition of the unconditional variance of selected observable variables into three categories of shocks. Markup shocks include price and wage markup innovations; technology shocks include labor- and capital-augmenting productivity disturbances. Demand shocks include investment, risk premium, intertemporal preference, government spending and monetary policy shocks.

as follows. First, given the estimated parameters, we obtain the sequence of exogenous disturbances that, according to the model, were behind the observed evolution of the economy since 1954:III. Then, we feed these “historical” disturbances into counterfactual versions of the model that feature different parameter values or equilibrium conditions from those used in estimation. The differences between the observed data and those produced in the counterfactual simulations summarize the importance for observed fluctuations of the particular feature of the model that is modified in the counterfactual.

Our first experiment gauges the role of inequality in marginal utility. It does so through a “perfect-consumption-insurance” counterfactual, in which the marginal utility of income of the two types of households is equalized in every period. Formally, we impose $\Lambda_{H,t} = \Lambda_{S,t}$ as one of the model’s equilibrium conditions, as a substitute for the mapping from pre- into post-tax income given by the transfer function. As a result, transfers are determined in equilibrium to deliver equal marginal utilities, subject to balancing the government’s budget. Therefore, the resulting equilibrium is the one that would be chosen by a planner who can transfer resources across agents in every period, although not over time. With the same initial conditions, and the same degree of habit formation, consumption for the two types would be the same in this equilibrium, resulting in a representative consumer counterfactual.⁶

An alternative, seemingly more direct approach to eliminating inequality in the model would be to set the value of θ , the fraction of H agents, to zero. This restriction would result in a model with a representative consumer similar to that featured in typical RANK medium-scale DSGEs. However, it would also have no aggregate output, since production requires a positive amount of labor of type H in our baseline economy. For this reason, we cannot pursue this approach to eliminating inequality in the model in our setup.

Figure 8 shows the results of the counterfactual with no inequality in marginal utilities. Both GDP growth and especially hours worked are significantly less volatile in this counterfactual economy. For instance, the standard deviation of GDP growth is 27 percent lower in the counterfactual than in the data. For de-trended hours, this number is 36 percent. More concretely, this means that hours worked would have fallen below trend by only two percentage points during the Great Recession, rather than more than six percentage points as in reality, if fiscal policy had enacted enough transfers to completely eliminate the impact of the recession on consumption inequality. According to these calculations, the Great Recession would have been of a similar

⁶With unequal initial conditions, this counterfactual leads to equal consumption only over time, as the influence of those initial conditions wanes. However, counterfactuals that equalize the level of consumption instead of the marginal utility produce almost identical results.




Figure 8: No-inequality counterfactual: GDP growth and de-trended hours.

order of magnitude as the much milder recession of the early 2000s without inequality, at least in terms of its impact on the labor market. Symmetrically, though, the boom of the late 1990s would have not been much of a boom if inequality had not declined below its trend over that period. On net, the observed movements in inequality appear as an important source of amplification of business cycles, leading to frothier booms and deeper recessions, especially on the labor market. These results provide an answer the first question that motivates our paper: Does inequality matter for business cycles? Yes, it does.

A complementary approach to assess whether the evolution of inequality in marginal utilities is important for business cycles is to compare the fit of the baseline model with household heterogeneity to that of an otherwise identical model but with complete markets (i.e. the model we have used for our counterfactual above). The log marginal likelihoods of the two models are equal to $-1,757$ and $-2,152$, respectively. Recall that the marginal likelihood is a measure of one-step-ahead density forecast accuracy, which makes clear that the model with complete

markets fits the data much worse than the baseline. This deterioration in fit is due to the fact that matching the data with the complete-market model requires substantially more volatile shocks, especially those typically labeled as “demand” shocks, such as the government spending and the investment shocks (Justiniano et al., 2010). This finding confirms that the endogenous wedge in the aggregate Euler equation of the baseline model (see section 2.5) enhances the propagation of demand shocks.

Our next set of experiments addresses our second motivating question, by exploring the mechanisms through which inequality amplifies aggregate fluctuations. In fact, the counterfactual of figure 8 conflates the role of four separate but interacting channels, which have all been shut down in the economy without inequality in marginal utilities: (i) cyclical inequality; (ii) steady-state (or permanent) inequality; (iii) cyclical idiosyncratic risk; and (iv) steady-state (or permanent) idiosyncratic risk. Our next simulations quantify the importance of these four channels by shutting them down one at a time, and then turning them on sequentially.

Figure 9 shows the results of eliminating cyclical inequality through transfers that keep the distance between the marginal utilities of the two types of households constant over time at its steady-state level. The rest of the model is unchanged, including the probability of switching from one type to the other, which varies over the cycle as in the baseline economy. The top panel of the figure shows that this simplified model produces essentially the same output fluctuations as the baseline specification. This exercise indicates that cyclical inequality, in and of itself, has a negligible quantitative role in the amplification of business cycles.

Surprisingly, hours are even more volatile in the counterfactual economy with constant inequality in marginal utilities than in the data, as shown in the bottom panel of the figure. The explanation for this result comes from the interaction between wealth effects on labor supply and heterogeneity. During booms, H consumption rises relative to S , leading to a stronger negative wealth effect on the labor supply of H workers. In turn, this produces an attenuation of the overall boom in hours and an increase in labor productivity, given that H workers are less productive. This mechanism, therefore, loosens the connection between the expansion in total hours and that in output over the cycle, making the latter relatively less volatile than the former. Of course, in the estimated model, that connection must be what we observe in the data, since total hours and output growth are both among the observable variables in estimation. But in the counterfactual that eliminates the counter-cyclicalities of inequality, and hence the relative wealth effect, this attenuation effect is no longer there. As a result, the volatility of hours tends to rise compared to that of output in that counterfactual, as we observe in the simulation.




Figure 9: No-cyclical-inequality counterfactual: GDP growth and de-trended hours




Figure 10: No-steady-state-inequality counterfactuals: GDP growth and de-trended hours.

Households in our model differ in steady state, and not only in the way in which their relative consumption and income move over time. And as in the data, this permanent component of inequality is large compared to its cyclical fluctuations. To explore its role in the transmission of business cycles, our next simulation studies the effect of eliminating only the differences between household types in steady state. This counterfactual is obtained through a redistributive scheme that equalizes the marginal utility of consumption of H and S households in steady state, but not their relative fluctuations. The results of this experiment are presented in figure 10. They demonstrate that steady-state inequality is crucial in the amplification of business cycles. When we eliminate the steady-state differences in households' marginal utilities, the volatility of both GDP growth and hours falls significantly. In fact, the counterfactual with no steady-state inequality is very similar to that without any inequality at all.

Figure 11 isolates the role of cyclical risk, by setting the switching probability in the estimated model equal to its steady-state value, $s_t = s$, and eliminating its dependence on the state of the




Figure 11: No-cyclical-risk counterfactuals: GDP growth and de-trended hours.

business cycle. This experiment indicates that cyclical risk plays a crucial quantitative role in amplifying output fluctuations. Without it, the standard deviation of GDP growth is 24 percent less than in the data, which is very close to the result of the no-inequality counterfactual. This is less true when looking at hours, for which the no-cyclical-risk counterfactual produces fluctuations that are in between those in the data and in the model of figure 8 without any inequality. The reason for this discrepancy in the output and hours counterfactuals is also related to the wealth effect on labor supply as discussed above.

Our next experiment studies the importance of the permanent component of idiosyncratic risk, by setting $s_t = 1$ in the estimated model. This value for the switching probability implies that households face no risk of changing type, as in two-agent New Keynesian models such as Galí et al. (2007), Bilbiie (2008), Debortoli and Galí (2018) and Bilbiie et al. (2022). The counterfactual behavior of GDP growth and hours under this assumption is essentially identical to the red lines in figure 11, which we had obtained by setting $s_t = s$. Put differently, idiosyncratic risk

induced by a constant probability of switching type has little impact on aggregate fluctuations. The effects of idiosyncratic risk are large only when this switching probability is counter-cyclical.

Taken together, these counterfactual experiments suggest that the main source of the notable effect of household heterogeneity on aggregate fluctuations is the combination of long-run inequality and the counter-cyclical risk. Savers reduce their consumption to insure themselves against the idiosyncratic risk of large income drops in the event that they become constrained—a risk that rises in recessions, as captured by the estimate of the parameter s_1 . This view is consistent with the counter-cyclical risk of the left skewness of the income growth distribution documented in section 4.

Table 2 reports the standard deviations of GDP growth and hours in a sequence of counterfactuals that are complementary to those shown in the figures above. The latter focus on the role of eliminating some of the sources of heterogeneity present in the baseline estimated model, one at a time. The table, instead, builds up to the baseline estimated model starting from the no-inequality counterfactual of figure 8, and adding back increasing layers of inequality and risk. The figures and the table provide complementary information because the decomposition of the role of heterogeneity into different channels is not additive, as these channels interact with each other.

The second column of the table adds steady-state inequality to the no-inequality counterfactual, but with no switching between agents types, and hence no risk. With no risk, even the large amount of inequality present in the steady state of the estimated model barely affects the dynamics of output, even though it has a large effect on the volatility of hours. At the margin, a constant risk of switching types, as in the third column, does not have a discernible effect on the volatility of either output or hours. But when the probability of switching varies with the cycle, as in the model of column four, the standard deviation of GDP growth jumps much closer to that of the estimated model that features all mechanisms, while that of de-trended hours is even higher than in the data. This is the basis of our conclusion that steady-state inequality, together with cyclical movements in the switching probability, do most of the heavy lifting in amplifying business cycle shocks. The addition of cyclical inequality in marginal utilities to the model (column five) has small effects on output volatility, but it helps to highlight the important quantitative role of the wealth effect on labor supply in the estimated model. When income inequality is counter-cyclical, as in the data, H agents become relatively richer than S agents in expansions. But with lower marginal utility of income, they also supply relatively less labor. In our model, this supply-side mechanism reduces the overall volatility of hours by a notable amount, even if

	no inequality	ss ineq no risk	ss ineq ss risk	ss ineq cyclical risk	cyclical ineq cyclical risk	no ss ineq
	Fig 8			Fig 9	Data	Fig 10
	$\Lambda_H = \Lambda_S$	$\Lambda_H \neq \Lambda_S$	$\Lambda_H \neq \Lambda_S$	$\Lambda_H \neq \Lambda_S$	$\Lambda_H \neq \Lambda_S$	$\Lambda_H = \Lambda_S$
	$\hat{\Lambda}_{H,t} = \hat{\Lambda}_{S,t}$	$\hat{\Lambda}_{H,t} = \hat{\Lambda}_{S,t}$	$\hat{\Lambda}_{H,t} = \hat{\Lambda}_{S,t}$	$\hat{\Lambda}_{H,t} = \hat{\Lambda}_{S,t}$	$\hat{\Lambda}_{H,t} \neq \hat{\Lambda}_{S,t}$	$\hat{\Lambda}_{H,t} \neq \hat{\Lambda}_{S,t}$
		$s_0 = 1$	$s_0 = 0.987$	$s_0 = 0.987$	$s_0 = 0.987$	$s_0 = 0.987$
		$s_1 = 0$	$s_1 = 0$	$s_1 = 0.36$	$s_1 = 0.36$	$s_1 = 0.36$
GDP growth	1.64	1.65	1.65	2.19	2.25	1.75
de-trended hours	1.66	2.82	2.82	3.30	2.61	1.65

Table 2: Standard deviations of annual GDP growth and de-trended hours in the counterfactual and baseline estimated economies.

its effect on the standard deviation of output growth is small.

6 Concluding Remarks

This paper developed and estimated a medium-scale New Keynesian model of the business cycle with heterogeneous households. Its starting point is complementary to that of the burgeoning HANK literature. Rather than enriching the Bewley/Ayagari model of idiosyncratic income risk with some of the basic elements of New Keynesian business cycle models, we took a standard empirical DSGE framework with a host of nominal and real rigidities, and extended it to accommodate household heterogeneity in a tractable way. Despite its stylized view of heterogeneity with only two household types, our model is flexible enough to feature four important mechanisms that affect the interaction between inequality and business cycles. The first one is the presence of inequality itself, as agents in the model have different incomes and consumptions. The second one is its cyclical nature, that is the fact that income and consumption differences across agents shrink during booms and expand in recessions. The third mechanism is idiosyncratic risk, namely the probability of experiencing large income and consumption drops. And the fourth is the fact that this probability is higher in recessions.

These four mechanisms are of course also present in economies with more complex dimensions of heterogeneity. But the tractability of our model makes it easier to design counterfactual experiments that isolate the role of each of these channels in shaping the effect of inequality on

business cycles. We have two main findings. First, taken together, these four channels provide meaningful amplification of business cycle shocks. Without any inequality and risk, our model would experience more moderate fluctuations in real activity. Second, most of this amplification comes from the interaction between the high permanent level of inequality and the cyclicity of idiosyncratic risk—the first and fourth mechanisms in the list above. The high level of inequality magnifies the potential losses of the unlucky households who are hit by economic hardship. And the cyclicity of idiosyncratic risk makes the risk of this hardship more acute in recessions. To protect themselves against this risk, unconstrained households engage in cyclical precautionary saving, which amplifies fluctuations.

A Data

This appendix describes the nine data series used for the estimation of the model. The source of our seven aggregate time series is the FRED dataset, available on the website of the Federal Reserve Bank of St. Louis. The series are available from 1955:III to 2019:IV, and defined as follows (series acronym in parenthesis).

1. Growth rate of real GDP per capita:

$$100 \cdot \Delta \log \left[\frac{\text{Gross Domestic Product (GDP)}}{\text{Population} \cdot \text{GDP Implicit Price Deflator (GDPDEF)}} \right]$$

2. Growth rate of real consumption per capita:

$$100 \cdot \Delta \log \left[\frac{\text{Personal Consumption Expenditure: Nondurable Goods (PCND)} + \text{Services (PCESV)}}{\text{Population} \cdot \text{GDP Implicit Price Deflator (GDPDEF)}} \right]$$

3. Growth rate of real investment per capita:

$$100 \cdot \Delta \log \left[\frac{\text{Gross Private Domestic Investment (GPDI)} + \text{Personal Consumption Expenditures: Durable Goods (PCDG)}}{\text{Population} \cdot \text{GDP Implicit Price Deflator (GDPDEF)}} \right]$$

4. Logarithm of hours per capita:

$$100 \cdot \log \left[\frac{\text{Total Economy: Hours of All Persons}}{\text{Population}} \right]$$

5. Growth rate of real hourly wages:

$$100 \cdot \Delta \log \left[\frac{\text{Total Economy: Compensation of Employees (W209RC1Q027SBEA)}}{\text{Total Economy: Hours of All Persons} \cdot \text{GDP Implicit Price Deflator (GDPDEF)}} \right]$$

6. Inflation rate:

$$100 \cdot \Delta \log [\text{GDP Implicit Price Deflator (GDPDEF)}]$$

7. Short-term nominal interest rate:

$$\frac{\text{Effective Federal Funds Rate (FEDFUNDS)}}{4},$$

where the population series used to compute the quantities per capita is the Hodrick-Prescott trend (estimated with smoothing parameter equal to 1600) of the logarithm of the Civilian Non-institutional Population (CNP16OV) series. The reason to use this smooth population series is to avoid the spikes in the original series that correspond to the census years. The series of GDP, PCND, PCESV, GDPDI, PCDG and W209RC1Q027SBEA are in current dollars, while GDPDEF is a chain-type price index that is equal to 100 in 2012. The series of hours worked comes from the Total U.S. Economy Hours & Employment data file, available on the Bureau of Labor Statistics website at www.bls.gov/lpc/special_requests/us_total_hrs_emp.xlsx.

The remaining two time series measure inequality in labor earnings and disposable income, and have been constructed using data from the Annual Social and Economic (ASEC) supplement of the Current Population Survey, as reported by IPUMS. Following [Heathcote et al. \(2020\)](#), we drop households with any member in the armed forces, or having zero or negative ASEC weight, or reporting no hours worked but non-zero earnings, and households with no reference person or without any member between 25 and 60 years of age. We then construct the series of household labor earnings as (IPUMS's series acronym in parenthesis)

$$\frac{\text{Salary (INCWAGE)} + (2/3) \cdot [\text{Non-Farm Business Income (INCBUS)} + \text{Farm Business Income (INCFARM)}]}{\text{Number of Adult Equivalents}}$$

and disposable income:

$$\frac{\text{Total Income (INCTOT)} - \text{Tax Liabilities}}{\text{Number of Adult Equivalents}},$$

where the Number of Adult Equivalents follows the definition in [Heathcote et al. \(2010\)](#). Tax liabilities after 1991 are calculated as the sum of Federal Income Tax Liability Before Credits (FED-TAX), State Income Tax Liability Before Credits (STATETAX) and Social Security Retirement Pay-

roll Deduction (FICA), minus Earned Income Tax Credit (EITCRED). Before 1991, IPUMS does not provide tax information, so we compute tax liabilities using [Taxesim](#), which estimates Federal Taxes (FIITAX), State Taxes (SIITAX, available starting in 1977) and Social Security Retirement Payroll Deduction (FICA). Accordingly, before 1991, tax liabilities are calculated as FIITAX + SIITAX + FICA. Section [A.1](#) describes how we obtain Taxesim estimates.

After assembling household labor earnings, we compute their inequality as follows. First, for each year in the sample, we discretize the distribution of labor earnings using its 15th (P15), 35th (P35), 45th (P45), 55th (P55), 65th (P65), 75th (P75), 85th (P85) and 95th (P95) percentiles. Each percentile is assumed to represent a specific pool of households in the population: P15 represents the 30 percent of households with labor earnings below the 30th percentile, P35 the 10 percent of households with labor earnings between the 30th and 40th percentiles, and so on. We then take the logarithm of these percentiles and compute their weighted standard deviation in each year.⁷ Given that the resulting time series has a well-known pronounced upward trend that our model is not equipped to match, we de-trend it using a band-pass filter that extracts fluctuations with periodicities lower or equal to 30 years. We use an identical procedure to calculate our yearly measure of disposable income inequality.

A.1 Taxesim estimates

To obtain tax liability estimates from Taxesim, we feed the system with CPS data following the guidelines provided by the [IRS](#). In particular, the inputs are:

1. **mstat**: Reporter's marital status. All reporters are single (mstat=1), unless their spouse is present in the household (mstat=2). For married couples, the reporter is the household's head (RELATE=101) unless their income is 0 (in this case, the spouse (RELATE=201) reports for the couple).
2. **state**: Reporter's state of residency. Whenever CPS identifies reporters within a region (union of states) rather than individual states, the reporter is assigned to the first state in the union.

⁷In the computation of the yearly standard deviation, the logarithm of P15 has a weight equal to 0.3, because it represents 30 percent of the population; the logarithms of P35, ..., P95 have all weights equal to 0.1, since they represent 10 percent of the population. Our discretization of the distribution starts with P15, because percentiles lower than P15 contain zero or negative values in at least one year of the sample. Therefore, we cannot take their logarithm.

3. **swages**: Wage of reporter's spouse (only for those filing together). Spouse's wage is $INCWAGE+INCBUS+INCFARM$.
4. **dep**: Reporter's number of dependents. All children below age 19 and all children below age 24 and going to school are dependents.⁸ Qualifying relatives who make less than \$4,300 (in 2020 \$) are also dependents.⁹ Finally, disabled people living within the household are dependents.
5. **dividends**: IPUMS's $INCIDR$ before 1975 and $INCDRT$ and $DIVID$ thereafter.
6. **intrec**: IPUMS's $INTREC$.
7. **otherprop**: IPUMS's $INCRENT$.
8. **nonprop**: IPUMS's $INCALIM$.
9. **pensions**: IPUMS's $INCRETIR$.
10. **gssi**: $INCSS + INCWELFR + INCGOV + INCALOTH + INCSSI + INCVET + INCSURV + INCDISAB + INCEDUC + INCCHILD$.
11. **ui**: IPUMS's $INCUNEMP$.

References

- ACHARYA, S., W. CHEN, M. D. NEGRO, K. DOGRA, A. GLEICH, S. GOYAL, D. LEE, E. MATLIN, R. SARFATI, AND S. SENGUPTA (2023): "Estimating HANK for Central Banks," Staff Reports 1071, Federal Reserve Bank of New York.
- ACHARYA, S. AND K. DOGRA (2020): "Understanding HANK: Insights from a PRANK," *Econometrica*, 88, 1113–1158.
- AGUIAR, M. AND M. BILS (2015): "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review*, 105, 2725–56.
- ASCARI, G., A. COLCIAGO, AND L. ROSSI (2017): "Limited Asset Market Participation, Sticky Wages, And Monetary Policy," *Economic Inquiry*, 55, 878–897.

⁸This is unless the children provide more than 50 percent of their support during the year, in which case they are independent taxpayers. To identify these cases, we compute the amount spent on a child as 90 percent of the household income excluding the child's income divided the number of adult equivalents in the household, times the adult equivalent weight of the child (the 90 percent follows the personal savings statistics provided by FRED). If the child makes more than half of this, he/she is an independent reporter.

⁹We disregard the IRS's second requirement of not providing for more than half of his/her support during the year, assuming that the first is more stringent.

- ATTANASIO, O. P. AND L. PISTAFERRI (2016): “Consumption Inequality,” *Journal of Economic Perspectives*, 30, 3–28.
- AUCLERT, A. (2019): “Monetary policy and the redistribution channel,” *American Economic Review*, 109, 2333–67.
- AUCLERT, A., M. ROGNLIE, AND L. STRAUB (2018): “The intertemporal keynesian cross,” .
- (2020): “Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model,” .
- BAYER, C., B. BORN, AND R. LUETTICKE (2020): “Shocks, Frictions, and Inequality in US Business Cycles,” CEPR Discussion Papers 14364.
- BAYER, C., R. LÜTTICKE, L. PHAM-DAO, AND V. TJADEN (2019): “Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk,” *Econometrica*, 87, 255–290.
- BERGER, D. W., L. BOCOLA, AND A. DOVIS (2019): “Imperfect Risk-Sharing and the Business Cycle,” NBER Working Papers 26032, National Bureau of Economic Research, Inc.
- BILBIIE, F. O. (2008): “Limited asset markets participation, monetary policy and (inverted) aggregate demand logic,” *Journal of Economic Theory*, 140, 162–196.
- (2018): “Monetary policy and heterogeneity: An analytical framework,” .
- (2020): “The new Keynesian cross,” *Journal of Monetary Economics*, 114, 90–108.
- BILBIIE, F. O., D. R. KAENZIG, AND P. SURICO (2022): “Capital and income inequality: An aggregate-demand complementarity,” *Journal of Monetary Economics*, 126, 154–169.
- BILBIIE, F. O. AND X. RAGOT (2017): “Optimal monetary policy and liquidity with heterogeneous households,” .
- BROER, T., J. DRUEDAHL, K. HARMENBERG, AND E. OBERG (2022a): “The unemployment-risk channel in business-cycle fluctuations,” Mimeo.
- BROER, T., N.-J. HARBO HANSEN, P. KRUSELL, AND E. ÖBERG (2020): “The New Keynesian transmission mechanism: A heterogeneous-agent perspective,” *The Review of Economic Studies*, 87, 77–101.
- BROER, T., J. KRAMER, AND K. MITMAN (2022b): “The Curious Incidence of Monetary Policy Shocks Across the Income Distribution,” Mimeo.
- CAMPBELL, J. Y. AND N. G. MANKIW (1989): “Consumption, income, and interest rates: Reinterpreting the time series evidence,” *NBER Macroeconomics Annual*, 4, 185–216.
- CARROLL, C., J. SLACALEK, AND M. SOMMER (2019): “Dissecting Saving Dynamics Measuring Credit, Wealth and Precautionary Effects,” Mimeo.

- CHALLE, E., J. MATHERON, X. RAGOT, AND J. F. RUBIO-RAMIREZ (2017): “Precautionary saving and aggregate demand,” *Quantitative Economics*, 8, 435–478.
- CHANG, M., X. CHEN, AND F. SCHORFHEIDE (2024): “Heterogeneity and Aggregate Fluctuations,” *Journal of Political Economy*, forthcoming.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): “Nominal rigidities and the dynamic effects of a shock to monetary policy,” *Journal of Political Economy*, 113, 1–45.
- COIBION, O., Y. GORODNICHENKO, L. KUENG, AND J. SILVIA (2017): “Innocent bystanders? Monetary policy and inequality,” *Journal of Monetary Economics*, 88, 70–89.
- CONSTANTINIDES, G. M. AND D. DUFFIE (1996): “Asset Pricing with Heterogeneous Consumers,” *Journal of Political Economy*, 104, 219–240.
- CURDIA, V., A. FERRERO, G. C. NG, AND A. TAMBALOTTI (2015): “Has U.S. monetary policy tracked the efficient interest rate?” *Journal of Monetary Economics*, 70, 72–83.
- DEBORTOLI, D. AND J. GALÍ (2018): “Monetary policy with heterogeneous agents: Insights from TANK models,” .
- (2023): “Idiosyncratic income risk and aggregate fluctuations,” *American Economic Journal: Macroeconomics*, forthcoming.
- DOMINGUEZ-DIAZ, R. (2022): “Precautionary savings and financial frictions,” Mimeo.
- EGGERTSSON, G. B. AND P. KRUGMAN (2012): “Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach,” *The Quarterly Journal of Economics*, 127, 1469–1513.
- FERRIERE, A. AND G. NAVARRO (2020): “The Heterogeneous Effects of Government Spending: It’s All About Taxes,” Mimeo, Paris School of Economics.
- GALÍ, J., J. D. LÓPEZ-SALIDO, AND J. VALLÉS (2007): “Understanding the effects of government spending on consumption,” *Journal of the European Economic Association*, 5, 227–270.
- GORNEMANN, N., K. KUESTER, AND M. NAKAJIMA (2016): “Doves for the rich, hawks for the poor? Distributional consequences of monetary policy,” .
- GRAVES, S. (2023): “Does Unemployment Risk Affect Business Cycle Dynamics?” Mimeo,.
- GUERRIERI, V. AND G. LORENZONI (2017): “Credit crises, precautionary savings, and the liquidity trap,” *The Quarterly Journal of Economics*, 132, 1427–1467.
- GUVENEN, F., F. KARAHAN, S. OZKAN, 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., S. OZKAN, AND J. SONG (2014): “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, 122, 621–660.
- 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.
- HAAN, W. J. D., P. RENDAHL, AND M. RIEGLER (2018): “Unemployment (Fears) and Deflationary Spirals,” *Journal of the European Economic Association*, 16, 1281–1349.
- HAGEDORN, M., I. MANOVSKII, AND K. MITMAN (2019): “The fiscal multiplier,” Tech. rep., National Bureau of Economic Research.
- HEATHCOTE, J. AND F. PERRI (2018): “Wealth and Volatility,” *Review of Economic Studies*.
- HEATHCOTE, J., F. PERRI, AND G. VIOLANTE (2020): “The Rise of US Earnings Inequality: Does the Cycle Drive the Trend?” *Review of Economic Dynamics*, 37, 181–204.
- HEATHCOTE, J., F. PERRI, AND G. L. VIOLANTE (2010): “More Unequal We Stand? Inequality Dynamics in the United States 1967–2021,” *Review of Economic Dynamics*, 13, 15–51.
- HEATHCOTE, J., F. PERRI, G. L. VIOLANTE, AND L. ZHANG (2023): “More Unequal We Stand? Inequality Dynamics in the United States 1967–2021,” *Review of Economic Dynamics*, forthcoming.
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2017): “Optimal Tax Progressivity: An Analytical Framework,” *The Quarterly Journal of Economics*, 132, 1693–1754.
- JUSTINIANO, A., G. PRIMICERI, AND A. TAMBALOTTI (2011): “Investment Shocks and the Relative Price of Investment,” *Review of Economic Dynamics*, 14, 101–121.
- JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2010): “Investment shocks and business cycles,” *Journal of Monetary Economics*, 57, 132–145.
- (2013): “Is There a Trade-Off between Inflation and Output Stabilization?” *American Economic Journal: Macroeconomics*, 5, 1–31.
- KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): “Monetary policy according to HANK,” *American Economic Review*, 108, 697–743.
- KAPLAN, G. AND G. L. VIOLANTE (2022): “The Marginal Propensity to Consume in Heterogeneous Agent Models,” *Annual Review of Economics*, 14, 747–775.
- KAPLAN, G., G. L. VIOLANTE, AND J. WEIDNER (2014): “The wealthy hand-to-mouth,” *Brookings Papers on Economic Activity*, 2014, 77–138.
- KRUSELL, P., L. E. OHANIAN, J.-V. RIOS-RULL, AND G. L. VIOLANTE (2000): “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68, 1029–1054.

- LENZA, M. AND J. SLACALEK (2021): “How Does Monetary Policy Affect Income and Wealth Inequality? Evidence from Quantitative Easing in the Euro Area,” CEPR Discussion Papers 16079, C.E.P.R. Discussion Papers.
- MCKAY, A., E. NAKAMURA, AND J. STEINSSON (2016): “The power of forward guidance revisited,” *American Economic Review*, 106, 3133–58.
- MCKAY, A. AND R. REIS (2016): “The role of automatic stabilizers in the US business cycle,” *Econometrica*, 84, 141–194.
- NISTICO, S. (2016): “Optimal Monetary Policy And Financial Stability In A Non-Ricardian Economy,” *Journal of the European Economic Association*, 14, 1225–1252.
- OH, H. AND R. REIS (2012): “Targeted transfers and the fiscal response to the great recession,” *Journal of Monetary Economics*, 59, S50–S64.
- PARKER, J. A., N. S. SOULELES, D. S. JOHNSON, AND R. MCCLELLAND (2013): “Consumer Spending and the Economic Stimulus Payments of 2008,” *American Economic Review*, 103, 2530–2553.
- PATTERSON, C. (2019): “The matching multiplier and the amplification of recessions,” .
- PFLUEGER, C., E. SIRIWARDANE, AND A. SUNDERAM (2020): “Financial Market Risk Perceptions and the Macroeconomy,” *Quarterly Journal of Economics*, 135, 1443–1491.
- RAVN, M. O. AND V. STERK (2017): “Job uncertainty and deep recessions,” *Journal of Monetary Economics*, 90, 125–141.
- (2020): “Macroeconomic fluctuations with HANK & SAM: an analytical approach,” *Journal of the European Economic Association*.
- SLACALEK, J., O. TRISTANI, AND G. L. VIOLANTE (2020): “Household balance sheet channels of monetary policy: A back of the envelope calculation for the euro area,” *Journal of Economic Dynamics and Control*, 115.
- SMETS, F. AND R. WOUTERS (2007): “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, 97, 586–606.
- SOULELES, N. S., J. A. PARKER, AND D. S. JOHNSON (2006): “Household Expenditure and the Income Tax Rebates of 2001,” *American Economic Review*, 96, 1589–1610.
- STORESLETTEN, K., C. TELMER, AND A. YARON (2004): “Cyclical Dynamics in Idiosyncratic Labor Market Risk,” *Journal of Political Economy*, 112, 695–717.
- WERNING, I. (2015): “Incomplete markets and aggregate demand,” *NBER Working Paper*, No. 21448.