The redistributive power of business cycle fluctuations*

Marcin Bielecki† Michał Brzoza-Brzezina‡ Marcin Kolasa§

Abstract

How do business cycle fluctuations redistribute welfare between different age cohorts? Is this redistribution quantitatively important? What are the channels through which it occurs and how is it affected by monetary policy? To answer these (and other) questions, we construct a New Keynesian life-cycle model and estimate it for the U.S. We find that business cycles are an important source of intergenerational redistribution. Years of large fluctuations can impact the remaining lifetime welfare of some cohorts by an amount equivalent to 30% of their annual consumption. These first-order effects do not net out over a typical life cycle: some cohorts (e.g., born in the 1950s) have been much less lucky than others. Life cycle aspects also significantly amplify second-order costs of business cycles, increasing them by half relative to a comparable setup with representative agents. Safety and monetary policy shocks play an overproportional role in driving intergenerational redistribution. Systematic monetary policy that responds more actively to inflation and output increases the amount of redistribution.

JEL: E31, E51, E52, H5, J11

Keywords: Business Cycles, Welfare Redistribution, Monetary Policy, Life-cycle Model

*We wish to thank Michał Reczek for providing excellent research assistance and participants of the ASSA meeting in Philadelphia, Computing in Economics and Finance conference, Summer Workshop on Macro & Finance, Czech National Bank Research Conference, EUI-UW conference, LMU-NYU workshop, and seminars at the ECB and NBP for useful comments and discussions. Comment received from Arpad Abraham, Richard Blundell, Morteza Ghomi and Mathias Klein are gratefully acknowledged as well. The views presented in this paper are those of the authors and not those of the institutions they are affiliated with. Parts of the project were done when Marcin Bielecki and Michał Brzoza-Brzezina were visiting the European Central Bank. Financing by the Polish National Science Center grant 2024/55/B/HS4/00276 is also gratefully acknowledged.

†University of Warsaw and Narodowy Bank Polski; Email: m.p.bielecki@uw.edu.pl.

‡SGH Warsaw School of Economics and Narodowy Bank Polski; Email: mbrzez@sgh.waw.pl.

§SGH Warsaw School of Economics and International Monetary Fund; Email: mkolas@sgh.waw.pl.

1 Introduction

Business cycle fluctuations do not affect all agents equally. As households differ, among other things, in their labor market status or asset holdings, they are differently affected by the same shocks. This observation is not new and there exists a large literature analyzing the redistributive impact of inflation and monetary or fiscal policy. In this paper, we look at the business cycles as a whole, but focus on how their impact differs between households belonging to different age cohorts. Taking this life-cycle perspective allows us to ask a number of important questions. How strongly do business cycle fluctuations redistribute across generations? Does their positive and negative influence on agents net out over a typical lifetime? Which business cycle shocks are mainly responsible for this type of redistribution and through which channels? Which stabilization policies dampen and which amplify the intergenerational redistributive impact of business cycles?

While a detailed presentation of our findings follows further down the road, we begin with a result that highlights the significance and novelty of our findings. We show that the redistributive impact of cyclical fluctuations is large and does not net out over typical lifetimes. Historically, in years of large fluctuations, some cohorts faced gains or losses in remaining lifetime welfare equivalent to over 30% of annual consumption. While such effects can be undone by subsequent shocks that act in the opposite direction, this averaging-out process is far from perfect in the context of a typical lifespan. In our sample, some cohorts were lucky to live through the period of mainly favorable shocks, which boosted their lifetime welfare even by the equivalent of an additional annual consumption. Others were less lucky and faced a lifetime welfare loss. We consider these results striking and giving a clear motivation to investigate the topic in more depth.

We do so by constructing a New Keynesian life-cycle model and estimating it with US economic data going back to 1960. We use the model to filter the sources of business cycles and peek into the drivers of intergenerational redistribution. We show that business cycle shocks typically considered in the literature differ strongly in their redistributive power. For instance, technology shocks affect all cohorts in the same direction, while monetary policy and bond preference shocks redistribute heavily by creating winners and losers. The latter group of shocks are also the main drivers of the redistribution. The prominence of monetary shocks is particularly intriguing as their role in driving business cycles (e.g., fluctuations in output) is usually estimated to be small. We also discuss the channels through which the redistribution occurs. It proves to work its way primarily via returns on nominal assets, which explains the disproportionately large role of shocks driving nominal interest rates and inflation, and via labor income. Contrary to popular beliefs, fluctuations in house prices play a relatively small role.

We next document how cyclical fluctuations redistributed welfare historically over the period covered by our data sample. We find some episodes during which all cohorts were affected in the same direction – selected booms raised and recessions lowered everyone's welfare. However, most business cycle swings created winners and losers at the same time. This was in particular the case during the Covid-19 pandemic, which, according to our model, was to a large extent driven by highly redistributive shocks. Our general conclusion is hence that not all cycles are the same in their power to redistribute between agents of different age.

As already mentioned, one of our key findings is that the welfare consequences of fluctuations can be large and do not net out over a typical lifetime. Let us explain the latter finding, which has two sources. First, in contrast to a typical model with infinitely lived agents, in our world positive and negative shocks do not net out, as agents lifeline is finite. Second, even if this was the case, shocks still might affect welfare to first order. To keep things simple, let us consider an agent who is affected by a contractionary monetary policy shock when young and an (equally sized) expansionary shock when old. The former lowers her/his welfare, as monetary contractions harm young agents. But the latter lowers her/his welfare as well, as expansionary policy harms old households. These effects - which are inherent to the life-cycle framework - explain why history matters and why some cohorts have been more lucky than others as we highlighted above – a typical lifespan of a household is just too short for the highly redistributive episodes to net out.

Last but not least, the role of monetary policy is particularly interesting. As already mentioned, monetary shocks (i.e., the unexpected component of central bank actions) belong to the most redistributive drivers of business cycles. But systematic policy can affect redistribution as well. Our counterfactual simulations show that, within a class of simple Taylor rules, redistribution is minimized under relatively inactive monetary policy that responds to inflation approximately one-for-one. This is because such a feedback rule helps stabilize the return on nominal assets, which is the main channel through which intergenerational redistribution occurs. Aggressive policy reactions, either with respect to output or inflation, can increase fluctuations-driven redistribution substantially.

We also want to stress here that the nature of redistribution in our model is different from that arising in increasingly popular heterogeneous agent New Keynesian (HANK) models that typically abstract away from life-cycle aspects. In these models, heterogeneity in asset holdings reflects differences in individuals' histories of idiosyncratic shocks. As a result, high (low) asset holdings indicate good (bad) luck and so, from the social welfare perspective, inequality is largely undesired as it is essentially a manifestation of financial market incompleteness. In contrast, in a model with realistic life cycle features, asset inequality between between households of different age is desired as it helps smooth consumption when transitioning between distinct stages of life. This distinction may have important normative consequences as HANK models typically call for aggressive interest rate adjustments to stabilize consumption inequality (Bhandari et al., 2021; Acharya et al., 2023). Our analysis shows that such policies, while helpful in shielding households from uninsurable income risk, may have undesired consequences in form of increased redistribution across generations.

Postponing a detailed literature review to the next section, we wish to underscore three dimensions of the original contribution of our work. First, to our knowledge, there exists no other study that analyses redistributive effects of business cycle fluctuations across generations. Second, we stress the crucial role of first-order (level) effects of cyclical fluctuations. As explained earlier, these arise for two reasons: (i) shocks do not net out over a typical lifetime and (ii) even two perfectly symmetric shocks, hitting a household at different stages of life, can significantly affect its lifetime welfare. While reason (i) is relatively trivial, we believe that (ii) offers a new and important insight into the welfare consequences of fluctuations. These features make aggregate fluctuations more risky from an individual perspective, and hence also significantly amplifies their second-order (volatility) costs. Compared with an otherwise identical representative-agent framework, our model implies that agents would be willing to give up nearly 50 percent more of their lifetime consumption to eliminate aggregate uncertainty. As a result, a realistically calibrated life-cycle model offers additional insights on redistribution and inequality over standard HANK models, in which households face idiosyncratic income shocks but are ex ante identical and have the same planning horizon. Third, while being a byproduct of our main goal, we provide a full-information Bayesian estimation of a model that combines fairly detailed life-cycle features with a standard set of real and nominal rigidities considered in the DSGE literature. While computationally intensive, this task proves feasible, thus opening the door to new quantitative research relying on aggregate models that account for age as an important dimension of household heterogeneity.

Relationship to the literature

Our paper is primarily linked to three streams of the economic literature. First and foremost, we connect to papers which analyze the distributional consequences of various shocks and policies. Several aspects have been discussed so far, both on theoretical and empirical grounds. Doepke and Schneider (2006) show that unexpected inflation transfers resources from old to young households. Similar evidence is provided by Albanesi (2007), who demonstrates the positive relationship between inflation and income inequality in OECD countries. Adam and Zhu (2016) use the Household Finance and Consumption Survey conducted in the euro area and show that surprise inflation generates a significant wealth redistribution. Pallotti et al. (2023) use the same data to measure the welfare effects of the post-pandemic inflation surge and identify young households as net winners.

Regarding monetary policy, Coibion et al. (2017) use micro data to show that a monetary contraction increases inequality. Dossche et al. (2021) document that a conventional monetary policy easing has an inequality-reducing impact mainly by lowering unemployment among the poorest households. Lenza and Slačálek (2021) come to a similar conclusion for unconventional monetary policy tools. Auclert (2019) discusses the channels via which monetary policy redistributes, stressing the role of unhedged interest rate exposure that naturally arises in a life-cycle framework like ours. Bielecki et al. (2022) use a life-cycle model to show that expansionary monetary policy redistributes from old to young generations. HANK models are also used in the context of distributional effects, showing i.a. that poor households are more affected by monetary policy shocks (see e.g., Kaplan et al., 2018; Guo et al., 2023). On the fiscal front, Bhattarai et al. (2023) analyze the distributional consequences of fiscal shocks while Brzoza-Brzezina et al. (2024) study the redistribution driven by the concerted monetary-fiscal stimulus during the Covid pandemic.

In contrast to the literature discussed above, which concentrates on a particular shock or policy, we analyze the distributional consequences of all business cycle fluctuations and their drivers. This allows us to compare the role of various shocks for redistribution, revealing, among other things, the relatively high importance of monetary policy shocks. To our knowledge, the only paper with such a scope is Bayer et al. (2024), who discuss the consequences of business cycle fluctuations for inequality in a HANK framework. In contrast, we concentrate on heterogeneity and redistribution of welfare along the life-cycle dimension.

The second related stream of the literature estimates the cost of business cycle fluctuations. In his seminal contributions, Lucas (1987; 2003) argued that this cost is negligible. Subsequent studies have pointed out that this result may be fragile once realistic frictions and heterogeneity are taken into account. Krusell and Smith (1999) show that allowing for incomplete markets and idiosyncratic income risk can raise the welfare cost of fluctuations as households face different exposure to aggregate shocks. Storesletten et al. (2001) and Krusell et al. (2009) demonstrate that the cost increases further once the impact of aggregate fluctuations on individual risk is taken into account. Other papers accentuated the role of job displacement (Krebs, 2007), persistence of the consumption process (Reis, 2009), or agents' risk aversion (Otrok, 2001; Alvarez and Jermann, 2004).

In general, these studies focus on the second-order effects of business cycles, i.e., on their impact on the volatility of individual choices and the associated precautionary reactions. This is justified if agents are assumed to have infinite planning horizons, as positive and negative shocks then ultimately cancel out, leaving their first-order welfare effects null. Deviations from this approach involve accounting for stochastic growth (Obstfeld, 1994), hysteresis in TFP or unemployment (e.g. Walentin and Westermark 2022; Tervala 2021), or occurrence of non-Gaussian recessions with permanent output losses (Barro, 2009; Jorda et al., 2024). Such mechanisms lead to first-order (i.e. level) welfare effects of economic fluctuations, which may be substantially higher than the second-order effects and can be helpful in explaining the risk premia observed in financial data.

Our paper shows that life cycle features greatly amplify the second-order effects of cyclical fluctuations, but also brings forward the role of first order effects in a different context than discussed in the existing literature. Since human life is finite, first order effects of business cycle fluctuations need not cancel out. As we already explained above, even if shocks average to zero over an individual's lifetime, their effect will not be null because the same shock affects people differently at different ages. We are not aware of papers that document and quantify this effect in the context of business cycle fluctuations, and we show that it can be very large. The most affected cohorts in our sample gained or lost the equivalent of over one annual consumption, even though our model does not include most of the mechanisms that have been found to increase welfare costs.

Finally, on the modeling front, our work is related to papers using an overlapping generations (OLG) setup with New Keynesian features. Such studies typically rely on the stylized Blanchard-Yaari framework and include, among others, Del Negro et al. (2012) on the forward guidance puzzle, Eggertsson et al. (2019) on secular stagnation, and Angeletos et al. (2023) on the possibility of fiscal deficits becoming largely self-financing due to their expansionary effect. As a key simplifying step that allows for tractable household sector aggregation, these models feature age-independent composition of assets, ensured by existence of complete annuity markets. In contrast, our framework features a detailed life-cycle setup in which agents are characterized by age-dependent asset portfolios that closely match their empirical distributions. Moreover, we estimate the model with full-information Bayesian techniques (in the spirit of Smets and Wouters, 2003), which, to our knowledge, have not yet been used in the context of full-scale life-cycle models with real and nominal frictions.

The rest of the paper is structured as follows. Section 2 presents our model and section 3 discusses its calibration and estimation. In section 4, we show our main findings about the impact of business cycle fluctuations on welfare and its redistribution across age cohorts. Section 5 focuses on the role of monetary policy. Section 6 concludes.

2 Model

Our model can be viewed as an extension of a standard medium-sized New Keynesian framework to include life cycle features. Compared to a typical DSGE model (see, e.g., Smets and Wouters, 2007), we change the model frequency to annual and replace the representative agent setup with a block of 80 age-specific cohorts of households that optimize over their life-cycle. These households have access to a fairly broad array of assets: housing, nominal assets, and real assets (physical capital and claims on firm profits). They dynamically solve their portfolio problems, readjusting asset holdings with age and in response to aggregate shocks. Otherwise, the model features a standard set of nominal and real rigidities that can be found in the DSGE literature. A full set of equations defining the model equilibrium can be found in Appendix A.

2.1 Demographics and notation

The model economy is populated by overlapping generations of households. We index their age with j, assuming that $j = 1$ corresponds to the age of 20 and marks the moment of entering the labor market, and $j = J = 80$ is the maximum lifespan of 100 years. The only uninsurable idiosyncratic risk facing households is age-dependent mortality, with the probability of dying denoted as ω_j. Then, the size of age cohort j in period t, denoted as $N_{j,t}$, evolves according to $N_{j+1,t+1} = (1-\omega_j)N_{j,t}$ and total population is given by $N_t = \sum_{j=1}^{J} N_{j,t}$. We assume that the size of the youngest cohort $N_{1,t}$ grows at a constant rate n. Together with time-invariant age-specific survival probabilities, this means that each age cohort, and hence the total population, also grows at rate n.

In what follows, and unless indicated otherwise, all cohort-specific variables (hence indexed by j) are defined at the household level. All other variables are expressed in per capita terms, i.e., they are implicitly divided by total population N_t. All real prices are expressed relative to the price of final goods P_t. Variables without time subscripts indicate the steady state.

2.2 Households

Individual households belonging to a particular age cohort are indexed by ι. Each household supplies differentiated labor services $\ell_{j,t}(\iota)$ to labor unions that pay real wage $z_j w_t(\iota)$ net of a proportional labor income tax at rate τ, where $z_j \geq 0$ denotes age-specific productivity. Due to staggered wage contracts, labor income is subject to idiosyncratic risk. However, we assume that this risk can be perfectly insured within the cohort so that all other allocations chosen by households of the same age are identical. Households can accumulate three types of assets: housing $h_{j,t}$, nominal bonds $b_{j,t}$, and claims on physical capital $k_{j,t}$. Capital can be rented to firms, earning real rental rate $r_{k,t}$. Households also hold shares of monopolistically competitive firms, proportionally to individual physical capital holdings, and hence receive firm profits $f_{j,t}$. Finally, households receive age-specific lump sum transfers $t_{j,t}$ from the fiscal authority, as well as unintentional bequests beq_t that consist of assets of deceased households. The real budget constraint of a household aged j can be written as

$$
\begin{aligned}
& c_{j,t} + p_{h,t}\left[h_{j,t} - (1-\delta_h)\, h_{j-1,t-1}\right] + p_{k,t}\left[k_{j,t} - (1-\delta_k)\, k_{j-1,t-1}\right] + b_{j,t} \\
&\quad = (1-\tau) z_j w_t \ell_{j,t} + r_{k,t} k_{j-1,t-1} + f_{j,t} + \frac{R_{t-1}}{\pi_t} b_{j-1,t-1} + t_{j,t} + beq_t,
\end{aligned} \tag{1}
$$

where $c_{j,t}$ is real consumption, $p_{h,t}$ and $p_{k,t}$ denote the real price of housing and physical capital, δ_h and δ_k are the rates at which these assets depreciate, R_t is the nominal interest rate paid on bonds, $\pi_t = P_t/P_{t-1}$ is the change in the aggregate price level P_t, and $w_t \ell_{j,t} \equiv \int w_t(\iota)\ell_{j,t}(\iota)\mathrm{d}\iota$.

Households derive utility from private consumption, exogenously supplied public consumption g_t, housing, and bond holdings, and experience disutility from working. Household ι chooses its consumption and asset holdings to maximize[1]

$$
U_{j,t}(\iota) = \mathbb{E}_t \sum_{i=0}^{J-j} \beta^i \frac{N_{j+i,t+i}}{N_{j,t}} \left(\begin{array}{c} (1-\varrho)\log(c_{j+i,t+i}(\iota) - \varrho c_{j+i,t+i-1}) + \frac{g}{c_{j+i}} \log(g_{t+i}) \\ -\phi_{j+i} \dfrac{\ell_{j+i,t+i}(\iota)^{1+\varphi}}{1+\varphi} + \psi_{j+i,t+i} \log(h_{j+i,t+i}) \\ +\zeta_{j+i,t+i} \log(1 + b_{j+i,t+i}) \end{array} \right), \tag{2}
$$

where φ is the inverse Frisch elasticity of labor supply, ϱ is the external habit persistence parameter, while $\psi_{j,t} \equiv \psi_j(1+\varepsilon_t^h)$, $\zeta_{j,t} = \zeta_j(1+\varepsilon_t^b)$ and ϕ_j are the age-specific weights that control the importance of housing, bond holdings, and leisure in utility. The weights on housing and bonds are additionally modified by stochastic shocks ε_t^h and ε_t^b that generate exogenous fluctuations in demand for housing and for bonds, respectively. The assumed weights on private and public consumption ensure that their marginal utility is equal in the steady state.

[1]Households do not directly choose their labor effort but supply it according to the demand from labor aggregators, at the wage rate set by labor unions. The problems solved by these two types of labor market intermediaries is described below.

2.3 Labor market intermediaries

Perfectly competitive labor aggregators representing households aged j purchase their differentiated labor services at rate $w_t(\iota)z_j$ and combine according to

$$\ell_{j,t} = \left[\int_0^1 (z_j \ell_{j,t}(\iota))^{1/\mu_w} \, \mathrm{d}\iota\right]^{\mu_w}, \tag{3}$$

where μ_w can be interpreted as the steady state wage markup. We assume that the thus defined bundles of labor services are homogeneous across j, so, when hired by firms, they are remunerated at the same real wage w_t per unit of $\ell_{j,t}$ for all j. Labor aggregators maximize their profits $w_t \ell_{j,t} - \int_0^1 w_t(\iota) z_j \ell_{j,t}(\iota) \mathrm{d}\iota$.

Wage setting is performed by monopolistically competitive labor unions. For tractability, we assume that they operate on behalf of all households, implicitly aggregating the marginal rate of substitution between consumption and leisure over the whole working-age population. Their revenue is taxed at a stochastic rate ε_t^w, which is then rebated back to all unions in a lump sum fashion, and which can be interpreted as a wage cost-push shock. While setting the nominal wage for each individual household, labor unions face a Calvo-style rigidity. The probability that the wage is not reoptimized equals θ_w, in which case it is automatically updated with steady state inflation.

2.4 Firms

Our model features a standard New Keynesian corporate sector, with two stages of production. Perfectly competitive final good producers purchase differentiated inputs $y(f)$ at price $P_t(f)$ and combine them according to the following Dixit-Stiglitz aggregator

$$y_t = \left[\int_0^1 y_t(f)^{\frac{1}{\mu_p}} \mathrm{d}f\right]^{\mu_p}, \tag{4}$$

where μ_p determines the steady state product markup. They maximize $P_t y_t - \int_0^1 P_t(f) y_t(f) \mathrm{d}f$.

At the upstream stage of production, risk-neutral intermediate goods producers rent capital $k_t(f)$ and labor $\ell_t(f)$ to produce according to the standard Cobb-Douglas technology

$$y_t(f) = (1+\varepsilon_t^z)k_t(f)^\alpha \ell_t(f)^{1-\alpha} - \Phi, \tag{5}$$

where ε_t^z is a stochastic total factor productivity (TFP) shock, α is the capital share in production and Φ is the fixed cost.[2] The real profit of firm f is
$p_t(f)y_t(f) - (1+\varepsilon_t^p)\left(w_t\ell_t(f) + r_{k,t}k_t(f)\right) + t_t^f$, where $p_t(f) = P_t(f)/P_t$ and ε_t^p is a stochastic price cost-push shock that we introduce in form of a tax on firms' costs that is rebated back to all firms in a lump sum fashion so that $t_t^f = \varepsilon_t^p\left(w_t\ell_t + r_{k,t}k_t\right)$. Intermediate goods producers set their prices subject to a Calvo rigidity. The probability of not receiving a reoptimization signal is θ_p, in which case the price is automatically updated with steady state inflation.

2.5 Capital goods producers

Perfectly competitive and risk-neutral capital goods producers purchase undepreciated capital from households at real price $p_{k,t}$ and combine it with final goods i_t, subject to quadratic investment adjustment costs and an investment specific technology shock ε_t^i. The thus produced new capital is then resold to households at real price $p_{k,t}$. The economy-wide capital stock per capita hence evolves according to

$$(1+n)k_t = (1-\delta_k)k_{t-1} + (1+\varepsilon_t^i)\left[1 - \frac{S_k}{2}\left(\frac{i_t}{i_{t-1}} - 1\right)^2\right] i_t, \tag{6}$$

where S_k controls the degree of adjustment costs.

2.6 Fiscal and monetary authority

The government has two types of expenditures. It purchases $g_t = g(1+\varepsilon_t^g)$ of final goods, where ε_t^g can be interpreted as a public consumption shock. These goods are then equally redistributed across all households. The government also pays lump sum transfers to house-

[2] We assume that $\Phi = (\mu_p - 1)y$ so that firms earn zero profits in the steady state.

holds, using this instrument to stabilize its debt $b_{g,t}$ in the long run. More specifically, we assume

$$\frac{t_t}{t} = \left(\frac{t_{t-1}}{t}\right)^{\gamma_t} \left(\frac{b_{g,t-1}}{b_g}\right)^{-\eta} (1+\varepsilon_t^t), \tag{7}$$

where ε_t^t is a stochastic component of transfers, η controls how fast the fiscal debt is brought back to the steady state, and γ_t is the transfer smoothing coefficient. Changes in aggregate transfers are proportionally reflected in transfers to individual cohorts so that we have $t_{j,t}/t_j = t_t/t$ for every j. Fiscal spending is financed by collecting labor income taxes from households and by issuing debt that pays a nominal interest rate R_t. The real government budget constraint can hence be written as

$$(1+n)\, b_{g,t} = \frac{R_{t-1}}{\pi_t} b_{g,t-1} + g_t + t_t - \tau w_t h_t. \tag{8}$$

The central bank follows a Taylor-like feedback rule with interest rate smoothing

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R}\right)^{\gamma_R} \left[\left(\frac{\pi_t}{\pi}\right)^{\gamma_\pi} \left(\frac{y_t}{y}\right)^{\gamma_y}\right]^{1-\gamma_R} (1+\varepsilon_t^r), \tag{9}$$

where ε_t^r is a monetary policy shock, γ_R is the interest rate smoothing coefficient, while γ_π and γ_y control the response to deviations of inflation from the target and to output deviations from the steady state.

2.7 Market clearing

The model is closed by imposing equilibrium on all markets. We also assume that the housing stock is fixed in per capita terms at a constant value h. The aggregate resource constraint can be written as

$$y_t \Delta_{p,t} = (1+\varepsilon_t^z) k_{t-1}^{\alpha} \ell_t^{1-\alpha} - \Phi, \tag{10}$$

$$y_t = c_t + i_t + \delta_h p_{h,t} h + g_t, \tag{11}$$

where $\Delta_{p,t} = \int \left(\frac{P_t(f)}{P_t}\right)^{\frac{\mu_p}{1-\mu_p}} \mathrm{d}f$ is a measure of price dispersion due to staggered price contracts.

The model features nine exogenous stochastic shocks, which are assumed to follow $AR(1)$ processes (except the monetary policy and transfer shocks, which are assumed i.i.d).

3 Calibration and estimation

We parameterize the model to reflect the US economy. As it is standard in the DSGE literature, we calibrate most of the parameters that determine the steady state proportions, using long-run averages as targets. Given our focus on redistribution across generations, we pay particular attention to the age profiles of key components of household income and assets. We then estimate the parameters that are crucial for the model dynamics, using Bayesian methods and nine macroeconomic time series.

3.1 Calibrated parameters

We start by characterizing the age profiles of selected income and wealth components for US households. To calculate them, we rely mainly on the Survey of Consumer Finance (SCF) data spanning the period 1989-2016. Details on data sources, definitions, and transformations can be found in Appendix B.

The first row of Figure 1 presents the age profiles of labor income and transfers to households. The labor income follows a well-documented pattern, growing steadily through the middle age, peaking in the early 50s, and then steadily declining. This shape reflects two major processes: (i) wages tend to increase over the life cycle as workers become more productive and gain experience, and (ii) labor market participation starts dropping as households begin moving to retirement. The profile for transfers is fairly flat through most of working age and then increases steeply, mainly as households start receiving pensions, so that transfers become the main source of income for households in their late 60s.

The remaining four panels in Figure 1 show the age profiles of household assets and their main components. The evolution of total net worth over the life cycle follows a well-known pattern: households accumulate assets until retirement, after which they start running them down. However, the composition of net worth changes markedly over the life cycle. While housing tends to be accumulated at a roughly constant rate through middle age, the increase in real financial wealth is slow, and households actually accumulate nominal financial liabilities when young, which mainly take the form of mortgage debt. It is only around the late 40s when household nominal assets become positive, mainly reflecting accumulation of bonds, either held directly or via financial intermediaries such as insurance or pension funds. Another important observation is that households typically run down their real financial wealth fairly quickly after retirement, which contrasts with a moderate decumulation of housing and even slower reduction of nominal assets.

We incorporate these patterns of income and asset evolution over the life cycle to calibrate the steady state of the model. In order to simultaneously remain consistent with aggregate national account statistics that have a wider coverage than the SCF, we proceed as follows. We take mortality rates ω_j from the US demographic data, averaged over the period 1989-2017. Next, we set the age specific productivity z_j and labor disutility ϕ_j parameters to jointly match the smoothed labor income profile from the first panel in Figure 1.[3] We also match exactly the steady state age profile of transfers, scaling t_j such that aggregate transfers are equal to 0.17 of GDP, which is the long-run average observed in the US. We then proceed similarly for the age profile of housing and nominal assets, using the utility function parameters ψ_j and ζ_j, and targeting the ratios of housing and government debt to GDP. The age profile of real financial assets, and hence of total net worth, is determined endogenously in the model.

[3]While doing it, we assume that hours worked are equal for all age cohorts. This is a good approximation of the SCF data for households aged between 20 and 65, but not older ones as they enter retirement. We nevertheless make this assumption to keep things simpler, but our results would not be affected in any important way if we instead matched the profile of hours exactly to the one observed in the data. In other words, it is key for our results to match the total labor income profile rather than separately the profiles for wages and hours worked.

Figure 1: Age profiles in the data

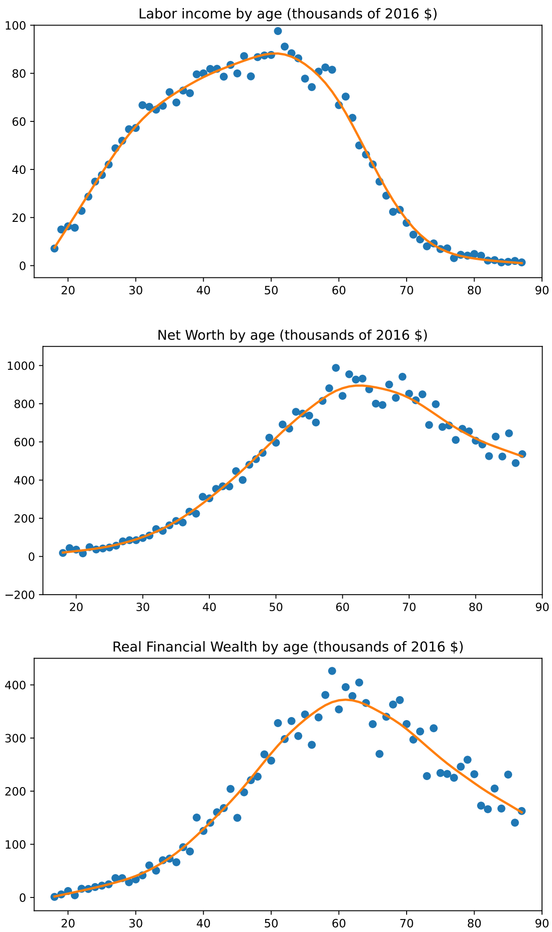

Note: The dots indicate the averages based on SCF spanning the period of 1989-2016. The lines show the fitted trends based on the HP filter with smoothing parameter 100.

The remaining parameters that we keep fixed in estimation are set to standard values in the literature or to match selected steady state proportions in the data. Table 1 summarizes our calibration.

Table 1: Calibrated parameters

Parameter	Value	Description	Source / Target
		A. Demography and households	
ω_j	profile	Mortality rates	Demographics
z_j	profile	Labor productivity profile	SCF
ϕ_j	profile	Hours worked profile	$\ell_j = \bar{\ell}$
ψ_j	profile	Housing profile	h_j as in SCF, $h/y = 1$
ζ_j	profile	Nominal assets profile	b_j as in SCF, $r^a - r = 0.01$
n	0.01	Population growth rate	Demographics
β	0.999	Discount factor	$r = 0.01$
φ	2	Inv. Frish elasticity	Literature
		B. Supply side	
α	0.4	Capital share in production	$k/y = 3$
δ_k	0.1	Physical capital depreciation	Literature
δ_h	0.035	Housing depreciation	$\delta_h h/y = 0.04$
μ_p	1.2	Price markup	Literature
μ_w	1.2	Wage markup	Literature
		C. Government	
g	0.2	Public consumption	$g/y = 0.2$
τ	0.53	Overall tax wedge	$b_g/y = 0.35$

3.2 Bayesian estimation

The remaining parameters of the model are estimated using Bayesian methods. We use nine macroeconomic time series spanning the period 1960-2024. Seven of them are typically used to estimate medium-sized DSGE models like Smets and Wouters (2007). These are: real GDP, real consumption, real investment, real wages (all expressed as log deviations from the HP trend), hours worked (expressed as a log deviation from the mean), inflation, and the nominal interest rate. Additionally, given the importance of nominal assets for redistribution and transfers for older agents, we also include the ratio of government debt to GDP and aggregate transfers to households. See Appendix B for more details on data sources, definitions, and transformations.

We center the prior distributions around standard values considered in the DSGE literature, after taking into account that our model operates at an annual rather than quarterly frequency.[4] We also make the prior distributions fairly diffuse so that their weight in estimation can be considered relatively low compared to most papers using Bayesian methods to estimate DSGE models. The characteristics of the prior and posterior distributions for the structural parameters are summarized in Table 2 and of the processes driving exogenous shocks in Table B.3 of Appendix 6.

Table 2: Estimated structural parameters

Parameter	Prior			Posterior		Description
	type	mean	std	mean	std	
ϱ	beta	0.75^4	0.1	0.19	0.05	Habit persistence
θ_p	beta	0.75^4	0.1	0.27	0.04	Calvo prob. for prices
θ_w	beta	0.75^4	0.1	0.25	0.05	Calvo prob. for wages
S_k	norm	1.0	0.5	0.32	0.06	Inv. adjustment cost
γ_R	beta	0.7^4	0.1	0.38	0.06	Interest rate smoothing
γ_π	norm	1.5	0.1	1.64	0.08	Response of interest rate to inflation
γ_y	beta	0.125	0.05	0.21	0.07	Response of interest rate to output
γ_t	beta	0.75^4	0.1	0.30	0.08	Tax smoothing
η	beta	0.1	0.05	0.11	0.05	Weight on debt in transfer rule

Although our model features a non-standard household block with substantial agent heterogeneity due to the life-cycle dimension, the estimated parameter values (after conversion between quarterly and annual data frequencies) are similar to those obtained in standard representative agent DSGE models. For example, the estimated Calvo probabilities for both prices and wages translate to around 0.66 in quarterly frequency terms, as often assumed in the literature describing the US economy. The estimated coefficients characterizing the interest rate reaction to output and inflation are also standard.

3.3 Business cycle properties

The model generates reactions of key macroeconomic variables that are well aligned with both economic intuition and existing results in the literature. This is evidenced by the impulse response functions presented in Appendix C. Of particular interest is the reaction of the

[4]More precisely, we raise the prior means typically used in quarterly DSGE models to the power of four for the parameters describing probabilities or persistence.

economy to a bond preference shock, which in the context of the US economy is sometimes referred to as a safety shock (Kekre and Lenel, 2024). This shock has the desirable feature of affecting consumption, investment and hours worked in the same direction, while at the same time generating procyclical reactions of inflation and the nominal interest rate. As such, it closely resembles the risk premium shock used by Smets and Wouters (2007) or the main business shock identified by Angeletos et al. (2020). As shown in Table 3, it plays a dominant role in explaining the fluctuations on the nominal side of the economy, overshadowing the relative impact of monetary shocks. The safety shock also accounts for about one-fifth of the unconditional variance of real variables (including real labor income), which are otherwise mainly driven by supply shocks. The supply shocks are key drivers of the return on real financial assets and house prices, the latter also significantly affected by the housing preference shock. Fiscal shocks play only a small role in explaining fluctuations in non-fiscal macroeconomic aggregates, but have substantial impact on consumption of old generations.

Table 3: Shock variance decomposition

Variable	Contribution of shock (in %)								
	ε^z	ε^i	ε^p	ε^w	ε^h	ε^b	ε^r	ε^g	ε^t
Output	10	7	4	45	0	19	14	1	0
Private consumption	10	8	5	41	0	22	13	1	1
Investment	9	15	2	46	0	14	13	1	0
Real labor income	4	6	22	27	0	21	18	1	0
Real house price	7	4	4	26	36	11	12	1	0
Inflation	10	2	7	8	0	50	21	1	2
Nominal interest rate	5	3	4	8	0	67	9	1	3
Real return on nominal assets	12	3	10	7	0	45	23	1	1
Real return on real assets	28	7	8	43	0	6	8	0	0
Public debt	2	3	3	16	0	4	13	3	57
Consumption of 20 yo	10	7	5	43	0	14	19	1	1
Consumption of 40 yo	10	7	5	43	0	12	22	1	0
Consumption of 60 yo	9	7	4	41	0	24	12	1	1
Consumption of 80 yo	5	5	4	20	0	29	8	2	27

4 Redistribution across age and time

This section documents our findings about the redistributive consequences of business cycles. First, we introduce the main concepts used in this section: unexpected welfare gains, their decomposition by channels of influence, and a metric of the total amount of redistribution. We then present our main results. We begin by showing how cyclical fluctuations have redistributed welfare in the US since 1960. We next examine the role of various business cycle shocks in this process—both on average and in historical contexts—and discuss the economic channels through which redistribution takes place. We proceed by checking to what extent the redistributive impact of cyclical fluctuations was netted out over the lifetimes of past generations, which allows us to assess whether business cycles can leave some cohorts significantly better-off than others, or, in other words, whether first order (level) effects of fluctuations matter for welfare. Last but not least, we discuss the impact of life-cycle features on the second order costs of business cycle fluctuations.

4.1 Main concepts

In the analysis that follows, we make extensive use of a central metric for welfare changes WG, referred to as unexpected remaining lifetime welfare gain (loss), or simply welfare gain. For any period t and cohort j, such defined measure will reflect the effects of changes in the economic environment resulting from unexpected shocks that arrive in that period and which, by definition, were unforecastable in period $t-1$. Since agents of different ages respond differently to changes in macroeconomic conditions, WG is cohort-specific, and business cycles therefore redistribute welfare across cohorts.

More formally, we define WG as

$$WG_{j,t} \equiv \mathbb{E}_t U_{j,t} - \mathbb{E}_{t-1} U_{j,t} \equiv \tilde{\mathbb{E}}_t U_{j,t}, \tag{12}$$

where

$$U_{j,t} = \mathbb{E}_t \sum_{i=0}^{J-j} \beta^i \frac{N_{j+i,t+i}}{N_{j,t}} \left(\begin{array}{c} (1-\varrho)\log(c_{j+i,t+i} - \varrho c_{j+i,t+i-1}) + \frac{g}{c_{j+i}} \log(g_{t+i}) \\ -\phi_{j+i} \dfrac{\left(\ell_{j+i,t+i}\right)^{1+\varphi}}{1+\varphi} \Delta_{w,t+i} + \psi_{j+i} \log(h_{j+i,t+i}) \\ +\zeta_{j+i} \log(1+b_{j+i,t+i}) \end{array} \right), \tag{13}$$

and operator $\tilde{\mathbb{E}}_t$ is such that, for any variable X_t, $\tilde{\mathbb{E}}_t X_{t+i} \equiv \mathbb{E}_t X_{t+i} - \mathbb{E}_{t-1} X_{t+i}$. There are two differences between individual household utility $U_{j,t}(\iota)$ given by equation (2) and $U_{j,t}$ defined above. First, we obtain the latter by aggregating across households belonging to the same cohort. As individual households work different hours due to staggered wage contracts, this aggregation introduces a wage dispersion term $\Delta_{w,t} = \int \left(\frac{w_t(\iota)}{w_t}\right)^{\frac{\mu_w}{1-\mu_w}(1+\varphi)} \mathrm{d}\iota$. Second, in formula (13) we omit shocks affecting weights on housing and bonds in utility, thus abstracting away from their direct impact on agents' utility. As these preference-type shocks are still allowed to affect aggregate quantities and prices that are relevant for household welfare, they remain important drivers of redistribution.

Another concept that we will frequently utilize is the amount of redistribution, or $AMOR$. It is supposed to capture any asymmetric movements in an individual cohort's welfare that are not aligned with changes in aggregate welfare. For any period t, the $AMOR$ statistic is calculated as follows:

$$AMOR_t = \sum_{j}^{J-20} \frac{N_{j,t}}{N_t} \left\| WG_{j,t} - \sum_j \frac{N_{j,t}}{N_t} WG_{j,t} \right\|. \tag{14}$$

Intuitively, if welfare of all cohorts changes exactly in unison, the value of $AMOR$ is 0. Otherwise, it becomes more positive the greater the differences in how individual cohorts are affected.[5]

To provide insights on the importance of channels through which business cycle fluctuations affect welfare gains $WG_{j,t}$, we offer a decompostition along the lines developed by

[5] As there is no data on asset holdings of the oldest cohorts, we present our results for households up to 80 years old. $AMOR$ is calculated in consistence with this convention.

Bielecki et al. (2022). More specifically, we focus on first-order effects, which allows us to apply the envelope theorem and express the welfare gains as a sum of components that describe the effects of changes in aggregate prices and quantities relevant for household optimization, while keeping the individual allocations fixed. The welfare gain of a j-aged cohort at time t can therefore be decomposed as follows:

$$WG_{j,t} = \Gamma^{h}_{j,t} + \Gamma^{b}_{j,t} + \Gamma^{kf}_{j,t} + \Gamma^{w\ell}_{j,t} + \Gamma^{tg}_{j,t} + \Gamma^{hab}_{j,t} + \mathcal{O}\left(2\right), \tag{15}$$

where the Γ's capture, respectively, the effects of changes in: house prices, return on nominal assets, return on real assets, labor income, transfers (including accidental bequests) and public consumption, and external habits. The $\mathcal{O}\left(2\right)$ term collects all second (and higher) order effects. As all the Γ's are first-order objects, they lend themselves to a straightforward linear aggregation. In particular, one can compute $WG_{j,t}$ and all of its components separately for each shock.

Since we will estimate our model after linearizing it around the non-stochastic steady state, it is natural to define the first-order accurate decomposition given by equation (15) also using this reference point. Then the respective Γ's can be expressed as follows:

$$\Gamma^{h}_{j,t} = u^{c}_{j} \sum_{i=0}^{J-j} \left(1+r\right)^{-i} \left[\left(1-\delta_{h}\right) h_{j+i-1} - h_{j+i}\right] \tilde{\mathbb{E}}_{t} p^{h}_{t+i} \tag{16}$$

$$\Gamma^{b}_{j,t} = u^{c}_{j} \sum_{i=0}^{J-j} \left(1+r\right)^{-i} b_{j+i-1} \tilde{\mathbb{E}}_{t} r_{t+i} \tag{17}$$

$$\Gamma^{kf}_{j,t} = u^{c}_{j} \sum_{i=0}^{J-j} \left(1+r\right)^{-i} \left\{\left[\left(1-\delta_{k}\right) k_{j+i-1} - k_{j+i}\right] \tilde{\mathbb{E}}_{t} p^{k}_{t+i} + k_{j+i} \tilde{\mathbb{E}}_{t} r^{f}_{t+i}\right\} \tag{18}$$

$$\Gamma^{w\ell}_{j,t} = u^{c}_{j} \sum_{i=0}^{J-j} \left(1+r\right)^{-i} z_{j+i} \left[\left(1-\tau\right) \ell_{j+i} \tilde{\mathbb{E}}_{t} w_{t+i} + \tfrac{\mu_{w}-1}{\mu_{w}} \left(1-\tau\right) w \tilde{\mathbb{E}}_{t} \ell_{j+i,t+i} \right] \tag{19}$$

$$\Gamma^{tg}_{j,t} = u^{c}_{j} \sum_{i=0}^{J-j} \left(1+r\right)^{-i} \tilde{\mathbb{E}}_{t} \left(t_{j+i,t+i} + beq_{t+i}\right) + \sum_{i=0}^{J-j} \beta^{i} \frac{N_{j+i}}{N_{j}} \frac{1}{c_{j+i}} \tilde{\mathbb{E}}_{t} g_{t+i} \tag{20}$$

$$\Gamma^{hab}_{j,t} = -\varrho \sum_{i=0}^{J-j} \beta^{i} \frac{N_{j+i}}{N_{j}} u^{c}_{j+i} \tilde{\mathbb{E}}_{t} c_{j+i,t+i} \tag{21}$$

where $u^{c}_{j,t} \equiv \left[\left(1-\varrho\right) c_{j,t}\right]^{-1}$ is the marginal utility of consumption at age j , $r_{t} = R_{t-1}/\pi_{t}$

is the ex post real interest rate, and $r_t^f \equiv r_t^k + f_t/k_{t-1}$ represents the return on real assets without changes in valuation. Recall that all variables without time subscripts denote their steady state values.

4.2 Welfare redistribution over business cycles

Let us begin by taking a look at the redistribution of welfare in our sample. For each year, Figure 2 presents welfare gains (green) or losses (red) of cohorts aged between 20 and 80, all expressed as percentage of annual steady state consumption of a given cohort.[6] Two important conclusions emerge. First, gains or losses arising from shocks materializing in just one year can be substantial, in several cases exceeding 30% of annual consumption. Second, the welfare consequences of business cycle fluctuations are fairly heterogeneous across time. There are years where all cohorts faced gains or losses. This pattern is true for several NBER dated recessions, for instance 1970, 1974, 1982 or 2001, when all cohorts faced declining welfare. However, such episodes are outnumbered by periods of outright redistribution, when some cohorts faced losses, while others gained. This is not only the case in non-recessionary years, like 2002 or 2005, but also during some recessions, like in 2009.

What is the reason for such heterogeneity and, in particular, for the mostly redistributive pattern of fluctuations? Given that business cycles in our model originate from stochastic shocks, it must be that at least some of them affect different cohorts differently. To inspect this aspect of redistribution, below we discuss the redistributive consequences of selected shocks.

4.3 Redistributive and egalitarian shocks

To present how different types of shocks can differ in their redistributive power, we first present how welfare of individual cohorts is affected by a one-off shock of a given type,

[6]Since WG measures the impact of shocks hitting the economy in one year, we express it in *annual* consumption units, as this ensures meaningful comparability across cohorts. We will use the concept of *lifetime* consumption equivalent when we study the effects of overall aggregate uncertainty, as it is typically done in the literature on the costs of business cycle fluctuations.

Figure 2: Welfare effects of historical business cycles (all shocks)

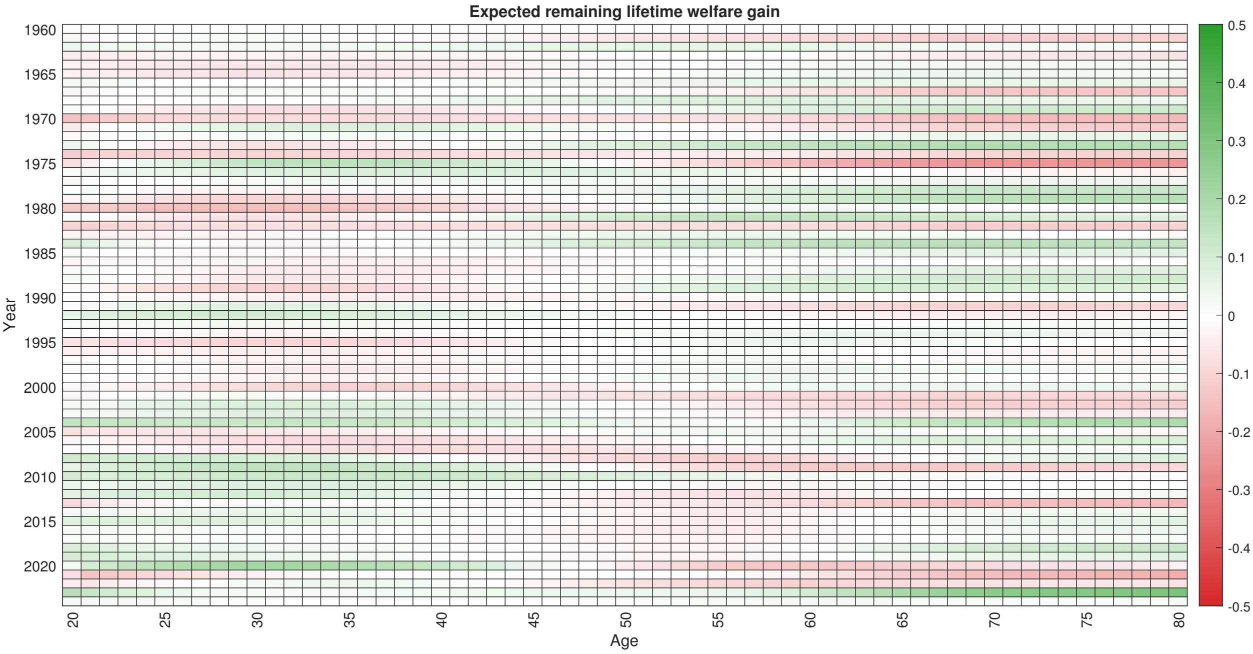

Note: Each rectangle represents the remaining lifetime welfare gain WG for households of different age, as a result of newly arrived shocks in a given year. The gains are expressed in percent of each cohort's steady state consumption.

hitting the economy when it is in the steady state. The ensuing explanation will utilize the decomposition defined in Section 4.1, the lifetime profiles of asset holdings presented in Figure 1, and the impulse responses of macroeconomic aggregates and prices reported in Appendix C.

Let us begin with a shock whose redistributive properties have already been discussed in the literature: an unexpected adjustment of the monetary policy stance. Panel (a) of Figure 3 presents the welfare effects of a monetary tightening and their decomposition across economic channels. As is known, this shock redistributes welfare from younger to older generations. In particular, a surprise fall in inflation and lower economic activity affect young cohorts negatively by increasing their real debt burden and reducing their labor income. Older agents, in contrast, benefit from higher interest rates and lower inflation, which increase the return on nominal assets they accumulated during their lifetime.

Panel (b) of Figure 3 depicts the effects of a safety shock. This shock, by shifting household preferences away from consumption and capital accumulation, generates strong comovement between main indicators of real activity (consumption, investment, output) and inflation. By lowering economic activity, it exerts a negative impact on labor income of working-age households. However, and in contrast to the monetary shock discussed above, young cohorts benefit via the nominal asset channel as the central bank reacts with a sharp monetary accommodation. The opposite is true for middle-aged and old owners of nominal assets. Finally, lower real asset prices benefit households that are in the process of increasing their holdings as part of their life-cycle wealth accumulation. The overall picture is a net welfare gain for households below 40 and a net loss for older agents.

The lower row of Figure 3 shows two shocks that, in contrast to those previously discussed, have relatively similar welfare consequences for all age cohorts. An investment-specific shock raises capital accumulation and hence improves labor market performance, thus increasing income of the young. However, it also raises stock prices, thus making their accumulation of real assets more expensive and lowering welfare of cohorts below 65. Slightly higher inflation benefits youngest, indebted households. Overall, as these forces balance out to a large degree, most cohorts end up with comparable gains.

A similar picture emerges for a total factor productivity shock. All cohorts benefit either from improved labor market performance or from higher transfers from the government (due to higher tax revenues). Interestingly, the effects associated with asset holdings mostly cancel out, with young agents gaining from lower nominal interest rates while losing from accumulation of more expensive real assets, and old agents facing lower returns on nominal assets balanced with higher dividends and stock prices. The overall picture is again a relatively flat redistribution profile, even though the gains clearly decrease with age for households that have reached the peak of their labor productivity.

Figure 3: Welfare effects of selected shocks

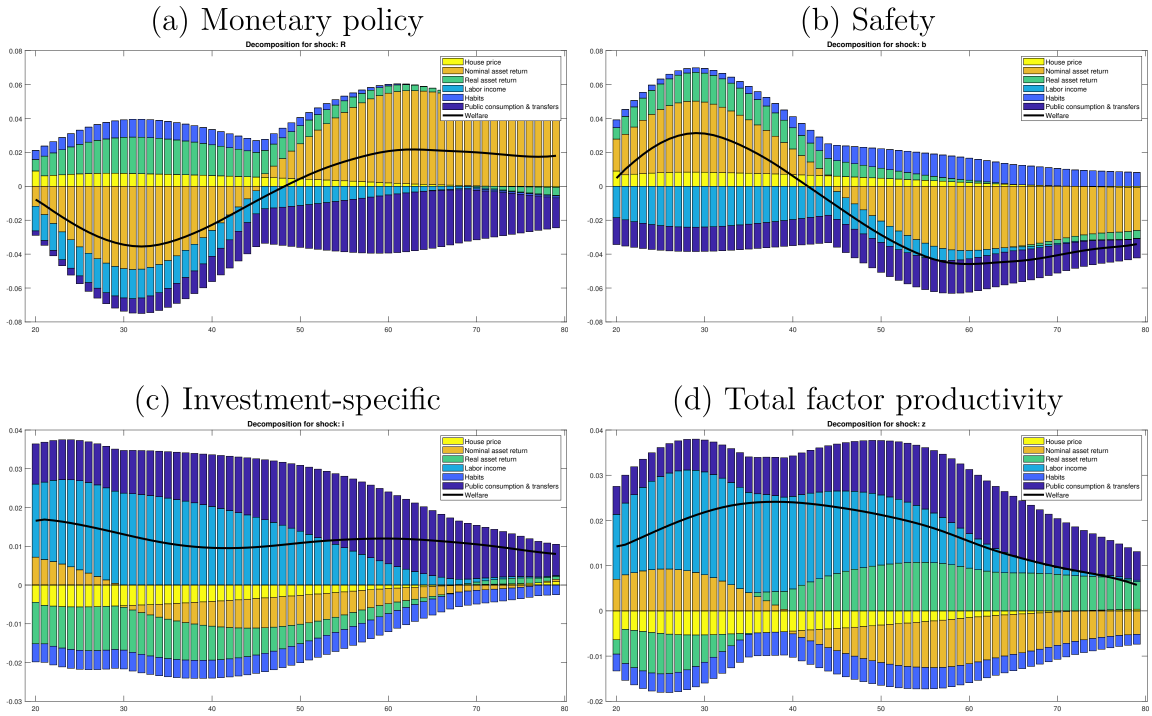

Note: The black lines represent the remaining lifetime welfare gain WG by age, as a consequence of a one-off innovation to a shock of a given type. The bars show the contributions of particular channels. The gains are expressed in percent of the cohort's steady state consumption.

Overall, the redistributive power can vary a lot between shocks. Some of them can be highly redistributive, creating winners and losers. Prime examples of "redistributive" shocks are safety and monetary policy shocks that we discussed above, but also price markup shocks shown in Appendix C. Some shocks, like public consumption and transfer shocks, move welfare of all households in the same direction but carry substantially different quantitative effects. The remaining shocks have a largely similar effect on welfare of different cohorts and hence can be classified as "egalitarian". The primary examples are the two productivity shocks that we discussed above.

The consequences of these differences in redistributive power can be seen in our historical decomposition. As a stark example, Figure 4 depicts the welfare effects for various cohorts resulting from bond preference shocks alone. A clear division runs between cohorts below or above age 42: the same shock benefited younger households at the expense of older ones or vice versa. Such patterns do not occur for “egalitarian” shocks. Figure 5 shows that the historical impact of investment specific shocks was almost symmetric across households of different age, either raising or lowering welfare for every cohort in a given year.

Figure 4: Welfare effects of historical safety shocks

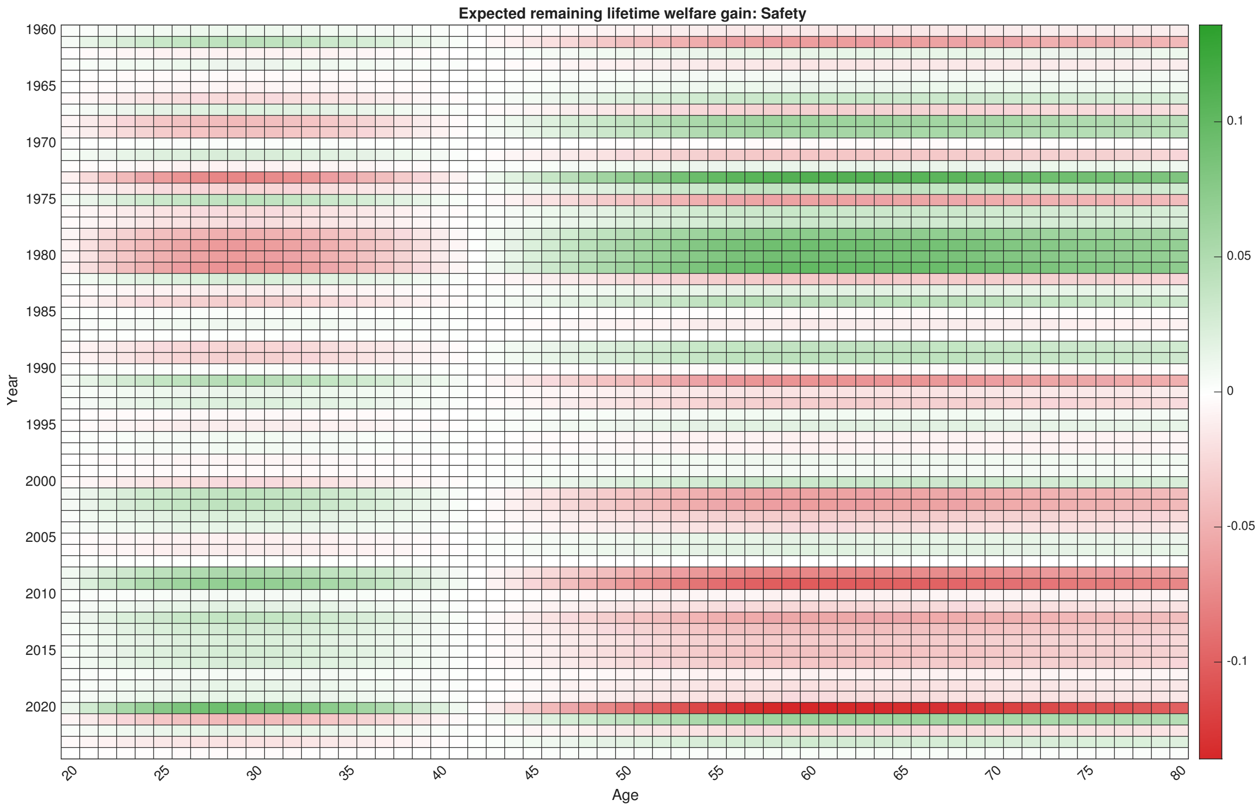

Note: Each rectangle represents the remaining lifetime welfare gain WG for households of different age, as a result of newly arrived safety shocks in a given year. The gains are expressed in percent of each cohort’s steady state consumption.

Figure 5: Welfare effects of historical investment specific shocks

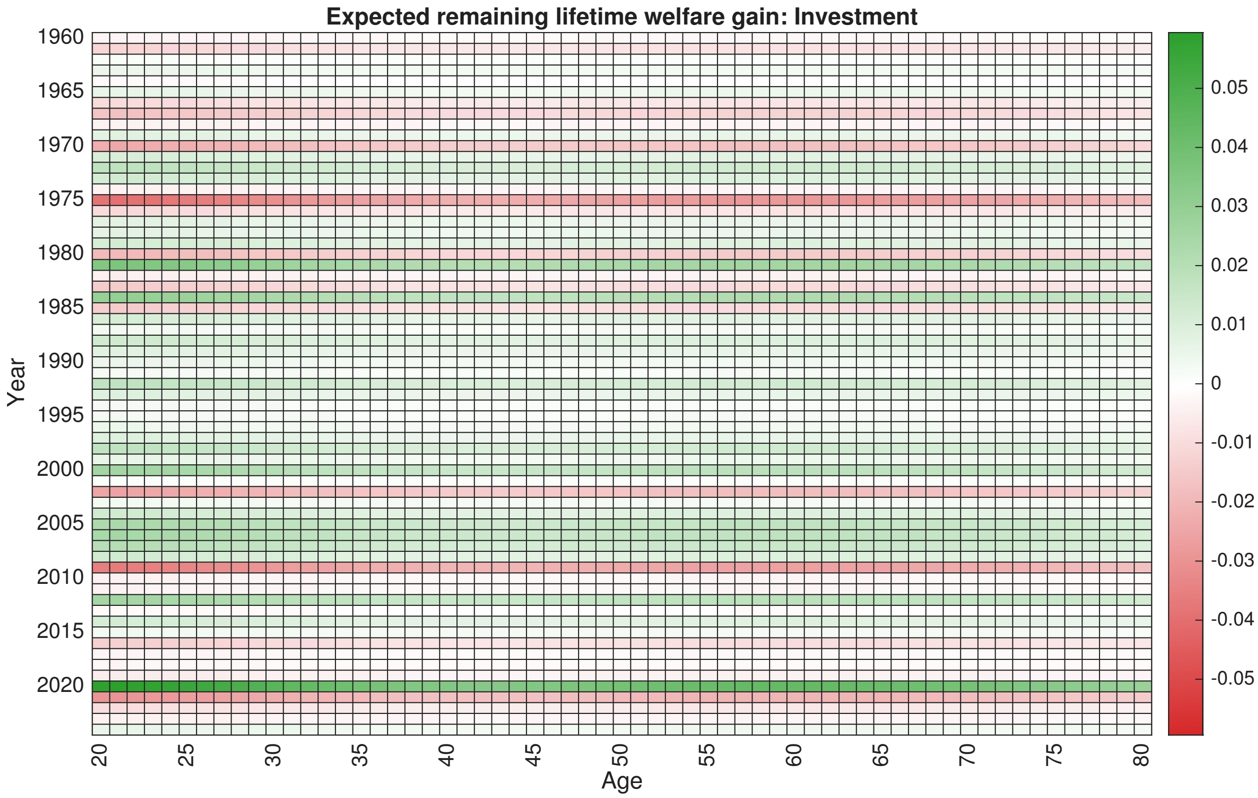

Note: Each rectangle represents the remaining lifetime welfare gain WG for households of different age, as a result of newly arrived investment-specific shocks in a given year. The gains are expressed in percent of each cohort's steady state consumption.

4.4 Drivers of redistribution

While the discussion above suggests the likely drivers of redistribution during cyclical fluctuations, the actual amount of redistribution that can be attributed to particular shocks or channels is a consequence of both their redistributive power and size. For instance, while we have demonstrated that monetary policy shocks are highly redistributive, most of studies find them to be relatively small as interest rate setting by central banks is typically dominated by a systematic response to inflation and other macroeconomic developments. As a result, monetary shocks are not among the key driving forces of aggregate cyclical fluctuations. Using a structural, estimated model, our framework enables us to sort the drivers of redistribution not only by source but also by their influence. To this end, we first calculate the $AMOR$ statistics defined above to gauge the overall amount of welfare redistribution for each year. We then focus on the effects of specific shocks and channels, to figure out the most important sources of welfare differences.

Figure 6 shows that highest redistribution occurred in 1975, 2009 and 2020.[7] All three dates coincide with US recessions. The one in 1975 was characterized by soaring energy prices and high interest rate volatility, and so our framework attributes a big role to price markup and monetary shocks. The other two recessions redistributed welfare mainly via the flight-to-safety behavior that is represented in our model by the safety shock. Figure 7 again shows the $AMOR$ statistics, this time breaking it down into quasi-contributions of macroeconomic channels. The 1975 recession-driven redistribution mostly worked through the nominal assets and labor income channels, highlighting the important role of surprise inflation in affecting the ex post return on nominal assets and real wages. These two channels were also key during the Great Recession, but this time the return on real assets also contributed heavily. Finally, the redistribution during the Covid recession was dominated by changes in asset returns with very little role of labor income channel, which is consistent with strong labor market performance during that period.

[7]The quasi-contribution of a given shock or channel in the $AMOR$ statistics is calculated by assuming that only this shock or channel is active. Due to the absolute value used in formula (14), these "contributions" do not add up to total $AMOR$ as the impact of shocks or channels can net out.

Figure 6: Overall amount of redistribution and quasi-contribution of specific shocks

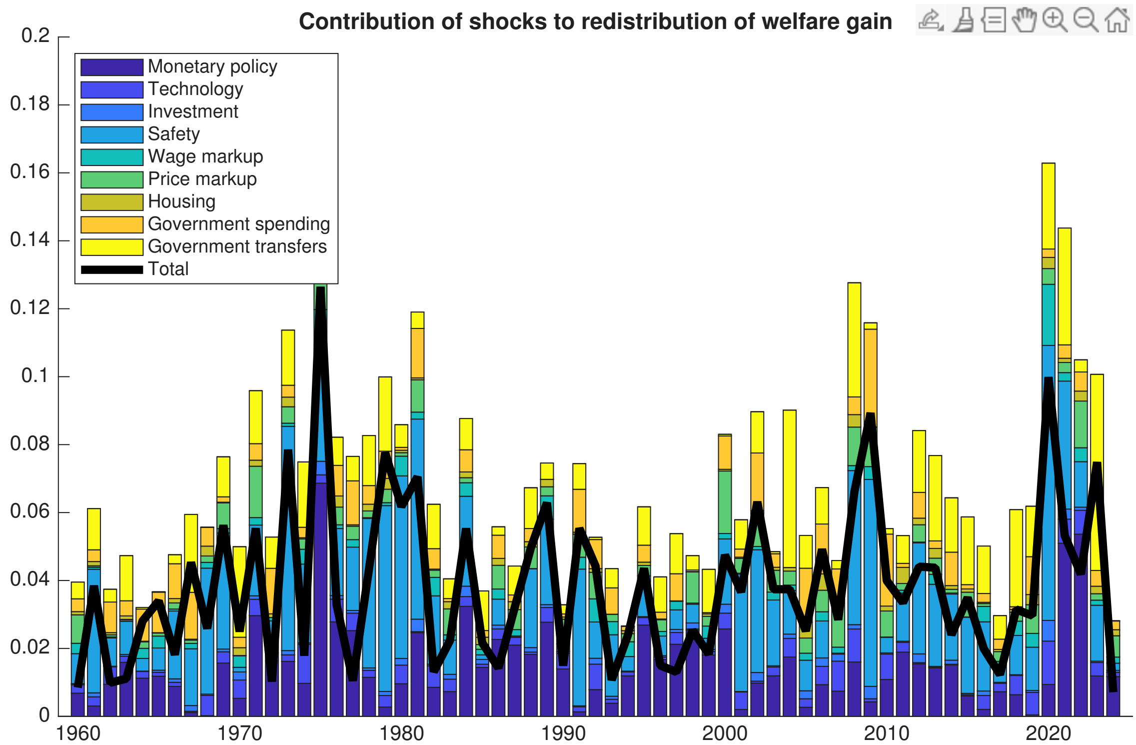

Note: The black line shows the total amount of redistribution $AMOR$ by year. The bars show the quasi-contributions of individual shocks. Due to the absolute value transformation in the $AMOR$ formula, the sum of these contributions is bigger than the combined effect of all shocks.

Figure 7: Overall amount of redistribution and quasi-contribution of specific channels

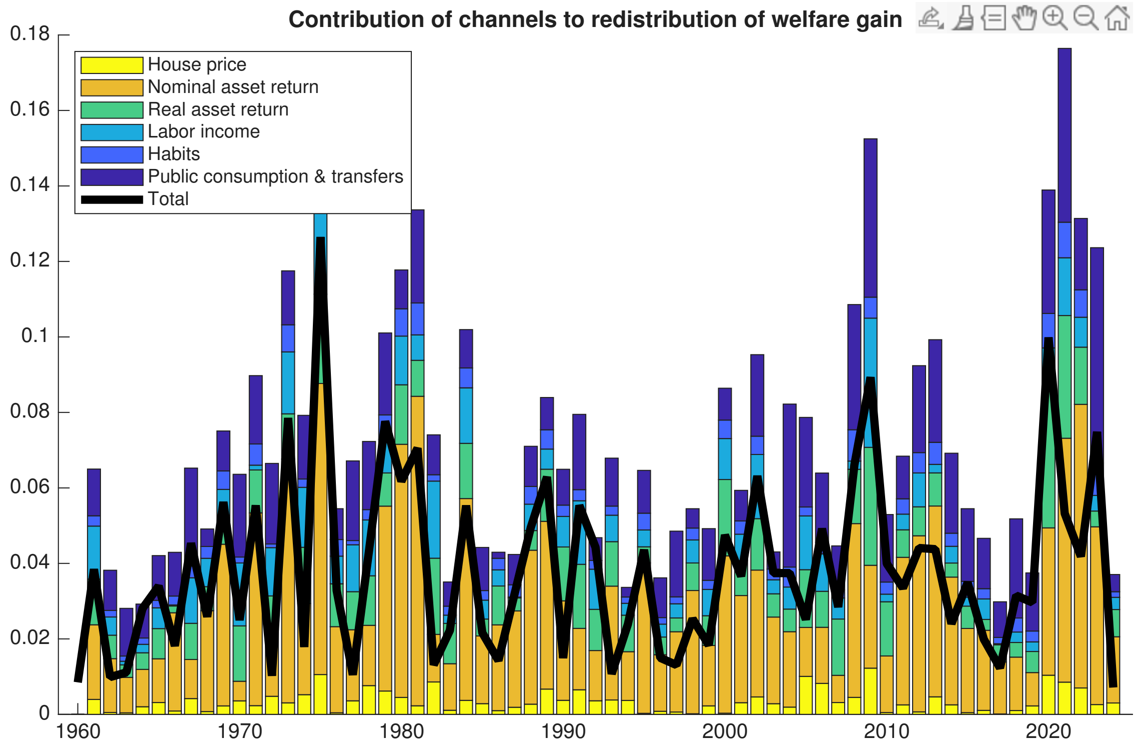

Note: The black line shows the total amount of redistribution $AMOR$ by year. The bars show the contributions of specific channels. Due to the absolute value transformation in the $AMOR$ formula, the sum of these contributions is bigger than the combined effect of all shocks.

Zooming out to encompass the entirety of our sample, Figure 8 presents the average importance of shocks and channels for redistribution in our sample. We find that the two shocks with a particularly high redistributive power – monetary policy and safety – played a key role in driving the intergenerational dispersion of welfare effects of the business cycles, together explaining over 50% of the observed redistribution. This number can be contrasted with the contribution of these shocks to output volatility, which is only about one-third. On the other end of the spectrum one finds – not surprisingly – the highly egalitarian investment shock, with an only 2% share in overall redistribution.

Moving to macroeconomic channels, almost 40% of overall redistribution can be attributed to the return on nominal assets, pointing to the key role of interest rates and inflation in affecting welfare. The labor income channel as well as real assets come next in importance among the channels related to the households budget constraint. Interestingly, and in contrast to common perception, the role of house prices is relatively minor. This is because housing is accumulated until fairly old age and then, unlike financial assets, run down relatively slowly. As a result, house price fluctuations only mildly affect welfare, since what matters in this context are not total asset holdings but their maturing part (Auclert, 2019).

Figure 8: Average contribution of specific shocks and channels to redistribution

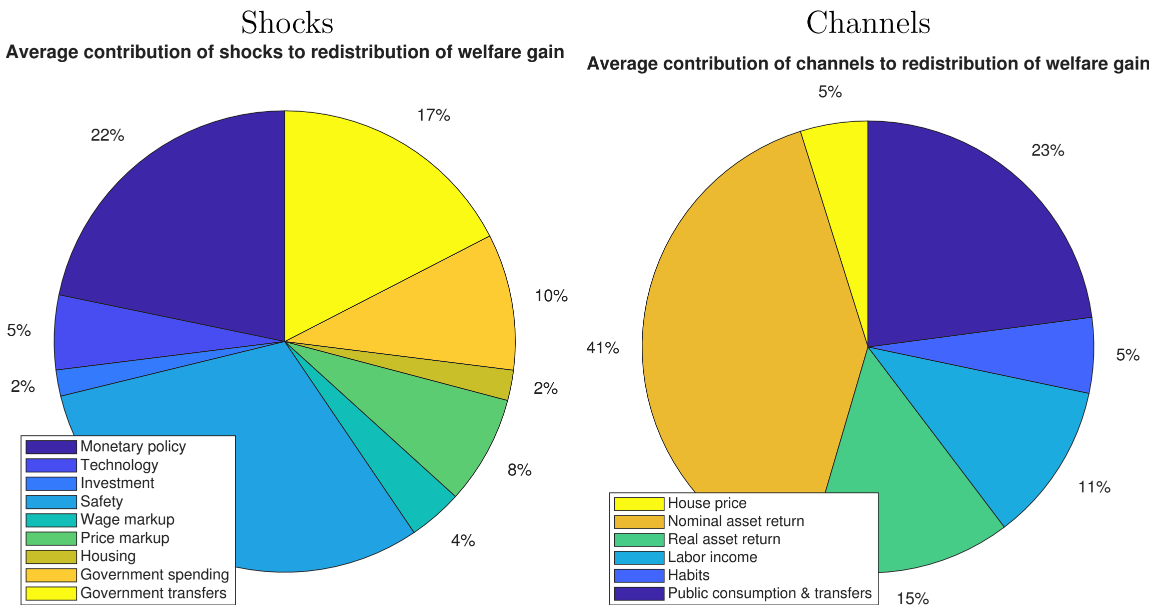

Note: The charts show the contributions of shocks (left panel) and specific channels (right panel) to the total amount of redistribution $AMOR$, averaged over our entire data sample 1960-2024.

4.5 First and second-order welfare effects of business cycle fluctuations

Let us now move to one of the most important findings of this paper. As already discussed in Section 1, there exists a large literature on the welfare consequences of cyclical fluctuations that concentrates on second-order effects of aggregate volatility. This focus is largely because of the standard assumption in the business cycle literature that households live infinitely. The natural measure of welfare is then the expected discounted sum of an infinite stream of period utility flows, as viewed at the beginning of the planning horizon before any uncertainty materializes. Together with the common assumption of zero-mean shocks, this implies that the first-order effects of business cycles are nil.

In our framework, life of an individual household is finite. This leads to two channels which may make cumulative first-order effects differ from zero, and the second-order effects bigger than in a representative agent framework. First, and quite trivially, shocks are not zero on average in a finite lifetime. Second, and less trivially, due to agent heterogeneity over the lifetime, even shocks that are zero on average might have non-zero welfare consequences. To illustrate this point, let us consider an agent who faces only two shocks during her life: a contractionary monetary policy shock when young and an equally sized expansionary shock when old. In line with the discussion in Section 4.3, both episodes will harm this agent's welfare. Figure 9, which presents the historical welfare gains due to monetary policy shocks alone, shows that this effect is not purely theoretical. The black, diagonal line is the lifeline of the cohort born in 1960 (and hence aged 20 in 1980). As is evident, monetary policy shocks affected its welfare almost always negatively. Households belonging to this cohort faced on average a contractionary monetary policy stance in the 1980s and 1990s (when they were young and indebted), and predominantly expansionary policy after 2009 (when they were old and with positive stock of nominal assets). As a result, they suffered welfare losses due to monetary shocks through almost all of their life.

Figure 9: Welfare effects of historical monetary policy shocks

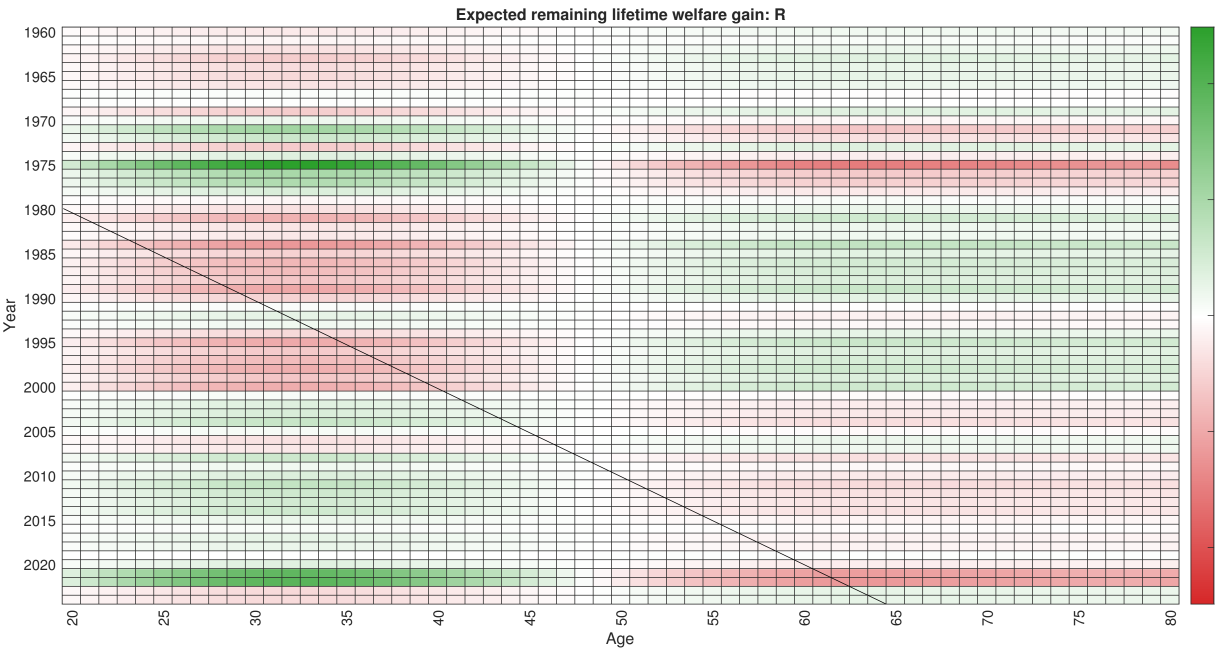

Note: Each rectangle represents the remaining lifetime welfare gain WG for households of different age, as a result of newly arrived monetary shocks in a given year. The gains are expressed in percent of each cohort's steady state consumption. The black line shows the lifeline of a cohort born in 1960.

How important are the first order welfare effects of business cycle fluctuations? Our framework allows us to assess the overall "luck" that a given cohort has experienced throughout its life so far. We are naturally limited by the end of the sample, so we assume no new shocks occurring after 2024 and simply track the paths of variables as they return to the steady state. We subject our cohorts to a "Nostradamus" exercise, in which we provide them with the actual (or forecasted if beyond 2024) stream of period utility flows experienced by surviving cohort members, and ask them to evaluate their welfare from the perspective of a 20-year old agent (which is the age at which agents enter the model). Formally, this boils down to evaluating formula (13) after dropping the expectations operator $\mathbb{E}_t$ and using the realizations (or forecasts) for all variables, so that the only residual uncertainty is the mortality risk.

The results in Figure 10 suggest that cohorts born between 1940 and 1968 drew a less lucky lot in life compared to their younger peers. Clearly, the effects of various shocks evidently do not net out over the life cycle, and some cohorts effectively enjoy or lose up to one extra year of consumption throughout their lives. Our approach also allows us to decompose the "Nostradamus" welfare measure into the contribution of shocks. As can be anticipated, the two most redistributive shocks – monetary and safety – play a prominent role again. In particular, one can observe the effects of unlucky realizations of monetary policy shocks for the cohort born in 1960, which we discussed above. The decomposition reveals that this cohort lost approximately 50% of annual consumption because it had to experience this particular sequence of shocks.

Figure 10: "Nostradamus" welfare at age 20

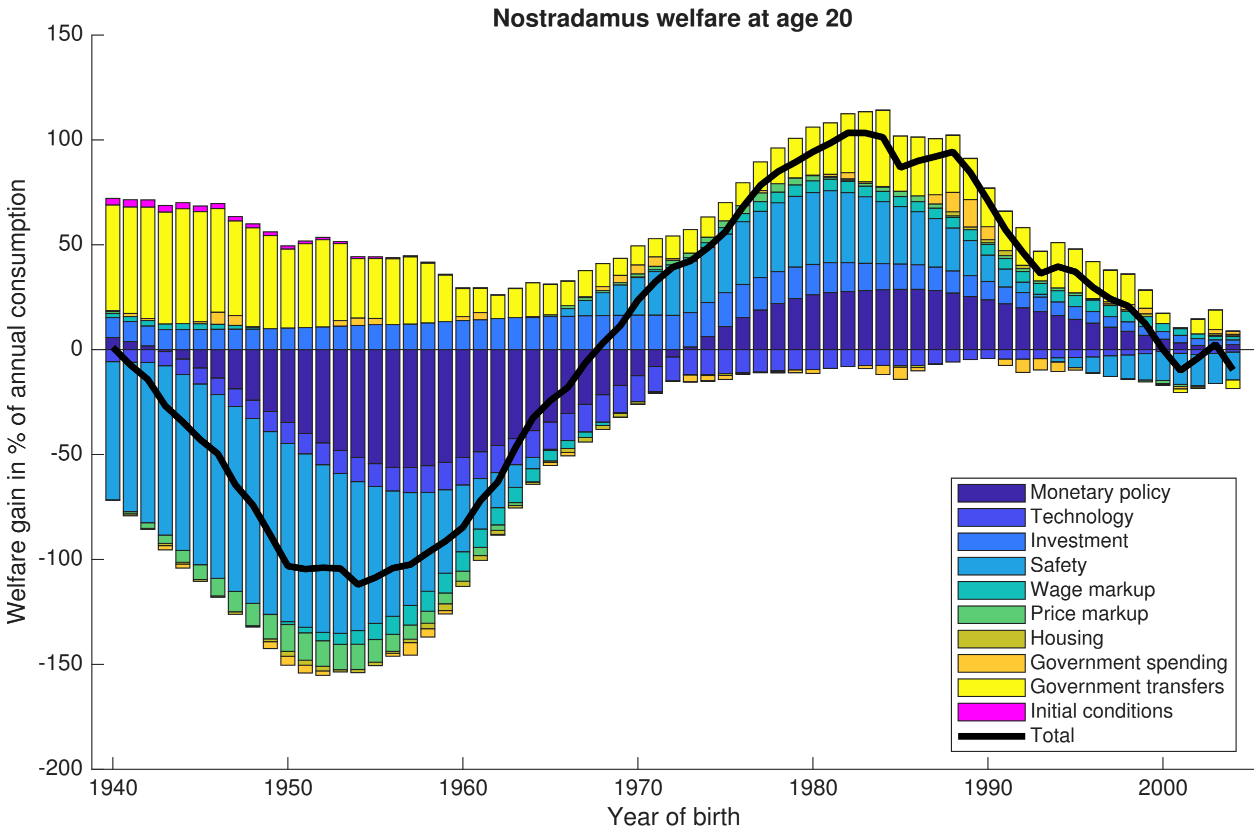

Note: The black line evaluates the remaining lifetime utility U for 20-year old agents born in different years, setting all its arguments equal to their realized (forecasted beyond 2024) values. The bars show the contributions of shocks. The gains are expressed in percent of lifetime steady state consumption.

In a quantitative sense, these first-order welfare effects stand out as high compared to the second-order effects typically found in the literature. Lucas (2003) famously reports losses associated with aggregate volatility worth approximately 0.05% of lifetime consumption. As we discussed in Section 1, other papers have come up with higher numbers by moving to richer models and taking into account precautionary reactions of risk-averse agents. In order to allow for a meaningful comparison with these results, we need to turn to second-order simulations and calculate how much of lifetime consumption a 20-year old individual would be willing to forego in order to avoid uncertainty. More precisely, we take the perspective of an individual who is aware of entering a stochastic world, but only knows the stationary distribution of aggregate variables and prices, as well as their laws of motion conditional on past states and current shocks. The second-order approximation of our model implies that such an individual would be willing to pay about 1.27% of lifetime consumption to live instead in a world without aggregate uncertainty. If we consider an otherwise identical model except that we replace the block of overlapping generations with an infinitely-lived household, the cost shrinks to 0.87%. Hence, life-cycle aspects amplify the second-order costs of business cycle fluctuations by nearly half. As discussed above, this finding reflects increased vulnerability to shocks with strong redistributive power—particularly safety shocks—whose effects do not cancel out over the life cycle, even when their mean is zero.

Figure 11 presents how the perceived second-order cost of aggregate uncertainty depends on household age. More precisely, we measure the willingness to pay of a household of age j, living in the ergodic state of a stochastic economy, for the permanent elimination of all current and future shocks.[8] Over the considered age range, households in their mid-20s are the most exposed to business-cycle fluctuations. These agents have not yet accumulated substantial net worth, and a large share of it consists of nominal liabilities, which we have shown to generate vulnerability. They also have limited ability to offset adverse shocks by working more hours as their labor productivity is still relatively low. After age 25, exposure to business-cycle risk decreases as households accumulate financial buffers, before rising again at very old ages (not shown) once assets have been largely depleted. It needs to be stressed that the implied variation in the exposure to business cycle fluctuations is sizable, with the most vulnerable cohorts willing to pay twice as much as those who are affected least.

[8]Note that this measure is conceptually different from assuming that a household (and all aggregate quantities and prices) could immediately move to the non-stochastic steady state (like implicitly assumed in the previous paragraph, where the cost for a 20-year old is 1.27%), but the qualitative implications of this alternative assumption would be similar to what we present.

Figure 11: Second-order welfare costs of business cycles by household age

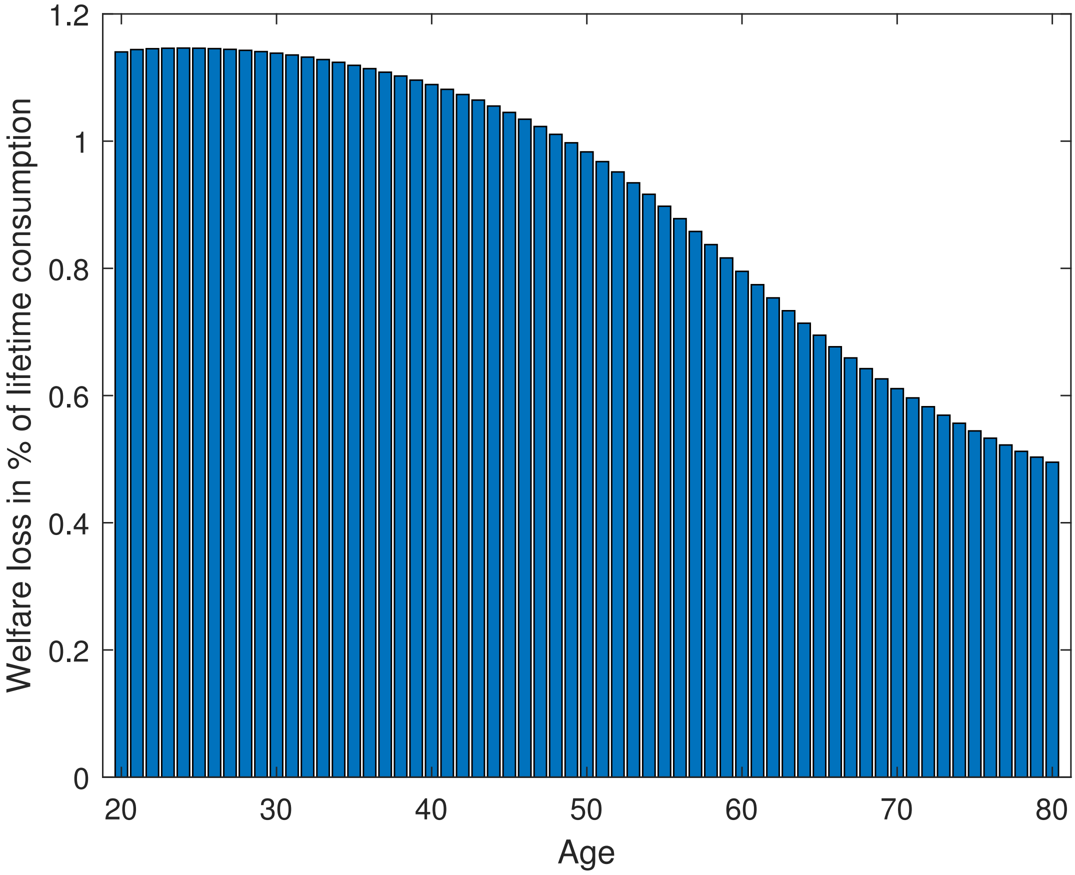

Note: The figure presents the gain of shutting down all current and future shocks as perceived by agents of different age living in a stochastic world. The gains are expressed in percent of lifetime steady state consumption.

5 The redistributive consequences of systematic monetary policy

All the results presented so far assumed that monetary policy follows a Taylor rule with parameters equal to the mean of the estimated posterior distribution. We have shown that monetary shocks, defined as deviations from this rule, can be among the important forces redistributing welfare across age cohorts. We now concentrate on the systematic component of monetary policy, i.e., its regular reaction to economic developments. Does this aspect of policy significantly affect the way business cycles redistribute welfare across different age cohorts?

There are many dimensions along which the effectiveness of monetary policy can be measured in this context. For consistency with the rest of the paper, we focus on the average amount of redistribution $AMOR$ as defined by equation (14). In our description, we are purely positive and abstract away from normative considerations. More specifically, we define a grid for monetary policy parameters γ_π and γ_y and, for each point on the grid, we simulate the model and calculate the average amount of redistributed welfare. This is implemented in two variants. The first takes the form of a counterfactual simulation, using the historical structural shocks obtained during the estimation procedure. The second has an ergodic flavor, relying on the estimated properties of the stochastic process. In both cases, we drop monetary policy shocks as their interpretation becomes dubious once the parameters of the policy rule are being changed, but their inclusion does not affect our conclusions significantly.

Our findings for the historical case are presented on Figure 12. Redistribution is highest when monetary policy responds strongly to inflation and output. For the highest points on the grid $\gamma_\pi = \gamma_y = 3$, $AMOR$ exceeds 4%, i.e., average redistribution equals 4% of annual steady state consumption. It needs to be stressed that this is clearly not an internal maximum and more hawkish responses would generate even more redistribution. More interestingly, we find an internal minimum of $AMOR$, which happens to be 2.5% for $\gamma_\pi = 1.1$ and $\gamma_y = 0$. It should be noted that these values imply a much weaker reaction of the policy rate to inflation and output than we estimated on the US data ($\gamma_\pi = 1.64$ and $\gamma_y = 0.21$, see Table 2), they are also far lower than the standard parameterization of the Taylor rule. The results for the ergodic variant are similar, so we do not report them here in detail.

Figure 12: Average amount of redistribution for alternative monetary policy rule parameters (historical shocks)

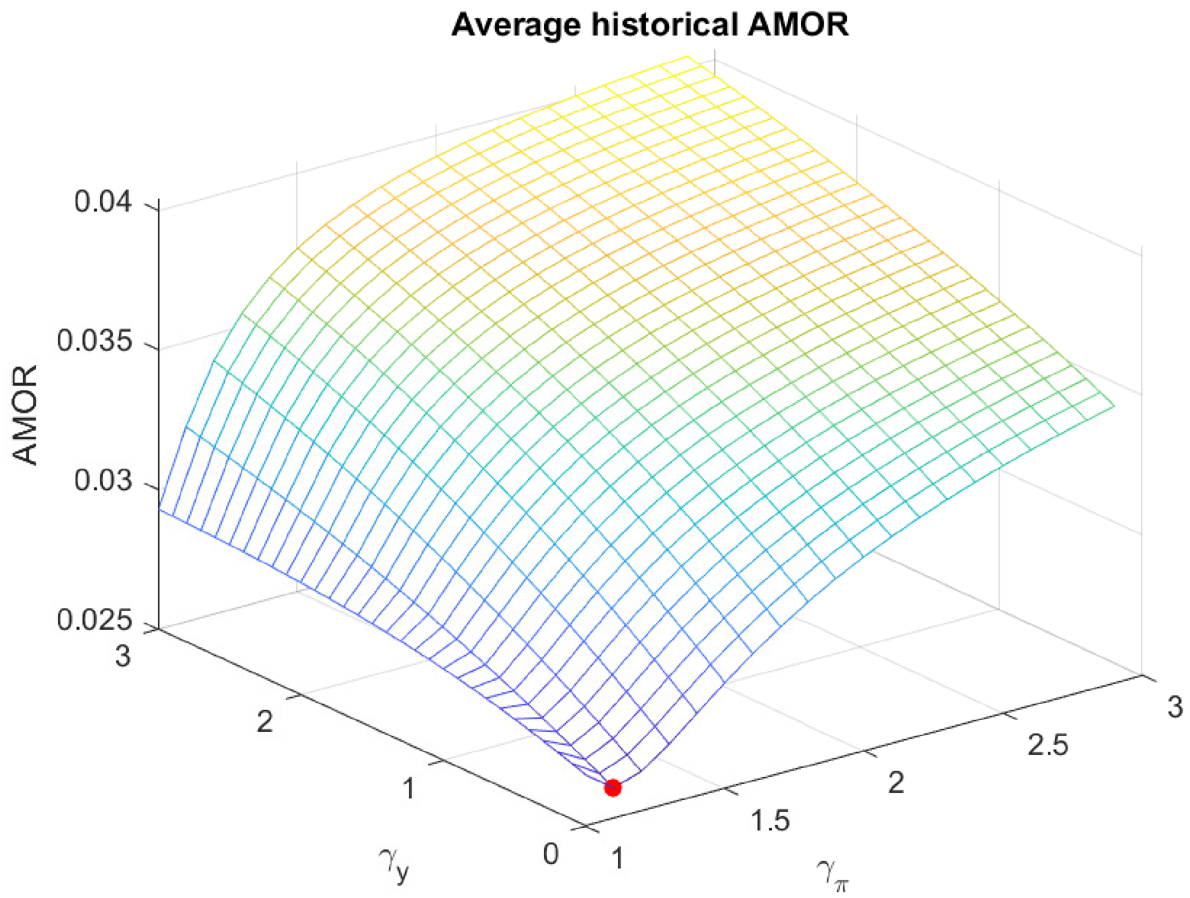

Note: The red dot denotes the minimum amount of redistribution. Monetary policy shocks have been excluded from the simulations.

Why is the amount of intergenerational redistribution minimized for relatively inactive monetary policy?[9] To understand this, it is instructive to concentrate on our main driver of redistribution, i.e., the safety shock. We first repeat the exercise presented above, but with all shocks turned off except this one. The results presented in Figure 13(a) are strikingly similar to those obtained for all shocks, confirming that the flight to safety shock is crucial to understand the mechanism. Figure 13(b) shows the welfare redistribution triggered by a one-off innovation to this shock under the assumption of inactive monetary policy parameters $\gamma_\pi = 1.001$ and $\gamma_y = 0$. If we compare it with the baseline case that uses the estimated

[9] We use the term "inactive" to denote the least responsive monetary policy which still fulfills the standard conditions for determinacy when fiscal policy is passive in the sense of Leeper (1991), as we assume throughout.

parameter values (Figure 3b), the overall redistribution is clearly smaller, and the main reason is that the nominal asset channel becomes much less pronounced.

What is so special about the safety shock? The key insight is that leaning against it drives the policy rate and inflation in opposite directions, which increases the volatility of the real ex-post return on nominal assets. This is demonstrated in Figure 13(c), which shows the reaction of this return to a contractionary safety shock for three variants of monetary policy: baseline ($\gamma_\pi = 1.64$, $\gamma_y = 0.21$), inactive ($\gamma_\pi = 1.001$, $\gamma_y = 0$), and hyperactive ($\gamma_\pi = \gamma_y = 3$). By strongly countering declining economic activity and inflation, the hyperactive monetary policy sharply lowers the ex post return on nominal assets. This boosts the welfare of indebted young cohorts and hurts older owners of nominal assets, thus increasing redistribution. By preventing a deeper recession, as illustrated in Figure 13(d), more active policy additionally benefits working-age cohorts, thereby further amplifying the relative gains of the young. These channels lead to increased redistribution between old and young. In contrast, under inactive policy positive returns on nominal assets on impact are balanced with losses in the following periods, which reduces the role of nominal asset returns in redistributing from old to young agents. Additionally, by allowing for a deeper decline in economic activity, inactive policy further lowers labor income of the young. The overall effect is, hence less redistribution from older to younger generations.

More generally, the safety shock is an instructive example of when policy that provides more macroeconomic stability can actually exacerbate intergenerational redistribution. By construction, this tradeoff is absent in models with representative agents. Indeed, since the safety shock shows up only in the household Euler equation, the canonical New Keynesian setup prescribes that monetary policy should adjust the policy rate to fully offset the impact of this shock on the output gap, thus also perfectly stabilizing inflation (divine coincidence). However, we have shown that such a policy can be highly redistributive in a model with realistic life cycle features.

Figure 13: Redistributive impact of safety shocks for alternative monetary policy

(a) Average historical amount of redistribution

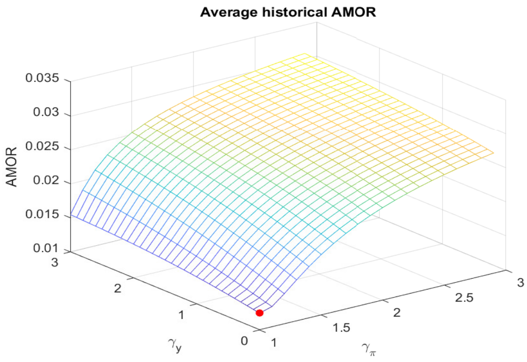

(b) Welfare redistribution under inactive policy

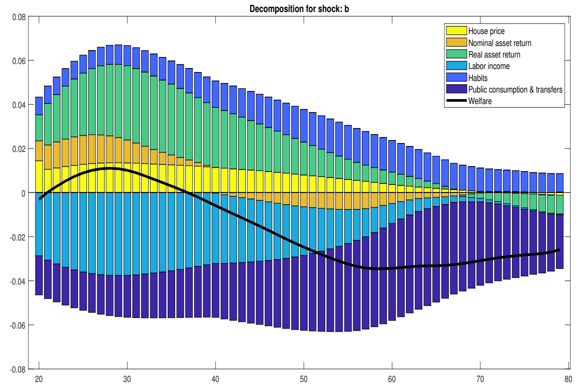

(c) Response of real ex-post return on nominal assets

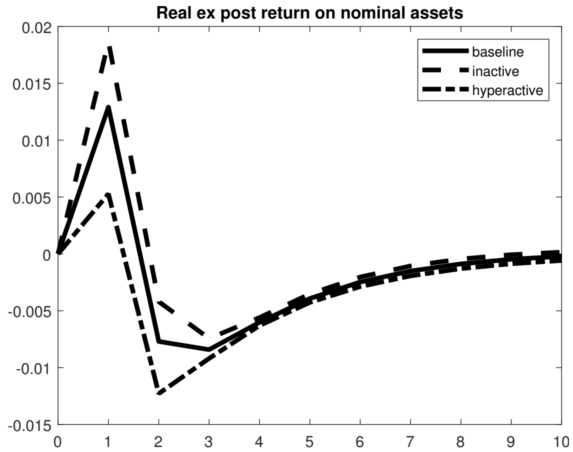

(d) Response of output

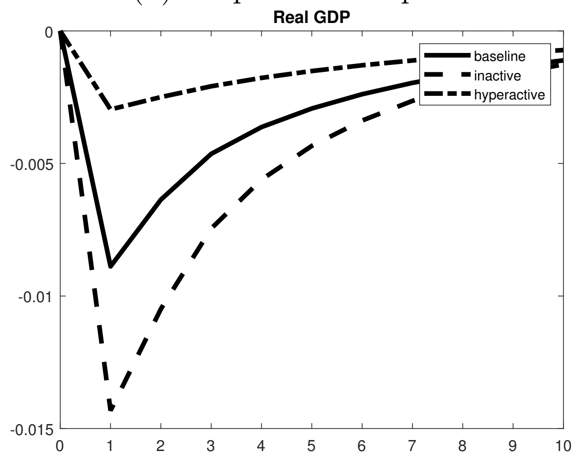

Note: The monetary policy variants are defined as follows: baseline ($\gamma_\pi = 1.64$, $\gamma_y = 0.21$), inactive ($\gamma_\pi = 1.001$, $\gamma_y = 0$), and hyperactive ($\gamma_\pi = \gamma_y = 3$). $AMOR$ calculation excludes monetary policy shocks.

Last but not least, we check how the alternative parameterizations of the monetary policy rule would have affected lifetime welfare of our agents. Figure 14 shows the Nostradamus welfare under the three policy variants that we have considered above. For reasons explained above and to preserve consistency throughout this section, we exclude monetary policy shocks from the calculations. The three lines clearly differ and the general finding is that less active policy would have generated less redistribution.

Figure 14: "Nostradamus" welfare at age 20 for alternative monetary policy parameterizations

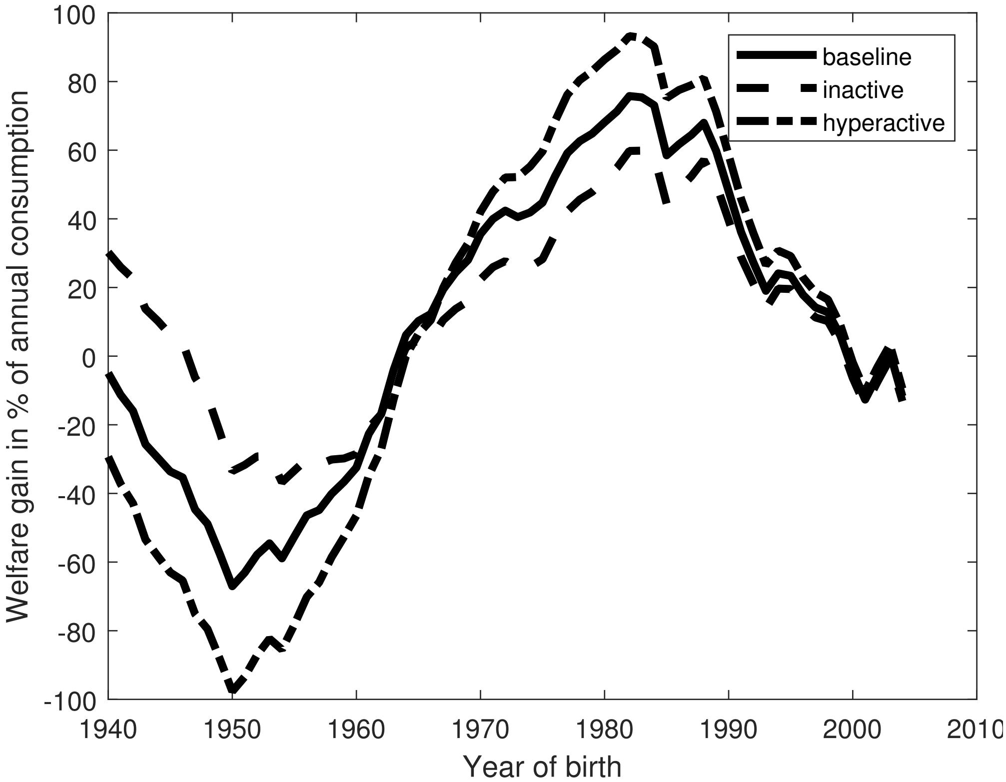

Notes: The monetary policy variants are defined as follows: baseline ($\gamma_\pi = 1.64$, $\gamma_y = 0.21$), inactive ($\gamma_\pi = 1.001$, $\gamma_y = 0$), and hyperactive ($\gamma_\pi = \gamma_y = 3$). The calculations exclude monetary policy shocks.

Moving beyond safety shocks, we also checked $AMOR$-minimizing parameters of the monetary policy rule for other shocks. Not surprisingly, the outcomes are shock-specific. For instance, the policy which minimizes redistribution conditional on price markup shocks calls for a similar response of the interest rate to inflation and output. Given that these shocks move output and inflation in opposite directions, such parameter combinations again guarantee relatively small variability of bond returns. Overall, while the Taylor rule parameters that minimize the amount of intergenerational redistribution are shock-specific, the general finding is that it is lower if monetary policy better stabilizes the ex post return on nominal assets.

All in all, our considerations about systematic monetary policy align with our earlier findings about the key drivers of redistribution across different age cohorts. First, the impact of systematic monetary policy on redistribution is largely shaped by how it affects the transmission of safety shocks. Second, the nominal asset channel is crucial in this context. In consequence policies that limit the volatility of returns on nominal assets also tend to minimize the amount of redistribution.

6 Conclusions

What are the distributional consequences of business cycle fluctuations across generations? Which shocks drive the redistribution? What economic channels are involved? How did cyclical fluctuations redistribute welfare in the historical context? Did their impact cancel out over human lifetimes or does it matter when an agent is born? Can monetary policy affect the amount of redistribution? We answer all these questions in the context of US business cycle fluctuations since the 1960s. To this end, we constructed a life-cycle model with nominal and real rigidities, estimated it on US data, and used it to run several counterfactual experiments.

What do we find? First and foremost, that cyclical fluctuations redistribute a lot between agents of different age and that their effects do not cancel out over typical lifetimes. This is not only a trivial consequence of a finite horizon (so that the mean of realized shocks need not be zero), but also the effect of a feature characteristic to life-cycle models, which is absent in representative agent or standard HANK models: the same shock may have very different consequences for agents of different age. In our data sample, some cohorts gained or lost an equivalent of one annual consumption due to shocks they drew during their life. In order to avoid such a risk, households would be willing to forego over 1% of their lifetime consumption, a number that is by half larger than implied by an otherwise identical representative agent model.

Second, we show that different macroeconomic shocks have different redistributive power. For instance, TFP shocks move all cohorts in the same direction, whereas safety and monetary policy shocks benefit some cohorts while leaving others worse off. As a consequence, the latter two shocks together explain over 50% of intergenerational redistribution over business cycles. The large contribution of monetary policy shocks is an interesting finding as their role in driving economic fluctuations is usually considered small. An important consequence of the differences in redistributive power between various shocks is that the drivers of a particular business cycle swing matter. Some recessions (expansions) affected all cohorts similarly by lowering (raising) welfare for everybody, but most of them heavily redistributed between households of different age.

Last but not least, we analyzed the role of systematic monetary policy. We showed that the amount of redistribution can be significantly affected by central bank policy. The general finding is that monetary policy dampens redistribution if it stabilizes returns on nominal assets. Given the main driving forces of redistribution in our sample, this translates into low parameters of the Taylor rule. We want to stress that this finding should be treated as entirely positive and not necessarily guiding the optimal policy behavior, a study of which we leave for future research.

References

Acharya, Sushant, Edouard Challe, and Keshav Dogra (2023) 'Optimal monetary policy according to hank.' *American Economic Review* 113(7), 1741–1782

Adam, Klaus, and Junyi Zhu (2016) 'Price-Level Changes and the Redistribution of Nominal Wealth Across the Euro Area.' *Journal of the European Economic Association* 14, 871–906

Albanesi, Stefania (2007) 'Inflation and inequality.' *Journal of Monetary Economics* 54(4), 1088–1114

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

Angeletos, George-Marios, Chen Lian, and Christian K. Wolf (2023) 'Can Deficits Finance Themselves?' NBER Working Papers 31185, National Bureau of Economic Research, Inc, April

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas (2020) 'Business-cycle anatomy.' *American Economic Review* 110(10), 3030–3070

Auclert, Adrien (2019) 'Monetary Policy and the Redistribution Channel.' *American Economic Review* 109(6), 2333–2367

Barro, Robert J. (2009) 'Rare disasters, asset prices, and welfare costs.' *American Economic Review* 99(1), 243–264

Bayer, Christian, Benjamin Born, and Ralph Luetticke (2024) 'Shocks, Frictions, and Inequality in US Business Cycles.' *American Economic Review* 114(5), 1211–1247

Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J. Sargent (2021) 'Inequality, business cycles, and monetary-fiscal policy.' *Econometrica* 89(6), 2559–2599

Bhattarai, Saroj, Jae Won Lee, and Choongryul Yang (2023) 'Redistribution and the monetary fiscal policy mix.' *Quantitative Economics* 14(3), 817–853

Bielecki, Marcin, Michał Brzoza-Brzezina, and Marcin Kolasa (2022) 'Intergenerational Redistributive Effects of Monetary Policy .' *Journal of the European Economic Association* 20(2), 549–580

Brzoza-Brzezina, Michał, Marcin Kolasa, Krzysztof Makarski, and Julia Jabłońska (2024) 'For whom the bill tolls: redistributive consequences of a monetary-fiscal stimulus.' Working Paper Series 2998, European Central Bank, Nov

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvial (2017) 'Innocent bystanders? monetary policy and inequality.' *Journal of Monetary Economics* 88, 70–89

Del Negro, Marco, Marc Giannoni, and Christina Patterson (2012) 'The forward guidance puzzle.' Staff Reports 574, Federal Reserve Bank of New York

Doepke, Matthias, and Martin Schneider (2006) 'Inflation and the redistribution of nominal wealth.' *Journal of Political Economy* 114, 1069–1097

Dossche, Maarten, Jiří Slačálek, and Guido Wolswijk (2021) 'Monetary policy and inequality.' *ECB Economic Bulletin, Issue 2/2021*

Eggertsson, Gauti B., Neil R. Mehrotra, and Jacob A. Robbins (2019) 'A model of secular stagnation: Theory and quantitative evaluation.' *American Economic Journal: Macroeconomics* 11(1), 1–48

Guo, Xing, Pablo Ottonello, and Diego J. Perez (2023) 'Monetary policy and redistribution in open economies.' *Journal of Political Economy Macroeconomics* 1(1), 191–241

Jorda, Oscar, Moritz Schularick, and Alan M. Taylor (2024) 'Disasters everywhere: The costs of business cycles reconsidered.' *IMF Economic Review* 72(1), 116–151

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante (2018) 'Monetary Policy According to HANK.' *American Economic Review* 108(3), 697–743

Kekre, Rohan, and Moritz Lenel (2024) 'The flight to safety and international risk sharing.' *American Economic Review* 114(6), 1650–1691

Krebs, Tom (2007) 'Job displacement risk and the cost of business cycles.' *American Economic Review* 97(3), 664–686

Krusell, Per, and Anthony A. Smith (1999) 'On the welfare effects of eliminating business cycles.' *Review of Economic Dynamics* 2(1), 245–272

Krusell, Per, Toshihiko Mukoyama, Aysegul Sahin, and Anthony A. Smith (2009) 'Revisiting the welfare effects of eliminating business cycles.' *Review of Economic Dynamics* 12(3), 393–402

Kuhn, Moritz, Moritz Schularick, and Ulrike I. Steins (2020) 'Income and wealth inequality in america, 1949-2016.' *Journal of Political Economy* 128(9), 3469–3519

Leeper, Eric M. (1991) 'Equilibria under 'active' and 'passive' monetary and fiscal policies.' *Journal of Monetary Economics* 27(1), 129–147

Lenza, Michele, and Jiří Slačálek (2021) 'How does monetary policy affect income and wealth inequality? Evidence from quantitative easing in the euro area .' *ECB Working Paper Series, No 2190/2018*

Lucas, Robert E. (1987) *Models of Business Cycles* (Oxford: Basil Blackwell)

Lucas, Robert E. (2003) 'Macroeconomic Priorities.' *American Economic Review* 93(1), 1–14

Obstfeld, Maurice (1994) 'Evaluating risky consumption paths: The role of intertemporal substitutability.' *European Economic Review* 38(7), 1471–1486

Otrok, Christopher (2001) 'On measuring the welfare cost of business cycles.' *Journal of Monetary Economics* 47(1), 61–92

Pallotti, Filippo, Gonzalo Paz-Pardo, Jiří Slačálek, Oreste Tristani, and Giovanni L. Violante (2023) 'Who bears the costs of inflation? Euro area households and the 2021-2022 shock.' *ECB Statistics Paper Series, No 18/2016*

Reis, Ricardo (2009) 'The time-series properties of aggregate consumption: Implications for the costs of fluctuations.' *Journal of the European Economic Association* 7(4), 722–753

Smets, Frank, and Rafael Wouters (2003) 'An estimated dynamic stochastic general equilibrium model of the euro area.' *Journal of the European Economic Association* 1(5), 1123–1175

Smets, Frank, and Rafael Wouters (2007) 'Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach.' *American Economic Review* 97(3), 586–606

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron (2001) 'The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk.' *European Economic Review* 45(7), 1311–1339

Tervala, Juha (2021) 'Hysteresis and the welfare costs of recessions.' *Economic Modelling* 95(C), 136–144

Walentin, Karl, and Andreas Westermark (2022) 'Learning on the job and the cost of business cycles.' *American Economic Journal: Macroeconomics* 14(4), 341–377

Appendix A: Model equilibrium conditions

Below we present the list of equations making up the model.

Demography

Population dynamics for $j = 1, ..., J$

$$N_{1,t+1} = (1+n)N_{1,t} \tag{A.1}$$

$$N_{j+1,t+1} = (1-\omega_j)N_{j,t} \tag{A.2}$$

Total population

$$N_t = \sum_{j=1}^{J} N_{j,t} \tag{A.3}$$

Households

Budget constraint for $j = 1, ..., J$

$$\begin{aligned} c_{j,t} + p_{h,t}\left[h_{j,t} - (1-\delta_h)\, h_{j-1,t-1}\right] + p_{k,t}\left[k_{j,t} - (1-\delta_k)\, k_{j-1,t-1}\right] + b_{j,t} \\ = (1-\tau) z_j w_t \ell_{j,t} + r_{k,t} k_{j-1,t-1} + \frac{f_t}{k_{t-1}} k_{j-1,t-1} + \frac{R_{t-1}}{\pi_t} b_{j-1,t-1} + t_{j,t} + beq_t \end{aligned} \tag{A.4}$$

Marginal utility of consumption for $j = 1, ..., J$

$$u^c_{j,t} = \frac{1-\varrho}{c_{j,t} - \varrho c_{j,t-1}} \tag{A.5}$$

Initial asset holdings

$$h_{0,t} = b_{0,t} = k_{0,t} = 0 \tag{A.6}$$

Optimal choice of assets for $j = 1, ..., J-1$

$$\frac{\psi_j(1+\varepsilon_t^h)}{h_{j,t}} = u^c_{j,t} p_{h,t} - \beta(1-\omega_j)(1-\delta_h)\mathbb{E}_t\left\{u^c_{j+1,t+1} p_{h,t+1}\right\} \tag{A.7}$$

$$\frac{\zeta_j(1+\varepsilon_t^b)}{1+b_{j+1,t+1}} = u^c_{j,t} + \beta(1-\omega_j)\mathbb{E}_t\left\{u^c_{j+1,t+1}\frac{R_t}{\pi_{t+1}}\right\} \tag{A.8}$$

$$p_{k,t}u^c_{j,t} = \beta(1-\omega_j)\mathbb{E}_t\left\{u^c_{j+1,t+1}\left(r_{k,t+1} + (1-\delta_k)p_{k,t+1} + \frac{f_{t+1}}{k_{t+1}}\right)\right\} \tag{A.9}$$

Optimal choice of assets for $j = J$

$$\frac{\psi_J(1+\varepsilon_t^h)}{h_{J,t}} = u^c_{J,t} p_{h,t} \tag{A.10}$$

$$\frac{\zeta_J(1+\varepsilon_t^b)}{1+b_{J,t}} = u^c_{J,t} \tag{A.11}$$

$$k_{J,t} = 0 \tag{A.12}$$

Transfers for $j = 1, ..., J-1$

$$\frac{t_{j,t}}{t_j} = \frac{t_t}{t} \tag{A.13}$$

Labor supply for $j = 1, ..., J-1$

$$\ell_{j,t} = \left(\frac{(1-\tau_t)z_j w_t}{\phi_j \mu_w} u^c_{j,t}\right)^{\frac{1}{\varphi}} \Delta_{\ell,t} \tag{A.14}$$

Aggregate wage

$$w_t = \left[\theta_w\left(w_{t-1}\frac{\pi}{\pi_t}\right)^{\frac{1}{1-\mu_w}} + (1-\theta_w)\left(\tilde{w}_t\right)^{\frac{1}{1-\mu_w}}\right]^{1-\mu_w} \tag{A.15}$$

Wage setting

$$\tilde{w}_t = \mu_w \frac{\Omega_{w,t}}{\Upsilon_{w,t}} \tag{A.16}$$

$$\Omega_{w,t} = \tilde{\phi} \left(\frac{w_t}{\tilde{w}_t}\right)^{\frac{\mu_w}{\mu_w - 1}(1+\varphi)} \ell_t^{1+\varphi} + \beta(1+n)\theta_w \mathbb{E}_t \left[\left(\frac{\pi}{\pi_{t+1}}\right)^{\frac{\mu_w}{1-\mu_w}(1+\varphi)} \Omega_{w,t+1} \left(\frac{\tilde{w}_{t+1}}{\tilde{w}_t}\right)^{\frac{\mu_w}{\mu_w - 1}(1+\varphi)} \right] \tag{A.17}$$

$$\Upsilon_{w,t} = (1+\varepsilon_t^w) u_t^c \left(\frac{w_t}{\tilde{w}_t}\right)^{\frac{\mu_w}{\mu_w - 1}} \ell_t + \beta(1+n)\theta_w \mathbb{E}_t \left[\left(\frac{\pi}{\pi_{t+1}}\right)^{\frac{1}{1-\mu_w}} \Upsilon_{w,t+1} \left(\frac{\tilde{w}_{t+1}}{\tilde{w}_t}\right)^{\frac{\mu_w}{\mu_w - 1}} \right] \tag{A.18}$$

where $\tilde{\phi} = \frac{w}{\mu_w \ell^{\varphi}} \sum_{j=1}^{JR-1} u_j^c \frac{N_{j+1,t+1}}{\sum_{i=1}^{JR-1} N_{i,t}}$, $u_t^c = \sum_{j=1}^{JR-1} u_{j,t}^c \frac{N_{j+1,t+1}}{\sum_{i=1}^{JR-1} N_{i,t}}$, and $JR = 45$ (notional retirement age)

Firms

Factor demand

$$w_t = (1-\alpha) mc_t \exp(\varepsilon_t^z) k_{t-1}^{\alpha} \ell_t^{-\alpha} \tag{A.19}$$

$$r_t^k = \alpha mc_t (1+\varepsilon_t^z) k_{t-1}^{\alpha-1} \ell_t^{1-\alpha} \tag{A.20}$$

Profits

$$f_t = y_t - w_t \ell_t - r_t^k k_{t-1} \tag{A.21}$$

Aggregate price

$$1 = \theta_p \left(\frac{\pi}{\pi_t}\right)^{\frac{1}{1-\mu_p}} + (1-\theta_p) \left(\tilde{P}_t\right)^{\frac{1}{1-\mu_p}} \tag{A.22}$$

Price setting

$$\tilde{P}_t = \mu \frac{\Omega_t}{\Upsilon_t} \tag{A.23}$$

$$\Omega_t = (1+\varepsilon_t^p) mc_t y_t + \theta \mathbb{E}_t \left[\left(\frac{\pi_{t+1}}{R_t} \right) \left(\frac{\pi}{\pi_{t+1}} \right)^{\frac{\mu}{1-\mu}} \Omega_{t+1} \right] \tag{A.24}$$

$$\Upsilon_t = y_t + \theta \mathbb{E}_t \left[\left(\frac{\pi_{t+1}}{R_t} \right) \left(\frac{\pi}{\pi_{t+1}} \right)^{\frac{1}{1-\mu}} \Upsilon_{t+1} \right] \tag{A.25}$$

Capital production

$$(1+n)k_t = (1-\delta_k)k_{t-1} + (1+\varepsilon_t^i) \left[1 - \frac{S_k}{2} \left(\frac{i_t}{i_{t-1}} - 1 \right)^2 \right] i_t \tag{A.26}$$

$$\begin{aligned} 1 = {} & p_{k,t}(1+\varepsilon_t^i) \left[1 - \frac{S_k}{2} \left(\frac{i_t}{i_{t-1}} - 1 \right)^2 - S_k \left(\frac{i_t}{i_{t-1}} - 1 \right) \frac{i_t}{i_{t-1}} \right] \\ & + E_t \left[p_{k,t+1}(1+\varepsilon_{t+1}^i) \frac{\pi_{t+1}}{R_t} S_k \left(\frac{i_{t+1}}{i_t} - 1 \right) \left(\frac{i_{t+1}}{i_t} \right)^2 \right] \end{aligned} \tag{A.27}$$

Government

Government budget constraint

$$(1+n)b_{g,t} = \frac{R_{t-1}}{\pi_t} b_{g,t-1} + g_t + t_t - \tau w_t \ell_t \tag{A.28}$$

Fiscal policy

$$\frac{t_t}{t} = \left(\frac{t_{t-1}}{t} \right)^{\gamma_t} \left(\frac{b_{g,t}}{b_g} \right)^{-\eta} (1+\varepsilon_t^t) \tag{A.29}$$

$$g_t = g \exp(\varepsilon_t^g) \tag{A.30}$$

Monetary policy

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R} \right)^{\gamma_R} \left[\left(\frac{\pi_t}{\pi} \right)^{\gamma_\pi} \left(\frac{y_t}{y} \right)^{\gamma_y} \right]^{1-\gamma_R} (1+\varepsilon_t^r) \tag{A.31}$$

Aggregation and market clearing

Aggregate allocations

$$c_t = \sum_{j=1}^{J} \frac{N_{j,t} c_{j,t}}{N_t} \tag{A.32}$$

$$\ell_t = \sum_{j=1}^{J} \frac{N_{j,t} z_j \ell_{j,t}}{N_t} \tag{A.33}$$

$$t_t = \sum_{j=1}^{J} \frac{N_{j,t} t_{j,t}}{N_t} \tag{A.34}$$

$$h = \sum_{j=1}^{J} \frac{N_{j,t} h_{j,t}}{N_t(1+n)} \tag{A.35}$$

$$k_t = \sum_{j=1}^{J} \frac{N_{j,t} k_{j,t}}{N_t(1+n)} \tag{A.36}$$

$$b_t = \sum_{j=1}^{J} \frac{N_{j,t} b_{j,t}}{N_t(1+n)} \tag{A.37}$$

Bequests

$$beq_t = \sum_{j=1}^{J} \frac{[N_{j-1,t-1} - N_{j,t}(1+n)]}{N_t(1+n)} \Bigg[(1-\delta_h) p_{h,t} h_{j-1,t-1} + \frac{R_{t-1}}{\pi_t} b_{j-1,t-1} + \left(r_{k,t} + (1-\delta_k) p_{k,t} + \frac{f_t}{k_{t-1}} \right) k_{j-1,t-1} \Bigg] \tag{A.38}$$

Aggregate production and resource constraint

$$y_t \Delta_t = (1+\varepsilon_t^z) k_{t-1}^{\alpha} h_t^{1-\alpha} - \Phi \tag{A.39}$$

$$y_t = c_t + i_t + \delta_\chi p_{\chi,t} \chi + g_t \tag{A.40}$$

Price and wage dispersion

$$\Delta_{p,t} = (1-\theta_p) \tilde{P}_t^{\frac{\mu_p}{1-\mu_p}} + \theta_p \left(\frac{\pi}{\pi_t} \right)^{\frac{\mu_p}{1-\mu_p}} \Delta_{p,t-1} \tag{A.41}$$

$$\Delta_{w,t} = (1-\theta_w) \left(\frac{\tilde{w}_t}{w_t} \right)^{\frac{\mu_w}{1-\mu_w}(1+\varphi)} + \theta_w \left(\frac{w_{t-1}}{w_t} \frac{\pi}{\pi_t} \right)^{\frac{\mu_w}{1-\mu_w}(1+\varphi)} \Delta_{w,t-1} \tag{A.42}$$

Average cohort welfare

$$U_{j,t} = (1-\varrho)\log(c_{j,t}-\varrho c_{j,t-1}) + \frac{g}{c_j}\log(g_t) - \phi_j \frac{(\ell_{j,t})^{1+\varphi}}{1+\varphi}\Delta_{w,t} \tag{A.43}$$
$$+ \psi_j \log(h_{j,t}) + \zeta_j \log(1+b_{j,t}) + \beta(1-\omega_j)\mathbb{E}_t U_{j,t+1}$$

Exogenous shocks

The model is driven by the following nine zero-mean exogenous shocks: ε_t^h, ε_t^b, ε_t^z, ε_t^i, ε_t^p, ε_t^w, ε_t^g, ε_t^t, ε_t^r. Each of them follows an independent linear AR(1) process, except for the transfer and monetary shocks ε_t^t and ε_t^r, which we assume to be white noise.

Appendix B: Calibration and estimation of the model

B.1 Income and asset profiles by age

Our primary data source for age-dependent income and asset holdings is the Survey of Consumer Finances (SCF), using waves from 1989 through 2016. We rely on the SCF+ dataset compiled by Kuhn et al. (2020). While we initially intended to incorporate earlier SCF waves (prior to 1989), the available information for those years was too incomplete for our purposes.

Another important data source is the Federal Reserve Board's Z.1 Financial Accounts. Specifically, we use Level Tables L.101, L.116, L.117, and L.122 to calculate the average annual shares of real versus nominal financial assets in the balance sheets of, respectively: (i) Households and Nonprofit Organizations, (ii) Life Insurance Companies, (iii) Private and Public Pension Funds, and (iv) Mutual Funds. Table B.1 provides the details of the transformations applied to the raw series.

Based on this data, we construct consolidated household-level holdings of housing, real financial assets, and net nominal financial wealth, in addition to labor income and transfer measures. Table B.2 summarizes the full procedure. The construction of labor income and

Table B.1: Construction of shares of real financial assets based on Z1

Share of real financial assets in Z1 balance sheets of ...	Operation on Z1 variables
Mutual funds [L.122]	= L.122 [LM653064100.Q] / L.122 [LM654090000.Q]
Pension funds [L.117]	= (L.117 [LM593064105.Q] + L.117 [LM593064205.Q] * Mutual funds share) / L.117 [FL594090005.Q]
Life insurance [L.116]	= (L.116 [LM543064105.Q] + L.116 [LM543064205.Q] * Mutual funds share) / L.116 [FL544090005.Q]
Households [L.101]	= (L.101 [LM153064105.Q] + L.101 [LM152090205.Q] + L.101 [LM153064205.Q] * Mutual funds share + L.101 [FL153050005.Q] * Pension funds share + L.101 [FL153040005.Q] * Life insurance share) / L.101 [FL154090005.Q]

transfers is straightforward, as the SCF+ dataset reports both variables directly. Asset variables require additional processing. Housing wealth is defined as the sum of owner-occupied housing and other real estate. Because some financial holdings in SCF+ are not clearly classified, we first compute a residual category of remaining mixed financial assets. We then allocate this category into real and nominal components based on the aggregate household-level shares derived from the Z.1 accounts.

Real financial wealth is defined as the sum of business assets, equity and other managed assets, plus the Z.1-based real shares of assets held via mutual funds, pensions, life insurance, and the mixed category. Nominal financial wealth is measured as the sum of liquid assets, certificates of deposit, and bonds, plus the Z.1-based nominal shares of assets held through mutual funds, pensions, life insurance, and the mixed category, net of nominal liabilities in the form of overall outstanding debt.

We exclude one category reported in SCF+: real non-financial wealth, which consists of durable goods such as vehicles. This component represents only a small share of total household net worth, and it is not obvious how it should be classified within our framework. Nevertheless, we continue to track it in the background, as doing so enables us to confirm that, after all transformations, aggregate household net worth is still correctly recovered.

Once household-level categories of income and asset holdings are constructed, we generate

Table B.2: Mapping between model objects and SCF+ objects

Category in the paper	Description	SCF+ code
Labor income	= income from wages, salaries and self-employment	incws
Transfers	= transfer income	inctrans
Housing stock	= asset value of house	house
	+ other real estate (net position)	oest
	+ other real estate debt	oestdebt
Remaining mixed financial assets	= financial assets (ffaequ, liqcer, bnd, mfun, ofin, life, pen)	ffafin
	- ffaequ - liqcer - bnd - mfun - ofin - life - pen	
Real financial wealth	= business wealth	ffabus
	+ equity and other managed assets	ffaequ
	+ Z.1 Mutual funds share * mutual funds	mfun
	+ Z.1 Pension funds share * pensions	pen
	+ Z.1 Life insurance share * life insurance assets	life
	+ Z.1 Households share * Remaining mixed financial assets	
Nominal financial wealth	= liquid assets and certificates of deposit	liqcer
	+ bonds	bnd
	+ (1 - Z.1 Mutual funds share) * mutual funds	mfun
	+ (1 - Z.1 Pension funds share) * pensions	pen
	+ (1 - Z.1 Life insurance share) * life insurance assets	life
	+ (1 - Z.1 Households share) * Remaining mixed financial assets	
	- total household debt (excluding other real estate debt)	tdebt
	- other real estate debt	oestdebt
Overall net worth	= net wealth	ffanw

life-cycle profiles for the average household. First, using SCF survey weights (wgtI95W95), we compute weighted averages within each wave by the age of the household head.[10] Second, we calculate age-specific averages across all waves combined. Finally, we apply a Hodrick-Prescott filter ($\lambda = 100$) to obtain smooth income and asset profiles by age, as shown in Figure 1. We verify that the HP-filtered average asset categories still sum to the HP-filtered average net worth, with only a negligible approximation error. These smoothed profiles serve as our calibration targets.

[10]To minimize sampling error, we restrict the sample to households with heads aged below 88.

B.2 Macroeconomic time series and estimation

We take the Case-Shiller Real Home Price Index from Robert Shiller's website. The series for general government US debt was taken from the Z.1 statistics. The remaining macroeconomic variables were taken from the St. Louis Fed's FRED database. Aggregate variables, including Gross Domestic Product, Personal Consumption Expenditures, Fixed Private Investment and Government Total Expenditures were divided by population aged 16 and above and divided by the GDP deflator. Real wages were constructed by dividing the Nonfarm Business hourly compensation by the GDP deflator. Hours worked are taken from total hours worked in the Nonfarm Business Sector and multiplied by the fraction of adults that are in the labor force. The inflation rate was calculated on the basis of annual GDP deflator dynamics, while the short-term rate is equal to the effective Federal Funds Rate. An HP filter with $\lambda = 6.25$ was applied to all trending macroeconomic time series, with the exception of US debt which was expressed in percent of GDP and demeaned. The model estimation was performed using standard Bayesian methods in Dynare 6.2.

B.3 Estimated properties of stochastic shocks

Table B.3 presents the estimation results for the processes driving the stochastic shocks in the model.

Table B.3: Estimated shock processes

Parameter	Prior			Posterior		Description
	type	mean	std	mean	std	
ρ_z	beta	0.75^4	0.1	0.28	0.07	TFP shock persistence
ρ_i	beta	0.75^4	0.1	0.28	0.07	Investment-specific shock persistence
ρ_p	beta	0.75^4	0.1	0.33	0.07	Price markup shock persistence
ρ_w	beta	0.75^4	0.1	0.36	0.08	Wage cost-push shock persistence
ρ_g	beta	0.75^4	0.1	0.37	0.07	Public consumption shock persistence
ρ_b	beta	0.75^4	0.1	0.68	0.05	Bond preference shock persistence
ρ_h	beta	0.75^4	0.1	0.45	0.08	Housing demand shock persistence
σ_z	inv. gamma	0.01	∞	0.009	0.001	TFP shock std. dev.
σ_i	inv. gamma	0.01	∞	0.015	0.002	Investment-specific shock std. dev.
σ_p	inv. gamma	0.01	∞	0.012	0.001	Price markup shock std. dev.
σ_w	inv. gamma	0.01	∞	0.117	0.024	Wage cost-push shock std. dev.
σ_g	inv. gamma	0.01	∞	0.027	0.002	Public consumption shock std. dev.
σ_b	inv. gamma	0.01	∞	0.545	0.082	Bond preference shock std. dev.
σ_h	inv. gamma	0.01	∞	0.197	0.033	Housing demand shock std. dev.
σ_t	inv. gamma	0.10	∞	0.156	0.014	Transfer shock std. dev.
σ_r	inv. gamma	0.0025	∞	0.0178	0.002	Interest rate shock std. dev.

Appendix C: Selected additional results

Figure C.1: Impulse response functions to a monetary shock

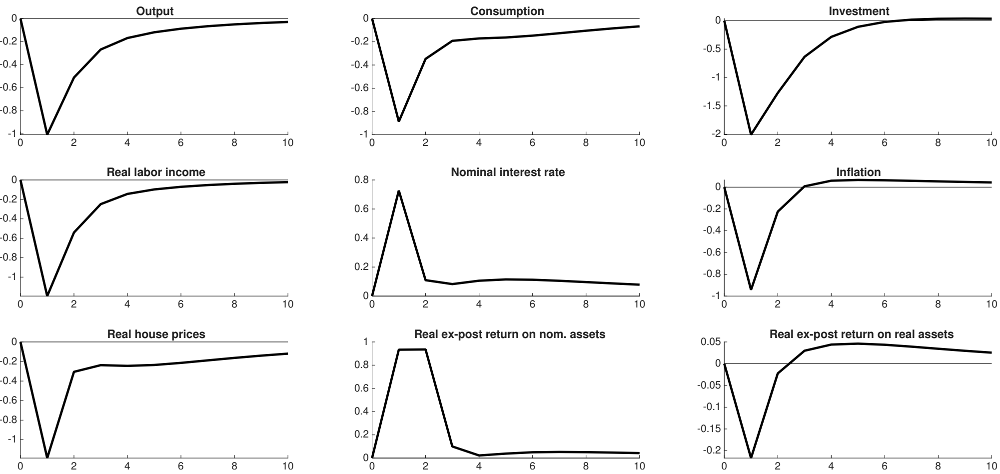

Figure C.2: Impulse response functions to a TFP shock

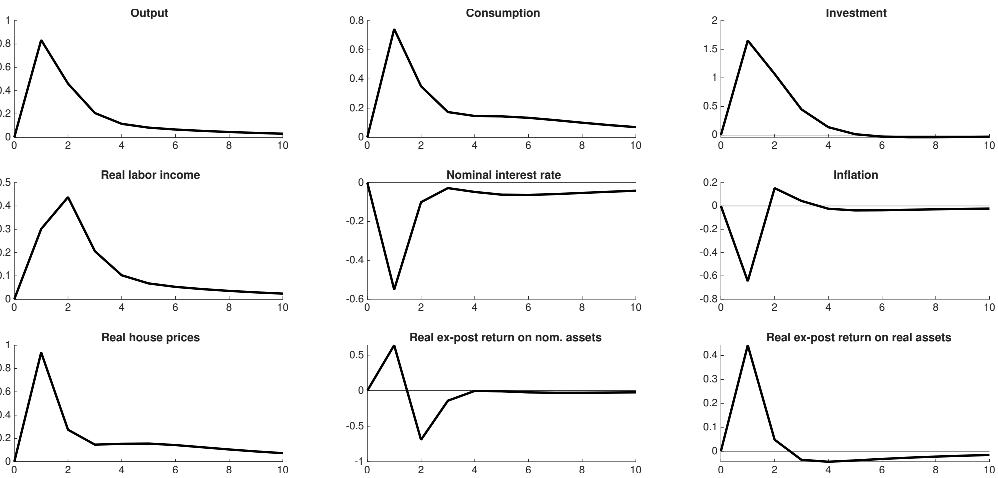

Figure C.3: Impulse response functions to an investment-specific shock

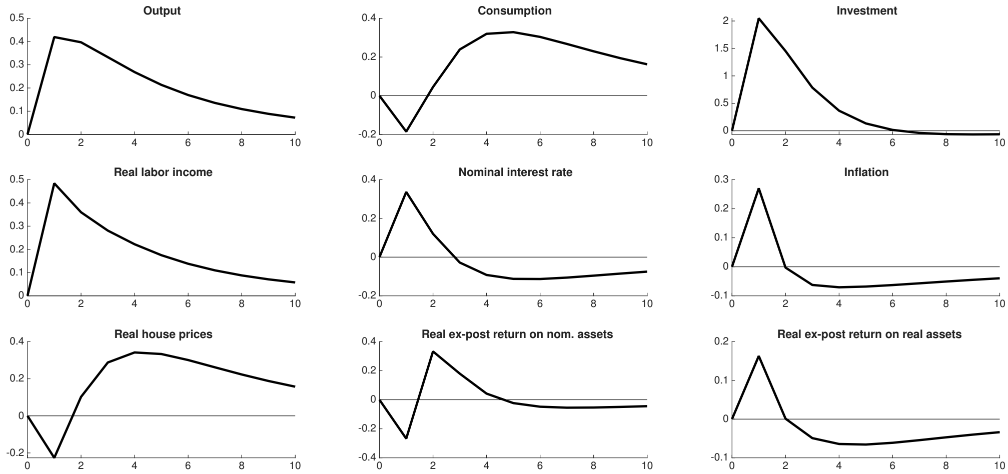

Figure C.4: Impulse response functions to a bond preference shock

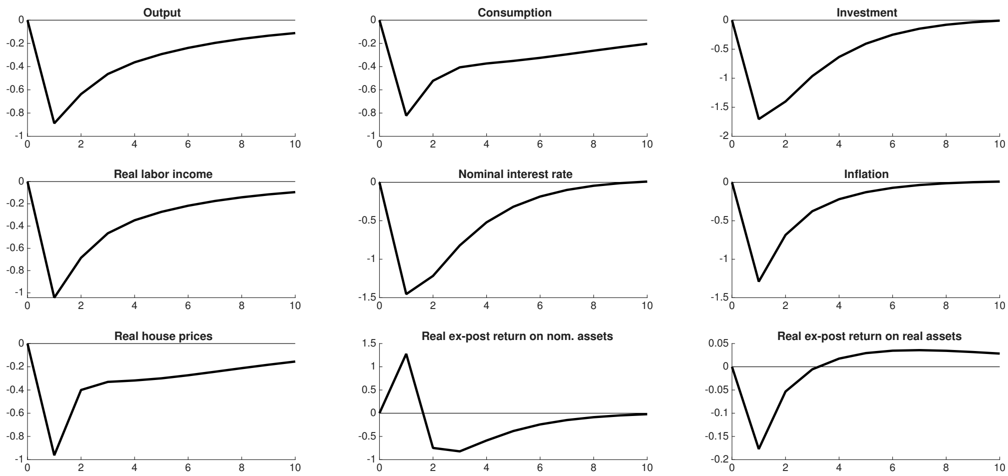

Figure C.5: Welfare effects of selected shocks and their transmission along specific channels

(a) Housing preference

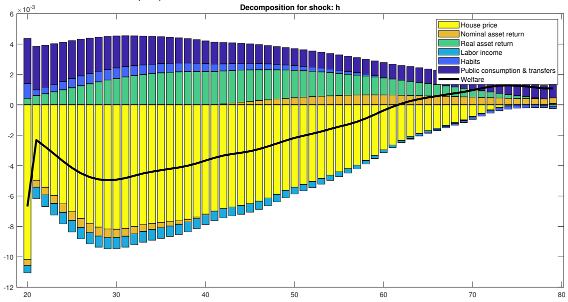

(b) Transfers

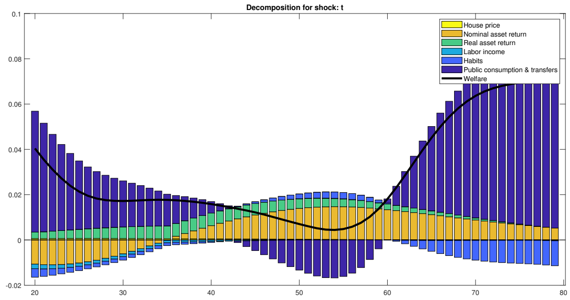

(g) Price markup

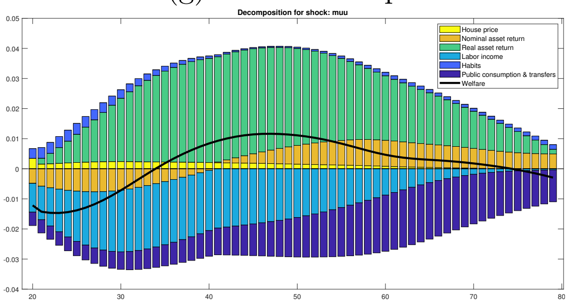

(h) Wage subsidy

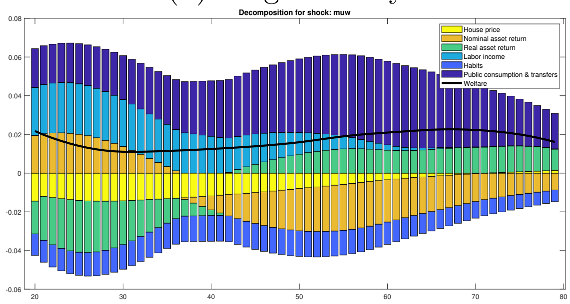

(g) Public spending