The redistributive power of business cycle fluctuations*

Marcin Bielecki† Michał Brzoza-Brzezina‡ Marcin Kolasa§

Abstract

How do business cycle fluctuations redistribute welfare between cohorts? Is this redistribution quantitatively important? What are the channels through which it occurs and how are they affected by monetary policy? To answer these (and other) questions, we construct a New Keynesian life-cycle model and estimate it for the US. We find that cyclical fluctuations are quantitatively important - for some cohorts and some years their impact on welfare can reach even 40% of annual consumption. More importantly, these effects do not tend to net out over a typical life cycle: some cohorts (e.g. born in the 1940s) have been historically substantially more lucky than others. Shocks that drive the cycle can be split into two groups: those that affect all cohorts' welfare in the same direction and those that create both winners and losers. Monetary policy shocks belong to the latter group and hence play an overproportional role in driving intergenerational redistribution, even if they do not contribute much to the variance of output. In contrast, more aggressive systematic monetary policy dampens the amount of redistribution.

JEL: E31, E51, E52, H5, J11

Keywords: Business cycles, Redistribution, Monetary policy, Life-cycle model

*We wish to thank Michał Reczek for providing excellent research assistance and participants of the Computing in Economics and Finance conference, Summer Worksop on Macro & Finance, and seminar at FPM ECB for comments and discussions. 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..

†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

1.1 Research problem and main findings

Business cycle fluctuations do not affect agents equally. As households differ, among other things, in their labor market status or asset holdings, they are differently affected by the same shocks. Some of these effects have already been tackled in the literature. For instance, the redistributive impact of inflation, monetary or fiscal policy has been widely analysed (more details below). In this paper, we ask a more general question: how do business cycles redistribute between agents of different age? If shocks driving the fluctuations are symmetric, is it enough to ensure that their influence on agents nets out over a typical lifetime? Which business cycle shocks are mainly responsible for redistribution? Which are the crucial channels through which cyclical fluctuations redistribute? And, last but not least, what is the role of systematic monetary policy - when does it dampen and when does it amplify the distributive impact of business cycle fluctuations?

While a detailed presentation of our findings follows further down the road, we begin with presenting a result that highlights the significance and novelty of our findings. The redistributive impact of cyclical fluctuations is large and does not net out over typical lifetimes. In case of large shocks, some cohorts may face gains or losses in remaining lifetime welfare equivalent to up to 40% of annual consumption. While these can be undone by subsequent shocks that act in the opposite direction, this averaging-out process is not perfect in the context of a typical lifetime. In our sample, some cohorts were lucky to encounter favorable shocks which boosted their lifetime welfare by up to 100% of annual consumption. Others were less lucky and faced a lifetime loss. This gives a clear reason to investigate the topic.

We do so by constructing and estimating a New Keynesian life-cycle model of the US economy. We use the model to peek into the drivers of redistribution. First, we show that business cycle shocks differ strongly in their redistributive power. For instance, markup shocks affect all cohorts in the same direction, while financial (risk premium) and monetary policy shocks redistribute heavily by creating winners and losers. Second, we identify the

main business cycle shocks that drive the redistribution. These are in particular price markup (i.e. cost) shocks, technology shocks and monetary policy shocks. While the former are well known to also drive cyclical (e.g. output) fluctuations, the large role of monetary policy shocks is an important finding, as their role in driving business cycles is rather small. Third, we also show via which channels redistribution occurs. It proves to work its way primarily via returns on nominal assets and via labor income. Contrary to popular beliefs, the revaluation of housing plays a negligible role.

Fourth, we document how cyclical fluctuations redistributed welfare historically. Typically, fluctuations affected all cohorts in the same direction - booms raise while recessions lower everyone's welfare. However, under specific conditions (when strongly redistributive shocks are the main drivers), fluctuations can also affect different cohorts differently. 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 (productivity and monetary policy). As a consequence, some cohorts gained at the expense of others.

Last but not least, the role of monetary policy is particularly interesting. On the one hand, as already mentioned, monetary policy shocks (i.e. the unexpected component of policy) belong to the the most redistributive drivers of business cycles. As a consequence, their role in redistributing welfare is much higher that their role in driving output fluctuations. On the other hand, systematic policy can contribute to dampening redistribution. Interestingly, there exists an interior solution for Taylor rule parameters which minimizes the amount of redistribution. This solution seems to be not very distant from typical parameters of central bank reaction to inflation, though the weight put on output is much higher than the one usually reported (and also estimated in this paper). Nevertheless, the amount of redistribution under the estimated policy rule is not dramatically higher than the minimized one.

Leaving a detailed literature review to the next section, we wish to underscore three dimensions of the original contribution of our work. First, to the best of our knowledge, there exists no other study that analyses the redistributive effects of cyclical fluctuations

accross generations. The closest paper to ours is Bayer et al. (2024), which discusses the impact of business cycle fluctuations on inequality. However, that study concentrates on the overall distribution of wealth and income, leaving the age dimension aside. We argue that this dimension is very important to understand redistribution, as life-cycle considerations drive agents to accumulate retirement assets, a process that critically affects the gains and losses that happen due to asset price changes (Auclert, 2019; Bielecki et al., 2022). Moreover, and this is our second key contribution, we show that even two symmetric shocks can affect agents' welfare when life cycle considerations are taken into account. This happens because the impact of shocks on household welfare depends crucially on their age. This is an important feature where life-cycle models offer insights on redistribution and inequality, which HANK models with infinitely-lived households cannot provide. Third, while being a side-product of our main goal, it is worth underscoring that we provide a full-scale first-order Bayesian estimation of a life-cycle model with real and nominal frictions. While computationally intensive, this task proved feasible. This achievement opens the door to new quantitative research in the spirit of the literature based on estimated representative agent DSGE models which, however, takes into account an important dimension of agent heterogeneity. We are not aware of other studies that would have done this.

1.2 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. Coibion et al. (2017) use micro data to show that contractionary monetary policy increases inequality. Dossche et al. (2021) document that a conventional monetary policy easing has an inequality-reducing impact mainly by reducing unemployment among the poorest households. Lenza and Slačálek (2021) come to a similar conclusion for unconventional monetary policy tools. 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 unexpected inflation generates a significant wealth redistribution. Pallotti et al. (2023) use the same data to measure the welfare effects of post-pandemic inflation and identify young households as net winners. Bielecki et al. (2022) use a life-cycle model to show that expansionary monetary policy redistributes from old to young generations and Bhattarai et al. (2023) analyze the distributional consequences of fiscal shocks. Heterogeneous agent New Keynesian (HANK) models are also increasingly gaining ground in the context of distributional effects, showing i.a. that poor households are more strongly affected by monetary policy shocks (see e.g., Kaplan et al., 2018; Guo et al., 2023).

In contrast to the literature discussed above, which concentrates on one particular source of redistribution, our paper offers a more holistic approach. We analyze the distributional consequences of a complete set of business cycle fluctuations and their drivers. To our knowledge, the only paper with such scope is Bayer et al. (2024), which discusses the consequences of business cycle fluctuations for inequality in a HANK framework. In contrast, we concentrate on the redistribution of welfare and focus on heterogeneity along the life-cycle dimension.

The second related stream of the literature calculates the cost of business cycle fluctuations. These have been broadly analyzed since the seminal contribution of Lucas (2003), who argued that in a world with representative agents this cost is negligible. However, this and many other related studies restrict their attention to second-order (i.e. volatility) effects of business cycles. This is because, for infinitely lived agents and symmetric business cycles, positive and negative fluctuations cancel out, leaving the first-order welfare effects to be zero. Our case is different as we concentrate on first order effects. First, clearly these need not cancel out if the horizon is finite. But, more importantly, there is another factor at play, as the same shock can affect agents differently at different age. A a result, even two symmetric shocks whose cyclical impact cancels out may leave the agent worse or better off if they hit

him/her at a different stage of life. This is a novel aspect in the literature on the cost of cyclical fluctuations and we show that, as a consequence, the effects of fluctuations can be several orders of magnitude higher then those calculated in the traditional fashion.

Finaly, 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 or 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 portfolio composition, 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, our model is estimated using Bayesian techniques (in the spirit of Smets and Wouters, 2003), which to our knowledge, have not 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 its calibration and estimation. Then, in section 4 we show our main findings about the impact of fluctuations on welfare (and its redistribution). Section 5 concentrates on the role of monetary policy and Section 6 concludes.

2 Model

Our model can be considered 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 for their

optimal portfolio allocations, readjusting their 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.

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 the 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. implicitly divided by total population N_t. All real variables are expressed relative to the price of final goods P_t. Variables without time subscripts indicate the steady state values.

2.2 Households

Individual households 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 tax at rate τ, where $z_j \geq 0$ denotes age-specific productivity. To allow for retirement, we assume that $z_j = 0$ for $j \geq J_R$ so that households reaching the retirement age J_R no longer work. 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 physical capital $k_{j,t}$. Capital can be rented to firms, earning the real rental rate r_t^k. 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 with age j can be written as

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

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 stocks 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 $g_{j,t}$, housing and bond holdings, and disutility from working. Household ι chooses its consumption and asset holdings to maximize[1]

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

where φ is the inverse Frisch elasticity of labor supply, ϱ is the external habit persistence parameter, and $\psi_{j,t} \equiv \psi_j \varepsilon_t^h$, $\zeta_{j,t} = \zeta_j \varepsilon_t^b$ and ϕ_j are the age-specific weights that control the importance of housing, bond holdings, and leisure in utility. The assumed weight on

[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.

private and public consumption ensure that their marginal utility are equal in the steady state. The weights on housing and bonds are additionally driven by stochastic shocks ε_t^h and ε_t^b, which are aimed to capture exogenous fluctuations in housing demand and convenience yields, respectively.

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 to 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 rate $\exp(\varepsilon_t^w) - 1$, where ε_t^w can be interpreted as a wage cost-push shock, which is then rebated back to all unions in a lump sum fashion. 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 $h_t(f)$ to produce according to the standard Cobb-Douglas technology

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

where ε_t^z is a stochastic productivity shock, α is the capital share in production and Φ is the fixed cost.[2] The real profit of firm f is $p_t(f) y_t(f) - \exp(\varepsilon_t^p)\left(w_t \ell_t(f) + r_{k,t} k_t(f)\right) + t_t^f$, 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 $t_t^f = \left(\exp(\varepsilon_t^p) - 1\right)\left(w_t \ell_t + r_{k,t} k_t\right)$. 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

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

where S_k controls the degree of adjustment costs.

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

2.6 Fiscal and monetary authority

The government has two types of expenditures. It purchases $g_t = g \exp(\varepsilon_t^g)$ of final goods, where ε_t^g can be interpreted as a public consumption shock. These goods are then equally redistributed across all households so that $g_{j,t} = g_t$ for all j. The government also pays lump sum transfers to households, using this instrument to stabilize its long-run debt. More specifically, we assume

$$\frac{t_t}{t} = \left(\frac{t_{t-1}}{t}\right)^{\gamma_t} \left(\frac{b_t^g}{b^g}\right)^{-\eta} \exp(\varepsilon_t^t), \tag{7}$$

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 housheols and by issuing debt that pays nominal interest rate R_t. The government real budget constraint can hence be written as

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

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

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

where ε_t^r is a stochastic monetary shock, γ_R is the smoothing coefficient while γ_π and γ_y control the long-run 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 housing stock is fixed in per capita terms at a constant value h. The aggregate resource constraint

can be written as

$$y_t \Delta_t = \exp(\varepsilon_t^z) k_t^\alpha \ell_t^{1-\alpha} - \Phi, \tag{10}$$

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

where $\Delta_t = \left(\frac{P_t(f)}{P_t}\right)^{\frac{\mu}{1-\mu}} \mathrm{d}f$ is a measure of the price dispersion due to staggered price contracts.

3 Calibration and estimation

We parametrize 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 key 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 key 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 A (TBA).

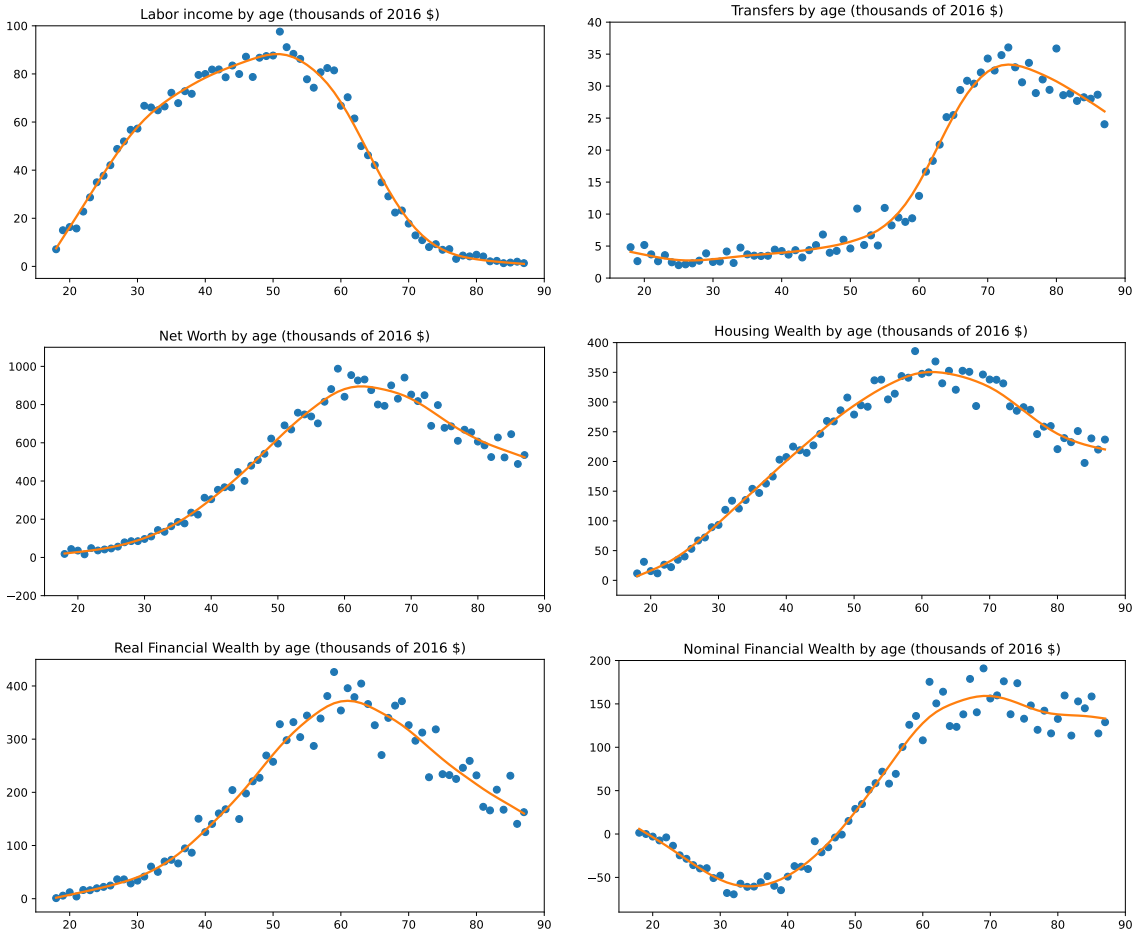
The first row of Figure 1 presents the age profiles of labor income and transfers to households. The labor income profile follows a well-documented pattern, growing steadily through the middle age, peaking in the early 50s, and then steadily declining. This shape reflects two major processes: (i) wages tend to increase over the life cycle as workers become more productive and gain experience, (ii) labor market participation starts dropping as households begin moving to retirement. In contrast, the profile for transfers is fairly flat through most of working age and then starts increasing 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 initially nominal financial liabilities, 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. 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 decumulation of nominal assets.

To incorporate these patterns of income and asset evolution over the life cycle in the model, at the same time keeping it 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 and ensure that hours worked are equal for all working age households.[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-term 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 housing-to-GDP ratio of 1 and the equity premium of 0.01, respectively. The age profile of real financial assets, and hence of total net worth, is determined endogenously in the

[3]A flat profile of hours per worker is a good approximation of the SCF data for households aged between 20 and 65, consistently with our model assumptions about the age of entering the labor market and retiring. We 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.

model.

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 3.1 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	Labor income tax rate	$b_g/y = 0.35$

3.2 Bayesian estimation

The remaining parameters of the model are estimated using Bayesian methods. We use nine macroeconomic time series spanning the period 1960-2023. 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 A (TBA) 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 3.2.

Table 2: Estimated structural parameters

Parameter	Prior			Posterior		Description
	type	mean	std	mean	std	
ϱ	beta	0.75^4	0.1	0.13	0.05	Habit persistence
θ_p	beta	0.75^4	0.1	0.49	0.03	Calvo prob. for prices
θ_w	beta	0.75^4	0.1	0.34	0.04	Calvo prob. for wages
S_k	norm	4.0	0.5	4.05	0.48	Inv. adjustment cost
γ_R	beta	0.7^4	0.1	0.13	0.06	Interest rate smoothing
γ_π	norm	1.5	0.1	1.51	0.07	Response of interest rate to inflation
γ_y	beta	0.125	0.05	0.13	0.06	Response of interest rate to output
γ_t	beta	0.75^4	0.1	0.34	0.08	Tax smooting
η	beta	0.1	0.05	0.12	0.04	Weight on debt in transfer rule

TBC

[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.

4 Redistribution across age and time

This section documents our findings about the redistributive consequences of business cycles. First, we show how fluctuations of the US economy redistributed welfare accross generations. Then, we explain how various business cycle shocks redistribute welfare. To this end, we discuss the role of economic channels via which this process happens. Next, we check how much these shocks and channels contributed to redistribution in the historical context. Last but not least, we check to what extent the redistributive impact of cyclical fluctuations was netted out over the lifetimes of past generations. This allows us to assess whether business cycles can leave some cohorts significantly better-off than others.

4.1 Main concepts

The subsequent analysis makes heavy use of a key metric of changes in welfare that we call "unexpected remaining lifetime welfare gain (loss)", or in shorthand, "welfare gain" (WG). Such defined welfare gain reflects the effects of changes in the agents' economic environment resulting from any unexpected shocks that arrive at the very beginning of period t and which, by definition, were unforecastable in period $t-1$.

For concreteness, imagine a $j-1$ aged household that optimized their entire lifecycle plan in period $t-1$ based upon all information that was available until period $t-1$. The implied expected remaining lifetime welfare of this now j-aged household in period t (but prior to t period shocks being revealed), is denoted with $\mathbb{E}_{t-1}U_{j,t}$. The arrival of new shocks at the beginning of period t necessitates reoptimization of the lifetime plan and results in a "shifted" expected remaining lifetime welfare with t period information taken into account $\mathbb{E}_t U_{j,t}$. We hence can define the welfare gain or loss of a household aged j due to shocks arriving at time t as:

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

To express the above metric in interpretable units, we present it as percentage of steady

state consumption of a given cohort.

Another concept that we will frequently refer to 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 specific period t, the $AMOR$ statistic is calculated as follows:

$$AMOR_t = \sum_j \left\| \mathcal{W}_{j,t} - \bar{\mathcal{W}}_t \right\| \quad \text{where} \quad \bar{\mathcal{W}}_t \equiv (1/J) \sum_j \mathcal{W}_{j,t} \tag{13}$$

Intuitively, if welfare of all cohorts changes exactly in unison, the value of $AMOR$ is 0. Otherwise, it is the more positive the more differently the individual cohorts are affected.

To provide insights on the importance of channels through which business cycle fluctuations affect welfare gains $\mathcal{W}_{j,t}$, we offer a decompostition along the lines developed by 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 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 period t due to a subset of shocks s can therefore be decomposed as follows:

$$\mathcal{W}_{j,t,s} = \Gamma^h_{j,t,s} + \Gamma^b_{j,t,s} + \Gamma^{kf}_{j,t,s} + \Gamma^{w\ell}_{j,t,s} + \Gamma^t_{j,t,s} + \Gamma^{hab}_{j,t,s} + \mathcal{O}(2) \tag{14}$$

where the Γ's capture, respectively, the effects of the resulting changes in: house prices, return on nominal assets, return on real real assets, labor income, transfers (including accidental bequests), and external habits. The $\mathcal{O}(2)$ term collects all second (and higher) order effects.

Since all the Γ's are first-order objects, they lend themselves to a straightforward linear aggregation. For example, one can track the overall effect of a single shock on a given cohort's total welfare gain $\mathcal{W}_{j,t,s}$. One can also aggregate the effects of all shocks along their respective channels, which often provides a convenient, direct link to the data objects $\mathcal{W}^c_{j,t} = \sum_s \Gamma^c_{j,t,s}$, with $c \in \{h, b, kf, w\ell, t, hab\}$.

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 14 also using this reference point. Then the respective Γ's can be expressed as follows (supressing shock index s for notational convenience):

$$\Gamma^{h}_{j,t} = u^{c}_{j} \sum_{i=0}^{J-j} (1+r^{a})^{-i} \left[\left(1-\delta^{h}\right) h_{j+i} - h_{j+i+1} \right] \left(\mathbb{E}_{t} p^{h}_{t+i} - \mathbb{E}_{t-1} p^{h}_{t+i} \right) \tag{15}$$

$$\Gamma^{b}_{j,t} = \mathbb{E}_{t} u^{c}_{j} \sum_{i=0}^{J-j} (1+r^{a})^{-i} b_{j+i} \left[\mathbb{E}_{t} \frac{R_{t+i}}{\pi_{t+i+1}} - \mathbb{E}_{t-1} \frac{R_{t+i}}{\pi_{t+i+1}} + (1+r) \left(\mathbb{E}_{t} \varepsilon^{b}_{t+i} - \mathbb{E}_{t-1} \varepsilon^{b}_{t+i} \right) \right] \tag{16}$$

$$\Gamma^{kf}_{j,t} = \mathbb{E}_{t} u^{c}_{j} \sum_{i=0}^{J-j} (1+r^{a})^{-i} \left\{ \left[\left(1-\delta^{k}\right) k_{j+i} - k_{j+i+1} \right] \left(\mathbb{E}_{t} p^{k}_{t+i} - \mathbb{E}_{t-1} p^{k}_{t+i} \right) + k_{j+i} \left(\mathbb{E}_{t} r^{f}_{t+i} - \mathbb{E}_{t-1} r^{f}_{t+i} \right) \right\} \tag{17}$$

$$\Gamma^{w\ell}_{j,t} = \mathbb{E}_{t} u^{c}_{j} \sum_{i=0}^{J-j} (1+r^{a})^{-i} z_{j+i} \left[\begin{array}{c} (1-\tau) \ell_{j+i} \left(\mathbb{E}_{t} w_{t+i} - \mathbb{E}_{t-1} w_{t+i} \right) - w \ell_{j+i} \left(\mathbb{E}_{t} \tau_{t+i} - \mathbb{E}_{t-1} \tau_{t+i} \right) \\ + \frac{\mu_{w}-1}{\mu_{w}} (1-\tau) w \left(\mathbb{E}_{t} \ell_{j+i,t+i} - \mathbb{E}_{t-1} \ell_{j+i,t+i} \right) \end{array} \right] \tag{18}$$

$$\Gamma^{t}_{j,t} = \mathbb{E}_{t} u^{c}_{j} \sum_{i=0}^{J-j} (1+r^{a})^{-i} \left[\mathbb{E}_{t} \left(t_{j+i,t+i} + beq_{t+i} \right) - \mathbb{E}_{t-1} \left(t_{j+i,t+i} + beq_{t+i} \right) \right] \tag{19}$$

$$\Gamma^{hab}_{j,t} = -\varrho \mathbb{E}_{t} \sum_{i=0}^{J-j} \beta^{i} \frac{N_{j+i}}{N_{j}} u^{c}_{j+i} \left[\mathbb{E}_{t} c_{j+i,t+i} - \mathbb{E}_{t-1} c_{j+i,t+i} \right] \tag{20}$$

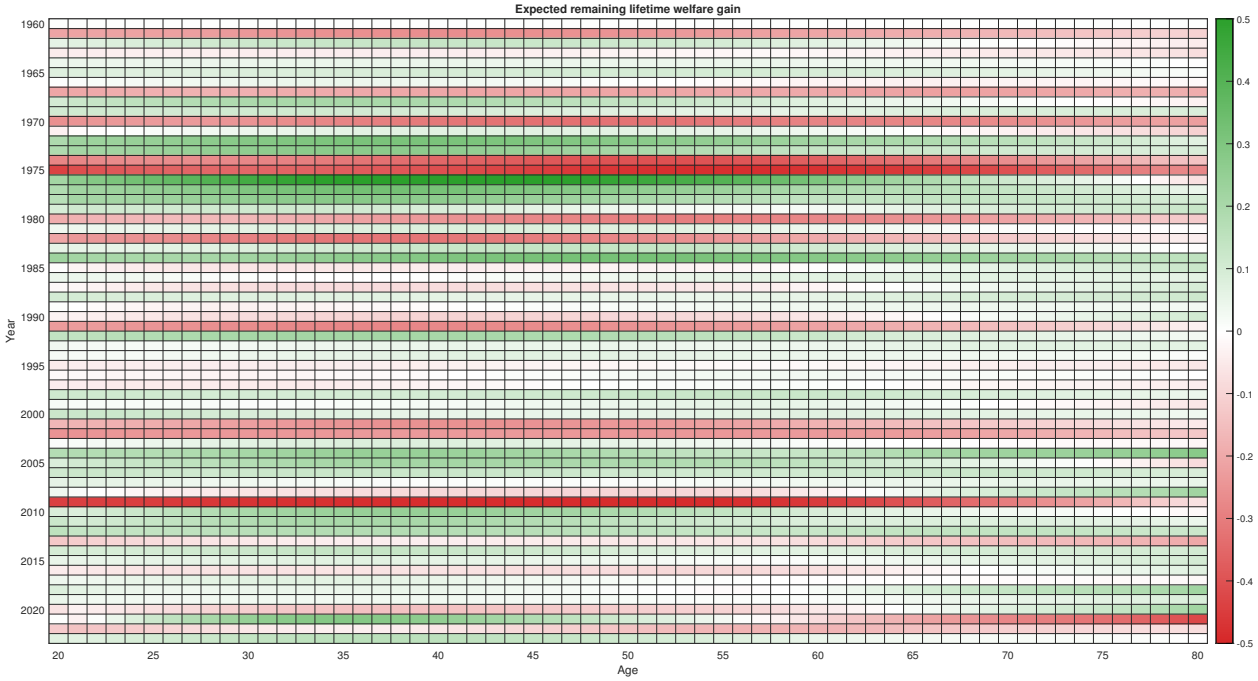
where $u^{c}_{j} \equiv \left[(1-\varrho) c_{j} \right]^{-1}$ is the steady-state marginal utility of consumption at age j and $r^{f} \equiv r^{k} + f/k - \delta^{k}$ represents the rate of return on real assets, which in our context is the appropriate discount rate. Note that all variables without time subscript denote their steady state values.

4.2 Welfare redistribution over business cycles

A typical recession involves declines in consumption and the worsening of the consumption-leisure tradeoff. As a result, for the majority of cases the direction of welfare gains (losses) due to newly arrived shocks is procyclical and largely uniform across cohorts. In Figure 2 one can easily discern the declines in welfare associated with e.g. 2001 and 2009 recessions. These unexpected changes to remaining lifetime welfare are non-trivial as due to finite planning

horizons they can correspond even to up to 50% loss of annual steady state consumption of any given cohort.

Figure 2: Overall welfare gain (loss)



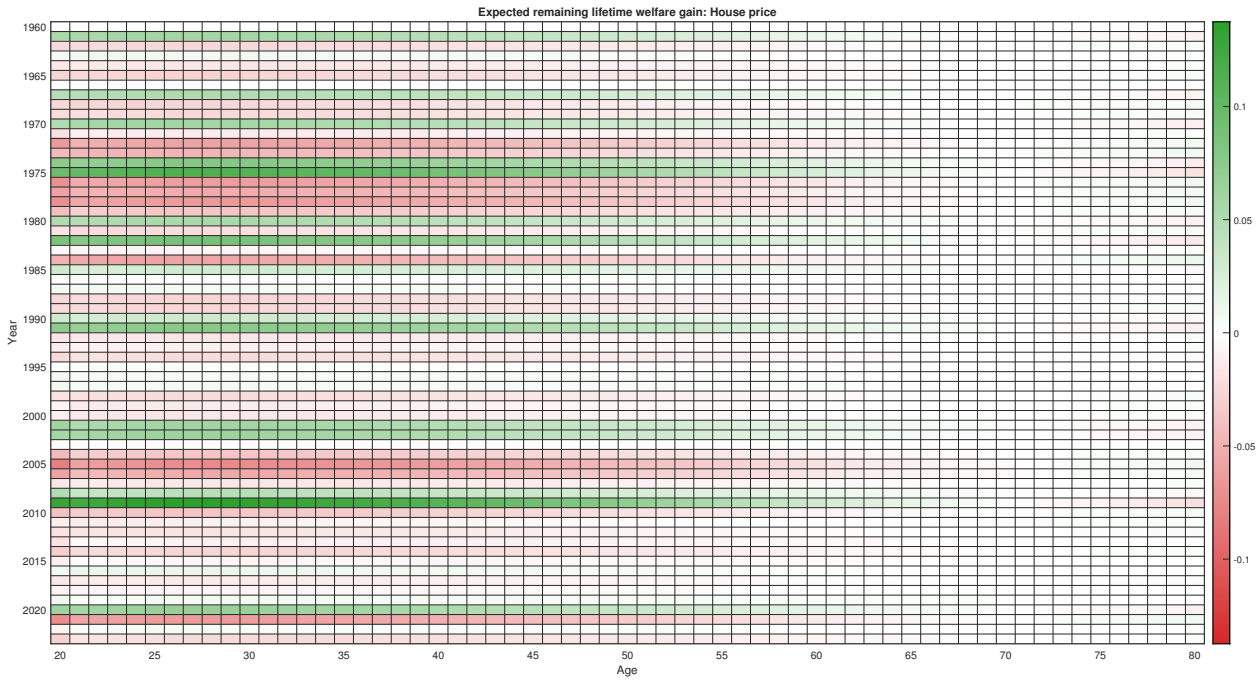
Note: Each rectangle represents a remaining lifetime welfare gain (loss) of a cohort aged j as a consequence of newly arrived shocks in year t, corresponding to a % gain (loss) of the cohort's steady state consumption.

Our decomposition of effects of shocks by channels can then be readily applied to (1) examine the transmission of shocks through specific items in the households' budget constraints to build additional intuition and (2) offer a glimpse in the heterogeneity of their impacts on specific household groups, being the primary source of redistribution across generations.

A good place to start the discussion of the various channels of shock transmission is the house price channel depicted in Figure 3. Although its quantitative importance is relatively minor in our model (due to housing being a continuous variable and households having access to within-cohort insurance), one can easily relate the historical house price movements to their effects on specific cohorts' welfare. Young households that have the steepest profile of housing accumulation gain in those periods where house price dynamics turned out to be lower than anticipated (e.g. in 2009 and 2020) and lose when housing quickly appreciates

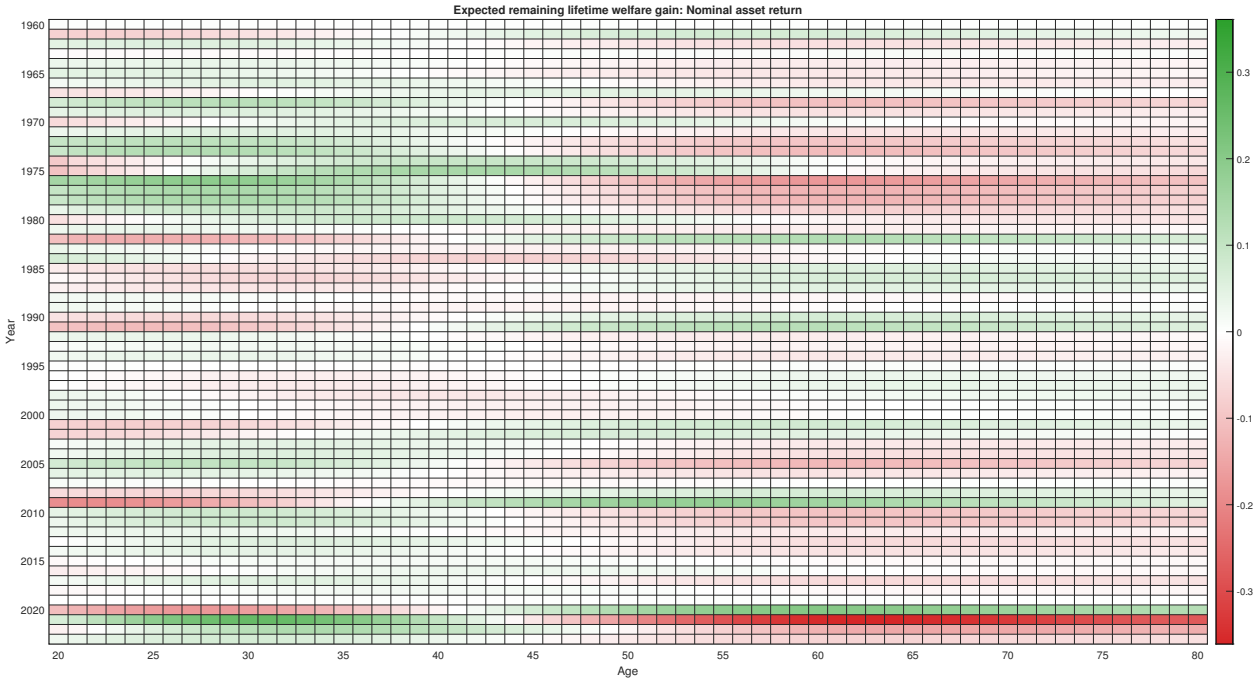
in value (e.g. in 2021). It is important to realize that the direction of welfare gain is not directly related to the already acquired ownership of housing. Even households aged 50 are net losers when house prices appreciate, despite owning already substantial "housing portfolio", as these households are still planning (on average) to acquire more housing due to their lifecycle considerations. The group that benefits from hikes in house prices are those that are reducing their housing ownership, which are mainly the already retired.

Figure 3: Welfare gain: house price channel



Nominal asset returns channel (in Figure 4) also lends itself to straightforward interpretation. Since young households typically hold nominal liabilities, while elderly households have positive net nominal assets position, unexpected changes to inflation or nominal interest rate carry first-order, heterogeneous welfare effects. Young households benefitted from low real interest rates on bonds and mortgages throughout the 1960s and 1970s, and bore the costs of Volcker disinflation in the first half of 1980s. More recently, mortgagors benefitted from post-GFC low interest rate environment and gained substantially from the post-pandemic surprise inflation surge, at the expense of the elderly bondholders.

Figure 4: Welfare gain: nominal asset return channel



The welfare effects along the real assets channel (Figure 5) are also easily related to the developments in the valuation of real financial assets. The very young households are in the accumulation phase while their portfolios are relatively small. Therefore, they strongly benefit from deep declines in the stock markets in e.g. 2009 and 2020, while their welfare deteriorates when the stock market rapidly booms, as in e.g. 2021. For households above the age of 35, and especially those already retired, the welfare effects run in the opposite

direction due to their already sizeable real asset holdings.

Figure 5: Welfare gain: real asset return channel

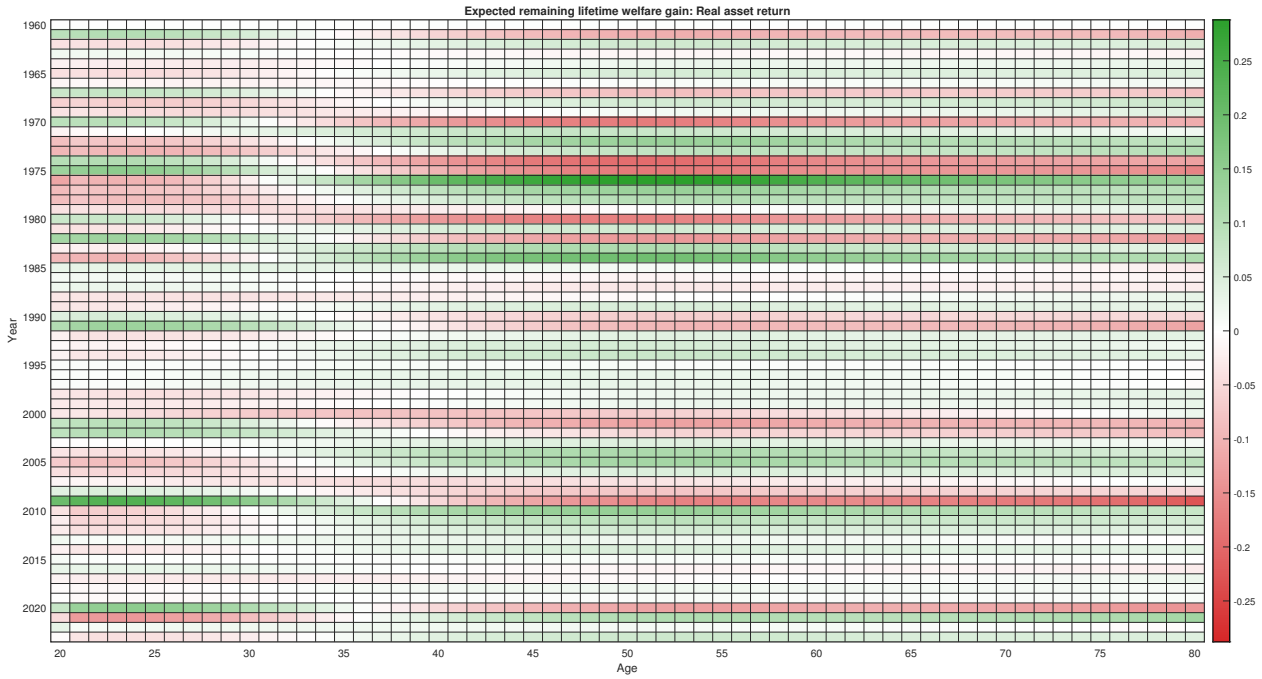
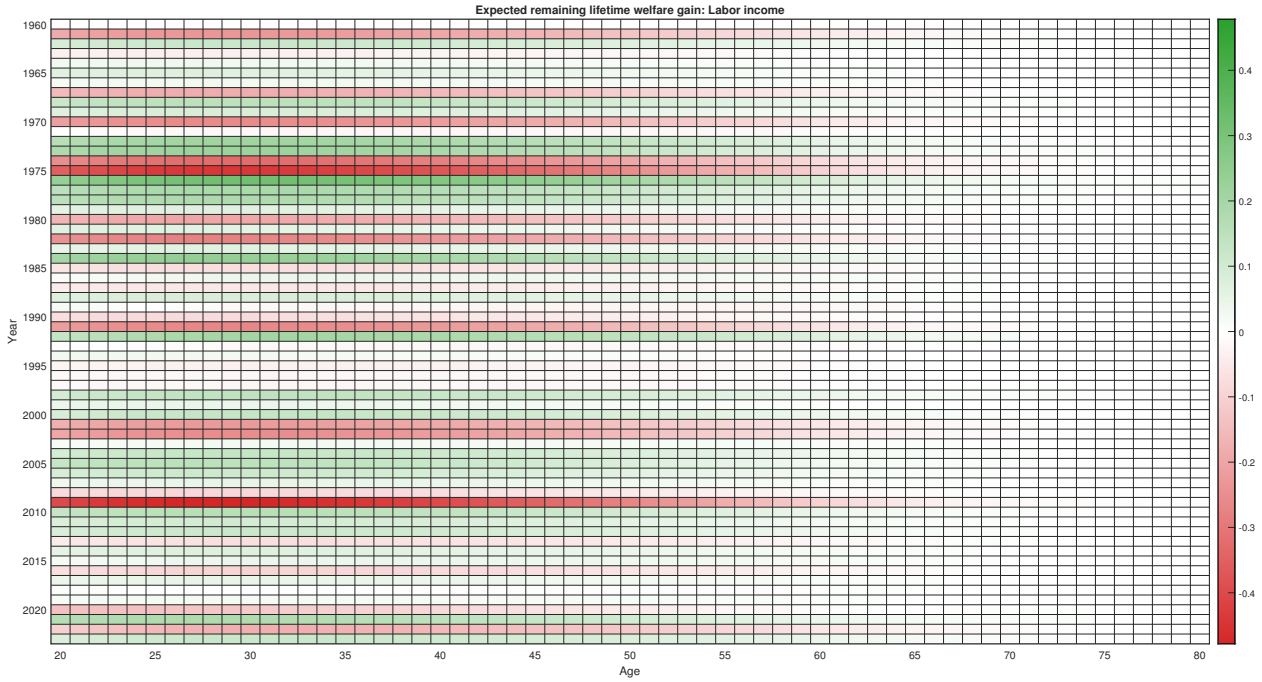


Figure 6: Welfare gain: labor income channel



The last easily interpretable channel captures changes in welfare due to labor income fluctuations in Figure 6. Naturally, the households that are the most exposed to labor market deteriorations are those still active in the labor market, and especially those that still plan to remain active for at least a couple of years. One can easily see the devastating effects of high-unemployment recessions of 1974-1975 and 2008-2009, or the pain of Volcker disinflation, the slowdown in the early 1990s and the dot-com recession of 2001-2002.

4.3 Redistributive and egalitarian shocks

Since the economic developments can evidently affect heterogeneous households' welfare differently, it is now instructive to examine the impacts of specific shocks that are the underlying sources of net gain or loss for specific households. In what follows we present first the overall welfare net gains for each cohort due to single shocks originating from the steady state, with their effects split into transmission via various channels. Second, we display the effects of historical sequences of shocks filtered out by our estimation procedure, focusing on their redistributive properties.

Figure 7: Welfare effects of monetary policy shock and its transmission along specific channels

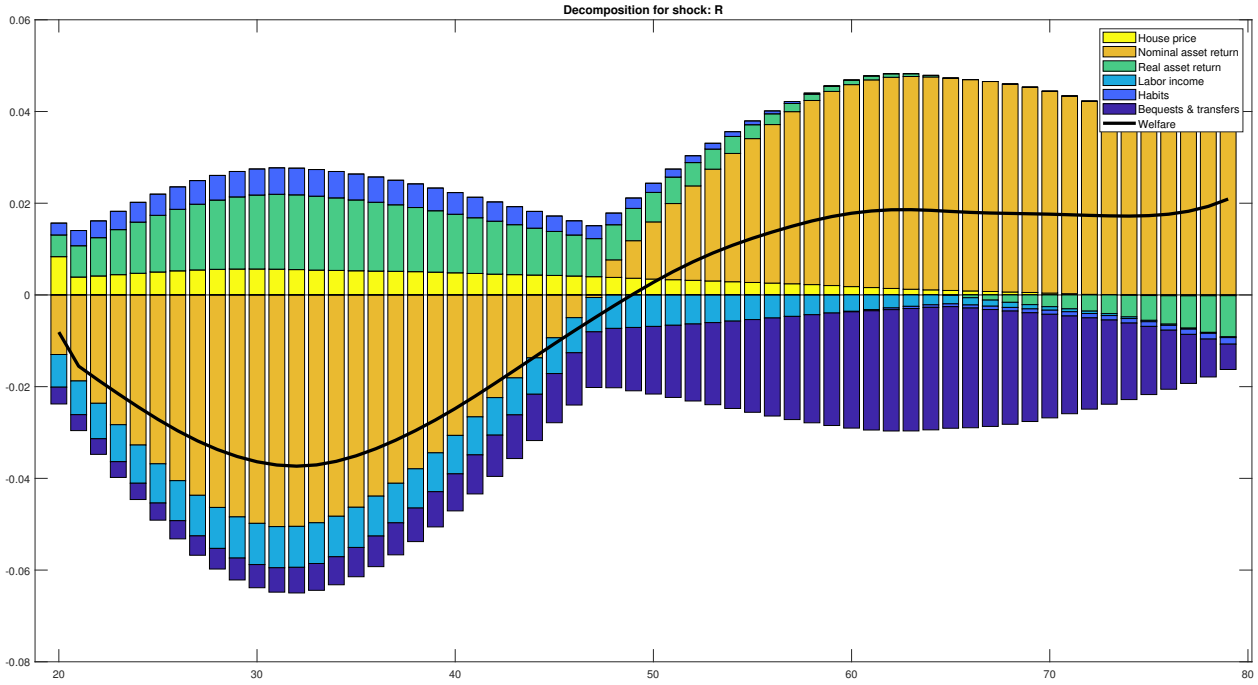
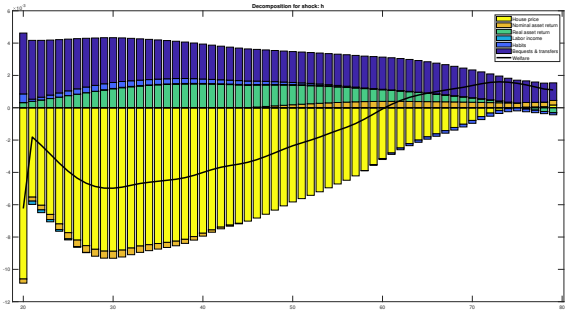
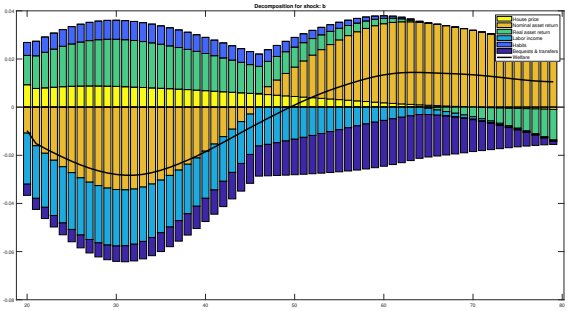


Figure 8: Welfare effects of shocks and their transmission along specific channels

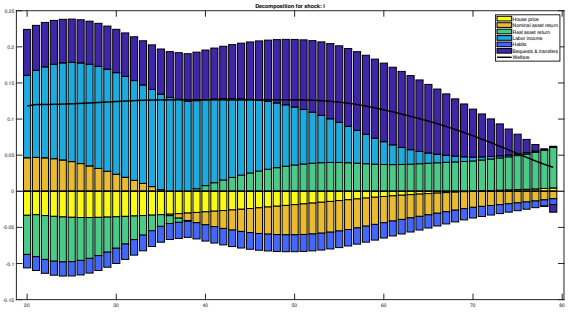
(a) Housing preference



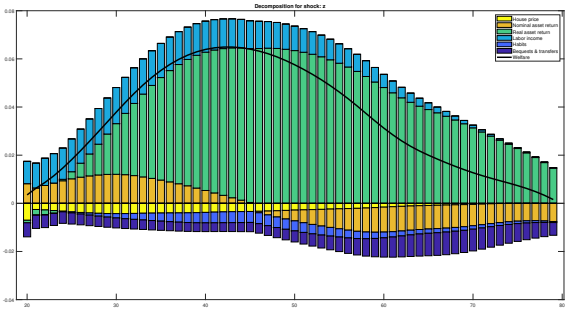
(b) Bond preference



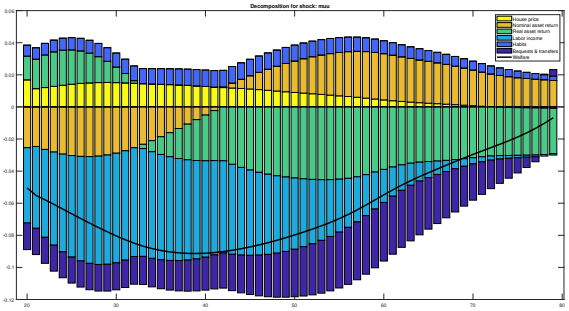
(c) Investment-specific



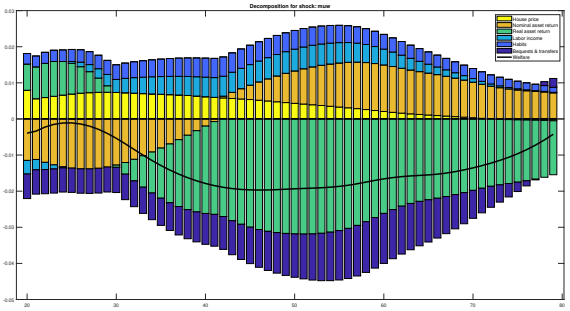
(d) TFP



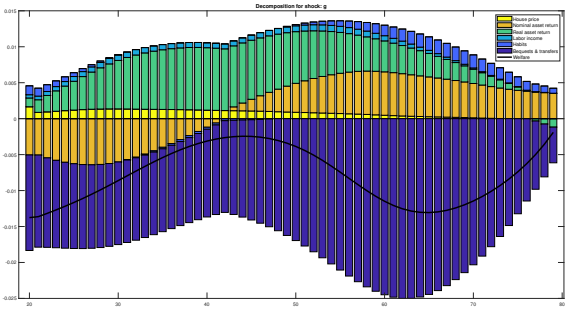
(e) Price markup



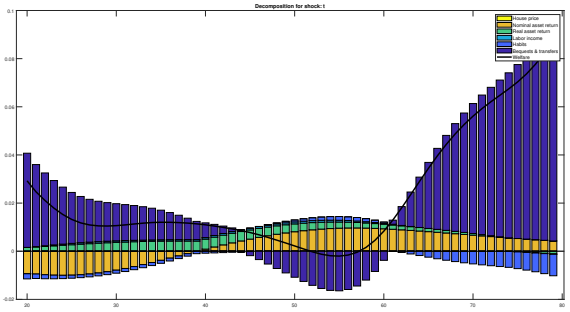
(f) Wage markup



(g) Public spending



(h) Transfers



As evidenced by Figures 7 and 8, all shocks affect households' welfare in a heterogeneous manner, but the welfare dispersion effect of certain shocks is much larger. We therefore adopt a following distinction: whenever a shock results is some households' welfare to increase while deteriorating the welfare of others, we call it "redistributive". Prime examples of redistributive shocks are both bond preference and monetary policy shocks, as they necessarily affect households differently due to their opposite unhedged interest rates exposure. We also decided to include in the "redistributive" group those shocks that despite affecting the welfare of the households in the same direction, nevertheless carry substantially different quantitative effects, as e.g. the TFP shock. The remaining shocks that have a largely similar quantitative effects for welfare, such as e.g. the investment-specific or price-markup shocks, we call "egalitarian".

This distinction can be readily seen in our historical decomposition. As a stark example, Figures 9 and 10 depict the welfare effects for various cohorts resulting from the monetary policy shocks. A clear division runs between cohorts below or above age 49: younger households benefit from monetary policy easing at the expense of older households. Moreover, when the monetary policy maintains a certain stance for a long period of time, as over the entire 2001-2021 period of accommodation, this may result in creating distinct "winners" and "losers" across the generations. To give a specific example (and focusing strictly on monetary policy), a household that entered the labor market in a few years prior to 1980 faced a decade of tight monetary policy that slowed its pace of asset accumulation. When monetary policy became accommodative after the turn of the millennium, the household - in their 40s and later 50s - again found itself on the losing side.

Such patterns are less likely - though not impossible - to occur for the more "egalitarian" shocks, like the price markup shock in Figures 11 and 12. Its welfare impact is almost symmetric across cohorts and the cumulated welfare effects result only from receiving disproportionately many shocks in a given direction throughout the lifecycle.

Figure 9: Welfare effects of new historical monetary policy shocks

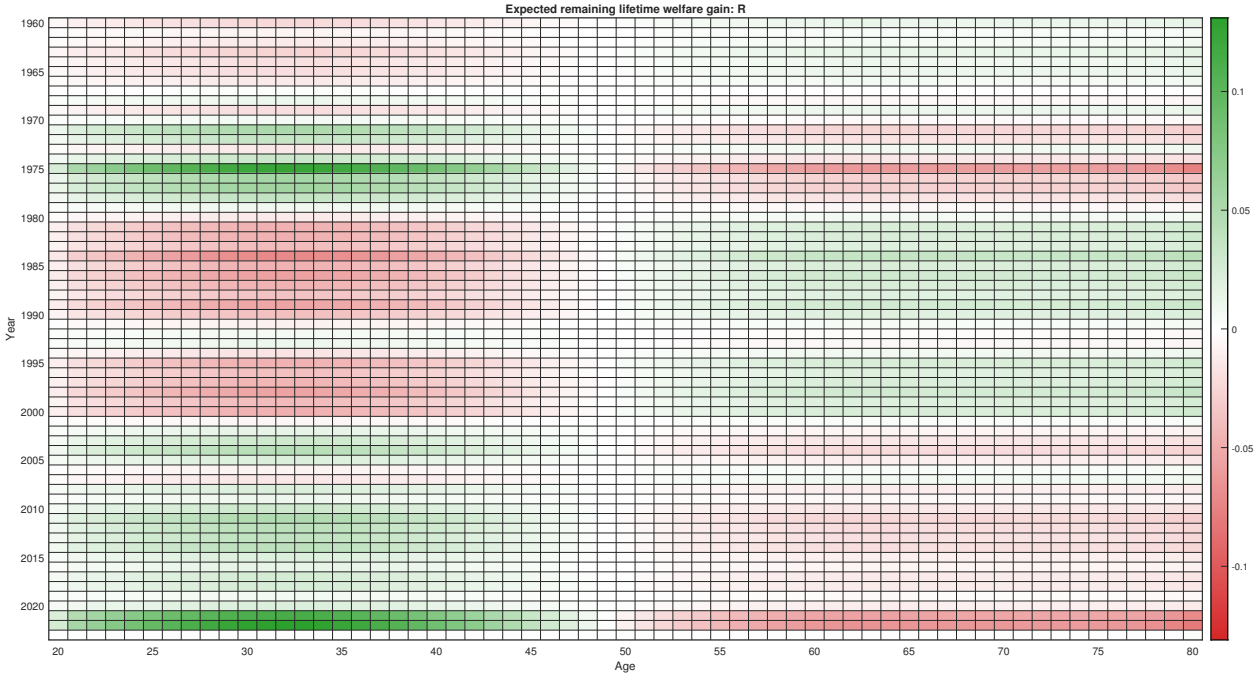


Figure 10: Cumulated welfare effects of historical monetary policy shocks

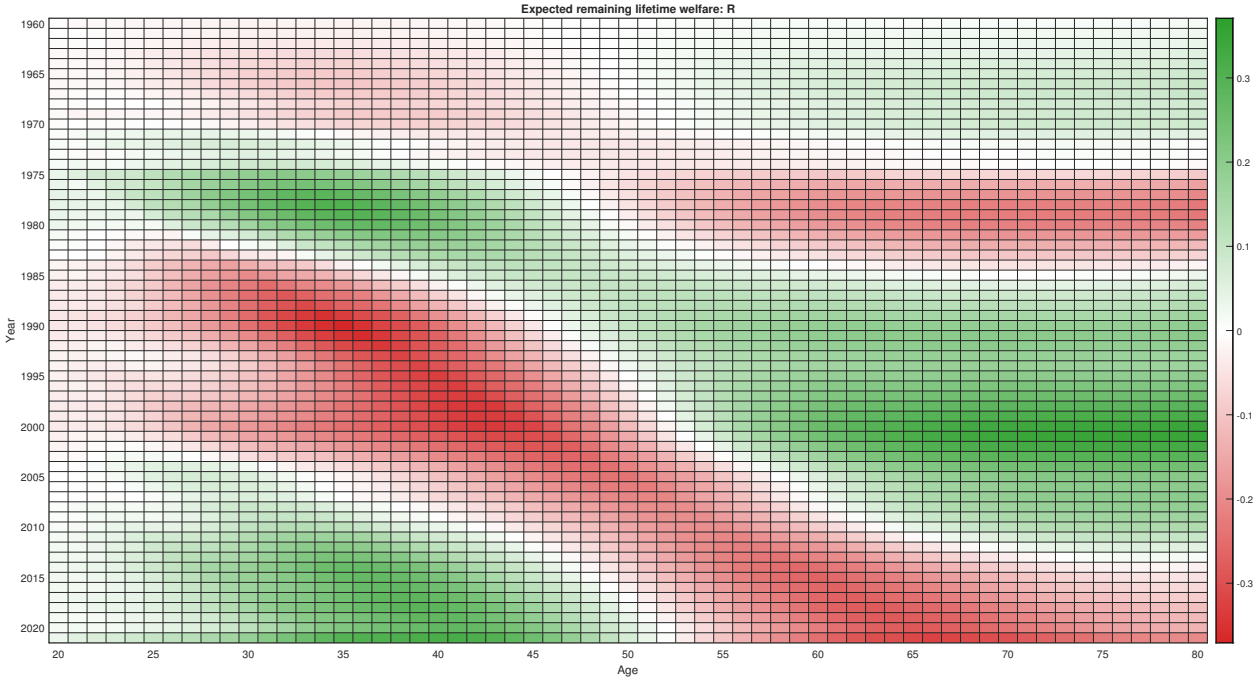


Figure 11: Welfare effects of new historical price markup shocks

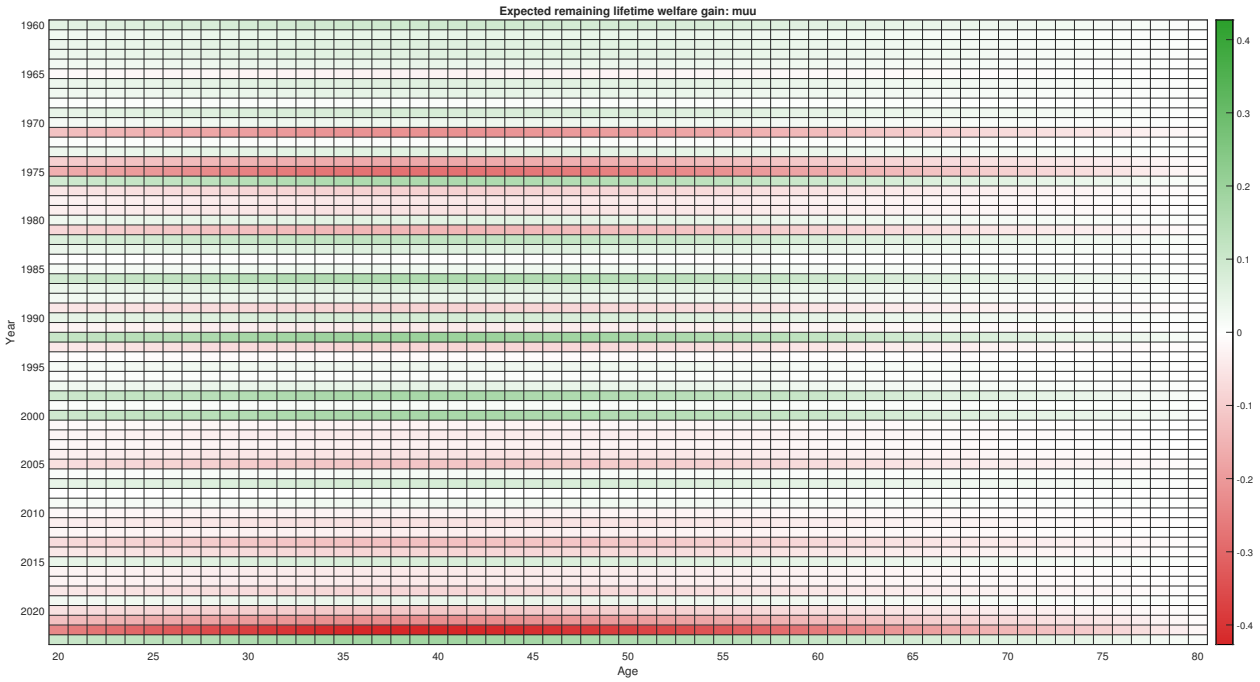
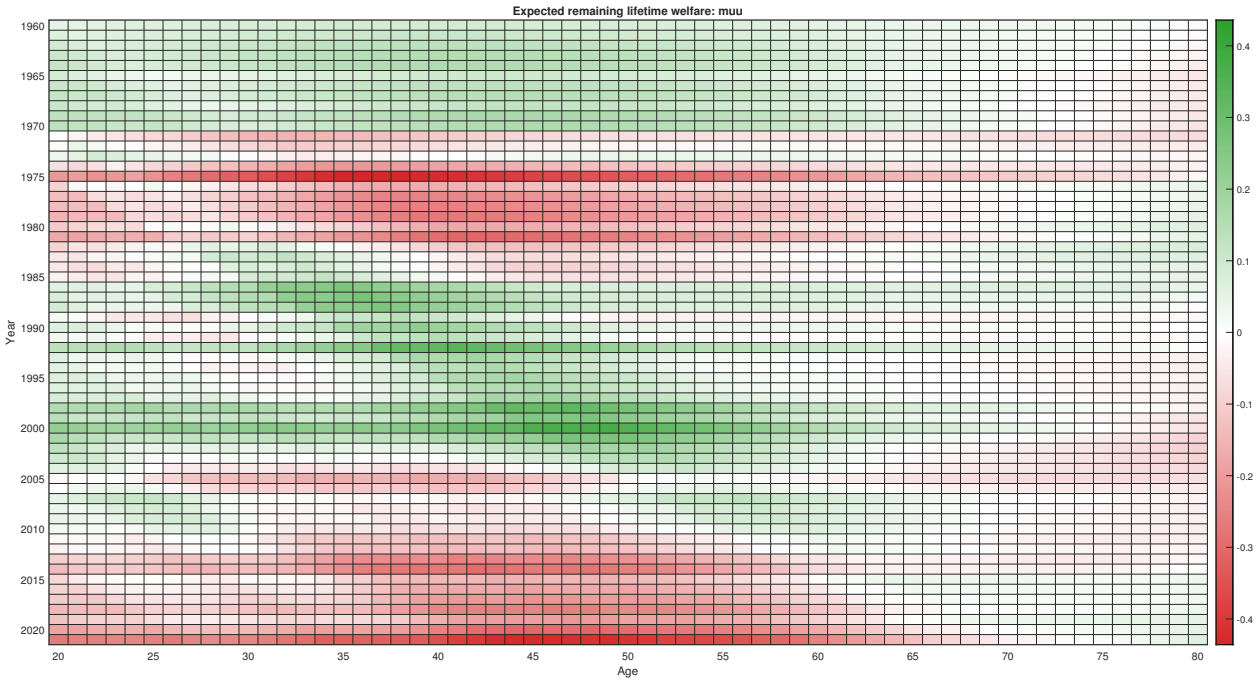


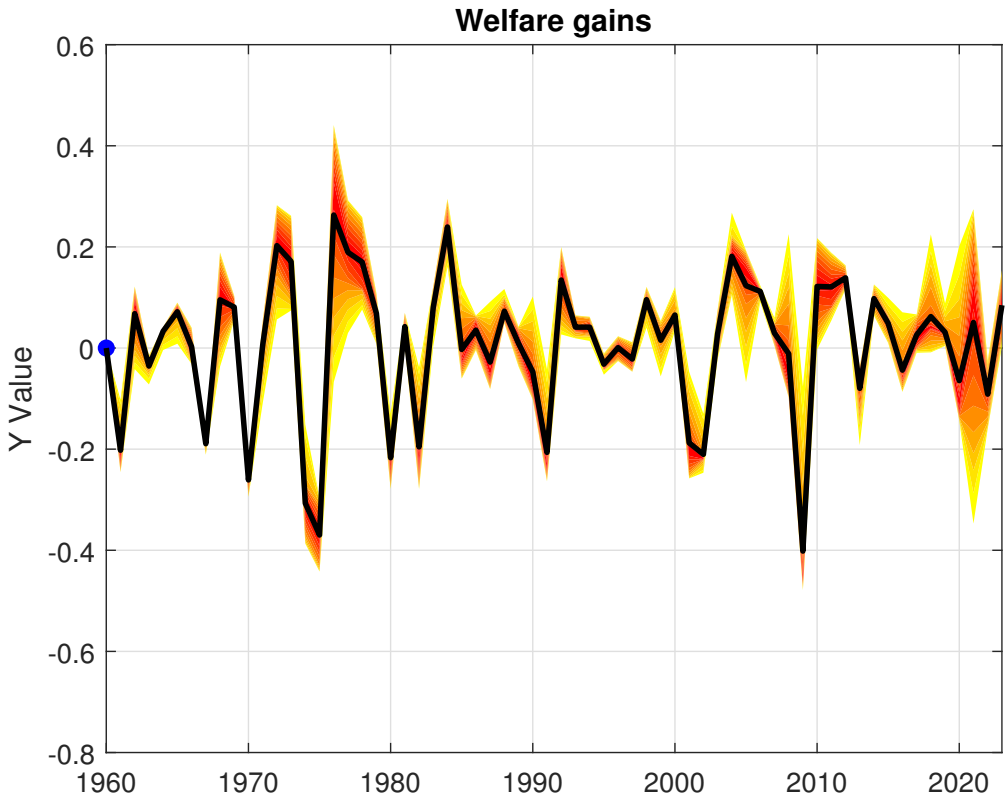
Figure 12: Cumulated welfare effects of historical price markup shocks



4.4 Amount of redistribution

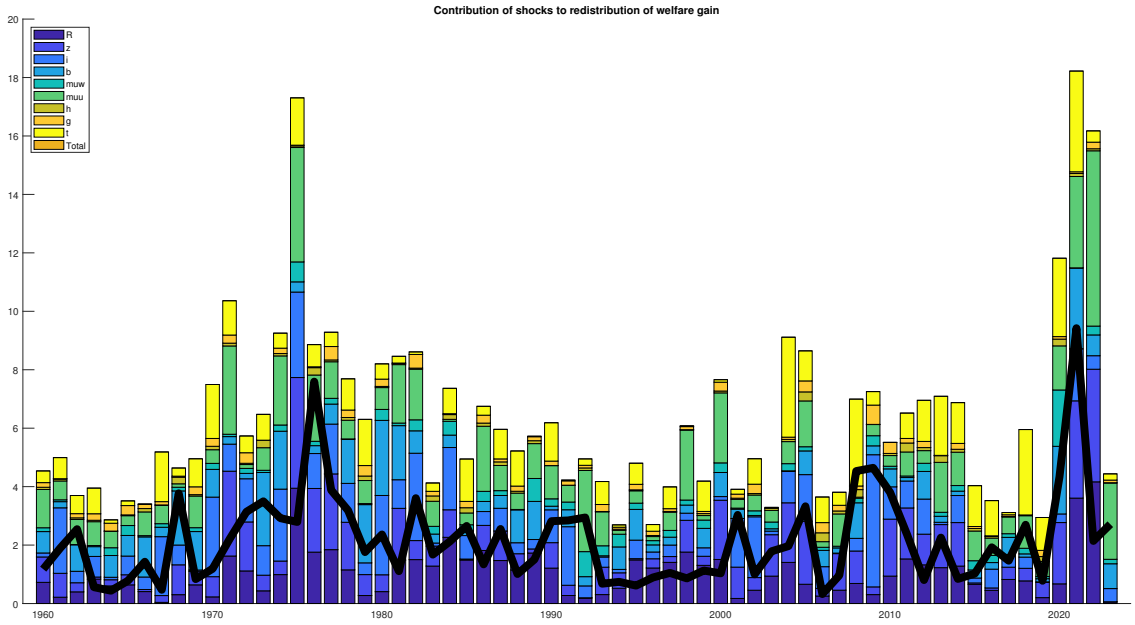
The documented heterogeneity in the welfare effects of various shocks calls for reconsidering the first-pass observation that most business cycles affect distinct households in a largely uniform manner. In Figure 13 we plot the average welfare gain across all cohorts in a given year (black line), surrounded by the "fanchart" depicting the range of changes in welfare across multiple household cohorts. While we can reconfirm that most of the time the welfare effects of any business cycle are rather symmetric, a few episodes distinctly stand out. The business cycles in the 1970s, involving large fluctuations in energy prices, resulted in a larger than usual dispersion in household welfare. A similar, but even more radical, occurence, was the set of shocks received since the 2020, involving the global Covid pandemic, a faster than expected reopening of the economies in 2021 and the Russian aggression on Ukraine in 2022, the latter two resulting in a significant energy and food prices spike at the end of our sample period. As a result our model infers a significant amount of redistribution across generations, involving also welfare effects of opposite signs.

Figure 13: Dispersion in historical welfare gains



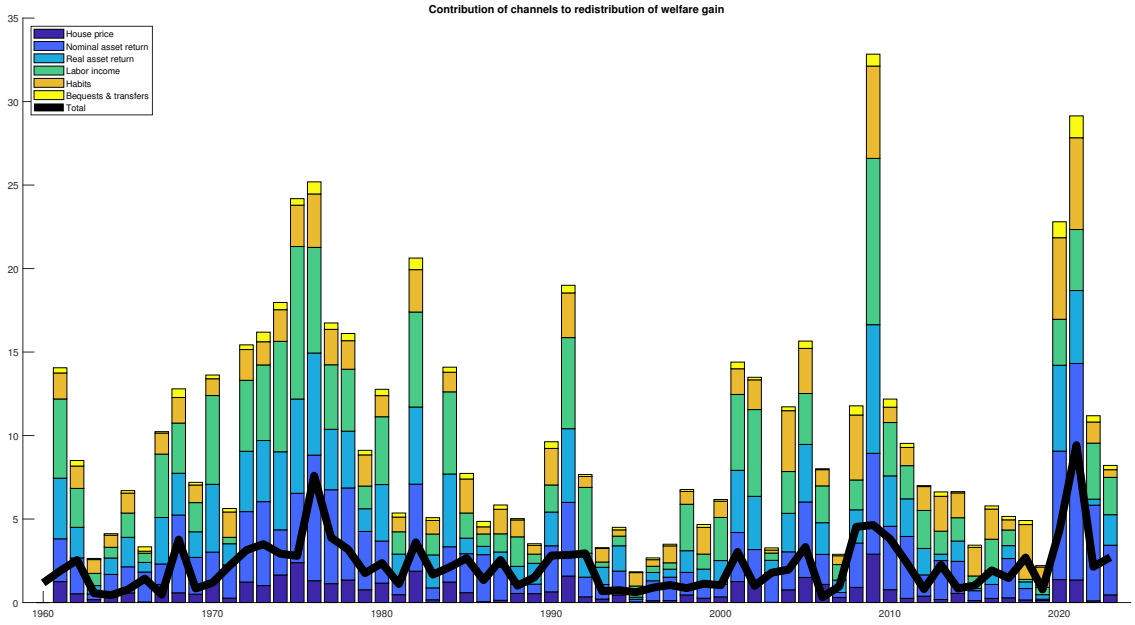
Our framework enables us to examine the sources of this unprecedented redistribution. First, we calculate the $AMOR$ statistic to gauge the overall amount of welfare redistribution for each year. We then focus on the effects of specific shocks and channels, to figure out the most important sources of welfare differences. Focusing on the years since 2020, we infer from Figure 14 that the key shocks driving the redistribution during and after the pandemic were price markup, TFP and monetary policy shocks, while the transmission of shocks was mostly relayed through the nominal assets channel (see Figure 15), although the other asset channels together with labor market channel also played a non-trivial role.

Figure 14: Overall redistribution and contribution of specific shocks



Note: Due to the absolute value transformation, the effects of specific shocks partially cancel out and they do not aggregate up additively to the overall $AMOR$.

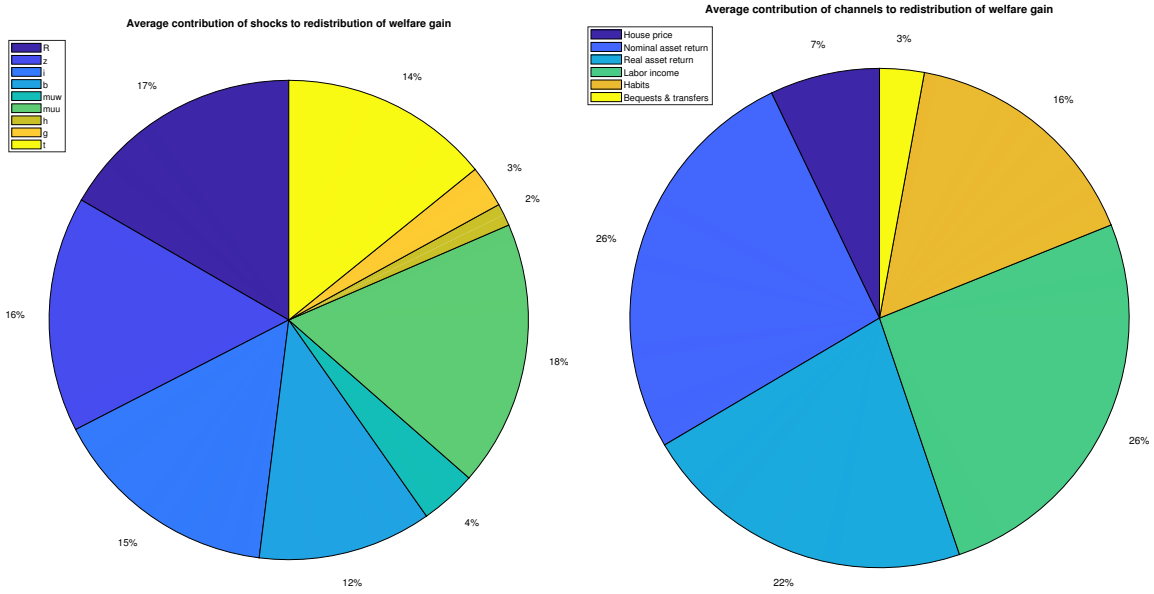
Figure 15: Overall redistribution and contribution of specific channels



Note: Due to the absolute value transformation, the effects along specific channels partially cancel out and they do not aggregate up additively to the overall $AMOR$.

Zooming out to encompass the entirety of our sample we assess the relative importance of shocks and channels toward the redistribution (Figure 16) We find that monetary policy shocks, although trivial in historical decomposition of e.g. output fluctuations, plays much more important role in driving the dispersion in household welfare over the business cycles. Other shocks that contribute much to welfare redistribution are: wage markup (active vs inactive in the labor market), TFP (similar), transfers (mostly toward retirees), investment-specific and bond-preference (both affecting rates of return on real and nominal assets, respectively). Almost 50% of overall redistribution occurs along the nominal and real asset return channels, the share increasing to almost 75% when inclusing the labor income channel as well. In contrast, the role of the house price channel is relatively minor.

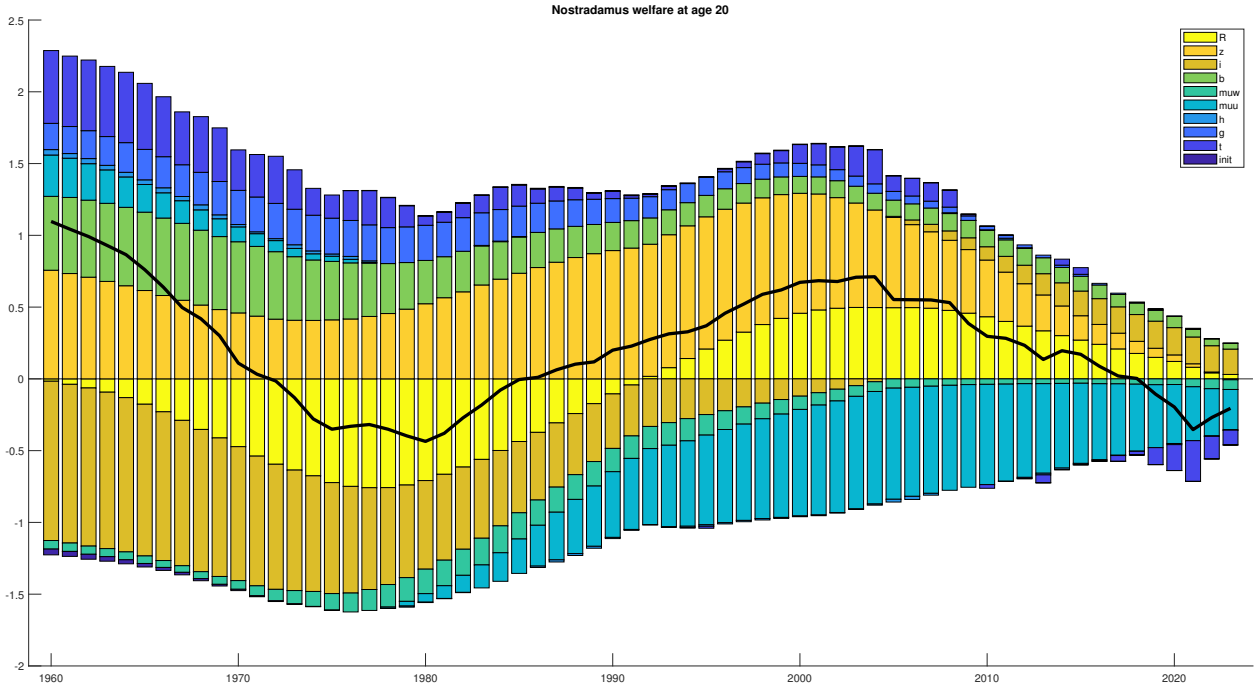
Figure 16: Contribution of specific shocks and channels to redistribution



4.5 Lucky "boomers" and doomed "gen-Zs"?

We can now assess the overall "luck" that a given cohort experienced throughout its life cycle. We are naturally limited by the available data, so we assume no new shocks occurring after 2023 and trace the path of endogenous variables as they return to their steady states. We subject our cohorts to a "Nostradamus" exercise in which we provide them with the realized (or post-2023, forecasted) allocations that they experienced throughout their lives, and ask them to evaluate their welfare from the 20-year old perspective (without the ability to change said allocations). This allows us to calculate the measure of an ex-ante lifetime welfare deviation from steady state under perfect foresight for agents entering model since 1960. The results in Figure 17 suggest, that (with perfect foresight and hindsight in hand), cohorts that entered the labor market in the 1960s, 1990s and 2000s drew a luckier lot in life compared to their 1970s and 2020s peers. Importantly, the effects of various shocks evidently do not net out across the lifecycle, and some cohorts effectively enjoy up to one extra year of consumption (or lose up to 50% of a year of consumption) throughout their lives.

Figure 17: Nostradamus welfare at age 20



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. Under this condition, we have shown that monetary shocks have the ability to relatively strongly redistribute welfare across age cohorts. Now we concentrate on the systematic component of monetary policy, i.e. its predictable reaction to economic developments. Does this component of policy affect the way welfare is redistributed over the business cycle?

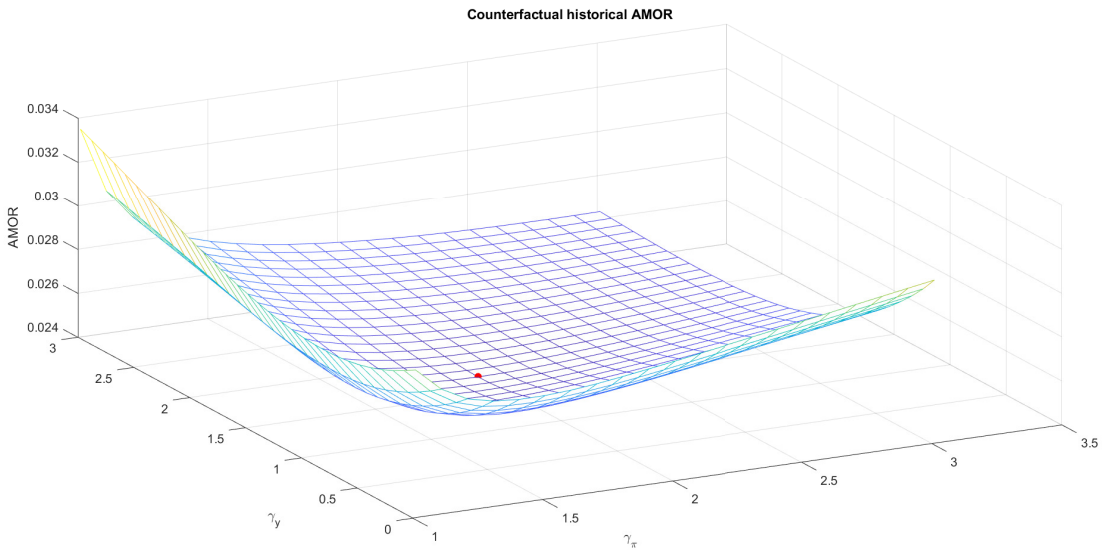
There are many dimensions along which the effectiveness of monetary policy can be measured in this context. For consistency with the rest of the paper, we concentrate on our measure of the average amount of welfare that is being redistributed defined in equation 13. In our description we are purely positive, and abstract away from normative questions, like which policy is desirable. After all, the goal of monetary policy is about macroeconomic stability, not about redistribution.

We do the following simulations. We specify a grid for monetary policy parameters γ_π and γ_y and simulate our model economy. For each simulation, we calculate the average amout of redistributed welfare. This is done in two variants: historical (using the historical filtered shocks) and stochastic (using the estimated shock distributions). In both cases, we drop monetary policy shocks as their interpretation becomes dubious once parameters of the policy rule are being changed. However, their inclusion does not affect our conclusions significantly.

Our findings for the historical case are presented on Figure 18. Redistribution is highest when either the reaction to inflation or to output is very weak. The highest $AMOR$ is almost 3.5%, though this is clearly determined by the grid cut-off, and not an internal maximum. More interestingly, we find an internal minimum, which happens to be 2.5% for $\gamma_\pi = 1.8$ and $\gamma_y = 1.3$. It should be noted that the reaction to inflation is not very far from our estimated value (1.53), though the reaction to output is much stronger (we estimated 0.15).

Nevertheless, the average amount of redistribution for the estimated parameters (2.8%) would not have been dramatically higher than the minimized one.[5]

Figure 18: Amount of redistribution conditional on monetary policy parameters (historical shocks)



Note: Red dot denotes the minimum

The results for the stochastic case are quite similar, so we do not report them in detail (there is an internal minimum as well, for $\gamma_\pi = 1.8$ and $\gamma_y = 1.6$). One, final observation is that, while both cases feature an internal minimum, from a quantitative perspective stronger reactions to both inflation and output are almost equally effective in reducing $AMOR$, which rises only slightly for higher values of γ_π and γ_y.

How does systematic monetary policy affect the redistribution, and why do we find an internal minimum? The problem is best understood by recalling the trade-offs faced by monetary policy. When cyclical fluctuations are (at least to some degree) driven by supply side shocks (e.g. markup, technology), monetary policy faces a well-known trade-off between stabilizing inflation and output. At the same time, inflation and output fluctuations are responsible for redistributing welfare. The most pronounced effects are related (though

[5]Note that a direct comparison with the historical AMOR (3.4%) would be misleading as it is affected by monetary policy shocks as well. The amount reported above (2.8%) is the counterfactual at the estimated policy parameters, but without monetary policy shocks.

not limited) to the role of nominal assets (surprise inflation redistributes between those who hold long and short positions) and the role of labor income (higher GDP growth boosts labor income of workers, but not of retirees). If the central bank puts more weight on stabilising inflation, it reduces the redistribution via the inflation channel. However, then volatility of output increases, raising the redistribution via the labor income channel. At some point the latter effect takes over and redistribution increases. A similar reasoning applies to increasing the weight put on output. Its higher value reduces the redistribution via the labor channel, but leads to more volatile inflation and related redistribution. At some point the latter effect prevails.

6 Conclusions

What are the distributional consequences of business cycle fluctuations? How do various shocks drive redistribution? what are the main drivers? 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 construct a life-cycle model with nominal and real rigidities, estimate it on US data and use it to run several counterfactual experiment.

What do we find? First and foremost, the distributional consequences of cyclical fluctuations do not cancel out over typical lifetimes. This is not only a trivial consequence of the finite horizon (so that the average shock need not be zero), but also the effect of a feature typical for life-cycle (but not e.g. for HANK) models. The same shock may have very different consequences for agents of different age. As a result even two shocks which perfectly cancel out may have non-zero welfare effects if they hit the agent at different stages of life.

We also show that different macroeconomic shocks have different redistributive power. For instance markup shocks affect all cohorts in the same direction, while financial (risk

premium) and monetary policy shocks leve some cohorts better, and others worse-off. The main driving forces of redistribution are price markup (i.e. cost) shocks, technology shocks and monetary policy shocks. While the former are well known to also drive cyclical (e.g. output) fluctuations, the large role of monetary policy shocks is an important finding, as their role in driving business cycles is rather small. This effect follows from their strong redistributive power.

Then, we check via which channels redistribution occurs. It proves to work its way primarily via returns on nominal assets and via labor income. Fourth, we document how cyclical fluctuations redistributed welfare historically. Typically fluctuations affected all cohorts in the same direction (booms raised while recessions lowered everyone's welfare).

Last but not least, the role of systematic monetary policy is analyzed. We show that it can contribute to dampening redistribution. Interestingly there exists an interior solution for Taylor rule parameters which minimizes the amount of redistribution. The reaction to inflation is similar to that found in empirical studies of Taylor rules. However, the reaction to output that minimizes redistribution is much stronger than usually encountered.

References

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

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

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

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

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

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

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

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

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

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 2021â€"2022 shock.' *ECB Statistics Paper Series, No 18/2016*

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