



EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Wouter J. Den Haan, Alessandro Ferrari,
Falk Mazelis, Annukka Ristiniemi

Battle of the ages: distributional and
aggregate effects of monetary policy
in a model with age demographics

No 3125

Abstract

We develop a model in which agents face unemployment risk, but also age and eventually retire. We study the impact of different retirement schemes on life-cycle consumption and the monetary transmission mechanism. Agents save because of a fall in income upon retirement, changes along the life-cycle wage profile, and unemployment risk. Changes in retirement policies affect the distribution of available assets (bonds) among the middle aged and the young, which in turn can have a strong impact on the ability of the young to insure themselves against unemployment risk. Interestingly, it is possible that an increase in retirement benefits leads to *higher* consumption levels during sustained unemployment spells even though the associated increase in taxes reduces unemployment benefits. The reason is that this policy induces the middle aged to save less which leaves more of the available asset supply to the young. A reduction in the interest rate has a bigger impact on those for whom labor market conditions improve the most and – due to a larger negative income effect – has a smaller impact on those who save more. In terms of the aggregate impact of monetary-policy shocks, our paper confirms conventional wisdom that the expansion is magnified in the presence of incomplete markets, since it is then accompanied by a fall in precautionary savings. The novel aspect of our analysis is that the extent of the incompleteness, i.e., the ability of those subject to unemployment risk to insure themselves, is endogenous. Specifically, it is reduced as the young (middle-aged) hold a larger (smaller) fraction of the available asset supply and this distribution is not only affected by retirement policies, but also by government bond supply and the life-cycle wage profile. Thus, understanding the distribution of assets across different age cohorts is not only important for understanding life-cycle consumption patterns, but also business cycles.

Keywords: New-Keynesian model, aging, monetary-policy shocks, unemployment risk, precautionary savings.

JEL Classification: E43, E52, E21, E24.

Non-technical summary

The modern macroeconomic literature has documented the importance of taking into account heterogeneity across different groups in society when studying the impact of monetary policy on the economy. This is important because monetary policy affects different people differently and because the presence of heterogeneity affects aggregate outcomes. One type of heterogeneity, however, has rarely been studied and that is age heterogeneity and in particular the presence of retirees and the need to prepare for retirement. This is especially important now that we observe an aging population and governments adapt pension systems. Before retirement systems are altered, it is important to understand how these changes interact with monetary policy.

This study shows that the consequences of changes in the policy rates by central banks are felt very differently by young job seekers, by young employed workers, and by workers who are closer to retirement. One area of concern is whether lowering policy rates entails a trade-off. That is, is it true that younger workers benefit more from the expansion in real activity and the increased job-finding probability, while older, wealthier savers are disadvantaged because the reduction in the return on their savings which could very well have a nontrivial impact on available funds during retirement. And are these outcomes affected by available transfers to the unemployed and retirees. This paper sheds light on how such demographic and policy factors affect both individual behavior and the broader economy. The paper explores this issue by developing a model that reflects key realities faced by households at different stages of life: youth, middle age, and retirement. Importantly, the model also accounts for the risk of becoming unemployed, which turns out to be an important factor in how different retirement systems affect the impact of monetary policy on the economy.

The paper uses a standard New-Keynesian framework, but modifies it to include aging and employment risk. The model also incorporates features such as retirement (benefits), unemployment (benefits), firm ownership, government bonds, capital investments, and an annuity available to retirees. It is calibrated using Euro-Area data and produces outcomes at a quarterly frequency, which means it is suitable to study the business-cycle impact of monetary contractions and expansions. This design makes it possible to study how different age groups respond to monetary policy changes *and* how (i) variations in retirement benefits, (ii) an increase in the supply of government bonds, (iii) a relaxation in the ability to borrow, and (iv) a steepening of the life-cycle wage profile influence both distributional and aggregate economic outcomes.

The main results are as follows. First, consumption of the middle aged increases substantially less than consumption of the young during a monetary expansion. Although, it is a theoretical possibility that consumption levels of savers decrease when interest rates falls, we find that this only happens for unusual parameter values for the average middle-aged person. Second, we find that consumption of the newly retired does robustly fall when the central bank reduces interest rates although the fall is always small (less than 0.3%). Third, the generosity of the pension system affects both the distributional impact of monetary policy *and* the impact on overall economic activity. There are two reasons for this. When pension benefits are less generous, then the middle aged save more which means that a smaller fraction of the supply of the safe asset is available to those who are subject to unemployment risk. Consistent with the literature, this increase in the inability to insure against unemployment risk amplifies business cycles including those induced by changes in the central bank's policy rate. This "battle" among different age groups on who gets to hold the available supply of the safe asset turns out to be important for several policy experiments considered. Another reason why the pension system matters is that it affects the level of the equilibrium real interest rate. The reason for this is that consumption choices depend on the Net Present Value of future income which is a nonlinear function of the interest rate.

The paper's findings indicate that demographic structure and pension policy not only affect the distributional impact of monetary policy but also the aggregate impact. For the ECB and other central banks, this means that the age and employment distribution of asset holders should be monitored in order to allow an evaluation of the likely impact of interest rate changes. Overall, this research contributes to a more nuanced understanding of monetary transmission and underscores the importance of integrating demographic and distributional factors into macroeconomic policy design.

1 Introduction

Following the development of advanced numerical methods in the nineties, a surge of papers has appeared that build business-cycle models with heterogeneous agents. One motivation is to understand whether there are differences in how groups in society are affected by business cycles. Another important motivation is that heterogeneity may very well affect aggregate outcomes.

One type of heterogeneity is understudied and that is age heterogeneity.¹ There are good reasons, however, to expect agents of different age are affected quite differently by business cycle fluctuations. For example, consider an expansion induced by a reduction in interest rates. Standard textbook reasoning tells us that this lowers the price of current compared to future consumption and this substitution effect stimulates current demand. But we know that there is also a negative (positive) income effect for savers (borrowers). And we can expect the negative income effect to become stronger as agents age and have had more time build up wealth.²

In theory, a sustained reduction in the interest rate – on its own – could induce middle-aged savers to *reduce* current spending and by doing so limit the negative impact of the low interest rate on available funds during retirement. But how important are these negative effects during typical monetary expansions? Specifically, could they be big enough to overturn the positive impact on these agents' current spending due to the associated increase in real activity? These are key questions that we address in this paper.

We also investigate whether age-related differential responses have aggregate effects? If a stronger negative income effect of the middle aged, i.e., those with large savings, is offset by a weaker negative income effect of the young, i.e., those with smaller savings or even debt, then this type of heterogeneity would not affect the aggregate outcome, at least not through this channel. However, these two effects may not cancel each other out. The reason is the following. Consider the comparison of an economy with a high level of pension benefits with one with lower benefits. The presence of lower pension benefits would induce an increase in desired savings for both the young and the middle aged as they aim to fill the gap through private pension savings. But the effect is stronger for the middle aged as they are closer to retirement. Consequently, there is a redistribution of wealth from the young to the middle aged which

¹Notable exceptions are [Bielecki et al. \(2022\)](#), [Bardóczy and Velásquez-Giraldo \(2024\)](#), and [Bielecki et al. \(2023\)](#).

²Such an adverse negative effect of low real interest rates on savers has received attention by policy makers (and the public) during the *sustained* period of low interest rates following the financial crisis. Instead, we focus on typical monetary expansions during which the central bank is able to affect the real interest rate temporarily due to the presence of sticky prices.

in turn implies that the negative income effect for the middle aged strengthens and that for the young falls. Do these effects cancel out? Recall that what matters for the income effect is how the change in the interest rate affects *lifetime* income. And as the young are expected to earn income for longer, their lifetime income is relatively less affected. This suggests that the re-distribution would increase the *average* negative income effect and dampen the aggregate impact of a monetary policy shock.

We show, however, that the opposite occurs. That is, the aggregate impact of monetary policy increases when retirement benefits are lower. The reason is that the increase in savings induced by a reduction in pension benefits goes together with a reduction in the steady-state real interest rate. And it turns out to be the case that consumption demand responds more strongly to changes in the interest rate at lower steady-state real interest rates, especially for the young.³ Since this effect is quite strong, the aggregate impact of monetary policy is stronger when retirement benefits are lower.⁴

There is another mechanism through which age heterogeneity matters for aggregate activity. The key element of that channel is how age heterogeneity interacts with the ability of workers to insure themselves against unemployment risk. It is well known that the latter magnifies business cycles.⁵ The reason for this is the following. When markets are incomplete, agents save to partially insure themselves against unemployment risk. During an expansion, unemployment risk falls, which leads to a reduction in this type of precautionary savings which stimulates demand, which in turn stimulates real activity.

The ability of those who are subject to unemployment risk to insure themselves against unemployment risk depends crucially on who else is competing for the supply of available assets. Age heterogeneity matters because agents also save for retirement and this is more important for those closer to retirement, i.e., the middle aged. Consequently, the inability of workers to insure themselves, and thus the impact of shocks on aggregate activity, is endogenous because changes in the economic environment

³We show this both with a simple analytical example and in the full model.

⁴Interestingly, both mechanisms provide a meaningful role for the real-interest rate in the monetary-transmission mechanism of the New-Keynesian model. [Rupert and Sustek \(2019\)](#) convincingly show that a monetary expansion in the New-Keynesian model is not driven by the traditional interest-rate channel. That is, an expansionary monetary-policy shock in the Taylor rule does lead to an expansion in real activity and increased inflationary pressure, but the real interest rate could either fall or increase. A simplified explanation is that key for the monetary transmission mechanism in the New-Keynesian model is the Phillips curve and the interest rate has no direct impact on this relationship. The distributional shifts described here operate like an aggregate preference shock affecting the intertemporal marginal rate of substitution, which we know affects aggregate activity.

⁵See [Kaplan et al. \(2018\)](#) on how the inability to insure against income shocks affects model properties. See [Den Haan et al. \(2017\)](#), [Alves et al. \(2020\)](#) and [Auclert et al. \(2020\)](#) for a quantitative analysis of the channels involved in the amplification of the business cycle. For an empirical analysis of the role played by this friction in amplifying the Great Financial Crisis see [Challe et al. \(2017\)](#). For a theoretical explanation in a tractable model see [Bilbiie \(2020\)](#) and its empirical counterpart [Bilbiie et al. \(2022\)](#).

affect who gets to hold most of the “insurance vehicle” in equilibrium. In our benchmark environment, the supply of available assets consists of government bonds and is determined by the government. As we consider a closed economy, we exclude the possibility that other countries supply assets. But another possibility is that firm owners trade assets with workers in financial markets and facilitate saving for retirement and/or unemployment spells. In this extension, workers can also hold (indirect) claims on capital and firm ownership.

We assume that only young agents are subject to unemployment risk. One important reason is that the model becomes quite challenging to solve having both age heterogeneity and employment risk. Consequently, some simplifying assumptions are necessary.⁶ But another advantage of our approach is that it allows us to much better understand the different mechanisms at play. For example, we find that the middle aged only lower their consumption during a typical monetary expansion for atypical parameter values. This indicates that the negative income effect is not that strong. This insight is more striking in a model like ours in which the middle aged do not increase spending because of a reduction in unemployment risk. That is, we give the negative income effect the best possible chance, but the positive impact of the substitution effect and beneficial aspects of the expansion such as higher wages dominate.

The following example illustrates the previous discussion. Consider an alternative economy with a higher level of retirement benefits financed by taxes. In equilibrium, the middle aged would need to save less privately, which in turn frees up part of the available bond supply to those who save to smooth consumption during unemployment spells. In other words, market incompleteness is reduced. But it does not only reduce consumption declines during unemployment spells, it also has a strong negative impact on the size of aggregate fluctuations.⁷

Workers in our model start out as young, after which they age and become middle aged, followed by retirement and then death. The three life phases are expected to last twenty years on average, with transitions from one age cohort to the next occurring randomly according to a fixed probability. Retirees do not work, but they receive retirement benefits, which are a fraction of previous wage income. Saving for retirement is important to limit the consumption drop upon retirement. A drop in the interest rate that is expected to persist lowers available resources during the retirement phase. In principle, this could

⁶Even with this simplification, we end up with a Dynare program of 830 equations.

⁷Since the increase in retirement benefits is financed by taxes, after-tax wage income are lower which in turn implies a fall in *average* consumption. Moreover, the calibration procedure is such that steady-state unemployment risk remains the same. Interestingly, we show that the ability to limit consumption drops during unemployment spells can improve by so much that consumption levels during extended unemployment spells actually increase even though the increase in taxes is assumed to lead to lower unemployment benefits in line with declining after-tax wages.

lead middle aged workers to lower consumption during a monetary expansion and increase saving for retirement. As mentioned above, we find that this is an unusual outcome for typical monetary expansions.

Given the importance of unemployment risk for setups like ours, we need to incorporate a serious model for unemployment into a New-Keynesian framework. To accomplish this, we build a model with two sectors, an intermediate-goods and a final-goods sector. The intermediate-goods sector uses labor to produce inputs for the final-goods sector. To model unemployment in this input-producing sector, we use a standard labor-market matching framework which introduces unemployment in a tractable way. The other sector combines the input good and capital to produce consumption and investment goods. This sector features standard New-Keynesian elements such as sticky prices.

We use this framework to study how key changes in the economic environment affect (i) savings and consumption over the life cycle and during unemployment spells, (ii) the distributional effects following monetary-policy shocks, and (iii) the aggregate impact of monetary-policy shocks.

We consider four different types of exercises. In the first, we consider the impact of a change in the level of tax-financed pension benefits. We also compare retirement benefits in the form of home production versus tax-financed benefits. Second, we consider changes in the supply of the available asset to smooth consumption during unemployment spells and upon retirement. In a model without aging, this would unequivocally improve the ability to insure against unemployment risk and reduce the severity of business cycles. In our framework, the question is whether the additional assets will be held by the young or the middle aged. We compare this model variation with one in which the borrowing constraint is relaxed. The two experiments described so far are analyzed in an environment in which worker productivity, and thus, the wage rate, does not increase with age. In reality, wages increase with age. But this introduces another savings motive making the model more complex. In the third experiment, we do analyze model properties when the middle aged earn a higher wage than the young. The fourth experiment also enriches the model. So far, saving to cover unemployment spells and to safeguard consumption during retirement could only be done through holding government bonds. Another possibility is that workers lend to or borrow from those who own the firms and the capital stock. In terms of steady-state properties, this possibility is similar to a change in the supply of government bonds. However, allowing for this interaction between different types of agents does turn out to be important during expansions and recessions. Thus, the idea of our setup is that we introduce additional features one by one. This has been helpful in terms of bringing to the surface and understanding relevant

mechanisms.

For all four experiments, the lesson is that knowing who holds the available supply of the asset is essential. When those who are subject to unemployment risk hold a larger fraction of the available asset, then (steady-state) consumption drops during unemployment spells are smaller, the impact of a monetary policy shock on consumption is more similar across agents, and the aggregate impact is smaller. In some sense, this result can be interpreted as a refined version of the point made in [Aiyagari and McGrattan \(1998\)](#). They point out that an increase in government bonds reduces market incompleteness. That is true in our environment as well, but with the caveat that the increased supply should end up in the hands of those who save to smooth consumption during unemployment spells.

Related literature. There are only a few papers that study the impact of monetary-policy shocks on the economy in the presence of age heterogeneity. [Bielecki et al. \(2022\)](#) and [Bardóczy and Velásquez-Giraldo \(2024\)](#) develop rich overlapping-generations (OLG) annual models to document the age-related heterogeneous response following monetary-policy shocks and they provide a detailed account for the reasons behind the different responses. We have chosen to adopt an extended perpetual-youth structure. It has the disadvantage of having a more simplified age structure. But it does allow us to reduce computational complexity and have a quarterly model, which makes the results more comparable to those in the literature dealing with monetary-policy shocks.

Key differences with the two papers are the following. First, we assume that workers only use a risk-free nominal asset to smooth consumption in response to changes in employment status and retirement.⁸ Capital and firms are owned by a separate group of agents which we refer to as firm owners. Recall that a key motivation is to explore the possibility that some agents reduce consumption during a monetary expansion and whether this would reduce the aggregate impact of monetary policy. This setup gives that possibility the best possible chance. Nevertheless, we find that the middle aged still increase their consumption, even under this “worst-case” scenario. Moreover, our framework is transparent enough to understand whether reduced impact on those with larger savings, i.e., the middle aged, cancels out against a stronger impact of those who borrow or have fewer savings. Second, our focus is different as we are interested in the question how changes in pension benefits affect distributional and aggregate responses to monetary-policy shocks.⁹ Third, we study both the impact of monetary-policy shocks *and*

⁸In addition, workers can buy an annuity upon retirement. And we allow workers to indirectly invest in capital and firm ownership in an extension.

⁹And we also consider the impact of other relevant changes such as the change in the supply of the risk-free asset. Retirees fully rely on their own savings in the model of [Bielecki et al. \(2022\)](#). [Bielecki et al. \(2023\)](#) do allow for transfers to retirees, but

the life-cycle fluctuations in consumption including consumption drops during unemployment spells for each environment. This turns out to be important because – consistent with the recent literature – the aggregate impact of monetary policy depends crucially on whether agents can insure themselves against unemployment spells, and changes related to pension benefits affect the availability of asset for those subject to employment shocks.

We have adopted a relatively simple framework and exclude features such as habits and housing. And – as mentioned above – we restrict which type of agent can own capital and firms. This has the advantage of transparency and allows us to uncover some important economic insights. Of course, it does mean that our paper is not suitable for generating accurate *quantitative* predictions.

Organization of the paper. The next two sections describe the model and the calibration. Section 4 considers how consumption responds to changes in the interest rate when those changes have no effect on economic activity. It considers some partial equilibrium models including a partial equilibrium version of our full model. Section 5 analyzes how key steady-state outcomes like the equilibrium real interest rate are affected by changes in pension benefits and other important model parameters. The last section studies how changes in those model parameters affect both the life-cycle consumption pattern and the aggregate and distributional impact of monetary-policy shocks.

2 The model

2.1 Demographics

A key objective of this paper is to study how demographics and idiosyncratic income risk interact with each other. Important lessons can already be learned when we abstract from aggregate shocks. As emphasized in the introduction, however, we are also interested in this question when the central bank changes interest rates to stimulate or cool down the economy. Understanding the consequences of unemployment risk as well as the impact of monetary policy requires a model in which the length of the time period is relatively short, as these occur mainly at business cycle frequencies. Quantitative business cycle models are typically quarterly, but introducing a realistic age distribution into a quarterly model comes with a challenge. If aging is deterministic and we include agents' economic life for $\bar{T}$ years, then we have $4 \times \bar{T}$ different age groups. Furthermore, within each cohort, we would have a distribution of different employ-

study how these affect the impact of changes in life expectancy on long-term outcomes like the natural rate of interest. Retirees in the model of [Bardóczy and Velásquez-Giraldo \(2024\)](#) do receive pension benefits, but their role in the monetary transmission mechanism is not the focus of the analysis.

ment histories. To keep heterogeneity tractable, we adopt an extended version of the perpetual-youth model of Blanchard (1985). This allows us to have realistic life horizons and agents in different age groups within a quarterly model.

In our economy, there are firm owners and workers.¹⁰ Workers begin their lives as “young” workers. While young, they face an idiosyncratic aging shock and become “middle aged.” This occurs with probability δ^y and is independent of the age of the worker. Similarly, a middle-aged worker becomes a retiree with probability δ^m at which point they can no longer supply labor. Thus, one crucial difference between the young and the middle aged is that the latter are closer to retirement at which point they would no longer receive wage income. Retirees have an idiosyncratic probability δ^r of dying, also independent of age. Deceased retirees are “replaced” by an equal number of young workers, so the population size remains constant. Age shocks occur at the beginning of the period before the realization of any other idiosyncratic shocks. The expected duration of a worker’s productive life is equal to $1/\delta^y + 1/\delta^m$ and the expected retirement period is equal to $1/\delta^r$. We model the behavior of firm owners using a representative infinitely-lived agent.¹¹

The mass of firm owners is denoted by γ^f . The mass of working-age young (middle-aged) workers is denoted by γ^y (γ^m) and the mass of retirees is indicated by γ^r . Given values for δ^y , δ^m , and δ^r , these three fractions can be solved from

$$\delta^y \gamma^y = \delta^r \gamma^r, \quad (1)$$

$$\delta^m \gamma^m = \delta^y \gamma^y, \text{ and} \quad (2)$$

$$\gamma^y + \gamma^m + \gamma^r = 1. \quad (3)$$

The first (second) equation establishes that the flow out of the pool of young (middle-aged) workers equals the inflow. The last equation indicates that we have normalized the total mass of workers to 1.

Extended perpetual-youth model and uncertainty. The assumption that all young workers have the same probability of aging, all middle-aged workers have the same probability of aging *and* retiring,

¹⁰The terminology is slightly confusing as firm owners also supply some labor as discussed below. With “workers” we refer to those who are not firm owners.

¹¹We abstract from life-cycle characteristics of firm owners, since our key objective is to study how monetary policy interacts with demographics in the presence of idiosyncratic employment shocks, which only affects workers. Moreover, some simplification is needed to keep the analysis tractable. Alternatively, one could assume that the stress of running a firm means you die in your prime before retirement and you leave your wealth to an heir whose utility is taken into account in the firm owner’s choices. By contrast, retired workers need their savings to finance their retirement.

and all retirees have the same probability of dying is extremely helpful in reducing the computational complexity. What matters for our model predictions is that the middle aged *as a group* face an expected drop in income and this affects their savings motives. In reality, there is heterogeneity with some older middle-aged workers for whom the drop in income will become sooner than for younger ones. This heterogeneity also introduces a nonlinearity in the problem which would affect *average* savings of the middle aged, just as uncertainty does. If there is a quantitative difference, then this may be another channel that is relevant for how retirement affects monetary policy. But we don't think it affects the logic of the mechanisms that we put forward in this paper.

Thus, these changes in a person's life cycle become stochastic. Specifically, retirement is often predictable. Moreover, workers have some control on when to retire.

2.2 Workers

In this section, we first discuss unemployment risk and how agents can partially insure themselves. Next, we discuss the economic environment for retirees. With these elements spelled out, we can describe the behavior of the young and middle aged.

2.2.1 Employment risk and limited insurance

Young workers are either employed (e) or unemployed (u). An employed worker remains employed with probability Ω^{ee} and an unemployed worker remains unemployed with probability Ω_t^{uu} . The former is constant, but the latter is endogenous and depends on conditions in the labor market. Both probabilities are taken as given by the workers. The distribution of the employment status among newly born is the same as that of other young workers.¹²

We assume that only young workers face the risk of unemployment. Their savings decision is, thus, affected by a precautionary savings component. The savings decision of the middle aged is also affected by a precautionary element, as investment returns fluctuate with economic conditions, which determines retirement income. The advantage of our approach is that we have one group of agents that is concerned about employment risk (the young) and one group of agents that is concerned about the return on their retirement savings (the middle aged). This makes it easier to analyze the roles of these economic phenomena for model properties like the distributional consequences of monetary policy.

¹²That is, some workers are born unemployed and they try to find a job in the matching market, whereas others are born employed.

Timing. At the beginning of the period, young (middle-aged) workers find out whether they have aged (retired). Immediately after that, the current employment status for young workers is determined. The mass of employed young workers is given by

$$\begin{aligned}
 n_t^y &= \frac{n_{t-1}^y}{\gamma^y} \delta^r \gamma^r \\
 &+ (1 - \delta^y) \Omega^{ee} n_{t-1}^y \\
 &+ (1 - \delta^y) (1 - \Omega_t^{uu}) (\gamma^y - n_{t-1}^y) \\
 &- \delta^y n_{t-1}^y,
 \end{aligned} \tag{4}$$

where $\delta^r \gamma^r$ is the mass of newborns which is equal to the mass of “departed” retirees and $\gamma^y - n_t^y$ equals the mass of unemployed young workers. The young age and become middle aged with probability δ^y ; with probability $(1 - \delta^y)$ they remain young and their employment status is determined by the probabilities Ω^{ee} and Ω_t^{uu} as defined in section 2.4.2.

Workers cannot buy contingent claims to fully protect themselves against unemployment risk. However, they can partially insure themselves by saving in a risk-free nominal asset and if needed they can take a short position which is equivalent to borrowing. One popular approach to limit borrowing is to assume that an individual’s position in the risk-free asset has to exceed $\bar{\eta}_d$ with $\bar{\eta}_d < 0$. Instead of such an inequality borrowing constraint, we assume that a variable premium limits the ability to take a short position in the risk-free asset, but does so in a more gradual manner than an inequality constraint.

Specifically, the interest rate on the risk-free asset is the sum of two components. The first is R_t^d , which is the same for all agents, and the second is a premium that depends on the position taken in the risk-free asset. The latter is given by

$$\tilde{R}^d(\bar{d}_{j,t}) = \eta_{d,0} (\bar{d}_{j,t} - \bar{\eta}_d)^{-\eta_{d,1}}, \tag{5}$$

where $\bar{d}_{j,t}$ is the position taken in the risk-free asset.

The shape of the premium is shown in figure 1 for our baseline parameters. The figure plots the sum of the “base rate,” R_t^d , plus the premium as a function of the position in the risk-free asset. The interest rate is fairly constant for positive deposit levels, but increases rapidly as the agents’ position approaches $\bar{\eta}_d$. This setup prevents agents from borrowing more than $-\bar{\eta}_d > 0$, since the premium they have to pay

goes to infinity as the short position approaches $\bar{\eta}_d$.¹³

To simplify first-order conditions (FOCs), we assume that the premium depends on the *average* amount saved/borrowed by agents of type j , $\bar{d}_{j,t}$, where type includes employment history.¹⁴ Hence, agents take the premium as given. It is still the case that borrowing is restricted when workers become unemployed. The reason is that other agents of the same type behave in the same way and as these agents borrow more to smooth consumption the interest rate they face increases.¹⁵

Financial intermediation implies positive profits, since the rate paid by borrowers exceeds the rate received by savers. Since the supply of the risk-free asset is provided by the government, we simply incorporate these intermediation profits into the government budget constraint, which means that they will return to households in the form of lower tax rates.¹⁶

2.2.2 Retirees

Retirees' income consists of government transfers and/or home production as well as the return on accumulated savings.¹⁷ The beginning-of-period wealth of worker j who retires in period $t = \bar{t}$, $\hat{x}_{\bar{t},j}$, is given by

$$\hat{x}_{\bar{t},j} = (r_t^d + \tilde{r}_{j,\bar{t}}^d(\bar{d}_{j,\bar{t}-1}))d_{j,\bar{t}-1}, \quad (6)$$

where r_t^d and $\tilde{r}_{j,t}^d$ are the real rates associated with R_t^d and $\tilde{R}_{j,t}^d$. This includes interest received at the beginning of period $\bar{t}$. This amount is used to buy annuities that cost q_t and pay out $1 - \delta^r$ each period. Thus, during their remaining retirement the retiree will receive $((1 - \delta^r)/q_{\bar{t}})\hat{x}_{\bar{t},j}$. Workers who retire in

¹³This specification implements the standard inequality constraint that $\bar{d}_{j,t} \geq \bar{\eta}_d$ when $\eta_{d,1}$ goes to infinity. This standard inequality constraint assumes that borrowing costs are constant for borrowing *any* amount less than $-\bar{\eta}_d$ and then jump discontinuously to infinity if the agents borrow just an epsilon more than $-\bar{\eta}_d$. Our approach not only seems more realistic, but also has computational advances by avoiding the discontinuity associated with the inequality constraint. As shown in [Den Haan and De Wind \(2012\)](#), the implied behavior of model variables with this type of smooth penalty function is similar to that generated with inequality borrowing constraints. Moreover, our simpler specification would actually be the correct one if lenders can observe the type of borrower, but cannot observe the *total* amount borrowed, for example, because borrowers can go to different lenders and by doing so hide the total amount borrowed.


¹⁴That is, we use $\bar{d}_{j,t}$ instead of $d_{j,t}$ to indicate it is an average across agents of the same type as agent j . Agents are of the same type if they have the same age and the same employment history. Agents of the same type would, of course, choose to hold the same amount of the risk-free asset. The essential aspect is that agents take this average – and thus the corresponding premium – as given.

¹⁵We have also considered the alternative where the premium does depend on the amount chosen by the individual agent. This does not seem to matter for the model properties but complicates the model equations.

¹⁶The alternative would be to return these profits to households in the form of dividends. This would be the same as long as the distributional impact is the same.

¹⁷For the retiree, it does not matter whether additional resources come in the form of a transfer or home production since home production (minus possible disutility) is assumed to be a perfect substitute for consumption produced in the market. They differ, however, because transfers have to be financed by the rest of the population whereas home production does not.

Figure 1: Annual interest rate received/paid on positions in the risk-free asset



Notes: The dotted horizontal line represents the base rate which is independent of the amount saved/borrowed. The solid line represents the total rate which includes the premium. A negative value for the position taken means that the agent is borrowing. The figure is based on our baseline parameter values. That is, $\eta_{d,0} = 3.9015e-04$, $\eta_{d,1} = 1$, and $\bar{\eta}_d = -0.3$.

different time periods receive different retirement benefits because q_t and $\hat{x}_{j,t}$ vary over time in response to aggregate conditions, but the benefits are constant for the remainder of a retiree's life.

The other source of income for retirees is μ_t^r which consists of a government transfer and/or home production. This term is also assumed to be constant during the life of the retiree but may vary across cohorts. Thus, a retiree's consumption level only depends on date- $\bar{t}$ variables and is equal to

$$c_{j,t}^r = (1 - \delta^r) \frac{\hat{x}_{j,\bar{t}}}{q_{\bar{t}}} + \mu_{j,\bar{t}}^r \quad (7)$$

for $\bar{t} \geq t$ until death.

Since retirement consumption is constant, it is possible to give an analytical expression for the value function of a retiree. Specifically, the value function of a worker retiring in period $\bar{t}$ is given by

$$V_{j,\bar{t}}^r = \frac{(c_{j,\bar{t}}^r)^{1-\phi} - 1}{1 - \beta^r(1 - \delta^r)}. \quad (8)$$

Having an analytical expression for the value function of the retiree greatly facilitates the analysis

because the first-order condition for savings by a middle-aged agent contains the derivative of this value function with respect to wealth brought into retirement.

2.2.3 Young and middle-aged workers

All agents face idiosyncratic ageing and aggregate shocks. Young agents also face idiosyncratic employment risk. Markets are incomplete and the only asset available to workers is a one-period liquid asset with a return that is risk-free in nominal terms. As discussed in section 2.2.1, short positions are limited by an interest rate premium that increases with the size of the short position.

A young household chooses consumption $c_{j,t}^y$, and risk free assets $d_{j,t}^y$ to solve the following problem:

$$\begin{aligned}
 & V^y(d_{j,t-1}^y, I_{j,t}^e; s_t) = \\
 & \max_{c_{j,t}^y, d_{j,t}^y} \left(\frac{(c_{j,t}^y)^{1-\phi}}{1-\phi} + \beta \mathbb{E}_t \left[\begin{aligned} & (1 - \delta^y) V^y(d_{j,t+1}^y, I_{j,t+1}^e; s_{t+1}) \\ & + \delta^y V^m(d_{j,t+1}^y; s_{t+1}) \end{aligned} \right] \right) \\
 & \text{s.t.} \\
 & c_{j,t}^y + d_{j,t}^y \\
 & = I_{j,t}^e (1 - \tau_t^y) w_t^y + (1 - I_{j,t}^e) \mu_t^y + \left(\frac{R_{t-1}^d + \tilde{R}^d(\bar{d}_{j,t-1}^y)}{\pi_t} \right) d_{j,t-1}^y,
 \end{aligned}$$

where $d_{j,t}^y$ denotes end-of-period t holdings of the risk-free asset, $c_{j,t}^y$ consumption, R_t^d the nominal base rate, π_t the gross inflation rate, w_t^y the wage rate, μ_t^y the level of unemployment benefits, τ_t^y the tax rate on wage income, $\bar{d}_{j,t}^y$ the amount of the risk-free asset held by agents of type j , s_t a vector that contains all aggregate state variables, and $I_{j,t}^e$ an indicator that is equal to 1 if the worker is employed and 0 otherwise.¹⁸ Whereas the nominal base rate on the risk-free asset chosen in period t , i.e., R_t^d , and the premiums, $\tilde{R}^d(\bar{d}_{j,t}^y)$ are known in period t , the corresponding real rates are not.

The optimization problem of the middle aged has the same structure except that the employment status is no longer a state variable, since the middle aged are always employed. Now, next period's value functions are V^m if they remain middle aged and V^r if they retire.

¹⁸The vector s_t includes the aggregate disturbances as well as the cross-sectional joint distribution of wealth, age, and employment status.

First-order conditions for the young and middle-aged workers. The associated FOCs are given by the two budget constraints and,¹⁹

$$(c_{j,t}^y)^{-\phi} = \lambda_{j,t}^y, \quad (10)$$

$$\lambda_{j,t}^y = \beta \mathbb{E}_t \left[\begin{aligned} &(1 - \delta^y) \left(\frac{(R_t^d + \tilde{R}^d(\bar{d}_{j,t}))}{\pi_{t+1}} \right) \lambda_{j,t+1}^y \\ &+ \delta^y \left(\frac{(R_t^d + \tilde{R}^d(\bar{d}_{j,t}))}{\pi_{t+1}} \right) \lambda_{j,t+1}^m \end{aligned} \right], \quad (11)$$

$$(c_{j,t}^m)^{-\phi} = \lambda_{j,t}^m, \quad (12)$$

$$\lambda_{j,t}^m = \beta \mathbb{E}_t \left[\begin{aligned} &(1 - \delta^m) \left(\frac{(R_t^d + \tilde{R}^d(\bar{d}_{j,t}))}{\pi_{t+1}} \right) \lambda_{j,t+1}^m \\ &+ \delta^m \left(\frac{(R_t^d + \tilde{R}^d(\bar{d}_{j,t}))}{\pi_{t+1}} \right) \frac{1 - \delta^r}{q_{t+1}} \frac{(c_{j,t+1}^r)^{-\phi}}{1 - \beta^r(1 - \delta^r)} \end{aligned} \right]. \quad (13)$$

The expectation is over aggregate shocks and (for the young) over employment status outcomes. The last FOC uses the analytical expression for the value function of the retirees derived above.

Wealth of the newborn. Heterogeneous-agent models are complex and obtaining a numerical solution requires some approximation. The key approximation in our numerical algorithm is that for each agent we keep track of the employment history for a finite number of periods. To limit the computational burden, we actually assume that newborns start with an employment status and wealth according to a distribution that is equal to those of the other young workers.²⁰

¹⁹As shown in section 2.5.1, the steady-state value of q_t turns out to be equal to $(1 - \delta^r) \frac{(1 + r^d)}{(r^d + \delta^r)}$ which means that the steady state value of

$$\frac{1 - \delta^r}{q_{t+1}} \frac{1}{(1 - \beta^r(1 - \delta^r))} \quad (9)$$

is equal to 1 when β^r times the steady-state value of the gross real interest rate is equal to 1. Thus, the two terms on the right-hand side of equation (13) would be completely symmetric for this parameter choice; that is, a transition probability times a gross return times a marginal utility of consumption. However, it would be a special choice for β^r . The value of the real interest rate depends on the precautionary saving motives, the penalty-premium, and the discount rates of the young and the middle aged. It would be quite specific if the discount rate for retirees is such that it is equal to this endogenous variable.

²⁰This is financed by a transfer of the other young agents. It is small since there are not that many newborns. So, although our model can be used to think of the transition from “young” to “middle aged” to “retired,” it cannot capture the transition from a newborn to a typical young worker, since our model assumes that the newly born start out in life resembling other young workers who have been “around” for a while.

2.3 Firm owners

Firm owners are entitled to firm profits.²¹ They also own the capital stock and they supply labor.²² To keep the analysis tractable, we model firm owners' behavior using a representative agent. This seems reasonable since these are the wealthier members of society and are less likely to be affected by binding borrowing constraints that prevent consumption smoothing. We also abstract from life-cycle aspects for this type of agent and assume that this is an infinitely-lived agent.²³

Although we assume that firm owners are not subject to idiosyncratic unemployment risk, we set the wage rate they receive equal to the average non-investment income received by *all* workers, including the unemployed and retirees. That is, the wage rate for firm owners is equal to²⁴

$$w_t^f = \frac{n_t^y w_t^y + n_t^m w_t^m}{\gamma^y + \gamma^m + \gamma^r}. \quad (14)$$

The representative firm owner solves the following optimization problem:

$$\begin{aligned} V_t^f &= \max_{c_t^f, k_t} \frac{(c_t^f)^{1-\phi}}{1-\phi} + \beta^f \mathbb{E}_t [V_{t+1}^f] \\ \text{s.t.} & \\ c_t^f + k_t &\leq w_t^f \gamma^f + (r_t^k + (1-\delta))k_{t-1} + \frac{1}{2} \kappa_k (k_t - k_{t-1})^2 + \hat{w}_t^f, \end{aligned} \quad (15)$$

where c_t^f denotes consumption, k_t beginning-of-period t capital, r_t^k the rental rate of capital, $\hat{w}_t^f$ firm profits. As is common in the literature, we assume that adjusting the capital stock is subject to adjustment costs. Firm owners do not pay taxes. The reason is that they also do not benefit from any of the transfers

²¹Our assumption that there is a representative agent who is entitled to firm profits means there is only one marginal rate of substitution to discount firm profits, which matters, for example, for price setting. Given the rich heterogeneity on the worker side, it would complicate the analysis if workers could also own firms as in Challe et al. (2017). In section 6.4, we consider financial interaction between workers and firm owners.

²²If firm owners do not supply labor, then their consumption drops following an expansionary monetary-policy shock due to the associated fall in the markup. This is a standard outcome in New-Keynesian models. By varying the mass of firm owners – and thus the relative importance of profits to their wage income – we affect consumption of firm owners following monetary-policy shocks.

²³The main motivation is tractability. But firm owners do not face the income drop that workers experience upon retirement as they could continue to receive income from their assets. Consequently, their need for retirement savings is likely to be less important than it is for workers.

²⁴Thus, wage income of firm owners is proportional to wage income of all workers and firm owners *as a group* are not treated differently. The determination of this wage rate does not matter that much. More important is the mass of firm owners relative to the mass of workers. For example, if the mass of firm owners becomes larger relative to the mass of regular workers, then issues related to idiosyncratic unemployment risk and life-cycle elements are likely to be less important for *aggregate* dynamics, since the income and wealth of the affected group are smaller relative to the size of the aggregate economy.

that the government finances out of taxes.

The associated FOCs are given by the budget constraint and,

$$\left(c_t^f\right)^{-\phi} = \lambda_t^f, \quad (16)$$

$$(\kappa_k(k_t - k_{t-1}) + 1)\lambda_t^f = \beta \mathbb{E}_t[\lambda_{t+1}^f(r_{t+1}^k + (1 - \delta) + \kappa_k(k_{t+1} - k_t))]. \quad (17)$$

2.4 Production

There are two types of producers. Retail-goods producers produce differentiated goods using capital and a homogeneous intermediate good and sell these goods to households. The final good consists of a basket of these differentiated retail-goods, and the standard Dixit-Stiglitz aggregator is used. This final good can be used for consumption and investment. The intermediate good is produced with a linear technology using labor only. Retail-goods producers are subject to price-adjustment costs and the intermediate-good producer is subject to a matching friction when hiring workers.²⁵

2.4.1 Retail sector and final-goods production

This part of the model is standard. Retail-goods producers use capital and intermediate goods as inputs. They sell their goods to buyers in markets that are characterized by monopolistic competition. The price of the retail good i is denoted by $P_{i,t}$. Price changes are subject to a quadratic adjustment cost $\frac{\kappa_\pi}{2} \left(\frac{P_{i,t} - P_{i,t-1}}{P_{i,t-1}} \right)^2 y_t$, as in Rotemberg (1982).²⁶ In a symmetric equilibrium, all retail-goods producers are identical and aggregate output is given by²⁷

$$y_t = \bar{z} k_t^\alpha n_t^{1-\alpha}, \quad (18)$$

where k_t denotes capital and n_t the amount of the intermediate good used. This setup results in the standard New-Keynesian Phillips curve:

$$1 - \varepsilon_t + \varepsilon_t MC_t = \kappa_\pi (\pi_t - 1) \pi_t - \kappa_\pi \mathbb{E}_t \left[\Lambda_{t,t+1}^f \frac{y_{t+1}}{y_t} (\pi_{t+1} - 1) \pi_{t+1} \right], \quad (19)$$

²⁵The advantage of splitting the production sector into two firms is that the equations for the retail-goods producers resemble those of the standard New-Keynesian model and those for the intermediate-goods producer resemble those of a matching model. The same setup is used in [Leduc and Liu \(2016\)](#) and [Den Haan et al. \(2021\)](#).

²⁶We assume that price-adjustment costs are paid as compensation to other members of the corporate sector and, thus, flow back into the coffers of the representative firm owner. Consequently, these do not affect aggregate resources.

²⁷We abstract from aggregate productivity shocks as our main interest is to study monetary-policy shocks.

where $\Lambda_{t,t+1}^f$ is the MRS of firm owners, ε is the demand elasticity of the Dixit-Stiglitz aggregator, and the marginal cost of producing an extra good, MC_t , is given by

$$MC_t = \frac{\left(\frac{r_t^k}{\alpha}\right)^\alpha \left(\frac{\tilde{w}_t}{1-\alpha}\right)^{1-\alpha}}{\bar{z}}, \quad (20)$$

where r_t^k is the rental rate of capital and $\tilde{w}_t$ the price of the intermediate good.

Cost minimization implies that

$$\frac{\tilde{w}_t n_t}{r_t^k k_t} = \frac{\alpha}{1-\alpha}. \quad (21)$$

2.4.2 Intermediate goods producers and the labor market matching market

Intermediate goods producers are one-worker firms that use a linear technology to transform effective labor units into intermediate goods which are sold to retail-goods producers on a competitive market. There are n_t^y (n_t^m) intermediate-goods producers, each one matched with a young (middle-age) worker. The total amount of intermediate goods produced by these firms is equal to

$$n_t^y + z^{m/y} n_t^m, \quad (22)$$

where $z^{m/y}$ represents the productivity of a middle-aged worker relative to a young worker. These intermediate-goods producers can also be thought of as employment agencies that search for workers who are then placed at retail-goods producers.

Intermediate goods producers are subject to a search friction in attracting workers. Search takes place at the beginning of the period and a successful match allows production in the same period. Free entry ensures that the expected benefit of posting a vacancy is equal to the posting cost, ξ . The search is undirected, that is, firms cannot search specifically for young or middle-aged workers. Since a match with both types of workers has a positive surplus, there is no reason for an entrepreneur who is matched with a young worker to reject the match and go back to the matching market where the expected profits are zero when the posting cost is taken into account. All firms face an exogenous separation probability equal to $1 - \Omega^{ee}$ and free entry implies that firm value then drops to zero. Similarly, retirement implies that the firm loses its worker, which again means a zero continuation value. The employment relationship is not affected when a young worker becomes middle aged, except it becomes more valuable when middle-aged workers are more productive than young ones, i.e., when $z^{m/y} > 1$.

Wages are given by

$$w_t = \omega_0((1 - \omega)\tilde{w}_{ss} + \omega\tilde{w}_t), \quad (23)$$

$$w_t^y = w_t \quad (24)$$

$$w_t^m = z^{m/y} w_t, \quad (25)$$

where $\tilde{w}_{ss}$ is the steady-state value of $\tilde{w}_t$. Recall that $\tilde{w}_t$ is what the intermediate-good producer receives for selling a unit of the intermediate good. Thus, the revenues for a firm with a young worker are equal to $\tilde{w}_t$ and equal to $z^{m/y}\tilde{w}_t$ for a firm with a middle-aged worker. And wages are simply a weighted average of steady-state and current-period revenues. By varying the parameter ω we can control how sticky wages are.

The life-cycle aspects of our model require a few modifications of the standard matching model. To keep the model tractable, we assume that only young workers are subject to unemployment risk. This means that middle-aged workers who face an exogenous separation and unemployed young workers who age find a job with probability 1.²⁸

The matching block is standard except for these modifications. The relevant equations are the follow-

²⁸That is, matches are first allocated to these workers and remaining matches go to young unemployed workers. Since matching occurs at the beginning of the period before production, these workers are employed and productive in period t even though they started the period as searchers in the matching market.

ing:

$$V_t^y = \tilde{w}_t - w_t^y + \beta \Omega^{ee} \mathbb{E}_t \left[\frac{\lambda_{t+1}^f}{\lambda_t^f} ((1 - \delta^y) V_{t+1}^y + \delta^y V_{t+1}^m) \right] \quad (26)$$

$$V_t^m = z^{m/y} \tilde{w}_t - w_t^m + \beta (1 - \delta^m) \Omega^{ee} \mathbb{E}_t \left[\frac{\lambda_{t+1}^f}{\lambda_t^f} V_{t+1}^m \right] \quad (27)$$

$$\tilde{u}_t = \tilde{\gamma}_t^{y,u} + \tilde{\gamma}_t^{m,u} \quad (28)$$

$$\tilde{\gamma}_t^{y,u} = (1 - \delta^y) (\gamma_{t-1}^{y,u} + (1 - \Omega^{ee}) \gamma_{t-1}^{y,e}) + \frac{\gamma^y - n_{t-1}^y}{\gamma^y} \delta^r \gamma^r \quad (29)$$

$$\tilde{\gamma}_t^{m,u} = \delta^y (\gamma_{t-1}^{y,u} + (1 - \Omega^{ee}) \gamma_{t-1}^{y,e}) + (1 - \delta^m) (1 - \Omega^{ee}) \gamma_{t-1}^{m,e} \quad (30)$$

$$m_t = \psi \tilde{u}_t^v v_t^{1-v} \quad (31)$$

$$h_t = \frac{m_t}{v_t} \quad (32)$$

$$f_t^y = \frac{m_t - \tilde{\gamma}_t^{m,u}}{\tilde{\gamma}_t^{y,u}} \quad (33)$$

$$f_t^m = 1 \quad (34)$$

$$\xi = h_t \left(\frac{m_t - \tilde{\gamma}_t^{m,u}}{m_t} V_t^y + \frac{\tilde{\gamma}_t^{m,u}}{m_t} V_t^m \right) = \frac{m_t - \tilde{\gamma}_t^{m,u}}{v_t} V_t^y + \frac{\tilde{\gamma}_t^{m,u}}{v_t} V_t^m \quad (35)$$

$$n_t^y = \frac{n_{t-1}^y}{\gamma^y} \delta^r \gamma^r + (1 - \delta^y) \Omega^{ee} n_{t-1}^y + f_t^y \tilde{\gamma}_t^{y,u} \quad (36)$$

$$n_t^m = \delta^y \Omega^{ee} \gamma_t^{y,e} + (1 - \delta^m) \Omega^{ee} n_{t-1}^m + f_t^m \tilde{\gamma}_t^{m,u} = \gamma^m \quad (37)$$

where V_t^i denotes the value of a match for a firm with a worker of type i , Ω^{ee} the exogenous separation rate, $\tilde{f}_t^i$ the job finding rate of a type i worker, and h_t the hiring rate. A $\tilde{\gamma}_t^i$ variable represents a mass of a particular type of worker who is *searching* whereas a γ_t^i variable indicates the employment status after matching has occurred. Similarly, $\tilde{u}_t$ denotes workers who are searching, not to the mass of unemployed, i.e., not-producing, workers.

As discussed above, we assume that firm owners supply labor to ensure that their consumption responses are procyclical in response to monetary-policy shocks. Specifically, we assume that their supply of labor is proportional to total employment of the combined worker households, that is, to $n_t^y + z^{m/y} n_t^m$. Thus, total supply of intermediate goods is equal to

$$n_t = (n_t^y + z^{m/y} n_t^m) \left(1 + \frac{\gamma^f}{\gamma^y + \gamma^m + \gamma^r} \right). \quad (38)$$

2.4.3 Transfers to firm owners

Net-revenues of the corporate sector are equal to

$$y_t - \tilde{w}_t n_t - r_t^k k_t + n_t(\tilde{w}_t - w_t) - \xi v_t \quad (39)$$

Recall that price-adjustment costs are paid as compensation to (other) members of the corporate sector.²⁹

The *net* transfer of the corporate sector to firm owners is equal to

$$\hat{w}_t^f = y_t - w_t^y n_t^y - w_t^m n_t^m - r_t^k k_t - \xi v_t + \frac{1}{2} \kappa_i (k_t - k_{t-1})^2, \quad (40)$$

which takes into account that wages paid to firm owners drop out. Using this expression, we can write the budget constraint for firm owners as

$$c_t^f + k_t \leq y_t - w_t^y n_t^y - w_t^m n_t^m + (1 - \delta) k_{t-1} - \xi v_t. \quad (41)$$

2.5 Government

2.5.1 Annuity

The annuity pays $1 - \delta_r$ units of consumption each period starting in the period the annuity is bought. The real price of an annuity is equal to q_t . Thus, for each unit of real wealth brought into retirement a retiree would get $(1 - \delta^r)/q_t$ units of consumption until death.

The only participants in the annuity market are the government and the newly retired. The retirees have a vertical demand function, which means that we cannot have a fixed supply. In order to have a sensible equilibrium, we assume that the government sets the annuity price such that it is equal to the Net Present Value (NPV) of expected payoffs using short-term real interest rates as discount rates and fulfills all demand at this price. Specifically,

$$q_t = 1 - \delta^r + \mathbb{E}_t \left[\frac{(1 - \delta^r)^2}{1 + r_{t+1}^d} + \frac{(1 - \delta^r)^3}{(1 + r_{t+1}^d)(1 + r_{t+2}^d)} + \dots \right],$$

²⁹Whether adjustment costs are transfers or loss of goods does not affect the first-order conditions. The advantage of the assumption that they are transfers is that the number of goods produced with capital and intermediate goods still adds up to consumption plus investment.

where r_{t+1}^d is the real interest rate on an investment from period t to period $t + 1$ which is not known in period t since it depends on the realized price level for period $t + 1$. This implies that

$$\frac{1 - \delta^r}{1 + r_{t+1}^d} q_{t+1} = \frac{(1 - \delta^r)^2}{1 + r_{t+1}^d} + \mathbb{E}_{t+1} \left[\frac{(1 - \delta^r)^3}{(1 + r_{t+1}^d)(1 + r_{t+2}^d)} + \frac{(1 - \delta^r)^4}{(1 + r_{t+1}^d)(1 + r_{t+2}^d)(1 + r_{t+3}^d)} + \dots \right].$$

Combining the last two equations gives

$$q_t = (1 - \delta^r) + \mathbb{E}_t \left[\frac{1 - \delta^r}{1 + r_{t+1}^d} q_{t+1} \right]. \quad (42)$$

In a steady state, we get

$$\begin{aligned} q &= (1 - \delta^r) + \frac{1 - \delta^r}{1 + r^d} q \\ &= (1 - \delta^r) \frac{1 + r^d}{r^d + \delta^r}. \end{aligned} \quad (43)$$

Key feature of our annuity. An important objective of this paper is to study how the desire to safeguard retirement consumption through savings is affected when a central bank lowers its policy rate and by doing so lowers returns on savings. The price of an annuity is high when the current and expected future short-term interest rates are low, which means that retirees will get less retirement income from their savings during periods with low interest rates. The same would happen if retirees would have to invest themselves during retirement. But the approach using an annuity does reduce the complexity of the analysis.³⁰

The law of motion for the *number* of annuities held by the government, x_t , is given by

$$x_t = \frac{\bar{\hat{x}}_t \delta^m \gamma^m}{q_t} + (1 - \delta^r) x_{t-1}, \quad (44)$$

where $\bar{\hat{x}}_t$ is the average of the beginning-of-period wealth of the newly retired.

The government's net revenue flow is equal to

$$g_t^x = x_t (1 - \delta^r) - \bar{\hat{x}}_t \delta^m \gamma^m, \quad (45)$$

³⁰In particular, the heterogeneity across retirement cohorts does not introduce additional state variables as is made clear in the remaining equations of this section.

where the first term represents current obligations and the second the sale of new annuities. In steady state, this cost would be equal to

$$g^x = \widehat{x}_t \gamma^m \delta^m \left(\frac{1 - \delta^r}{q \delta^r} - 1 \right). \quad (46)$$

This cost would be equal to zero when the interest rate is equal to zero.

2.5.2 Transfers and tax rates

Unemployment insurance. Unemployment benefits, μ_t^y , and associated expenditures are given by

$$\mu_t^y = (\mu^y + \mu_{hp}^y) w_t^y, \quad (47)$$

$$g_t^u = \mu_{\tau}^y w_t^y (\gamma^y - n_t^y), \quad (48)$$

where $\mu^y w_t^y (\gamma^y - n_t^y)$ is the total transfer to unemployment workers and this amount needs to be financed through taxes, whereas $\mu_{hp}^y w_t^y$ is home production which does not. ³¹

Social security. These transfers are constant in real terms during retirement, but vary between cohorts. The two relevant terms are μ_t^r which affect the level of individual retirement benefits and g_t^r which represents total outlays related to social security.

$$\mu_t^r = (\mu^r + \mu_{hp}^r) w_t^m, \quad (49)$$

$$g_t^r = (1 - \delta^r) g_{t-1}^r + \mu_t^r w_t^m \delta^m \gamma^m, \quad (50)$$

where $\mu_t^r w_t^m$ denotes post-retirement transfer for retirees who retire in period t and this will remain constant for this cohort.

Tax rate for workers. The government budget constraint is given by

$$\tau_t (n_t^y + z^{m/y} n_t^m) w_t = g_t^x + g_t^r + g_t^u + \int_{d_{t-1}} \frac{R_{t-1}^d + \widetilde{R}^d(d_{t-1}) - 1}{\pi_t} d_{t-1} dF_{d_{t-1}}(d_{t-1}),$$

where $F_{d_{t-1}}$ is the cumulative distribution function of risk-free asset holdings brought into period t . ^{32, 33}

³¹With some abuse of notation, we refer to μ_t^y as the (time-varying) amount of resources available to an unemployed worker and to μ^y and μ_{hp}^y as parameters indicating how μ_t^y relates to the wage level.

³²Repayment of principal and new asset issuance cancel out because the supply of assets is fixed. Thus, only interest payments show up.

³³Firm owners are not taxed. The reason is that they also do not benefit from any of government expenditures.

2.5.3 Monetary policy

The nominal risk-free base rate, that is, the rate excluding the premium, is determined by a standard Taylor rule

$$R_t^d = \left(\bar{R}^d\right)^{1-\rho_r} (R_{t-1}^d)^{\rho_r} \left[\left(\frac{\pi_t}{\bar{\pi}}\right)^{\theta_\pi}\right]^{1-\rho_r} \left[\left(\frac{y_t}{\bar{y}}\right)^{\theta_y}\right]^{1-\rho_r} \varepsilon_t^R, \quad (51)$$

where R_t^d is the rate of return for an investment from period t to $t + 1$ and is known in period t . The standard deviation of the innovation, ε_t^R , is denoted by σ_{ε^R} .

2.5.4 Supply of the risk-free asset

Supply is fixed and equal to $\bar{d}$.

2.6 Equilibrium

Goods-market equilibrium is left out of the system because it is implied by the other equations according to Walras law. Equilibrium on the rental market for capital is imposed by using the same symbol (namely k_t) for the supply of capital by firm owners and the demand for capital by final-goods producers. The same is true for equilibrium on the market for the intermediate good, where we use n_t for both demand and supply.

The equilibrium condition for the market of the risk-free asset is given by

$$\int_{d_t} d_t dF_{d_t}(d_t) = \bar{d}. \quad (52)$$

Cashless economy. There are two ways to motivate a cashless NK economy. The first is Woodford's cashless limit. The other is to assume that agents can (i) invest in money and (ii) agents hold money because money gives direct utility and the utility component is additive. The central bank can set the interest rates on the risk-free asset in which they do not trade by controlling the return on a close substitute, i.e., money, and letting the money supply adjust. Because utility is additive, all the equations related to money drop out of the model. In the representative-agent version of this money-in-the-utility (MIU) model, money holdings do not have to be small. Strictly speaking, this would not be a cashless economy as money drops out of the equations, but is still present in the economy. But money would NOT drop out in the corresponding version with heterogeneous agents. In the representative-agent model, money is not

a source of wealth, because there are offsetting tax terms. Although that is true *in the aggregate* with heterogeneous agents, money does create another way to smooth consumption for *individual* agents. This means that the additive MIU justification with heterogeneous agents *also* requires that money balances are driven to zero, which would happen when the assumption is made that the marginal utility of money drops very fast as money balances increase (starting at zero), so that real money balances are tiny and can play no role in insuring against unemployment risk.

3 Calibration

The objective of this paper is to study economic outcomes when the economy is characterized by the presence and interaction of both precautionary savings to deal with unemployment spells *and* saving for retirement. To highlight this interaction clearly, we assume that the ones for whom retirement savings are important (i.e., the middle aged) are not subject to unemployment risk.³⁴ Although stylized, it allows us to analyze whether changes in policies and parameters affect agents with these different types of savings motives differently.

Saving and retirement. The average retirement duration and time spent in middle-age saving for retirement are key parameters. We assume that each life experience, i.e., young, middle aged, and retired, is expected to last 20 years. For computational reasons we adopt an extended perpetual-youth model. This means that we do not have a range of middle-aged agents varying in how close they are to retirement. Instead, we mainly have a representative middle-aged agent capturing *aggregate* behavior of this age group.³⁵ What matters for our analysis is that agents *expect* to age and lose their labor income at some point. This, together with the desire to smooth consumption, creates a savings motive and we study how this is affected by changes in the interest rate. The added stochastic element to aging and retirement is unlikely to be important for our main results as uncertainty does not play an important role in this type of model when standard utility functions are adopted.³⁶

Key parameters that affect retirement savings by the middle-aged are of course the retirement transfer parameter, μ^r , and retirement home production, μ_{hp}^r . In our baseline calibration, we assume that non-investment retirement income is equal to 50% of after-tax income earned while middle aged in the form

³⁴Middle-aged workers are, of course, also subject to unemployment spells. But displacement rates are substantially higher for younger workers as is documented in figure 3 of Davis and Von Wachter (2011).

³⁵As explained in appendix B, our numerical approximation also includes two “transition” types of middle-aged agents for those who recently became middle aged.

³⁶And it is of course to realize that the behavior of our middle-aged agents should be considered as representative of a range of middle-aged persons, some close to retirement and some not.

of a tax-financed transfer and home production is equal to zero. The generosity of pension benefits varies quite a bit, both across and within countries. Given the importance of non-investment retirement income for savings of the middle aged and given our interest in that variable, we consider several alternatives to the 50% benchmark value. These robustness exercises, will help us to understand how differences in the generosity of retirement benefits across countries affect economic outcomes, both at the aggregate level and for different groups in society.

Another key aspect that affects retirement savings is the supply of available assets, $\bar{d}$, and the demand for these assets by other agents. In our model, these are the young who save to insure against unemployment spells and – although to a lesser extent – for retirement.³⁷ In our benchmark, we assume that saving for retirement occurs through risk-free government bonds and these are in fixed supply. Indeed, a big part of retirement savings does occur through government bonds and transaction costs and risk aversion will limit the desire to rely on riskier alternative investments. In an extension, we allow workers to invest in capital and firm ownership, but with the firm owners as intermediaries. That is, workers and firm owners can interact in the market for the liquid asset which would affect the investment (and consumption) decision of firm owners.

Saving and unemployment spells. A model period corresponds to a quarter. In terms of unemployment risk, we closely follow the calibration of [Den Haan et al. \(2017\)](#), which also focuses on the Euro Area. Specifically, the target steady-state unemployment rate is equal to 10.7% and expected unemployment duration is equal to 3.62 quarters. The latter implies that the job finding probability is equal to 26.71%.³⁸ As is common in the literature, we set the curvature of the matching function equal to 0.5, based on the empirical work of [Pissarides and Petrongolo \(2001\)](#). This assumption, together with the two targets, pins down the exogenous separation rate, $1 - \Omega^{ee}$, and the ratio of the job creation cost over the scaling coefficient of the matching function, ξ/ψ .^{39,40}

A key parameter affecting the importance of unemployment risk is, of course, the level of unemploy-

³⁷We also consider the case in which the wage rate increases with age which would add another aspect to the savings decision of the young.

³⁸This takes into account that an unemployment spell ends immediately when a young worker becomes middle aged.

³⁹It is possible to pin down the two values separately if one makes an assumption about the success probability for a firm posting a vacancy. Given our targets, employment would not be affected. Moreover, a different assumption about this success probability would also leave the steady-state value of job creation costs, ξv_t unaffected.

⁴⁰We set the separation rate for a middle-aged agent equal to that of a young agent. This value does not matter much for the model properties we study. It does not matter at all for the middle-aged worker, since they will find a job without any interruption in employment status. From the firms' perspective it makes a job with a young and a middle-aged worker more similar. Moreover, if separation is exogenous, then it is not unreasonable to assume that age is not a key factor.

ment benefits. We set this equal to 50% of after-tax wages. As pointed out in [Den Haan et al. \(2017\)](#), this level is a bit less than what is received on average in most European countries, but in our model unemployment benefits are received indefinitely which simplifies computational aspects.⁴¹

The ability to insure against unemployment spells is not only affected by unemployment benefits, but also by the amount of available liquid assets. How the supply of this risk-free asset is divided among the different types of agents is an equilibrium outcome that depends on several model parameters. But if the supply increases, all agents will hold more. We set per capita bond supply, $\bar{d}$, equal to 3 which means that supply per worker is equal to 4.5, which is equal to 65% (130%) of before (after-) tax annual wage income. In our baseline calibration, the young hold 16.3% of these assets. The much higher holdings by the middle aged means that interest-rate-sensitive investment income is a nontrivial part of retirement savings. Specifically, the fraction of retirement consumption financed by investment income is equal to 30.1%.

The average amount of risk-free assets held by the young in the steady state is equal to 42.4% of their annual after-tax income. This is enough to cover 3.4 quarters of the post-displacement income drop. [Den Haan et al. \(2017\)](#) report that the liquid financial wealth to annual labor income ratio is equal to 31.2% and 21.3% when wealth is defined as gross and net wealth, respectively. Thus, our baseline values are generous, especially if one takes into account that the middle aged actually hold substantially more of the risk-free asset. But they imply reasonable consumption drops upon displacement as discussed below. If we would lower the supply of the risk-free asset, however, then the middle-aged would also hold less which would mean that savings are less important relative to non-investment income. Given our objective to study how the presence of this type of savings affect economic outcomes and, in particular, the impact of monetary policy, we prefer to have a benchmark in which this possibility gets a good chance of being quantitatively important.

The parameters of the penalty function are also important for the ability to smooth consumption. When negative, the parameter $\bar{\eta}_d$ acts as a borrowing constraint, since borrowing costs sharply increase as the holdings of the asset approach this level. We set this parameter equal to -0.3 which is around 34.7% of annual after-tax wage income of an employed agent. We set $\eta_{d,1}$ equal to 1 and $\eta_{d,0}$ equal to $3.9015e-04$. As shown in figure 1, these two parameter values imply that the premium is small when long positions are chosen, but increases sharply when the agent enlarges its short positions especially close to

⁴¹In the baseline version, this is in the form of a tax-financed transfer and home production is equal to zero.

$\bar{\eta}_d$. In section 6.2, we discuss model properties for alternative values for $\bar{\eta}_d$.

For our purpose, it is important that we get empirically plausible post-displacement consumption drops. An agent who is displaced after a long employment spell faces a 15.2% consumption drop after being unemployed for one year and a 26.0% drop after a year and a half at which point the consumption level no longer drops. These numbers are in line with the average 26.3% drop after one year of unemployment for individuals and 15.8% for households reported in Kolsrud et al. (2015). However, these observations are for Swedish data which has a relatively high replacement rate of 72%. In the experiments in which we lower the supply of the risk-free liquid asset, we obtain larger consumption drops.

Utility function. The elasticity of intertemporal substitution, $1/\phi$, is set equal to 1, a standard value in the literature, but we do consider other values. We set the discount factor equal to 0.985. This is somewhat lower than the usual 0.99. Moreover, the effective discount factor is even lower because of the probability of dying. But there is a nontrivial downward effect on the real interest rate in our model because of the precautionary savings motive and with $\beta = 0.985$ we get a real quarterly interest rate equal to 1.07%, which is close to the usual 1% value. We set the firm owner's discount factor such that the firm owner's demand for the liquid asset is equal to zero in the steady state.⁴² This ensure that the steady state of the version in which firm owners do interact with workers in the market for the liquid asset over the business cycle is equal to the steady state of the version in which we shut down such interaction, which is the benchmark.

Taylor-rule coefficients. The Taylor-rule coefficient related to inflation, θ_π , is set equal to 1.99, the one related to real activity, θ_y , equal to 0.24, and the smoothing parameter, θ_r , equal to 0.84. These values are estimated in a medium-scale DSGE model for the euro area by Mazelis, Motto and Ristinieni (2023)

An important objective of this paper is to study the monetary transmission mechanism when changes in the real interest rate affect retirement savings. Thus, it is key that a monetary-policy shock generates a change in the real interest rate that is realistic in terms of its persistence. With these parameter values, the model predicts a peak response for the real interest rate in the fourth quarter and that peak response is halved seven quarters later. Thus, the monetary-policy shock reduces the real interest rate for a nontrivial

⁴²That is, the firm owner's discount rate is equal to the real interest rate which is accomplished when $\beta^f = 0.98937$.

amount of time.⁴³

Parameters to get reasonable aggregate responses. We set the parameter controlling the cost of changing prices, κ_π , equal to 1,500. Price stickiness ensures that nominal shocks have an impact on real demand for goods. To ensure that this leads to a quantitatively realistic impact on employment, we follow [Hagedorn and Manovskii \(2008\)](#) and choose the profit margin, $1-\omega_0$, and the degree of wage flexibility, ω , low enough such that expected profits display sufficient volatility which in turn ensures that vacancy posting and job creation are sufficiently volatile.

Specifically, a 25 basis point reduction in the real interest rate leads to a peak response in employment equal to 0.85%, a peak response in output equal to 0.55%, and a peak increase of year-on-year inflation equal to 3.5 basis points.

Mass of firm owners. The mass of workers is normalized to be equal to 1. The value for γ^f , i.e., the mass of firm owners, is actually not that important for many model properties, since we use a representative firm owner. Specifically, the size of the aggregate capital stock and the economy-wide employment rate are not affected by the value of γ^f . In a New-Keynesian (NK) framework, firm profits decline during an expansion that is induced by a monetary stimulus because of the fall in the markup. This goes together with a reduction in consumption by firm owners. As explained in [section 2](#), we allow firm owners to benefit from a monetary expansion by letting them supply labor and earn wage income. The value of γ^f does not affect aggregate profit income but it does increase wage income earned by firm owners. Consequently, wage income becomes relatively more important as γ^f increases which makes it more likely that consumption of firm owners would not decrease following a monetary expansion.

For the questions addressed in this paper, the value of γ^f mainly matters when we let firm owners facilitate consumption smoothing of workers by buying or selling the risk-free asset to them. It is, of course, easier for workers (as a group) to borrow more from firm owners without a significant dampening effect on the interest rate when firm owners are richer, that is, when a larger share of wage income goes to firm owners. In our baseline calibration, we set γ^f equal to 0.3.

⁴³For the estimated value of θ_y , the smoothing parameter barely matters for the persistence. That is, higher persistence of the shock is undone by the response to changes in the level of economic activity. The parameter θ_y is also important for the quantitative impact of monetary-policy shocks on economic activity, together with parameters controlling the stickiness of prices and wages.

Allowing workers to borrow or lend from firm owners. We consider the case when firm owners do participate in the bond market and when they do not. In the second case, workers who save to insure against unemployment risk and those who save for retirement have to do with a fixed supply of the risk-free asset. This case allows us to cleanly study how these savings motives are affected during the monetary transmission mechanism. In the first case, we can address the question whether a group of agents that is not subject to unemployment risk and does not have to save for retirement can act as a shock absorber.

When we do allow firm owners to invest in the bond market, then we assume that they also face a trade restriction. One possibility is to assume that they face a penalty function like workers do. But in terms of allowing for more interaction, i.e., larger trades, without running into Blanchard-Kahn indeterminacy issues, a simple quadratic adjustment costs on bond holdings works better. That is, a position in bonds, d_t^f comes with a (symmetric) cost equal to $\eta_d^f (d_t^f)^2$ and we set η_d^f equal to 0.0001.⁴⁴

4 Interest rate changes and retirement savings in simple models

A key objective of this paper is to study the *differential* impact of monetary policy on the economic behavior of agents who differ in age and unemployment risk. Many aspects are relevant which complicates the analysis. Thus, we first consider some simpler models before turning to the full model. A common element in all models considered in this section is that income levels are kept fixed when interest rates change. We start by discussing two very simple models in which agents live for a finite number of periods. Next, we discuss a partial equilibrium version of our full model.

4.1 Simple finite-lifetime models

In this section, we consider simple models in which agents live for a finite number of periods. This allows us to derive some analytical results and more easily illustrate the income and substitution effects. The first is a standard textbook partial-equilibrium model in which agents live for two periods. Despite its simplicity and popularity, we find it useful to review it as its predictions are at odds with those of our much more complex environment. The second is an overlapping-generations (OLG) model. It mimics our full model in that there are three generations, the young, the middle-aged, and retirees. Key in our analysis of this model is that the interest rate shock occurs at the steady-state *equilibrium* interest rate.

⁴⁴Recall that the discount factor of firm owners is such that the steady-state value of d_t^f is equal to zero. Thus, this represents a cost for deviations from the steady-state value. The approach with the quadratic-adjustment cost is perhaps less elegant than the one with the penalty function. Both approaches lead to similar results, however, as long as the case with the penalty function does not lead to indeterminacy issues.

It turns out that income level during retirement affect the value of this equilibrium interest rate which in turn affect the responses to an interest rate shock.

A simple two-period model. The textbook setup to study the impact of a change in the interest rate on the consumption/saving decision of agent i is the following:

$$\begin{aligned} \max_{c^m, c^r, b} \quad & \frac{(c^m)^{1-\phi} - 1}{1-\phi} + \beta \frac{(c^r)^{1-\phi} - 1}{1-\phi} \\ \text{s.t.} \quad & c^m + b^m = y^m, \end{aligned} \tag{53a}$$

$$c^r = \mu^r + (1+r)b^m. \tag{53b}$$

Here, c^m denotes consumption in the first period, c^r consumption in the second period, b^m savings, and r the interest rate. If we think of the first period as the one when the agent is middle aged and the second when the agent is retired, then this model captures in a very crude and simplified manner the savings problem of a middle-aged agent in our model. After substituting out c^m and c^r , we get the following first-order condition for b^m :

$$(y^m - b^m)^{-\phi} = \beta(1+r)(\mu^r + (1+r)b^m)^{-\phi}. \tag{54}$$

In appendix A, we show that the solution for consumption is given by

$$c^m = \frac{[\beta(1+r)]^{-\frac{1}{\phi}}(1+r)}{1 + [\beta(1+r)]^{-\frac{1}{\phi}}(1+r)} \left[\frac{\mu^r}{1+r} + y^m \right] \tag{55}$$

and the solution for bond holdings, b^m , is equal to $y^m - c^m$.

Microeconomics teaches us that the total effect of a change in the interest rate can be decomposed into a substitution and an income effect. The income effect induces a decrease (increase) in current consumption for savers (borrowers) as it shifts in (out) the budget constraint when the interest rate falls. The substitution effect of a decrease in the interest rate induces an increase in current consumption for both borrowers and savers, since it decreases the price of current to future consumption.⁴⁵

⁴⁵See, for example, section 6.2 of [Kurlat \(2020\)](#).

The income effect of a middle-aged worker. One of the questions we want to address in this paper is whether a strong negative income-effect associated with a monetary-policy stimulus could lead them to *reduce* consumption. This does seem to be a possibility. Consider the (extreme) version when retirees have only investment income, that is, $\mu^r = 0$. Equation (55) implies that consumption in the first period, i.e., consumption before retirement, would increase in response to a decline in the interest rate when $0 < \phi < 1$, decrease when $\phi > 1$, and remain unchanged when $\phi = 1$. Thus, it is – in theory – possible that the income effect for the middle aged dominates the substitution effect and their demand would *fall* following a monetary expansion. However, there are several reasons why the result from this simple model has limited applicability to the full model.

One reason to be careful is that setting μ^r equal to zero is special and extreme. If $\mu^r > 0$, then lifetime income increases and would be relatively less affected by the change in the current interest rate and the income effect would become less important. Also, a change in the interest rate is a permanent change in this 2-period model, which is not representative for the temporary changes in the real interest rate induced by central bank monetary policy. Finally, the middle aged are also likely to be affected by increased economic activity and/or higher wages in the more comprehensive setting captured by the full model.

Pension benefit levels and impact of interest rate changes on demand. Next, we turn to the question how the sensitivity of demand to changes in the interest rate is affected by changes in the level of pension benefits. We will look at consumption demand by the young, consumption demand by the middle aged, and aggregate demand. We will see that differential changes across age cohorts are important.

A simple OLG model with three generations. We consider an overlapping-generations (OLG) model with three generations: the young, the middle aged, and the retirees. In contrast to the full model described in section 2, generations live for only one (long) period. Income of the young is equal to y^y , the middle aged get y^m , and the retirees get μ^r . The latter is a transfer from the young and the middle aged. When young, the agent solves the following optimization problem:

$$\begin{aligned} \max_{c^y, c^m, c^r} \quad & \ln(c^y) + \ln(c^m) + \ln(c^r) \\ \text{s.t.} \quad & c^y + \frac{c^m}{1+r} + \frac{c^r}{(1+r)(1+\bar{r})} = y^y + \frac{y^m}{1+r} + \frac{\mu^r}{(1+r)(1+\bar{r})}. \end{aligned} \quad (56)$$

The solution satisfies:

$$c^y = \left[y^y + \frac{y^m}{1+r} + \frac{\mu^r}{(1+r)(1+\bar{r})} \right] / 3, \quad (57)$$

$$c^m = c^y(1+r), \quad (58)$$

$$c^r = c^y(1+r)(1+r^f). \quad (59)$$

The values of y^y and y^m are assumed to be equal to each other. Since aggregate resources are kept constant at Y , the transfer to the retirees satisfies $\mu^r = Y - y^y - y^m$. The *current* interest rate at which the young and middle aged can borrow and lend is denoted by r and the interest rate at which the current young can trade in the future when they are middle aged is denoted by r^f . We study an unexpected change in the current interest rate, r keeping the future interest rate, r^f fixed at its steady-state equilibrium value, r_{eq} . For the young this means a change in r keeping r^f fixed. Since the change in the interest rate is unexpected, the middle aged made decisions when young based on the assumption that $r = r^f = r_{eq}$. The economy is in equilibrium when

$$(y^y - c^y) + (y^m - c^m) = 0, \quad (60)$$

that is, total demand for one-period bonds adds up to supply which is assumed to be equal to zero. Faced with the unexpected change in the interest rate, the middle aged can adjust their consumption/savings decision. Specifically,

$$c^m = \frac{1}{2} \left[y^m + (y^y - c^y(r_{eq})(1+r_{eq})) + \frac{\mu^r}{1+r} \right]. \quad (61)$$

Note that both the savings the current middle aged brought with them and the return earned are based on r_{eq} , that is, the risk-free interest rate that was in place when they were young before the shock occurred.

We are interested in the following two elasticities.

$$\frac{\partial c^y}{\partial r} = - \left(y^m + \frac{\mu^r}{1+r_{eq}} \right) \frac{1}{3(1+r)^2}, \quad (62)$$

$$\frac{\partial c^m}{\partial r} = -\mu^r \frac{1}{2(1+r)^2}. \quad (63)$$

The formulas reveal that these elasticities are highly nonlinear and demand is more sensitive to changes in the interest rate at low levels of r . To understand this consider the extreme case when the

interest rate is infinite. Then the NPV of future income is equal to zero and changes in the interest rate would have no effect on consumption.

Another important insight is that consumption when young is in general more interest sensitive.⁴⁶ For example, suppose that $y^y = y^m = \mu^r$. Then $r_{eq} = 0$ and $c^y = c^m = c^r$. Then we get⁴⁷

$$\frac{\partial c^y}{\partial r} = -\frac{2}{3}, \quad (64)$$

$$\frac{\partial c^m}{\partial r} = -\frac{1}{2}. \quad (65)$$

The intuition is fairly simple. The young have more future income than the middle aged which means that an interest rate change has a bigger impact on its discounted value. Note that this is true even though the middle aged save as much as the young before the shock, namely none at all.

We are now ready to consider our main objective, namely the sensitivity of these elasticities to permanent changes in pension benefits, i.e., changes in μ^r . The values of y^y and y^m have to fall if μ^r increases, since we keep resources fixed and we assume that $y^y = y^m$.

First, there is a direct effect in that a decrease in μ^r reduces $|\partial c^y / \partial r|$ and $|\partial c^m / \partial r|$, that is, reduces the response of consumption to changes in the interest rate. This is quite intuitive. A reduction in μ^r lowers future income which means that the NPV of future income is less sensitive to changes in the interest rate. Since we are considering permanent changes in μ^r , we also have to take into account changes in the equilibrium interest rate, r_{eq} , that is, the value of r at which we have to evaluate the elasticities. A reduction in μ^r would induce working agents to save more which would lower the equilibrium interest rate. And this would *increase* the sensitivity of consumption to changes in the interest rate. Figure 2 plots the values of these elasticities. The left column plots $\partial c^y / \partial r$ and $\partial c^m / \partial r$. This is the column to consider if we are interested in the impact on aggregate consumption, i.e., $c^y + c^m$, and which component is driving it. The right column plots $\partial \ln c^y / \partial r$ and $\partial \ln c^m / \partial r$. And these will be more useful when discussing impulse response functions of the full model below.

The following conclusions can be drawn from the figure. Most importantly, the equilibrium effect


⁴⁶The value of y^m would have to be substantially below μ^r for this not to be true.

⁴⁷Throughout this section, we will evaluate these elasticities at $r = r_{eq}$. That is, they measure demand pressure in response to an exogenous shock when the economy starts out in its steady-state equilibrium. These changes in c^y and c^m are not equilibrium outcomes. But that is not what we are interested in. We want to understand how notional demand changes which then could affect real activity in a full model with sticky prices.

always dominates the direct effect for consumption of the young as well as the aggregate effect.⁴⁸ Even for consumption of the middle aged, there is a large range of values for μ^r at which the indirect effect dominates. But this cannot be true over the whole range since the elasticity must approach zero as μ^r goes to zero.

The reason for the strong indirect effect is that the demand for consumption is a nonlinear function of the interest rate and the drop in demand accelerates as the interest rate falls.

Figure 2: Interest rate elasticity consumption demand



Notes: This figure plots the responsiveness of consumption to changes in the interest rate as a function of pension benefits in the OLG model with three generations. The top (bottom) row gives the results for the young (middle aged) and the left (right) column gives the absolute (percentage) numbers. The change in r is evaluated at a level for r at which the bond market is in equilibrium and this equilibrium value for r falls as μ^r falls.

Re-emergence of an interest-rate channel in a New-Keynesian environment. [Rupert and Sustek](#)

⁴⁸Note that the range of values taken on by $\partial c^y / \partial r$ is about one order of magnitude larger than the one for $\partial c^m / \partial r$.

(2019) document that a negative monetary-policy shock in the representative-agent New-Keynesian model would induce an increase in inflation and economic aggregates like consumption and GDP. The paper also shows that such an inflationary expansion happens when the *real* interest rate increases and that a fall in the real interest rate – if it does happen – is not the driving force behind the expansion.⁴⁹ In the New-Keynesian model, the key equation for the monetary-policy transmission is the Phillips curve, which establishes a positive relationship between real activity and inflation and the standard Phillips curve does *not* contain an interest rate. In other words, there is no interest-rate channel in a representative-agent New-Keynesian model. The analysis of our simple OLG model with retirement indicates that changes in the interest rate do play a role for aggregate activity due to heterogeneous responses. Specifically, if the consumption/savings responses to a change in the interest rate do not average out to the one of the representative-agent version of the model, then the presence of heterogeneity would be similar to a preference shock affecting the Euler equation. And such a shock is known to have real consequences for the economy. Specifically, it would lead to increased inflation when it is expansionary which leads to increased economic activity as predicted by the New-Keynesian Phillips curve.⁵⁰

4.2 Partial-equilibrium outcomes in our infinite-horizon model

As discussed in the introduction, an important question is whether those whose decisions are affected by retirement issues respond differently to monetary-policy changes and in particular whether their consumption could *decrease* during a monetary expansion. In our full model, agents' decisions are affected by the direct effect of the interest rate as well as changes in income levels. To understand the role of both, we consider a partial equilibrium version of our model in which income levels are kept fixed.

We first consider the case when parameters take on their baseline values. At these parameter values, savings of the middle-aged are substantial and these finance 20.5% of retirement consumption in the steady state. Moreover, they face a 37.7% drop in consumption upon retirement, again in the steady state. Next, we consider more extreme values for key parameters.

The real interest rate is modeled as an exogenous AR(1) process and the shock is a 25 basis points reduction on impact. Wage income, non-investment retirement income, as well as the savings accumulated while young and carried into the middle-aged phase of life are kept fixed at their steady state values.

⁴⁹By contrast, in the standard (undergraduate) IS-LM (or IS-MP) model a reduction in the real interest rate plays a key role as it directly affects aggregate demand and, thus, aggregate activity.

⁵⁰Another possible role for the real interest rate is through how it affects households' exposure to mortgage debt as shown in Garriga et al. (2021).

In other aspects, the partial-equilibrium environment is identical to that of the full general-equilibrium model.

A simple calculation. Suppose that the quarterly real interest rate drops by 25 basis points permanently. This leads to an 11.8% increase in the price of the annuity. If middle-aged savings remain fixed, then this implies an equal percentage reduction in retirement consumption financed out of the annuity and a 2.2% drop in total retirement consumption. Persistence matters a lot. If the change in the interest rate is not permanent, but follows an AR(1) process with a coefficient equal to 0.95, then the annuity price increases by only 3.4% on impact.

What can we expect? A reduction in the interest rate implies that expected consumption growth must fall. That is, the consumption IRF for an *individual* middle-aged worker must be downward sloping as long as the interest rate including the premium follows the same pattern.⁵¹ Moreover, the reduction in the interest rate does not stimulate the economy because all aggregate variables other than the interest rate and the annuity price are kept constant. At the same time, the drop in the interest rate does imply a reduction in investment income which is the cause of the negative income effect. If this income effect is strong enough, then consumption would already fall on impact.⁵² It is possible that consumption does increase on impact because of the substitution effect, but the IRF *must* turn negative at some point because a positive consumption response in every period is impossible when the lifetime budget constraints shifts inward due to a reduction in the return on savings. The IRFs we plot basically follow these rules except in the limit. The reason is that they display the response of the consumption level *averaged* across middle-aged workers. Over time, young people who age will be added to this group and they will become middle aged with wealth equal to the full model's steady-state values in this partial-equilibrium experiment. This means that by construction, the IRFs will eventually go back to zero.

Results when the AR(1) coefficient for the interest rate, r is equal to 0.8. We first consider the case when the drop in the real interest rate is fairly persistent. An AR(1) process with a coefficient equal to 0.8 implies that the half-life is slightly more than three quarters and it takes fourteen quarters for the change in the interest rate to be less than 5% of the change on impact.

Figure 3 plots the associated Impulse Response Functions. The following observations can be made.

⁵¹The rate including the premium is endogenous and depends on positions taken in the risk-free asset, but always follows the same pattern as the exogenous part of the interest rate.

⁵²This means that on impact savings would increase and this would dampen consumption drops in subsequent periods.

Consumption of the middle aged increases at first and only falls below its pre-shock value after ten quarters. The amount held in the risk-free asset displays an immediate and persistent drop. Post-retirement consumption must fall, since non-investment retirement income is kept fixed, savings fall, and the fall in the interest rate increases the annuity price. The IRF in the bottom-middle panel plots consumption of the *newly* retired. Recall that the newly retired lock in post-retirement income and consumption levels by buying an annuity. This means that those who retired before the drop in the interest rate are not affected. Those who retire in the period when the shock hits are affected the most by the increase in the annuity price, but not at all by the reduction in savings since their savings are still based on the pre-shock interest rate value. Those who retire later, however, are affected more by the reduction in the level of savings brought into retirement. The IRF for consumption upon retirement reaches a maximum drop of almost 0.4%. Aggregate retirement consumption drops by less, since only the newly retired are affected as the annuity insures those who retired before the shock. The drop in aggregate retirement consumption is extremely persistent and takes sixty-six quarters to reach its lowest value as the asset holdings of the middle aged take a long time to revert to normal which means that newly retired enter the retirement pool with lower savings for quite a while.


In partial equilibrium, it is possible that the asset holdings of *all* agents drop. In general equilibrium, however, this would put upward pressure on the real interest rate until – on average – there is no change in asset holdings. Whether the middle aged will be able to reduce asset holdings in our general-equilibrium model will depend on what the young would like to do when the real interest rate falls and on how economic activity affects the demand for the risk-free asset.

The fact that consumption rises initially and falls later reveals a useful quantitative insight. That is, the substitution effect of the interest-rate reduction dominates the income effect even though the middle aged hold a large amount of the risk-free asset and face a non-trivial consumption drop upon retirement.⁵³

Results when AR coefficient for r is equal to 0.95. In the next experiment, we consider a very persistent drop in the real interest rate. That is, the AR(1) coefficient is equal to 0.95 which implies a half life of around thirteen quarters. When the drop in the interest rate is more persistent, then the annuity price increase is larger putting more pressure on retirement consumption. Moreover, with a more per-


⁵³This outcome is quite different from the two-period model considered above. But there are important differences between our life-cycle model and the two-period model. In our life-cycle model, the fact that there is a drop in non-investment income upon retirement income is only relevant for those middle-aged agents who retire in the next period and they are only a small fraction of all middle-aged agents. In the two-period model considered, all agents face a drop in non-investment income and this happens with perfect certainty. Also, the drop in the interest rate considered here is temporary whereas it is for the remainder of the agent's life in the two-period model.

Figure 3: Interest-rate shock IRFs: Partial-equilibrium; $\rho_{PE} = 0.8$




Notes: This figure plots the IRFs for an interest-rate shock when the real interest rate follows an AR(1) with autoregressive coefficient equal to 0.8. This is for the partial-equilibrium version of our model and all aggregate variables besides the real interest rate and the annuity price remain equal to their steady-state values of the full model.

Figure 4: Interest-rate shock IRFs: Partial-equilibrium; $\rho_{PE} = 0.95$




Notes: This figure plots the IRFs for an interest-rate shock when the real interest rate follows an AR(1) with autoregressive coefficient equal to 0.95. This is for the partial-equilibrium version of our model and all aggregate variables besides the real interest rate and the annuity price remain equal to their steady-state values of the full model.

Figure 5: Interest-rate shock IRFs: Partial-equilibrium; $\phi = 4, \beta = 0.96296$




Notes: This figure plots the IRFs for an interest-rate shock when the real interest rate follows an AR(1) with autoregressive coefficient equal to 0.8 and the intertemporal elasticity of substitution, $1/\phi$, is equal to $1/4$. This is for the partial-equilibrium version of our model and all aggregate variables besides the real interest rate and the annuity price remain equal to their steady-state values of the full model. The value of β is adjusted so that the steady-state value of the real interest rate remains equal to its baseline value.

Figure 6: Interest-rate shock IRFs: Partial-equilibrium; $\phi = 5.5, \beta = 0.94436$




Notes: This figure plots the IRFs for an interest-rate shock when the real interest rate follows an AR(1) with autoregressive coefficient equal to 0.8 and the intertemporal elasticity of substitution, $1/\phi$, is equal to $1/5.5$. This is for the partial-equilibrium version of our model and all aggregate variables besides the real interest rate and the annuity price remain equal to their steady-state values of the full model. The value of β is adjusted so that the steady-state value of the real interest rate remains equal to its baseline value.

Figure 7: Interest-rate shock IRFs: Partial-equilibrium; $\mu^r = 0, \beta = 0.960176$



Notes: This figure plots the IRFs for an interest-rate shock when the real interest rate follows an AR(1) with autoregressive coefficient equal to 0.8 and investment income is the only source of income for retirees. This is for the partial-equilibrium version of our model and all aggregate variables besides the real interest rate and the annuity price remain equal to their steady-state values of the full model. The value of β is adjusted so that the steady-state value of the real interest rate remains equal to its baseline value.

Figure 8: Interest-rate shock IRFs: Partial-equilibrium; $\bar{d} = 6, \beta = 0.98749$



Notes: This figure plots the IRFs for an interest-rate shock when the real interest rate follows an AR(1) with autoregressive coefficient equal to 0.8 and the supply of the riskfree asset, $\bar{d}$ is twice its baseline value. This is for the partial-equilibrium version of our model and all aggregate variables besides the real interest rate and the annuity price remain equal to their steady-state values of the full model. The value of β is adjusted so that the steady-state value of the real interest rate remains equal to its baseline value.

sistent interest rate drop, the loss in investment income will be larger relative to lifetime income. But a more persistent interest rate drop also means that the price of current consumption relative to future consumption is lower for longer which would strengthen the substitution effect. As shown in Figure 4, consumption of the middle aged still increases on impact. That is, the substitution effect still dominates. In fact, savings of the middle aged now display an even larger peak drop, namely 9.2% compared to an 1.7% drop when the persistence parameter is equal to 0.8.⁵⁴ Consistent with this finding, we see that the percentage increase in middle-aged consumption is roughly triple the increase observed for the less persistent drop in the interest rate. And the peak percentage consumption drop for the newly retired increases from 0.38% to 2.1%.

Lower intertemporal elasticity of substitution. Next, we consider the results when ϕ , i.e., the inverse of the intertemporal substitution, is equal to 4 instead of 1. We reduce the discount factor of the workers to ensure that the steady-state values of the real interest rate and thus the annuity price are not affected by the change in ϕ . Without this adjustment, the annuity price would increase sharply which would mean that the same amount of savings would imply lower post-retirement consumption. With this adjustment, steady-state savings level of the middle aged as well as the importance of investment income for steady-state retirement remain roughly the same. Moreover, the steady-state consumption path is also very similar to the baseline outcome, both when employment status changes and when an agent ages.⁵⁵

The re-calibration of ϕ increases the desire of the middle aged to avoid a consumption drop when the interest rate falls. As shown in Figure 5, savings of the middle aged now drop by less and retirement consumption drops by less. The initial increase in consumption of the middle aged is still positive, but substantially smaller and is followed by a persistent decline.⁵⁶ Figure 6 displays the results when ϕ is increased even further to 5.5 and the discount factor is again adjusted to keep steady-state properties similar. Now consumption of the middle aged displays a uniform drop.⁵⁷ Thus, it is indeed possible that some agents in the economy respond to a reduction in the interest rate by consuming less not more, although a value of ϕ equal to 5.5 is substantially higher than what is typically used in the literature.

⁵⁴Note that savings levels when the shock occurs are identical in both cases as shock persistence does not affect the steady-state.

⁵⁵That is, the effects of the higher discount factor and the lower intertemporal substitution offset each other.

⁵⁶On impact, the response in consumption and the response in savings must add up to zero. This is not quite clear from the IRFs since they express the responses as a percentage from their steady-state values.

⁵⁷And whereas the peak drop of consumption upon retirement is equal to 0.38% when $\phi = 1$, it is now only 0.24%. The change in the value of ϕ has a bigger impact on the IRFs for middle-aged consumption and savings than on the IRF of consumption upon retirement. The reason is that the IRF of the change in retirement consumption represents a change that is permanent (until death) and will continue long after the interest rate has gone back to its steady-state value.

Moreover, we want to remind the reader that this reduction in the real interest rate is not necessarily representative for a monetary stimulus as all other prices and aggregate variables are kept constant.

No non-investment retirement resources. The two-period model illustrates that the income effect can be very strong when period-2 income depends solely on investment income. In the examples considered so far, retirees have non-investment income since $\mu^r > 0$. Setting μ^r equal to zero without a change in any of the other parameter values would lead to such a large drop in the real interest rate that the steady-state interest rate plus the death rate would turn negative which implies an infinite annuity price. To keep the steady-state environment comparable, we lower the discount factor to equalize the steady-state real interest rate to its baseline value. When we earlier changed ϕ , the adjustment of the discount factor to keep the real interest rate the same also ensured that other steady-state properties remained similar. That is not quite true now. For example, middle-aged steady-state savings increase from 7.02 to 8.78.⁵⁸ The re-calibration of the discount factor ensures that the annuity price remains constant, but the increase in savings is not enough to offset the drop in μ^r ; consumption now drops by 87.1% instead of 38.1% upon retirement when the economy is in the steady state.

Figure 7 displays the responses to an interest rate reduction. The key observation is that consumption of the middle aged still increases on impact, even in this extreme case. Thus, the behavior of the middle aged would still stimulate demand in response to a monetary stimulus.⁵⁹ Consumption increases on impact even though the percentage drop in consumption of the newly retired is now substantially higher and this is undoubtedly unpleasant given that retirement consumption is already quite low.

Although consumption still increases on impact when $\mu^r = 0$, the increase is smaller and less persistent as it already turns negative after five periods.⁶⁰ Thus, there may be a *delayed* dampening effect of a monetary stimulus through reduced demand by the middle aged.

⁵⁸That is, the reduction in retirement benefits and the reduction in the discount factor have exactly offsetting effects on the interest rate (by construction), but do redistribute holdings of the risk-free asset from the young to the middle aged.

⁵⁹Recall that in these partial-equilibrium exercises, the IRF of consumption of the middle aged *must* turn negative at some point, since non-investment income has not changed and investment income falls for savers when the interest rate drops. Thus, key is what happens on impact.

⁶⁰In our extended perpetual-youth model, the consumption and savings variables for the middle-aged capture the cross-section of middle-aged workers. Of course, there are middle-aged workers close to retirement for whom the predictions of the two-period model are quite relevant. But then there are also quite young middle-aged workers for whom the probability of retirement is much lower than the retirement probability in our model. With our setup we capture the *average* across these different middle-aged workers and with this simplification we vastly reduce the computational burden. Another aspect to keep in mind is that retirement income is assumed to depend solely on the returns from the risk-free asset and not on the returns from other assets like equity or real estate which may increase during an expansion induced by a reduction in the interest rate.

Increased deposit supply. In the last experiment of this section, we consider an increase in the supply of the risk-free asset from 3 to 6. This will increase the steady-state wealth value of the middle aged and thus strengthen the income effect. However, the analysis with the two-period model shows that the magnitude of the substitution effect may change as well, leaving the total effect unchanged.

Once again, we adjust the discount factor to ensure that the real interest rate, and thus the annuity price, remain equal to their general-equilibrium steady-state values. Savings of the middle aged more than double in steady state to 15.54 and the drop in post-retirement consumption is now only 19.6% and the fraction of post-retirement consumption financed by investment income is equal to 37.9%. Figure 8 displays the responses of a monetary-policy shock. The percentage drop in retirement consumption is now substantially higher which can be explained by the fact that investment income is now more important for retirees. But the IRF for consumption of the middle aged still increases on impact. In fact, the response is even slightly larger.

Summary. Although it is possible for consumption of the middle aged to drop upon impact following a reduction in the real interest rate, this only happens for more extreme parameter values. Thus, the insight from the two-period model where this could easily happen does not generalize to our life-cycle model.

5 Steady-state asset holdings and equilibrium rates of return


In this section, we discuss steady-state equilibrium holdings of the risk-free asset across the three worker types, i.e., employed young, unemployed young, and the middle aged.⁶¹ Moreover, we discuss how these holdings – as well as the level of the equilibrium interest rate – change in response to changes in key parameter values. These findings will be helpful when we consider the full model with monetary-policy shocks.

Changes in non-investment retirement benefits. Figure 9 shows how the equilibrium interest rate and average holdings of the risk-free asset by worker type vary when we adjust the magnitude of the tax-financed transfer to retirees from $\mu^r = 0.05$ to $\mu^r = 1$.⁶² An intuitive prediction of our model is that

⁶¹Recall that those that retire sell their holdings of the risk-free asset and use the proceeds to buy an annuity.

⁶²When μ^r and μ_{hp}^r are both equal to zero, then retirement consumption is financed solely by investment income. How much retirees can consume depends not only on the savings they bring into retirement, but also on the annuity price. In general equilibrium, the reduction in the interest rate associated with changes in μ^r , which controls the level of non-investment retirement income, implies an increase in the annuity price. With log utility, the real interest rate converges to $-\delta^r$ as μ^r goes to zero, which implies that the annuity price goes to infinity. Consequently, retirement income would go to zero. For this reason we have chosen a small positive number as the lower bound for μ^r to construct the figures. Setting μ^r equal to 1 is a natural upper bound, since it implies that pre-retirement after-tax wage income is equal to the retirement transfer which is defined such that it abstracts from taxes.

Figure 9: Steady-state asset-market outcomes: Role of tax-financed retirement benefits




Notes: This figure plots the average for the indicated type of agent of the steady-state per capita holdings of the risk-free asset as well as the equilibrium interest rate for different values of μ^r , i.e., the level of the tax-financed retirement transfer, which is defined as a fraction of the after-tax wage rate of middle-aged workers. The interest rate is the base rate which excludes the level-dependent premium.

an increase in μ^r reduces retirement savings by the middle-aged which has an upward effect on the real interest rate. The increase in the real rate is quite large. Specifically, the annual real interest rate increases by 8.9 percentage points when μ^r increases from 0.05 to 1, i.e., from virtually no retirement benefits to a level that equals pre-retirement wage income (see RHS panel). This increase in μ^r leads to a reduction in average per capita holdings of the risk-free asset by the middle aged from 8.53 to 6.27 (see dash-dotted line in LHS panel). Relative to *after-tax* annual wage income, however, holdings increase from 1.25 to 1.53 (not shown). The reason it moves in the opposite direction is that the increase in μ^r leads to a large fall in after-tax wage income because it goes together with a substantial increase in the tax rate.

Even when retirees face no drop in non-investment income, i.e., when $\mu^r = 1$, it is still true that the middle aged hold most of the available risk-free asset, namely 70%. With no drop in non-investment income and a substantial amount of savings, retirees would then enjoy a substantial *increase* in consumption upon retirement. Thus, the question arises why they hold such a high fraction of available assets. Under our baseline calibration, wage rates do not increase with age and with a 50% replacement rate and a steady state unemployment rate equal to 10.7%, average income of the young including the unemployed is only 11% lower than that of the middle aged. Thus, this cannot explain the higher wealth level of the middle aged. Precautionary saving to cover unemployment spells cannot explain the difference either, since that would imply that those subject to employment risk, i.e., the young, should hold more of the asset. The reason for the wealth difference is that the young have to accumulate savings and that takes time.

Figure 10: Steady-state asset-market outcomes: Role of home-production retirement income



Notes: This figure plots steady-state average per capita holdings of the risk-free asset for the indicated type of agent and the equilibrium interest rate varying μ_{hp}^r between 0.5 and 1. μ_{hp}^r is the level of retirement home production which is defined as a fraction of the after-tax wage rate of middle-aged workers. The tax-financed retirement transfer is set equal to zero. The interest rate is the base rate which excludes the level-dependent premium.


Figure 10 considers the case when retirement benefits are in the form of home production, i.e., $\mu_{hp}^r > 0$, instead of a tax-financed transfer. In this economy, an increase in μ_{hp}^r implies that overall resources in the economy increase. The qualitative impact of increases in μ_{hp}^r on the interest rate is identical to the impact when μ^r increases. The main difference between the tax-financed and home-production retirement system is that asset holdings of the middle aged drop much faster with μ_{hp}^r than with μ^r . Instead of the drop from 8.53 to 6.27, we now find a drop from 8.52 to 4.88. Moreover, an increase in μ_{hp}^r is not associated with an increase in government expenditures and risk-free asset holdings of the middle aged now also drop *relative* to the after-tax wage income. It is key to take into account that we keep the supply of the risk-free asset constant. As μ^r increases, the after-tax wage rates of both the young and the middle-aged workers decrease substantially. But this means that it becomes harder for the young to accumulate wealth when μ^r increases than when μ_{hp}^r increases. Consequently, asset holdings by the middle aged drop by less when increases in retirement benefits are accompanied by a tax increase.

Changes in the supply of the risk-free asset. The available quantity of the risk-free asset plays a key role in models with heterogeneous agents. For example, consider the case when agents cannot borrow. Consumption smoothing during unemployment spells is still possible when the supply is positive because workers would accumulate assets while employed and sell these when unemployed. But this would not be possible if the supply of the asset is equal to zero. Similarly, the supply of the risk-free asset affects the ability of the middle-aged to save for retirement. Section 6.2 discusses in detail how the supply of the

risk-free asset, $\bar{d}$, affects individual consumption patterns when agents face changes in their employment status and when they age. Here we discuss how $\bar{d}$ affects steady-state values of some aggregate outcomes, i.e., averages for different worker types.

Figure 11 shows how the average holdings of the risk-free asset and the equilibrium interest rate vary with the supply of the risk-free asset, $\bar{d}$. The impact on these key variables is again substantial. A thirty-three percent increase in $\bar{d}$ from our baseline value of 3 to 4 leads to an increase in the annual real interest rate of 39 basis points. The figure documents that the increased supply ends up mainly in the hands of the middle aged. We will see in section 6.2, however, that the somewhat higher holdings by the young do help significantly in reducing consumption drops upon displacement. The higher supply ends up mainly in the hands of the middle aged for the following reason. An increase in $\bar{d}$ goes together with an increase in the tax rate, both because higher government debt directly increases interest payments and because of the general-equilibrium increase of the interest rate. For example, the increase in $\bar{d}$ from 3 to 4 leads to an increase in the tax rate from 29.3% to 31.2%. If the after-tax wage rate drops relative to the supply of available assets, then it will take longer for agents to acquire the equilibrium supply which in our environment means that the young will hold relatively fewer assets than the middle aged.⁶³

Figure 11: Steady-state asset-market outcomes: Role of asset supply quantity




Notes: This figure plots average per capita holdings of the risk-free asset for the indicated type of agent and the equilibrium interest rate varying the total level of risk free assets in the economy, $\bar{d}$ between 1 and 5. Other parameter values are set equal to their baseline parameter values. Specifically, $\mu^r = 0.5$ and $\mu_{hp}^r = 0$. All values are steady-state values.

Relaxation of the borrowing constraint. Figure 12 displays the results when we relax the borrowing constraint and considers values for $\bar{\eta}_d$ that are below our baseline value of -0.3 . If the borrowing

⁶³Given our calibration strategy, the unemployment rate remains fixed and is not affected by changes in equilibrium changes in variables like the tax rate and the interest rate.

constraint is an inequality constraint, then a relaxation simply means that agents can borrow more. In our more general specification, a reduction in $\bar{\eta}_d$ not only means that agents can borrow more, but also that the interest rate for levels of debt below the previous limit decline.⁶⁴ A reduction in $\bar{\eta}_d$ allows young agents to borrow more during unemployment spells which in equilibrium implies that the middle aged hold more of the risk-free asset and the real interest rate increases, that is, the gap between the real interest rate and the discount rate gets smaller. Thus, both an increase in the supply of the risk-free asset, $\bar{d}$, and a relaxation of the borrowing constraint help unemployed agents to smooth consumption. However, there are some key differences. An increase in $\bar{d}$ affects the tax rate both directly because of increased borrowing by the government and because of an increase in the equilibrium interest rate, whereas a decrease in $\bar{\eta}_d$ increases the tax rate only because of the change in the interest rate.

Figure 12: Steady-state asset-market outcomes: Role of borrowing restrictions



Notes: This figure plots average per capita holdings of the risk-free asset for the indicated type of agent and the equilibrium interest rate for values of the borrowing constraint, $\bar{\eta}_d$, between -3 and -0.3.


Changes in the life-cycle income profile. Figure 13 plots the equilibrium interest rate and the per capita average equilibrium amount of the risk-free assets for different types of households as a function of the steepness of the life-cycle income profile. A value of $z^{m/y}$ equal to 1 implies that the middle aged are as productive as the young and therefore have the same wage rate. By contrast, workers face an increase in labor income with age $z^{m/y} \geq 1$. We lower the overall productivity, $\bar{z}$, when we increase $z^{m/y}$ to keep the overall amount of resources in the economy the same.⁶⁵

The figure documents that asset holdings of the young fall when $z^{m/y}$ increases. The reason is that

⁶⁴That is, the penalty term displayed in figure 1 would not only shift to the left allowing for larger short positions, but would also take on lower values at positions that are possible at the previous value of $\bar{\eta}_d$.

⁶⁵If an increase in $z^{m/y}$ increases resources, then this would imply a reduction in the supply of the risk-free asset *relative* to available resources and the results above indicate that changes in the supply of the asset have large effects on the interest rate.

Figure 13: Steady-state asset-market outcomes: Role of lifetime income profile



Notes: This figure plots average per capita holdings of the risk-free asset for the indicated type of agent and the equilibrium interest rate when we vary productivity of the middle-aged relative to the young, $z^{m/y}$ from 1 to 1.7. All other parameter values are equal to their baseline values. All values are steady-state values.

workers save less when young in anticipation of higher wages later in life. The effect on the interest rate is small. Specifically, the annual interest rate increases by only 6 basis points when $z^{m/y}$ increases from 1 to 1.7. The reduced desire to save by the young should put upward pressure on the interest rate. However, lower savings by the young mean that when middle aged they have less wealth build up relative to their higher income when middle aged. Thus, to safeguard their retirement income, increased saving by the middle aged would push the interest rate in the opposite direction.

6 The “battle” for the safe asset

Financial frictions limit the ability of workers to insure themselves against post-displacement consumption drops which necessitates the need for precautionary savings by employed workers. A large literature has emerged to study how changes in the level of precautionary savings – in response to changes in (perceived) unemployment risk – amplify fluctuations over the business cycle.⁶⁶ But the unemployed are not the only ones who face potentially large consumption drops and use financial markets to soften the blow. The middle aged save to compensate for the drop in income upon retirement.⁶⁷ Desired changes in

⁶⁶See footnote 5.

⁶⁷The drop in income is associated with a drop in spending. As discussed in [Olafsson and Pagel \(2018\)](#), there is a debate in the literature on whether this implies a reduction in the actual amount of “goods and services” enjoyed. After all, retirement may reduce the opportunity costs for efficient shopping and home production. Indeed, some older studies focusing on “nutrient intake” found that this does not drop upon retirement. But [Stephens and Toohey \(2025\)](#) argue that those previous studies rely on biased estimators and their alternative unbiased estimator does indicate a drop. Another related aspect issue is that there would be less need for “work-related” consumption. Finally, health issues and associated precautionary motives may effect savings. Our model allows for increased home production by setting $\mu_{hp}^r > 0$. What matters for our purpose is that the middle aged have substantial savings and, thus, face a nontrivial income effect when the central bank changes interest rates. Our framework allows us to investigate whether that effect is smaller when μ_{hp}^r is larger and the consumption drop upon retirement smaller.

such savings may very well affect the effectiveness of monetary policy, just as changes in precautionary savings related to employment risk have been shown to be important.

In this section, we focus on the general-equilibrium *interaction* between precautionary savings to deal with unemployment spells and saving for retirement. Specifically, we consider different environments to study two key questions. First, how do changes in key parameters affect consumption variation across unemployment spells and over the life cycle, i.e., consumption levels when young, middle aged, and retired. Second, how do these changes in the environment affect the distributional and aggregate consequences of monetary policy.

Specifically, we will study these two questions for the following experiments: (i) changes in retirement benefits in section 6.1, (ii) changes in the supply of the risk-free asset and relaxation of financial constraints in section 6.2, (iii) changes in the life-cycle income profile in section 6.3, and (iv) a change in the financial environment such that workers have (indirect) access to investment in capital and firm ownership by trading with firm owners in the risk-free asset in section 6.4. We will document that there can be a “battle” in the sense that improved consumption smoothing for one group implies larger consumption fluctuations for another. But that is not necessarily the case. For example, it is possible that policy changes that explicitly lower the consumption decline upon retirement also improve the ability for the young to smooth consumption.

6.1 Varying retirement benefits

In section 6.1.1, we study how the level of retirement benefits (and associated changes in tax rates and holdings of the risk-free asset) affect consumption variation over the life cycle and during unemployment spells. In section 6.1.2, we study how the level of retirement benefits affects the monetary transmission mechanism.

6.1.1 The role of retirement benefits for individual consumption outcomes

To understand how retirement benefits affect individual consumption levels, we analyze the following three cases. First, we compare the case where retirement benefits are transfers and financed with taxes ($\mu^r = 0.5$) with the case where they are at the same level, but consist of home production ($\mu_{hp}^r = 0.5$). Next, we check how individual consumption fluctuations are affected when we change the level of retirement benefits. Finally, we consider a particular policy experiment that introduces a tax-financed retirement transfer that ensures that retirement consumption no longer falls.

An important result is that retirement benefits affect consumption variability for all agents, not just the retirees and the middle aged. In particular, they also affect consumption drops during unemployment spells that the young experience. To calculate individual lifetime consumption paths, we focus on the case when the aggregate economy is in its steady state.

Tax-financed retirement benefits versus home production. Figure 14 plots a particular individual consumption path for the economy with a tax-financed retirement system, $\mu^r = 0.5, \mu_{hp}^r = 0$ and also for the economy with retirement home production, $\mu^r = 0, \mu_{hp}^r = 0.5$. Specifically, the agent starts out as a young agent and has been employed long enough to be at the highest possible consumption level. Next, the agent becomes unemployed for a long enough time period so that consumption reaches its lowest observed level.⁶⁸ Next, the agent ages and becomes a middle-aged agent for a while and eventually retires. In the economy with home production, there are additional resources which means that the consumption time path is uniformly higher.⁶⁹ To make the two consumption time paths comparable, we scale them so that the initial value is equal to 1 in both cases.

The figure shows that individual consumption volatility is higher when non-investment retirement benefits are in the form of home production. Specifically, with a tax-based retirement system consumption drops by 31% during a sustained unemployed spell and by 33% in the economy with home production. Upon retirement, consumption drops by 38% in the economy with tax-financed transfers and with 41% in the economy with retirement home production. Thus, having retirement home production instead of tax-financed retirement transfers increases consumption volatility for the unemployed, as well as the retirees. What explains these findings? Although the differences are not that big quantitatively, the answer to this question reveals an aspect of the model that turns out to be important throughout this paper. When retirement benefits are a transfer, then this increases the tax rate. Retirement benefits are specified as a fraction of *after-tax* pre-retirement wages, which means that non-investment retirement benefits drop proportionally when the tax rate increases. Moreover, the supply of the risk-free asset in the two economies is the same in absolute terms. Relative to total resources, however, it is less in the economy with the additional home-production resources. Consequently, investment income provides *relatively* less of a buffer in the richer economy with home production, which explains the larger percentage post-retirement consumption drop. The larger consumption drop during unemployment spells in the home-production economy can be explained with the same reasoning. If retirement benefits are in the


⁶⁸And the amount borrowed is close to the implicit borrowing constraint implied by our penalty function.

⁶⁹The *relative* drop in non-investment income upon retirement is the same in both economies, because the retirement-benefit parameter determines non-investment retirement income as a fraction of *after-tax* income of the middle aged in both cases.

form of home production instead of transfers, then *after-tax* wages are higher. This means, however, that the fixed supply of the risk-free asset will provide less insurance. Throughout this paper, we will see that the amount of available assets to workers relative to other aspects of the economy is key for several outcomes, including distributional ones.

As resources in the economy with home production are higher, one can expect all agents to demand more of the risk-free asset. That is not possible in general equilibrium with a fixed supply of the asset. So the question arises whether the *distribution* of this fixed supply remains the same or changes. It turns out that average per capita holdings of the risk-free asset by the middle aged (young) are 6.0% lower (21.3% higher) in the economy with home production. Thus, the increase in demand in the economy with home-production retirement benefits is stronger for the young than for the middle aged. One relevant aspect is that the lower taxes in the home-production economy imply a larger (absolute) increase in the Net Present Value of labor income for the young than for the middle aged as they have a longer expected remaining life span.

Figure 14: Individual consumption time paths: Role of type of retirement benefits




Notes: This figure plots the life-cycle path for consumption of an agent when $\mu^r = 0.5$ and $\mu_{hp}^r = 0$ (tax-financed retirement system) and when $\mu^r = 0$ and $\mu_{hp}^r = 0.5$ (home production). The starting point is a young agent who has been employed for a while and has reached their highest consumption level. Then the agent has a sustained unemployment spell during which they reach the lowest observed consumption level. After another employment spell, the agent becomes middle aged and then retires. Aging is not deterministic, since we use an extended perpetual-youth model. Expected duration is the same for being young, middle aged, and retired, but we have chosen to give the middle-aged agent the retirement shock rather soon, since not much happens when middle aged and there are no aggregate shocks.

Changing the level of retirement benefits. Next, we consider changes in the level of retirement benefits. In section 5, we document that an increase in the level of retirement benefits lowers the demand for

the risk-free asset by the middle aged, which leads to a nontrivial increase in the real interest rate. Here we investigate how such changes affect individual consumption patterns, again in general equilibrium.

We first analyze how changes in retirement benefits that are in the form of home production affect individual consumption paths. This leads to more focused insights since changes in retirement benefits do not lead to direct changes in tax rates (which can be substantial).⁷⁰ Specifically, we consider values for μ_{hp}^r equal to 0.1, 0.5, and 1. In the next exercise, we consider a change in retirement benefits financed by taxes.

Figure 15: Individual consumption time paths: Role of level of home-production retirement benefits



Notes: The figure displays individual consumption time paths for the indicated level of retirement benefits which are in the form of home production. For additional notes see Figure 14.


Figure 15 displays a particular lifetime consumption profile for the three values of μ_{hp}^r considered. Not surprisingly, retirement consumption increases as μ_{hp}^r increases. In fact, consumption actually increases upon retirement when $\mu_{hp}^r = 1$. The reason is that there is no drop in post-retirement non-investment income, but savings taken into retirement are still positive and are used to buy the annuity which leads to an additional source of income. A more intriguing model result is that higher retirement benefits sharply reduce consumption fluctuations for the young, that is, the ones in our model who are subject to unemployment risk. When $\mu_{hp}^r = 0.1$, a long-term unemployed worker faces a consumption drop of 50.4%, whereas this number is only 9.2% when $\mu_{hp}^r = 1$. The reason is that when μ_{hp}^r increases, a

⁷⁰Since we have an equilibrium model, the tax rate still changes slightly because the interest rate the government pays on its debt varies. However, aggregate activity as well as the unemployment rate are guaranteed to remain the same by our calibration strategy.

larger fraction of the supply of the risk-free asset is available to help insure against unemployment spells of the young. Specifically, for $\mu_{hp}^r = 0.1$ average per capita bond holdings of the middle aged are equal to 8.32 and that of the young only 0.68. By contrast, these numbers are equal to 4.88 for the middle aged and 4.12 for the young, respectively, when $\mu_{hp}^r = 1$. Thus, even when the retirees face no drop in non-investment income upon retirement, it is still the case that the middle aged are substantially richer than the young. An important factor is that the young have to build up their savings.⁷¹ Another relevant element is that average income of the young is lower because of unemployment spells.

The retirees benefit substantially from the higher interest rate when μ_{hp}^r is higher. When μ_{hp}^r increases from 0.1 to 1, the savings brought into retirement drop by 41.9%, but the associated increase in the annual real interest rate is equal to 6.9 percentage points. Consequently, consumption financed by the purchased annuity actually increases with 55% for this change in μ_{hp}^r .

Figure 16: Individual consumption time paths: Impact of a change in tax-financed retirement benefits



Notes: The figure displays individual consumption time paths for the indicated level of tax-financed retirement benefits. The positive value for μ^r is chosen such that consumption upon retirement is equal to that when middle aged. In addition, retirees benefit from home production in both cases ($\mu_{hp}^r = 0.5$). For additional notes see figure 14.

Higher retirement benefits financed with taxes. Would these results be different when higher retirement benefits are in the form of a tax-financed transfer instead of home production? Figure 16 shows two consumption lifetime profiles for two different values of μ^r . The first is when retirement income is solely based on home production and μ_{hp}^r is equal to 0.5. For this level of retirement income, there is a substan-

⁷¹Individual newborns are assumed to resemble existing newborns in terms of wealth which is finance as a transfer from other young agents. These newborns replace those who age who keep their assets. Thus, the young as a group have to invest to get assets at desired levels.

tial drop in consumption upon retirement. The second profile shown is for the case when a tax-financed transfer to the retirees is introduced and home production benefits remain the same. The value of μ^r is set equal to 0.4175. At this value of μ^r , there is no longer a consumption drop for the retirees. The associated increase in the tax rate causes the income of both the employed and the unemployed to drop.⁷² Thus, it is not surprising that this tax-financed transfer to the retirees leads to a reduction in *average* consumption for young and middle-aged workers. Given the reasoning above, it is also not surprising that the young face reduced relative consumption fluctuations during unemployment spells. As retirement income is increased, the middle aged save less which means that a larger share of the available supply of the risk-free asset is available for the young which makes it easier for them to buffer consumption upon unemployment. The increase in retirement benefits, however, implies an increase in the tax rate and a 20% reduction in unemployment benefits. But the higher level of savings improves the ability to smooth consumption by so much that the consumption level of a long-term unemployed agent is actually higher even though unemployment benefits have fallen.

Importance of general equilibrium. The key conclusion that can be drawn from these exercises is that the level of retirement benefits also affect those who are quite distant from retirement. This is only partly because the young are forward looking. General-equilibrium consequences are important as well. As those closer to retirement change their savings behavior, then this changes holdings of the financial asset by the young and the associated equilibrium interest rate. This in turn affects their ability to buffer consumption drops during unemployment spells.

6.1.2 The role of retirement benefits for the monetary transmission mechanism

In this section, we address the question whether the source of non-investment retirement income and its level matters for the impact of monetary-policy shocks on aggregate-level variables and/or in terms of distributional consequences.

General patterns of responses. Figure 17 displays the impact of a monetary-policy shock when non-investment income is a tax-financed transfer (our baseline) and when it is in the form of home production.⁷³ For both retirement regimes, we get qualitatively similar results that are consistent with conventional views on the monetary transmission mechanism. That is, a monetary-policy shock leads to a drop

⁷²Also, unemployment benefits are defined as a fraction of the *after-tax* wage, which implies that the income of the unemployed drops (and by the same percentage).

⁷³The size of the shock is set to generate a maximum twenty-five basis point drop in the nominal rate, which is obtained in either the third or the fourth period.

in both the nominal and the real interest rate. The shock is fairly persistent and the nominal interest rate is still around ten basis points from its pre-shock level after three years. The increase in inflation is only a few basis points, but remains positive for around three years. Wage inflation is also moderate. Output, employment, and aggregate consumption increase as is standard. Magnitudes are possibly at the upper end of empirical estimates, but this has the advantage that we can study our channels of interest in an environment in which we do not underestimate the role of monetary policy.⁷⁴ Movements in the average premium component of the interest rate are always small. For example, for the baseline case the average premium drops by 0.62 basis points on an annual basis, then increases and eventually turns positive with a peak positive response of 0.66 basis points.⁷⁵

When analyzing the heterogeneous impact of monetary policy on different types of workers, we average the results by type, that is, for young employed workers, young unemployed workers, the middle aged, and the (newly) retired.⁷⁶ Average consumption increases for both the employed and unemployed young workers as well as the middle aged, but there are substantial quantitative differences. When non-investment pension income is a tax-financed transfer, then we find that the increases in consumption on impact are equal to 1.77%, 1.34%, and 0.80% for the young unemployed, the young employed, and the middle aged, respectively.⁷⁷ For the home-production case, we find a similar qualitative pattern. Quantitative differences across the two retirement schemes will be discussed below. One important reason for the differential responses across worker types is that the increase in the job-finding probability benefits the young directly, especially the unemployed.^{78,79} Another relevant aspect is that – as shown in

⁷⁴Figure 3 in [Miranda-Agrippino \(2017\)](#) reports estimated peak responses for industrial production between roughly 2 and 3 percent for a 1 percentage point change in the policy rate, which means a range between 0.5 and 0.75 percent for a 25 basis point drop. Standard-error bands, however, are quite large. Thus, our results are in line with these estimates. Of course, identification issues make it nontrivial to determine the true empirical impact of a monetary-policy shock.

⁷⁵Recall that the interest paid/received depends on the position taken and grows nonlinearly when the magnitude of the short position increases.

⁷⁶The consumption level for retirees depends only on aggregate conditions holding in the period when they retired. This does introduce heterogeneity among current retirees, but the heterogeneity among those that retired in the past is not affected by subsequent monetary-policy shocks as the annuity income of retirees is fixed in *real* terms.

⁷⁷[Bardóczy and Velásquez-Giraldo \(2024\)](#) find that during the first year of a monetary expansion, more than half of the aggregate consumption response is due to households below the age of 40. For us, that number is 57.2%.

⁷⁸The middle aged benefit directly from the increase in the pre-tax wage rate, but this effect is quite small since wages are quite sticky. If we assume that wages are fully sticky, then the initial increase in consumption of the middle aged drops from 0.80% to 0.76%. If – in addition – the tax rate on the wage rate of the middle aged is kept equal to its steady-state value, then the initial increase drops to 0.48%, but the response now turns negative after 12 quarters with a peak negative response of -0.093%. We focus on the consumption response averaged across all existing middle-aged workers. But as time goes by, this will include some workers who were young when the shock started. Thus, even though income of the middle-aged is kept fixed in this latest exercise, it is not quite similar to the partial equilibrium exercise of section 4.

⁷⁹One disadvantage of our focus on the average impact across worker types is that these are affected by composition effects. For example, the *share* of recently-separated workers among the unemployed increases (because the separation rate is fixed and the job-finding rate increases) and the recently separated have a higher level of assets than the average unemployed. This leads to a further increase in the consumption response for this group.

section 4.1 – consumption demand is more interest sensitive for young agents. Or in terms of the usual terminology, the middle aged hold more assets relative to their remaining lifetime income. In fact, they hold most of the supply. Specifically, per capita deposit holdings of the middle aged are equal to 138.0% of after-tax annual income in the steady state compared to 39.0% for the young. This means that the negative income effect associated with a reduction in the real interest rate is larger for the middle aged than the young. The reduction in the interest rate will put downward pressure on retirement investment income. One possible response would be for the middle aged to save more. Since the supply of the risk-free asset is fixed, this would mean that the young would have to hold less. We see the opposite, that is, risk-free asset holdings of the middle aged display a persistent drop which means that the holdings of the young employed and young unemployed *combined* must increase.⁸⁰ The reason is that the young benefit more from the expansion since they benefit from the increase in the pre- and after-tax wage rate like the middle aged do, but they also benefit from the increase in the job-finding rate. This means that in equilibrium in which the supply of the asset is fixed, the young buy more of the asset whereas the middle aged sell. Finally, the young are closer to the constraint on borrowing which puts upward pressure on the marginal propensity to consume.

On impact, consumption of the newly retired must fall because the reduction in the interest rate makes the annuity more expensive and the savings they brought into retirement were chosen before the shock occurred. In subsequent periods, consumption is also negatively affected by the lower savings levels chosen by the middle aged.

Consumption of the firm owner typically declines in New-Keynesian models following a monetary expansion because lower markups imply lower firm profits. We avoided this outcome by assuming that firm owners also earn wages and parameters are such that the increase in wage income more than offsets the drop in profits.

Incomplete markets and the impact of monetary policy. The literature has pointed out that business cycles are magnified if consumers are less well insured against employment risk.⁸¹ Moreover, the previous subsection showed that the generosity and the type of pension benefits could have a strong impact on the ability of young workers to smooth consumption during unemployment spells. And this turns

⁸⁰This may seem inconsistent with the plotted IRFs for the first few periods, which indicate that deposit holdings of both types of young workers drop. The explanation is a composition effect. The increase in the job-finding rate means that the fraction of young employed workers who were recently unemployed increases. Since these are less wealthy, the average savings level across young employed workers falls.

⁸¹See footnote 5.

out to be important to understand how changes in the retirement scheme in place do or do not affect the monetary transmission mechanism.

To quantify the importance of incomplete markets on the IRFs describing the impact of a monetary-policy shock, we consider the case when unemployment benefits are set equal to zero, that is, the replacement rate, μ^y , is equal to zero instead of its benchmark value of 0.5. Then, the consumption drop upon a long unemployment spell increases from 30.7% to 43.7%.⁸² And, the peak employment response following an expansionary monetary-policy shock increases by 11.8% from 0.853% to 1.010%. This may underestimate the role of incomplete markets for the following reason. In our economy, unemployment benefits are a fraction of the after-tax wage rate which means that they increase during an expansion. This will dampen the reduction in the tax rate induced by the fall in the level of unemployment. This matters since the unemployed have a higher marginal propensity to consume. No such mechanism is in place when μ^y is equal to zero. Thus, to better access the role of incomplete markets, we assume that the level of unemployment benefits is kept fixed to its steady-state value. Then the peak increase in employment is equal to 0.779% when $\mu^r = 0.5$. This would indicate that incomplete markets adds 1.010% minus 0.779%, i.e., twenty-three basis points, to the employment response.

In the remainder of this section, we first discuss whether the *type* of retirement benefits matters for the responsiveness of economic variables to monetary-policy shocks. Next, we focus on the relevance of the *level* of retirement benefits.

Aggregate responses across different retirement schemes: transfers versus home production. To understand the role of the retirement scheme, it may be useful to recall what is different and what is not across the two economies. Our calibration procedure is such that the steady-state values of key economic aggregates such as GDP, employment, and *market-produced* consumption are exactly the same in both economies. Although market-produced consumption is the same in both economies, part of this is consumed by retirees under the transfer scheme whereas this is much less so with home production.⁸³ When retirement benefits consist of home production, it makes sense to assume that they do not fluctuate with economic conditions. To make the two economies comparable, we assume that non-investment retirement benefits under the transfer system are of the same magnitude and also do not vary in response

⁸²The increase in the consumption drop during unemployment spells is dampened in general equilibrium since the reduction in the replacement rate induces a large shift in who holds the risk-free asset. Specifically, whereas the young hold 22.0% of the available supply under our baseline calibration, this increases to 65% when the replacement rate is equal to zero.

⁸³With home production, retirees augment their consumption obtained through home production with market-produced consumption financed with the annuity.

to aggregate shocks. Thus, the transfer-based economy differs from the home-production economy because (i) steady-state consumption of workers is lower since a larger fraction of the consumption goods produced in the market goes to retirees and (ii) part of aggregate market-produced consumption is fixed, namely the part that goes to retirees.⁸⁴ The lower value for worker's consumption level is brought about by higher taxes, thus a related difference is that *after-tax* wages are lower in the economy with tax-financed pension benefits. With a fixed supply of the risk-free asset, this means that *relative* to after-tax wages, workers hold more savings.⁸⁵

Figure 17 documents that there are substantial differences between the IRFs of a monetary-policy shock in the two economies. For example, the peak increase in employment is equal to 0.853% in the transfer economy and 1.232% in the home-production economy.⁸⁶ The previous paragraph makes clear that this cannot be due to different steady-state values for key economic activity measures like GDP or employment. To disentangle our results, we first consider the case when the value of the monetary-policy shock is the same in both economies and equal to the value in the baseline economy. Then, we find that the peak employment increase is equal to 0.902% with home production and the drop in the nominal interest rate is only 18.3 basis points, compared to 25 basis points in the baseline. Although the response in aggregate real activity is reduced, it is still larger than the 0.853% observed in the baseline economy with tax-financed retirement transfers. This stronger response in real activity is responsible for the smaller drop in the nominal interest rate since our Taylor rule assumes some tightening when economic activity increases. Thus, when we increase the value of σ_{eR} to get the same drop in the nominal interest rate in both economies, the difference in the employment responses across the two economies increases.

The question that remains to be answered is why economic activity responds less strongly in the economy with tax-financed retirement benefits to begin with, that is, when the monetary-policy shocks are of the same size. Above, we pointed out that an increase in the inability to smooth consumption magnifies the impact of monetary-policy shocks on the economy. Thus, one possible reason is that agents in the economy with tax-financed retirement benefits face smaller relative changes in consumption during unemployment spells. As pointed out above, it is indeed the case that during an extended unemployment

⁸⁴Non-investment retirement benefits are not only fixed for those who retired in the past, but also for newly retired workers. Of course, non-investment retirement income is also fixed with home production, but this is not part of market-produced consumption (and not part of GDP either).

⁸⁵This has a positive effect on the ability of young workers to smooth consumption during unemployment spells.

⁸⁶Throughout this section, we focus on employment, but the story is basically the same for output. Specifically, the corresponding numbers for GDP are 0.554% and 0.801%.

spell consumption drops by 30.7% with tax-financed benefits and with 32.9% with home production.⁸⁷ Although there is a bit more incompleteness in the economy with home-production benefits, the difference is too small to be significant for the substantial difference in the aggregate responses in the two economies.⁸⁸ Since this channel does turn out to be quantitatively important in subsequent experiments, we do want to point it out here as well even though its quantitative role is small.

This lower responsiveness of aggregate activity with tax-financed benefits is consistent with the fact that an important demand component of GDP is fixed in this environment, namely the market-produced consumption that retirees receive as a transfer.⁸⁹ But is this consistent with the behavior of interest rates? Is it not the case that the Euler equation implies that at least the real interest rate is less volatile when consumption is less volatile. To understand the following recall that the economy with tax-financed transfers does not only differ from the home-production one in that a share of market-produced consumption is fixed, but also that consumption levels are lower since there is no home production. The latter is important since the Euler equation tells us that *relative* changes in consumption matter for the volatility of interest rates.⁹⁰ And relative to its lower steady-state value consumption actually increases by more with tax-financed benefits when the same value for σ_{eR} is used.⁹¹ Specifically, for aggregate consumption, the two numbers are 0.556% and 0.513%.⁹² The Euler equation implies that a larger *percentage* increase in consumption goes together with a larger drop in the real interest rate. This is indeed what we find, namely 27.4 versus 21.1 basis points. And it turns out that the inflation response is very similar, which means that this difference in the response in the real rate is consistent with the observed difference in the responses of the nominal interest rate.

Distributional responses across different retirement schemes: transfers versus home production.

Above, we already discussed why the young – and especially the unemployed – experience the largest consumption increases following a monetary expansion. Here, we discuss the quantitative differences

⁸⁷This is the general equilibrium outcome and this difference is dampened by the endogenous change in who holds the deposits. Specifically, the young hold 22.0% of the supply in the baseline economy and 26.7% in the economy with home production.

⁸⁸Specifically, when we decrease the replacement rate from 0.5 to 0.429 in the baseline economy, then the consumption drop upon displacement is as big as the one in the home-production economy. But the peak employment response only increases from 85.3% to 87.0%, so it is still substantially lower than the 1.232% that is observed in the home-production economy.


⁸⁹Recall that we assume that retirement benefits are fixed in both economies to make the two economies comparable. But retirement consumption financed by taxes is part of GDP whereas home-production benefits are not.

⁹⁰That is, how strongly $c_{t+\tau}$ with $\tau \geq 1$ responds relative to c_t .

⁹¹Another way to think about this is the following. The fixed obligation to finance retirement benefits becomes relatively smaller as the economy expands. This leads to an additional reduction in the tax rate and means that the consumption response relative to its steady value is higher with tax-financed benefits.

⁹²And we find the same ranking for the consumption responses for all agent types.

Figure 17: Monetary-policy shock IRFs: Tax-financed (—) versus home-production (---) retirement benefits



Notes: This figure displays the IRFs for a monetary-policy shock for the economy with tax-financed retirement benefits (solid and black) and for the one in which retirement benefits consist of home production (dashed and red). For the tax-financed regime, we have $\mu^r = 0.5$ and for the home-production case, we have $\mu_{hp}^r = 0.5$. Thus, retirement benefits are the same as a fraction of pre-retirement after-tax wages. Magnitude of the monetary-policy shock is normalised to generate a 25 basis point decline in the nominal interest rate. Specifically, the values for σ_{ε^r} are equal to 0.0019165 and 0.0026169 for the case with tax-financed benefits and the one with home production, respectively. Other parameters are set according to the baseline parameterization.

across the two pension regimes. To correct for the different impact on aggregate activity, we consider how the *relative* consumption responses across worker types depend on the pension regime. With tax-financed pension benefits, the consumption response of the young employed is 1.68 times as big as the one for the middle aged and for the young unemployed this number is equal to 2.22. For the home-production pension benefits, the two corresponding numbers are equal to 1.91 and 2.88. The reason for the difference is that the employment response is stronger with home-production pension benefits, which benefits the young.⁹³

Role of the level of retirement benefits for monetary policy. Next, we discuss whether the level of μ^r matters for the impact of monetary policy.⁹⁴ Figure 18 displays the IRFs for a monetary-policy shock. The size of the shock is again rescaled to ensure that we get a maximum twenty-five basis point drop

⁹³Note that the observed small increase in steady-state holdings of the risk-free asset by the young in the home-production economy goes in the opposite direction as this would increase the negative income effect.

⁹⁴In terms of the lessons learned, the same results are obtained when we vary the level of post-retirement home production.

in the nominal interest rate for each of the three values of μ^r considered. The values considered for μ^r are 0.2685, 0.5, and 1. A lower value for μ^r implies a larger drop in non-investment retirement income but this drop is limited in general equilibrium because the middle aged save more when μ^r is lower. Consequently, a smaller fraction of the supply of the risk-free asset is available for the young, which in turn means that they are less capable of smoothing consumption during unemployment spells. And we know that the latter can magnify business cycles. When μ^r is equal to 0.2685 (and μ^y equal to its benchmark value of 0.5), then the drop in consumption during an extended unemployment spell is equal to 43.7%, which is exactly equal to the case discussed above when unemployment benefits are equal to zero. This compares to a drop of 30.7% in our baseline scenario with $\mu^y = \mu^r = 0.5$.

The key observation is that the generosity of the pension system has a large effect on the impact of monetary policy on the economy. Specifically, the peak response of employment is equal to 1.30%, 0.85%, and 0.54% when μ^r is equal to 0.2685, 0.5, and 1, respectively.⁹⁵

The results in the figure are based on the case when the size of the shock is adjusted to generate the same twenty-five basis point drop in the nominal interest rate. If the size of the shock is kept the same as in the baseline calibration and μ^r is equal to 0.2685, then the peak increase in employment is only equal to 0.91% and the drop in the nominal interest rate only 17.5 basis points. Thus, the increase in the employment rate is still bigger than in the baseline scenario, but less so. Whenever parameters are such that the impact on real activity is stronger, then monetary policy will make the difference smaller because the Taylor rule is such that there is upward pressure on the nominal interest rate when output increases. And that mechanism is quantitatively important as the endogenous response of the nominal interest rate is only 17.5 instead of twenty-five basis points.

One reason why the impact on economic activity is larger when μ^r is lower is that workers are less well insured which happens in equilibrium because at lower values of μ^r a larger fraction of the supply of the asset goes to the middle aged. As mentioned above, the drop in consumption during an extended unemployment spell when μ^r is equal to 0.2685 (and μ^y equal to 0.5) is the same as when μ^y is equal to zero (and μ^r equal to 0.5). Thus, these two economies are comparable in terms of the magnitude of idiosyncratic unemployment risk. Above, we calculated the impact of this level of incompleteness to be twenty-three basis points. This is a substantial fraction of the difference between the baseline peak response of 0.85% and the 1.30% peak response observed here, but not the complete story.

⁹⁵ And the peak response would be even bigger when even smaller values for μ^r are considered. For example, it is equal to 1.59% when μ is equal to 0.1, the value considered in the previous subsection.

Another relevant part of the story is the analysis of section 4.1 which makes clear that lower values of μ^r are associated with lower values of the steady-state level of the interest rate which in turn imply a higher sensitivity of demand to interest rate changes. And as shown in figure 10, the level of non-investment retirement benefits have a strong impact on the steady-state real interest rate in the full model as well. Moreover, the analysis of the simple OLG model indicates that changes in μ^r mainly affect demand elasticities of the young. And that is exactly what we find here. When μ^r is equal to 0.2685, the peak impact of a monetary-policy shock on aggregate consumption is equal to 0.82% compared to 0.56% in the baseline, that is, an increase of 52.2%. But the peak consumption response of the middle aged only increases to 0.83% from 0.80%. By contrast, the peak responses for the young increase from 1.35% to 1.75% and from 1.77% to 3.10% for the employed and unemployed, respectively.⁹⁶

The model is very rich and there may be other mechanisms that strengthen or offset the ones discussed here. But the result that the type and generosity of the pension system affect the impact of monetary policy is clear.

6.2 Higher bond supply versus relaxation borrowing constraint


A relaxation of the borrowing constraint makes it easier for individuals to smooth consumption in response to income fluctuations. Aiyagari and McGrattan (1998) argue that an increase in the supply of government bonds can do the same.⁹⁷ The reason is quite intuitive. The increased bond supply introduces an asset, i.e., a stock, and an offsetting tax liability consisting of a flow of future tax payments. On average, these would offset each other. When *individual* agents face an unemployment spell, however, they can liquidate (part of) the asset to smooth their consumption.

In this section, we analyze how an increased bond supply compares with a relaxation of borrowing constraints in our general-equilibrium framework. In particular, we check whether the mechanism described in the previous paragraph also holds in our framework with rich heterogeneity, that is, when agents do not only differ in employment status, but also in age. Since we have a general-equilibrium framework, we can study whether the interest and tax rates respond differently in response to these two changes in the economic environment. Age heterogeneity may matter because the question arises who

⁹⁶The analysis of section 4.1 indicates that a drop in μ^r from 0.5 to 0.2685 corresponds with a change in the interest rate elasticity for consumption of the young from -1.08 to -1.44 and for consumption of the middle aged from -0.60 to -0.61. Thus, the change for the young is much bigger than the one considered here. However, the interest rate change considered here does not last a generation. Nevertheless, the results of the full model are consistent with the remarkable outcome that the IRFs consumption responses of the young do depend on the value of μ^r and the ones for the middle aged to not.

⁹⁷As long as government expenditures are kept constant.

Figure 18: Monetary-policy shock IRFs: Dependence on the level of retirement benefits



Notes: This figure displays the IRFs for the indicated variables associated with a monetary-policy shock for three different levels of tax-financed retirement benefits. Magnitude of the monetary-policy shock is normalized to generate a 25 basis point decline in the nominal interest rate. Specifically, the values for σ_{ε^R} are equal to 0.0027320, 0.0019165 and 0.0013114 for μ^r equal to 0.2685, 0.5 and 1, respectively. Other parameters are set according to the baseline parameterization.

will hold the additional supply of the risk-free asset; those for whom additional supply could reduce consumption fluctuations caused by unemployment spells or those for whom it would imply a smaller consumption drop upon retirement.⁹⁸

Section 6.2.1 compares the impact of these two modifications on the steady-state time paths of consumption and asset holdings when employment status changes as well as when the agent ages and retires. Section 6.2.2 studies how the IRFs of monetary-policy shocks are affected.

6.2.1 Steady-state results

In our first modification, we increase the bond supply, $\bar{d}$, by 67% from 3 to 5. This increase mainly ends up in the hands of the middle aged whose average bond holdings increase by 81.6%, whereas average bond holdings of the young increase by only 13.7%. It makes sense that most of the increase goes to the middle aged as the young have less time to build up savings. Nevertheless, the relatively small increase in bond holdings for the young does help to stabilize their consumption during unemployment spells. In our baseline version of the model, the consumption of a long-term unemployed worker drops by 30.7%, whereas it drops by only 23.8% in this alternative environment.

The increase in $\bar{d}$ would be associated with higher tax rates to finance interest payments even without a change in the interest rate. In general equilibrium, the increased supply causes the interest rate to go up, which leads to a further increase in the tax rate. Specifically, when the bond supply, $\bar{d}$, increases from 3 to 5, then the model predicts the following steady-state increases: the annual interest rate increases from 4.37% to 5.13% and the tax rate from 29.3% to 33.4%.

To study the relaxation of the borrowing restriction, we lower $\bar{\eta}_d$. A reduction in $\bar{\eta}_d$ not only increases the maximum amount that agents can borrow, it also lowers the premium paid by borrowers at levels that were possible under the original baseline parameter value. To ensure that the percentage drop in consumption during a long unemployment spell is lowered by the same percentage as when bond supply increases, we lower $\bar{\eta}_d$ from -0.3 to -5.26 , which means that the maximum amount an agent can go short increases substantially. Aggregate resources are the same in both model modifications considered here.⁹⁹ We find that the young, including the employed, become borrowers. That is, their steady-state bond holdings fall from a positive level equal to 39% of their annual income to a negative 35%.


⁹⁸This is related to Samuelson (1958), who pointed out that introducing an asset like money or government bonds affects the ability to save for retirement.

⁹⁹Key is that our calibration procedure targets the same unemployment rate and is not affected by changes in equilibrium outcomes like the interest or tax rate.

With this modification, the interest rate also increases. The reason is that the relaxation of borrowing restrictions reduces the precautionary component of the demand for the risk-free asset. Specifically, the interest rate increases to from 4.37% to 5.02%, which is a substantial increase, but less than the one associated with the increase in bond supply. The steady-state tax rate increases from 29.3% to 30.6%. It makes sense that the tax rate increases by less for this second modification considered, because the interest rate increases by less *and* the government does not change the number of bonds issued.¹⁰⁰

Figures 19 and 20 plot time paths for individual consumption and bond holdings, respectively. Both for the baseline version of our model and the two modifications introduced here. Specifically, the starting point is the situation where a young worker has been employed for quite a while and his wealth has reached a steady-state value. Next, they face an unemployment spell of ten quarters, followed by a sustained period of employment before they become middle aged and eventually retire.


Figure 19: Individual consumption time paths: Higher supply of the risk-free asset versus relaxation financial constraint



Notes: This figure plots individual consumption paths for three environments: the baseline, the economy with a higher supply of the risk-free asset, $\bar{d}$, and the economy with more relaxed borrowing constraints, i.e., lower $\bar{\eta}_d$. The parameter changes for the two alternatives are such that the *relative* consumption drop during a sustained unemployed spell are equal. For additional notes see figure 14.

¹⁰⁰Because we have combined the government and the financial sector, there is one other minor factor that affects the tax rate. That is, the tax rate is also affected by profits made by the financial sector due to households paying more when borrowing than when saving. Although the effect is small, the higher level of borrowing associated with the reduction in $\bar{\eta}_d$ increases these intermediation profits which dampens the increase in the tax rate slightly. Also, see footnote 16.

Figure 20: Individual asset-holdings time paths: Higher supply of the risk-free asset versus relaxation financial constraint



Notes: This figure plots individual holdings of the risk-free asset for three environments: the baseline, the economy with a higher supply of the risk-free asset, $\bar{d}$, and the economy with more relaxed borrowing constraints, i.e., lower $\bar{\eta}_d$. The parameter changes for the two alternatives are such that the *relative* consumption drop during a sustained unemployed spell are equal. For additional notes see figure 14.

To interpret the figures, it is important to realize that an increase in the interest rate and the associated increase in the tax rate operates like a transfer to those who save more from those who borrow (or save less), because the increase in the tax rate affects workers equally, but the additional interest income is not distributed equally.¹⁰¹ And this mechanism is in place for both modifications. Keeping this in mind, we can make the following observations.

In the new steady state, this transfer from borrowers and small savers to big savers is larger when $\bar{d}$ increases than when $\bar{\eta}_d$ drops, since the former not only results in extra income for those holding most of the increased supply, it also leads to a larger increase in the interest rate. Given that there is such a transfer and given that aggregate resources remain the same, we can expect the consumption of the young to be lower under both modifications relative to the benchmark and the drop to be bigger for the increase in $\bar{d}$. As documented in the figure, however, this is only true for young employed workers. In fact, the increased capability to smooth consumption is so large that consumption during a sustained unemployment spell is actually higher than the benchmark value for both modifications. Perhaps not surprisingly, this is especially true when the borrowing restrictions are relaxed. As documented in figure 19, however, the “big winners” in terms of steady-state consumption are the retirees. Their consumption

¹⁰¹Unemployment benefits and retirement income are fixed as a fraction of after-tax wage income, so an increase in the tax rate also affects the non-investment income of the unemployed and the retirees proportionally.

increases by 15.2% when $\bar{d}$ increases and by 11.9% when $\bar{\eta}_d$ drops. The increase in $\bar{d}$ is more beneficial for the retirees as it not only increases their wealth when they retire, but also comes with the larger increase in the interest rate. For the middle aged, we see a drop in steady-state consumption although a smaller percentage decrease than for the young employed. The middle aged are of course also affected by the higher tax rate, but by holding more assets than the young, the middle aged benefit more from the increase in the interest rate. One aspect that lowers steady-state consumption of the middle aged is that the middle aged increase their asset holding by more when they age under both modifications, which – of course – depresses consumption.¹⁰²

6.2.2 Higher average bond supply versus relaxed borrowing constraint and monetary policy

Figure 21 displays the IRFs of a monetary-policy shock for the baseline environment and the two alternatives. The innovation standard deviations are chosen such that the peak response of the nominal interest rate is the same in all three environments and equal to 25 basis points.

Both the higher bond supply and the relaxation of the borrowing constraints lead to a smaller expansion. Specifically, whereas the peak employment response is equal to 0.853% in the baseline case, it is equal to 0.740% when borrowing restrictions are relaxed, and 0.750% when bond supply is increased. Both model modifications imply that agents are better insured which would dampen the magnification channel due to a reduction in precautionary savings as unemployment risk falls during the expansion.¹⁰³ For other aggregate variables we find relatively minor differences across the three different economies. Larger differences are observed across the responses for the population subgroups.

As discussed above, the young benefiting relatively more from the monetary expansion induces a relative shift of ownership of the fixed supply of the risk-free asset towards the young because of consumption smoothing motives. That is, the middle aged reduce their bond holdings and the young to increase theirs. The two changes considered induce a shift in steady-state wealth from the young to the old. This increases the negative income effect of the interest rate increase for the middle aged, but reduces it for the young.¹⁰⁴ Consequently, the mechanism just described becomes stronger because the benefits

¹⁰²Since we look at the steady state, we do not take into account the transition from the baseline to the alternative environments during which the middle aged have to sacrifice consumption to build up the larger asset holdings. We do take into account that the middle aged have to accumulate a larger increase in asset holdings when they become middle aged.

¹⁰³This is the well-known channel emphasized in the literature. See footnote 5. Quantitatively, the reduction in the impact of the shock is comparable to the one discussed in section 6.1.2. There we found that a reduction in the consumption drop during an extended unemployment spell from 43.7% to 30.7% was associated with a reduction in the employment response of twenty-three basis points. Here we are considering a reduction in the consumption drop from 30.7% to 23.8% and a reduction in the employment response of around ten basis points.


¹⁰⁴These effects may not cancel out because of nonlinearities in the model. In that case, this could affect the aggregate impact

of the monetary expansion for the middle aged are reduced relative to the benefits of the young. The result that the expansion and reduction in unemployment risk is smaller goes in the opposite direction. It turns out that the first effect dominates as we find that the reduction in bond holdings by the middle aged is larger for the two modifications considered and biggest for the increase in bond supply which also is associated with the biggest relative shift in steady-state wealth from the young to the old. The magnitude of the consumption responses of the young relative to the consumption responses of the old are also consistent with this mechanism. Specifically, we find the following. On impact, the baseline response of per capita consumption of young employed workers is 68% percent higher than the corresponding one for the middle aged. When we increase the bond supply, then this increases to 72% and increases further to 203% for the case with a relaxation of the borrowing constraint.

The larger drops in bond holdings of the middle aged lead to substantial larger drops in the consumption levels of the newly retired.

of a monetary expansion. But given that the aggregate responses are quite similar, this does not seem to matter.

Figure 21: Monetary-policy shock IRFs: High bond supply versus relaxation borrowing constraints



Notes: For the three different economies considered, this figure plots the IRFs for the indicated variables associated with a monetary-policy shock that leads to a maximum twenty-five basis points drop in the nominal interest rate. For the baseline calibration this meant setting σ_{ε^R} equal to 0.0018661. And the values for σ_{ε^R} for the case with a higher bond supply and looser borrowing constraints are equal to 0.0016723 and 0.0016380, respectively. Other parameter values are set according to the baseline calibration.

Comparing with previous experiment. How do the distributional differences documented here compare with the ones due to changes in μ^r ? There are some notable differences. When μ^r was lowered from 0.5 to 0.2685, the peak employment response increased from 0.556% to 0.825%, but the peak consumption response of the middle aged increased only a bit, namely from 0.800% to 0.830%. When $\bar{\eta}_d$ is lowered from -0.3 to -5.26, the peak employment response drops from 0.556% to 0.469%.¹⁰⁵ Thus, a much smaller change at the aggregate level. Nevertheless, whereas the consumption response of the middle aged was relatively constant when we varied μ^r , it now drops from 0.800% to 0.637%. But then the change in $\bar{\eta}_d$ led to a 53% increase in the steady-state holdings of the risk-free asset by the middle

¹⁰⁵The discussion is qualitatively similar for the increase in bond supply, but quantitatively the results are even stronger for the change in $\bar{\eta}_d$.

aged, which strengthens the negative income effect. The weaker employment response would imply that the consumption responses of the young *relative* to the one for the middle aged should fall. However, as the young now hold less of the risk-free asset their negative income effect gets smaller which would go in the opposite direction. In fact, the peak consumption response of the employed young basically remains the same.

To summarize, the comparative statics of this section show that changes in some key model parameters related to the ability of workers to insure themselves against unemployment fluctuations affect not only the impact of monetary-policy shocks, but also have important distributional consequences. And the key insight is that it is important to understand how changes in the economic environment affect who holds the risky asset.

6.3 Alternative life-cycle income profile

Our baseline results are based on a flat labor-income profile. This helps in understanding the mechanisms highlighted in this paper. But wages typically increase with age. This would increase the heterogeneity between the young and the middle aged? This section addresses the question whether this enhanced heterogeneity affects our results. Specifically, we consider the case when $z^{m/y} = 1.5$, that is, when productivity of a middle-aged agent is 50% higher than that of a young agent, a considerable increase. At the same time, we adjust the overall productivity so that total resources in the economy are not affected.


6.3.1 Steady state results

For this increase in the value of $z^{m/y}$, we find that average deposits of the middle aged increase by 14.6% and they decrease by 51.5% for the young. Since we keep aggregate resources fixed, the increase in $z^{m/y}$ implies an increase in the after-tax wage of the middle aged and a decrease for the young. This also means that the changes in bond holdings are a bit different when expressed relative to annual wage income. Relative to annual after-tax wage income, the holdings of the young drop by 36.1% (from 0.390 to 0.249), whereas for the middle-aged the fraction displays a mild increase (from 1.378 to 1.388). The real interest rate increases which is consistent with the young wanting to reduce their bond holdings to smooth consumption. But the increase is quite modest; the annual steady-state real interest rate increases from 4.37% to 4.42%.

Figure 22 plots a consumption time path for our baseline parameterization and the one in which the middle aged earn a much higher wage income. Again, it displays the time path of a young agent who has

been employed for a while, then becomes unemployed, followed by another sustained employment spell after which the agent “ages” and becomes middle aged and eventually retires. A striking observation is that consumption drops by more during unemployment spells for the higher value of $z^{m/y}$. Specifically, it drops by 30.7% when $z^{m/y} = 1$ and by 40.9% when $z^{m/y} = 1.5$. This is, of course a direct consequence of the young holding less of the risk-free asset relative to their wage income. Consumption increases when the agent becomes middle aged and their income displays a sudden large increase. The *percentage* decrease in consumption upon retirement is very similar, namely 37.7% when $z^{m/y} = 1$ and 36.8% when $z^{m/y} = 1.5$. This finding is consistent with the observation made earlier that *relative* to wage income holdings of the risk-free asset by the middle aged are roughly similar for the two economies considered.

Figure 22: Individual consumption time paths: Role of increasing life-cycle wage profile



Notes: This chart shows the lifetime consumption paths when the productivity of young and middle-aged is the same, $z^{m/y} = 1$, and when wages increase with age, $z^{m/y} = 1.5$.


6.3.2 Increasing life-cycle wage income profile and monetary policy

Figure 23 plots the IRFs for a monetary-policy shock for the baseline calibration and the alternative considered here when the middle aged earn substantially higher wages. Although the change in $z^{m/y}$ is substantial, it has only a small impact on the aggregate response. Specifically, the peak response of employment increases from 0.853 to 0.922 when we increase $z^{m/y}$ from 1 to 1.5. This is somewhat surprising. In anticipation of the expected increase in income, the young now save less. This implies that they are less able to insure against unemployment spells and in other experiments we found this to magnify *aggregate* responses. What is different here is that the young become quantitatively less

important for aggregate activity when $z^{m/y}$ increases, that is, when the middle aged earn more than the young. Moreover, the middle-aged have not only become more important in terms of wage income, but also as owners of the risk-free asset which means that the negative income effect associated with the decrease in the real interest rate becomes more relevant at the aggregate level.

Under the baseline calibration, the response of per capita consumption on impact is equal to 0.80%, 1.35%, and 1.77%, for the middle aged, the employed young, and the unemployed young, respectively. For the case considered here, those numbers are 0.84%, 1.73%, and 2.78%. So it is still true that the stronger expansions mainly affects the young. Although the expansion is not that much bigger at the higher value for $z^{m/y}$, the fact that the young are quantitatively less important means that it is easier to change bond holdings to increase consumption without any dampening general equilibrium effects through changes in the interest rate to kick in.

Figure 23: Monetary-policy shock IRFs: Role of increasing life-cycle wage profile



Notes: The figure considers the case when the wage profile is flat and increasing with age. For the two different economies, the figure plots the IRFs for the indicated variable associated with a monetary policy shock that leads to a maximum twenty-five basis points drop in the nominal interest rate. For the baseline calibration this meant setting σ_{ER} equal to 0.0019165 and to 0.0020874 for the case with a higher middle aged productivity. Other parameter values are set according to the baseline calibration.

6.4 Financial trades between workers and firm owners

The results presented up to now are based on a version of the model in which firm owners do not have access to the market for risk-free assets. This restriction allows us to cleanly highlight the differences in the demand for the risk-free asset across the different types of workers. This relates both to differences in steady-state holdings and in response to shocks or parameter changes. For example, we learned that an increase in the bond supply ends up mainly in the hands of the middle aged. In this section, we relax this assumption. This has the potential of introducing additional mechanisms. For example, if *all* workers would like to save more, then that is now possible by lending funds to firm owners. And this does still allow firm owners to keep their savings level the same (or change) as the additional funds can be invested in capital and/or firm creation (or not).

6.4.1 Model modification to allow for additional financial interactions

The baseline results presented in previous sections are based on parameter values that ensure that firm owners freely choose to take neither a long nor a short position in the risk-free asset *in the steady state*.¹⁰⁶ Thus, not allowing financial transactions between workers and firm owners was only binding in the previous analysis when considering aggregate shocks. Without this assumption, firm owners may very well hold a non-zero position in the risk-free asset in the steady-state as well, which would mean that workers indirectly hold a position in the capital stock and the corporate sector. This would resemble a change in the supply of the risk-free asset, $\bar{d}$.

We impose a quadratic adjustment cost, $\eta_d^f (d^f)^2$, to ensure that Blanchard-Kahn (BK) conditions are satisfied in the stochastic economy in which firm owners can borrow and lend at the risk-free rate. The scale parameter of the penalty term, η_d^f , is chosen to a value that is so small that choosing a lower value no longer affects the results. That is, the magnitude of the adjustment-cost parameter prevents instability due to Blanchard-Kahn problems, but allows for the largest possible amount of financial interaction.¹⁰⁷

¹⁰⁶Specifically, the firm owners' discount rate is set equal to the equilibrium real interest rate. This requires solving a nonlinear equation, because the firm owners' discount rate does affect the equilibrium interest rate. For example, it affects capital accumulation, which affects wage levels, which in turn affects the demand for the risk-free asset.

¹⁰⁷The derivative of this penalty term is zero at $d_f = 0$, which implies that the steady state of our model remains exactly the same when we allow for these additional financial trades. But some curvature is needed in the first-order condition to ensure that the Blanchard-Kahn conditions are satisfied. We also considered a penalty function like the one that workers face. This would require choosing more parameter values. More importantly, we found that the quadratic penalty term allowed for more trades before running into BK problems.

6.4.2 Additional financial market interaction and monetary policy

Figure 24 shows the IRFs for a monetary-policy shock, both when d_f is restricted to be zero and when it is allowed to vary. There are two questions that we are interested in. First, does this modification affect the differential responses across worker types presented above? Second, are aggregate results affected? We first discuss the second one. A striking observation is that the aggregate employment response is higher, namely 1.06% instead of 0.85%. Note that we have not re-scaled the standard deviation of the monetary-policy shock and the stronger economic response induces some dampening since $\theta_y > 0$, resulting in a maximum decrease of the nominal interest rate of only 12.7 instead of 25 basis points. If the shock would have been such that the drop in the nominal interest rate would have remained the same, then the output response would have more than doubled.

The larger increase in employment is made possible by a higher level of vacancy posting which in turn is made possible by firm owners taking a short position in the risk-free asset, that is, $d_{f,t} < 0$. A monetary stimulus increases the profits of the intermediary firm providing inputs to the final-goods producer. This will lead to more vacancy posting and this channel is strengthened when firm owners can benefit from lower interest rates by borrowing in the bond market. Overall profits of the monopolistic final-goods producers fall, but they increase when capital rental costs are excluded from the profit measure. It makes sense to exclude rental costs to study behavior of firm owners, since rental costs are paid by this group as owners of firms, but are also received by this group as they own the capital stock.¹⁰⁸

What is intriguing about the much larger aggregate response is that this is accomplished with firm owners borrowing relatively little. Namely, the peak increase in the amount borrowed is only 0.57% of the supply of the risk-free asset. Although beyond the scope of this paper, this does indicate the quantitative importance of the financing of job creation for the strength of the monetary-policy transmission.

How does this model modification affect different worker types? Workers as a group are lending to firm owners since the IRF for d_f is uniformly negative.¹⁰⁹ Moreover, the IRF of per capita holdings of the risk-free asset shifts up for all three worker types. In fact, the IRF for asset holdings of the middle aged changes sign and is now uniformly positive. These increases in their asset holdings put pressure on their consumption levels. However, workers benefit from the stronger aggregate response and lower

¹⁰⁸In our baseline calibration, the mass of firm owners and, thus, the wages earned are relatively small. Consequently, the results discussed in this section are virtually unchanged when firm owners do not provide labor. Recall that the main reason we allowed firm owners to earn wages is that it ensures that the IRF for firm owners' consumption is positive, also on impact.


¹⁰⁹Recall that with a fixed supply of the risk-free asset, a positive response for one group must necessarily mean a negative response for another.

unemployment rates. In terms of average per capita consumption levels, the peak response for young employed falls from 1.35% to 1.08%. But this is partly due to a composition effect in that the increased flow out of unemployment into employment lowers average wealth of the employed. The corresponding numbers for the middle aged are 0.80% and 0.47%. The response of per capita consumption for the unemployed is roughly unchanged. However, if we would have scaled the size of the shock to get the same drop in the nominal interest rate (i.e., 25 basis points), then consumption of all three worker types would have increased, especially for the unemployed.

A remarkable observation is that the consumption response for the newly retired switches from uniformly negative to mainly positive except for an initial drop. This is due to the switch in the sign of the response for asset holdings by the middle aged and the lower drop in the interest rate relative to the increase in resources.

In terms of how monetary policy affects consumption levels of the different types of workers, the results are *qualitatively* not affected much except for the newly retired. For example, the consumption response of the young employed is still much bigger than the consumption response of the middle aged.

Figure 24: Monetary-policy shock IRFs: With and without firm-owner access to risk-free asset market



Notes: The figure considers the case when firm owners can and cannot take positions in the risk-free asset. The solid line correspond to our baseline case when they cannot and the dashed line when they can. The figure plots the IRFs for the indicated variable associated with a monetary policy shock that leads to a maximum twenty-five basis points drop in the nominal interest rate in the baseline economy. Other parameter values are set according to the baseline calibration.

7 Concluding comments

One motivating factor in writing this paper was to see whether the negative income effect of reductions in the interest rate could be so strong that consumption for some would actually go down during a monetary expansion. Indeed, we find that this happens for the newly retired, although the drop is quantitatively small. It must be noted, however, that their only variable income is an annuity whose value is related to the interest rate. By contrast, we find that consumption of the middle aged robustly increases following a monetary expansion although the increase is always substantially less than the increase for the young. One caveat is that our computational setup allows us to only focus on the middle aged as a group. It may very well be the case that the consumption response of those who are close to retirement is close to the small but negative response of the newly retired.

Finding age-related heterogeneous responses is not surprising. More surprising, at least to us, is the result that the generosity of the pension system also affects the impact of monetary expansions on overall economic activity. There are two reasons for this. When pension benefits are less generous, then the middle aged save more which means that a smaller fraction of the supply of the safe asset is available to those who are subject to unemployment risk. Consistent with the literature, this increase in the inability to insure against unemployment risk amplifies business cycles including those induced by changes in the central bank's policy rate. This "battle" among different age groups on who gets to hold the available supply of the safe asset turns out to be important for several policy experiments considered, not just the one related to the pension system. Another reason why the pension system matters is that it affects the level of the equilibrium real interest rate. The reason for this is that consumption choices depend on the Net Present Value of future income which is a nonlinear function of the interest rate.

Incorporating aging and retirement (benefits) necessarily complicates the model. Since these modifications are rarely included, it has been important to us to keep the environment relatively simple so that we could highlight relevant mechanisms. Of course, computational constraints also require modesty in how rich the model can be. But these considerations mean that our analysis is incomplete and should be considered as a stepping stone for further research. There are several desirable additions to the model. The first is a richer investment portfolio including real estate. The second is the introduction of pension funds. The third is a refinement of the demographic structure so that the model can study not just the middle aged on average but also those who are really close to retirement. These modifications are likely to affect the interaction between pension systems and monetary policy. But we suspect that our main mechanisms – and in particular the battle for the safe asset – will remain relevant.

References

- Aiyagari, S.Rao, and Ellen R. McGrattan (1998) 'The optimum quantity of debt.' *Journal of Monetary Economics* 42(3), 447–469
- Algan, Yann, Olivier Allais, Wouter J. Den Haan, and Pontus Rendahl (2014) 'Solving and simulating models with heterogeneous agents and aggregate uncertainty.' In *Handbook of Computational Economics*, ed. Kenneth L. Judd and Karl Schmedders (Amsterdam, The Netherlands: Elsevier)
- Alves, Felipe, Greg Kaplan, Benamine Moll, and Giovanni L. Violante (2020) 'A further look at the propagation of monetary policy shocks in hank.' *Journal of Money, Credit and Banking* 52(S2), 521–

- Auclert, Adrien, Matthew Rognlie, and Ludwig Straub (2020) ‘Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model.’ NBER Working Papers 26647, National Bureau of Economic Research, Inc, January
- Bardóczy, Bence, and Mateo Velásquez-Giraldo (2024) ‘Hank comes of age.’ Finance and Economics Discussion Series 2024-052. Washington: Board of Governors of the Federal Reserve System
- Bielecki, Marcin, Michal Brzoza-Brzezina, and Marcin Kolasa (2022) ‘Demographics, monetary policy, and the zero lower bound.’ *Journal of the European Economic Association* 20(2), 549–580
- (2023) ‘Demographics, monetary policy, and the zero lower bound.’ *Journal of Money, Credit, and Banking* 55(7), 1857–1887
- Bilbiie, F., G. E. Primiceri, and A. Tambalotti (2022) ‘Inequality and Business Cycles.’ Cambridge Working Papers in Economics 2275, Faculty of Economics, University of Cambridge, December
- Bilbiie, Florin O. (2020) ‘The New Keynesian cross.’ *Journal of Monetary Economics* 114(C), 90–108
- Blanchard, Olivier (1985) ‘Debt, Deficits, and Finite Horizons.’ *Journal of Political Economy* 93(2), 223–247
- Challe, Edouard, Julien Mathéron, Xavier Ragot, and Juan F. RubioRamirez (2017) ‘Precautionary saving and aggregate demand.’ *Quantitative Economics* 8(2), 435–478
- Davis, Steven J, and Till M Von Wachter (2011) ‘Recessions and the cost of job loss.’ Technical Report, National Bureau of Economic Research
- Den Haan, Wouter J., and Joris De Wind (2012) ‘Nonlinear and stable perturbation-based approximations.’ *Journal of Economic Dynamics and Control* 36, 1477–1497
- Den Haan, Wouter J., Lukas B. Freund, and Pontus Rendahl (2021) ‘Volatile hiring: uncertainty in search and matching models.’ *Journal of Monetary Economics* 123, 1–18
- Den Haan, Wouter J, Pontus Rendahl, and Markus Riegler (2017) ‘Unemployment (Fears) and Deflationary Spirals.’ *Journal of the European Economic Association* 16(5), 1281–1349
- Garriga, Carlos, Finn E. Kydland, and Roman Sustek (2021) ‘MoNK: Mortgages in a New-Keynesian model.’ *Journal of Economic Dynamics and Control*

- Hagedorn, Marcus, and Iourii Manovskii (2008) ‘The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited.’ *American Economic Review* 98(4), 1692–1706
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante (2018) ‘Monetary Policy According to HANK.’ *American Economic Review* 108(3), 697–743
- Kurlat, Pablo (2020) *A Course in Modern Macroeconomics* (Stanford: Independently published)
- Leduc, Sylvain, and Zheng Liu (2016) ‘Uncertainty shocks are aggregate demand shocks.’ *Journal of Monetary Economics* 82, 20–35
- LeGrand, F., and X. Ragot (2022) ‘Managing Inequality over the Business Cycle: Optimal Policies with Heterogeneous Agents and Aggregate Shocks.’ *International Economic Review* 63, 511–540
- Mazelis, Falk, Roberto Motto, and Annukka Ristiniemi (2023) ‘Monetary policy strategies for the euro area: optimal rules in the presence of the ELB.’ Working Paper Series 2797, European Central Bank, March
- Miranda-Agrippino, Silvia (2017) ‘The transmission of monetary policy shocks.’ Bank of England Staff Working Paper No. 657
- Olafsson, Arna, and Michaela Pagel (2018) ‘The retirement-consumption puzzle: New evidence from personal finances.’ Working Paper 24405, NBER
- Pissarides, Christopher A., and Barbara Petrongolo (2001) ‘Looking into the Black Box: A Survey of the Matching Function.’ *Journal of Economic Literature* 39(2), 390–431
- Reiter, Michael (2009) ‘Solving Heterogenous Agent Models by Projection and Perturbation.’ *Journal of Economic Dynamics and Control* 33, 649–665
- Rupert, Peter, and Roman Sustek (2019) ‘On the mechanics of New-Keynesian models.’ *Journal of Monetary Economics* 102, 53–69
- Samuelson, Paul A. (1958) ‘An exact consumption-loan model of interest with or without the social contrivance of money.’ *Journal of Political Economy*
- Stephens, Melvin Jr., and Desmond Toohey (2025) ‘Changes in Nutrient Intake at Retirement.’ *American Economic Journal: Applied Economics* 17, 501–526

A Two-period model

Household problem:

$$\max_{c^m, c^r, b} \frac{(c^m)^{1-\phi} - 1}{1-\phi} + \beta \frac{(c^r)^{1-\phi} - 1}{1-\phi}$$

s.t.

$$c^m + b^m = y^m, \quad (66a)$$

$$c^r = \mu^r + (1+r)b^m. \quad (66b)$$

Here, c^m denotes consumption in the first period, c^r consumption in the second period, b^m savings, and r the interest rate. If we think of the first period as the one when the agent is middle aged and the second when the agent is retired, then this model captures in a very crude and simplified manner the savings problem of a middle-aged agent in our model. The usual Euler condition is equal to

$$(c^m)^{-\phi} = \beta (1+r) (c^r)^{-\phi}$$

Using the second-period budget constraint and rearranging gives

$$\begin{aligned} (c^m)^{-\phi} &= \beta (1+r) (\mu^r + (1+r)b^m)^{-\phi} \\ c^m &= [\beta (1+r)]^{-\frac{1}{\phi}} (\mu^r + (1+r)b^m) \\ c^m &= [\beta (1+r)]^{-\frac{1}{\phi}} (\mu^r + (1+r)(y^m - c^m)) \\ c^m &= [\beta (1+r)]^{-\frac{1}{\phi}} [\mu^r + y^m(1+r) - c^m(1+r)] \\ c^m \left(1 + [\beta (1+r)]^{-\frac{1}{\phi}} (1+r) \right) &= [\beta (1+r)]^{-\frac{1}{\phi}} [\mu^r + y^m(1+r)] \\ c^m &= \frac{[\beta (1+r)]^{-\frac{1}{\phi}}}{1 + [\beta (1+r)]^{-\frac{1}{\phi}} (1+r)} [\mu^r + y^m(1+r)] \\ c^m &= \frac{[\beta (1+r)]^{-\frac{1}{\phi}} (1+r)}{1 + [\beta (1+r)]^{-\frac{1}{\phi}} (1+r)} \left[\frac{\mu^r}{1+r} + y^m \right] \end{aligned}$$

That is, the optimal solution is to consume a given fraction of the lifetime income.

B Numerical implementation

The true rational expectations solution depends on the cross-sectional joint distribution of wealth, age, and employment status which is an infinite-dimensional object. Moreover, it is time-varying object in the presence of aggregate shocks. One possibility is to summarize this cross-sectional distribution with a set of moments.¹¹⁰ Agents' wealth heterogeneity is due to agents having different histories of idiosyncratic shocks. The idea of the approximation method developed in [LeGrand and Ragot \(2022\)](#) is to keep track of agents' histories over only a finite number of previous periods. Another key feature of the LeGrand-Ragot (LR) algorithm is that it follows [Reiter \(2009\)](#) and adopts a projection method to deal with idiosyncratic risk and uses perturbation to deal with aggregate risk. Idiosyncratic risk such as unemployment risk is nontrivial and the approximation method may require sufficient flexibility to deal with possible nonlinearities. By contrast, aggregate risk is much more moderate at least along typical business cycles. We find that this algorithm works well for our particular environment in which the young face idiosyncratic employment risk and the middle aged do not.

As mentioned above, the key approximating step is to keep track of only a finite number of past employment status outcomes. We will refer to this truncation parameter as N . If N is equal to infinity, then one captures the complete history of all agents just as is the case in the true rational expectations solution. The idea of the algorithm is set N high enough so that the results of interest remain unchanged when N is increased.¹¹¹

The economic interpretation of the solution for a particular value of N . In principle, one can describe the approximation for a fixed value of N as a particular recursive economy.¹¹² Some readers may find the environment a bit weird, as it involves agents living on different islands and pseudo-social planners. Those readers can simply ignore the “story” and simply consider the algorithm as mathematical approximation that keeps track of a finite number of past histories instead of the infinite past as is done with a rational expectations solution. The economy we are interested in is the one describe in section 2 and what matters is whether the approximation algorithm can accurately capture the key features of the original model for the chosen *finite* value of N .

What does this economy with a truncation parameter equal to N look like? The idea is that agents

¹¹⁰See [Algan et al. \(2014\)](#).

¹¹¹For our model, that is the case when N is equal to six.

¹¹²See the online appendix of [LeGrand and Ragot \(2022\)](#) for details.

with the same history during the last N periods are living on the same island. To explain the algorithm, we assume that N is equal to 2, that is, the algorithm keeps track of the individual state variables of the current and previous period for both the young and the middle-aged agents. Two model assumptions make it possible to deal with the heterogeneity across retirees. The first is that middle-aged agents who retire sell their holdings of the risk-free asset and buy annuities. The second is that their non-investment income is also fixed and equal to a fraction of the wage rate of middle-aged agents in the period they retire. This means that we have a rich heterogeneity of retirees with different consumption levels that depend on economic conditions in the period when they retire. But as shown in section 2, these two assumptions allow us to capture retirees without the need for numerical approximations.

Setting N equal to 2 simplifies the exposition; the results reported in our paper are based on a numerical approximation with N equal to 6 for young agents and 2 for middle-aged agents.

There are thirteen different groups when N is equal to 2. The vector containing the mass of each group is given by

$$\gamma_t = \{\gamma_t^{yyee}, \gamma_t^{yyeu}, \gamma_t^{yyue}, \gamma_t^{yyuu}, \gamma_t^{myee}, \gamma_t^{myeu}, \gamma_t^{myue}, \gamma_t^{myuu}, \gamma_t^{mnee}, \gamma_t^{mneu}, \gamma_t^{mmue}, \gamma_t^{mmuu}, \gamma\}.$$

The first (second) superscript denotes the age in the current (previous) period.¹¹³ The third (fourth) superscript denotes the employment status in the current (previous) period. As explained above, we can capture the heterogeneity across the retirees with one aggregate.

In terms of the economic environment for fixed N , the idea is that agents in the same group are living on the same island where they pool resources. Moreover, all agents who are on the same island consumer the same amount and will be given the same amount of savings to carry over into the next period (when they will either stay on the same island or move to another one). This is important to capture the incomplete markets feature of our model. That is, one cannot allocate more savings to agents who will be unemployed in the next period and less to those who will be employed.

It is important that agents remain price takers even though heterogeneity has been reduced to a finite number of types. To introduce this more formally one should think of the model as one with an infinite number of island groups each consisting of a set of islands. The idea of the approximation is that as one increases N the approximation converges to the original model with a continuum of different agents.

¹¹³Recall that we use an extended perpetual-youth setup according to which an agent ages with fixed probability.

Note that two agents who have the exact same infinite history of employment status outcomes would be identical in the original model and one could think of these as being on the same island making the same choices.

Let Ω_t denote the matrix with transition probabilities for γ_t . Thus,

$$\gamma_t = \Omega_t \gamma_{t-1}. \quad (67)$$

Also, let $\hat{\delta}^y$ the probability of a young agent remaining young, i.e. $\hat{\delta}^y = 1 - \delta^y$ and $\hat{\delta}^m$ the probability of a middle-aged agent not retiring, i.e. $\hat{\delta}^m = 1 - \delta^m$. The first eight columns of Ω_t are given by

$$\Omega_t = \begin{bmatrix} \hat{\delta}^y \pi_t^{y,e|e-1} & \hat{\delta}^y \pi_t^{y,e|e-1} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \hat{\delta}^y \pi_t^{y,e|u-1} & \hat{\delta}^y \pi_t^{y,e|u-1} & 0 & 0 & 0 & 0 \\ \hat{\delta}^y \pi_t^{y,u|e-1} & \hat{\delta}^y \pi_t^{y,u|e-1} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \hat{\delta}^y \pi_t^{y,u|u-1} & \hat{\delta}^y \pi_t^{y,u|u-1} & 0 & 0 & 0 & 0 \\ \delta^y \pi_t^{m,e|e-1} & \delta^y \pi_t^{m,e|e-1} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \delta^y \pi_t^{m,e|u-1} & \delta^y \pi_t^{m,e|u-1} & 0 & 0 & 0 & 0 \\ \delta^y \pi_t^{m,u|e-1} & \delta^y \pi_t^{m,u|e-1} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \delta^y \pi_t^{m,u|u-1} & \delta^y \pi_t^{m,u|u-1} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \hat{\delta}^m \pi_t^{m,e|e-1} & \hat{\delta}^m \pi_t^{m,e|e-1} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \hat{\delta}^m \pi_t^{m,e|u-1} & \hat{\delta}^m \pi_t^{m,e|u-1} \\ 0 & 0 & 0 & 0 & \hat{\delta}^m \pi_t^{m,u|e-1} & \hat{\delta}^m \pi_t^{m,u|e-1} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \hat{\delta}^m \pi_t^{m,u|u-1} & \hat{\delta}^m \pi_t^{m,u|u-1} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

and the last five columns are given by

$$\begin{bmatrix}
 0 & 0 & 0 & 0 & \delta^r \gamma^{yyee} / \gamma_y \\
 0 & 0 & 0 & 0 & \delta^r \gamma^{yyeu} / \gamma_y \\
 0 & 0 & 0 & 0 & \delta^r \gamma^{yyue} / \gamma_y \\
 0 & 0 & 0 & 0 & \delta^r \gamma^{yyuu} / \gamma_y \\
 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 \\
 \widehat{\delta}^m \pi_t^{m,e|e-1} & \widehat{\delta}^m \pi_t^{m,e|e-1} & 0 & 0 & 0 \\
 0 & 0 & \widehat{\delta}^m \pi_t^{m,u|e-1} & \widehat{\delta}^m \pi_t^{m,u|e-1} & 0 \\
 \widehat{\delta}^m \pi_t^{m,u|e-1} & \widehat{\delta}^m \pi_t^{m,u|e-1} & 0 & 0 & 0 \\
 0 & 0 & \widehat{\delta}^m \pi_t^{m,u|u-1} & \widehat{\delta}^m \pi_t^{m,u|u-1} & 0 \\
 \delta^m & \delta^m & \delta^m & \delta^m & 1 - \delta^r
 \end{bmatrix}$$

The i^{th} column of Ω_t gives the probabilities to which island an agent living on island j in the previous period can move to. These probability add up to 1. Consider agents living on the first island which consists of agents who were employed for two periods. If they do not age and remain employed, then they stay on the first island which happens with probability $\widehat{\delta}^y \pi_t^{y,e|e-1}$. If they do not age but become unemployed they move to the third island, i.e., the “yyue” island. If they do age and become middle aged, then they move to either the fifth (myee) or the seventh (myue) island. Note that agents from the retirement island are reborn as a young agent and are distributed among the islands with young agents according to the current distribution. But these newborns do not bring any wealth as the right to the annuity ends with death. The j^{th} row of Ω_t gives the flows into the j^{th} island. For example, one can stay on the first island if one does not age and remains employed, one could move to the first (yyee) island from the second (yyeu) island, and move their as a newborn from the retirement island.

Filling in the matrix with the matching block. This general specification for Ω_t allows for all probabilities across employment status to be endogenous. However, several are exogenous and constant through time. For example, the probability of a young agent becoming unemployed is equal to the constant separation rate, that is, $\pi_t^{y,e|e-1} = 1 - \Omega^{ee}$. By contrast, conditions in the matching market determine job finding rates which are endogenous. For example, $p_t^{y,e|u-1} = f_t^t$. This general exposition of the al-

gorithm also does not enforce the restriction we impose in the model that middle-aged agents are not subject to unemployment risk.

Deriving the agents' first-order conditions. To derive the correct first-order conditions of the numerical approximation, [LeGrand and Ragot \(2022\)](#) construct a hypothetical economy. Specifically, it specifies a particular optimization problem to derive a set of first-order conditions. [LeGrand and Ragot \(2022\)](#) refer to this as a quasi-planner problem. “Quasi” because prices are taken as given as is the case for individual agents in the underlying model and “planner” because it pins down the choices for all types of workers. Both the term quasi-planner and the environment may seem unusual. But the reader will see that we end up with a set of first-order conditions that will be identical to those of our underlying model if the truncation parameter, N , goes to infinity. The advantage of specifying this hypothetical economy is that it helps in getting the right equations for the numerical approximation.

The quasi-planner's flow utility function is given by

$$\sum_{i=1}^I \gamma_t^i \frac{(c_t^i)^{1-\phi} - 1}{1-\phi}, \quad (68)$$

where γ_t^i is the fraction of workers of island i , i.e., those who have the same history of idiosyncratic shocks for N periods and I is the number of islands. Thus, agents with a larger population get more weight. It will become clear below why this is important to get the correct first-order conditions. In addition to the i superscript to indicate the employment status history (island), we also could have added a j subscript to indicate that there are many (identical) versions of each island and each is small and takes prices as given.

Constraints The constraints of the optimization problem are the following:

$$c_t^i + d_t^i = I_t^{e,i} (1 - \tau_t) w_t + (1 - I_t^{e,i}) \mu_t + \sum_{i^* \in \mathcal{J}(i)} (r_t^d + \tilde{r}^d(\bar{d}_{t-1}^{i^*})) \frac{\Omega_t^{i,i^*} \gamma_{t-1}^{i^*}}{\gamma_t^i} d_{t-1}^{i^*}, \quad (69)$$

where $\mathcal{J}(i)$ is the set of islands from which agents can move to island i . For example, these would be islands 1 and 2 for island 1. The last term then adds up the resources brought into island i from the islands in $\mathcal{J}(i)$ and then divides this sum by γ_t^i since variables are expressed as per capita values.¹¹⁴ It may seem that moving to an island with a larger population is a disadvantage as you have to share the

¹¹⁴Since all islands of the same type have identical asset holding, we have that $\bar{d}_t^i = d_t^i$, but we can only impose that after FOCs have been written.

resources you bring to the island with more. However, this is exactly why the flow utility gives a larger weight to the bigger islands.

This optimization generates a set of first-order conditions for each island.¹¹⁵ For example, the FOC conditions for the choice variables of island *yyee* are the constraint given above and

$$\gamma_t^{yyee} (c_t^{yyee})^{-\phi} = \lambda_t^{yyee}, \quad (70)$$

$$\lambda_t^{e^y e^y} = \beta \mathbb{E}_t \left[\begin{aligned} & \Omega_{t+1}^{yyee, yyee} \frac{\gamma_t^{yyee}}{\gamma_{t+1}^{yyee}} (r_{t+1}^d + \tilde{r}^d(\bar{d}_t^{yyee})) \lambda_{t+1}^{yyee} \\ & + \Omega_{t+1}^{yyee, yyue} \frac{\gamma_t^{yyee}}{\gamma_{t+1}^{yyue}} (r_{t+1}^d + \tilde{r}^d(\bar{d}_t^{yyue})) \lambda_{t+1}^{yyue} \\ & + \Omega_{t+1}^{yyee, myee} \frac{\gamma_t^{yyee}}{\gamma_{t+1}^{myee}} (r_{t+1}^d + \tilde{r}^d(\bar{d}_t^{myee})) \lambda_{t+1}^{myee} \\ & + \Omega_{t+1}^{yyee, myue} \frac{\gamma_t^{yyee}}{\gamma_{t+1}^{myue}} (r_{t+1}^d + \tilde{r}^d(\bar{d}_t^{myue})) \lambda_{t+1}^{myue} \end{aligned} \right]. \quad (71)$$

If we substitute out the Lagrange multiplier, then we get

$$(c_t^{yyee})^{-\phi} = \beta \mathbb{E}_t \left[\begin{aligned} & \Omega_{t+1}^{yyee, yyee} (r_{t+1}^d + \tilde{r}^d(\bar{d}_t^{yyee})) (c_t^{yyee})^{-\phi} \\ & + \Omega_{t+1}^{yyee, yyue} (r_{t+1}^d + \tilde{r}^d(\bar{d}_t^{yyue})) (c_t^{yyue})^{-\phi} \\ & + \Omega_{t+1}^{yyee, myee} (r_{t+1}^d + \tilde{r}^d(\bar{d}_t^{myee})) (c_t^{myee})^{-\phi} \\ & + \Omega_{t+1}^{yyee, myue} (r_{t+1}^d + \tilde{r}^d(\bar{d}_t^{myue})) (c_t^{myue})^{-\phi} \end{aligned} \right]. \quad (72)$$

This Euler equation has the same structure as the one of our underlying model.¹¹⁶ Specifically, the right-hand side is a weighted sum taking into account idiosyncratic shocks, both in terms of the two possible outcomes for the employment status and the two aging possibilities. And even if there is no aggregate risk, then the consumption values at these different outcomes will be different which means that the solution will incorporate precautionary savings.

Although, the Euler equation has the same structure as the one of the original model, the budget constraint does not. Whereas the budget constraint in the original model simply has the resources from last period's investment for each agent, the budget constraint of the quasi-planner problem contains a weighted sum of different types of agents. However, as N goes to infinity it will also be identical to the original budget constraint. Consider an agent on island 1 which is for this discussion the island on which agents have been employed for N periods. And island 2 is the island with workers that have the same

¹¹⁵ Similarly, a grid-based method would have a set of first-order conditions for each grid point. When adopting a grid-based method, it may take some trial and error to find the appropriate range. That is not an issue with the LR method as the range of possible outcomes endogenously increases with N .

¹¹⁶ If one increases the value for N and the number of islands increases, then the number of first-order conditions increases, but this is also true for a grid-based method.

history except that these workers were unemployed $N - 1$ periods ago. As N increases, these islands will become more similar and the two d^{i*} values will converge to the same value. Moreover, the value of $\Omega^{i,i*} \gamma_{t-1}^* / \gamma_t^i$ will converge to 1.¹¹⁷

LR algorithm and the penalty function. A big advantage of the LR algorithm is that one accurately captures the nonlinearities associated with idiosyncratic uncertainty which is thought to be important because idiosyncratic risk such as unemployment is quantitatively important but uses low-order perturbation to deal with aggregate uncertainty which is much smaller. LR implement their algorithm in a model with inequality constraints. That complicates the algorithm since one has to find out on which island the constraint is binding and more importantly one has to assume that this does not change in response to aggregate shocks. The reason is, the LR requires knowing which equation are used to solve choices on each island, the Euler equation or the borrowing constraint. The advantage of penalty functions is that choices are always determined by the same set of equations.

Truncation parameter used for results in our paper. The results in our paper are based on the case when we set N equal to six for young agents. Since the young are subject to unemployment risk, there is more heterogeneity among the young than the middle aged. This means that we have sixty-four islands for the young. We assume that N is equal to 2 for the middle aged. This together with the assumption that the middle aged are always employed means that we have three islands for the middle aged, namely one for those that were young and employed last period, one for those that were young and unemployed last period, and one for those who were middle aged (and thus employed) last period. Thus, our agent quickly become a “typical” agent in our model and our approximation is not suited to capture the transition from young to middle aged well.¹¹⁸ But that is not the objective of this paper. What we are interested in is

¹¹⁷To understand this, consider islands 1 and 2, that is islands with workers who have been employed for quite a while and only differ in the last employment status we keep track of. That is, the $eee \dots eee$ and the $eee \dots ee u$ island. To simplify the argument ignore the possibility of aging. Also, to simplify the expressions, we assume that agents that separate cannot immediately get matched. The steady-state employment rate, n , is then equal to $f / ((1 - \Omega^{ee}) + f)$. The steady-state unemployment rate, u , is equal to $1 - \Omega^{ee} / ((1 - \Omega^{ee}) + f)$ and the mass of unemployed divided by the mass of employed is equal to $u/n = (1 - \Omega^{ee})/f$. Now consider the ratio of the mass of those who are employed this period and unemployed last period relative to those who are employed this and the last period. This is equal to $f u / (\Omega^{ee} n)$. Using the expression for u/n , we get that this ratio is equal to $(1 - \Omega^{ee})/\Omega^{ee}$. If we now reduce both groups each period by only considering those who remain employed then both groups are reduced by the same fraction, namely $1 - \Omega^{ee}$. Thus, the ratio remains the same. But this means that the ratio γ^2/γ^1 is equal to $(1 - \Omega^{ee})/\Omega^{ee}$ which in turn implies that

$$\Omega^{1,1} \frac{\gamma^1}{\gamma^1} + \Omega^{1,2} \frac{\gamma^2}{\gamma^1} = 1.$$

This – together with the observation that these two islands will have made the same choice if N is large – implies the term in the budget constraint simply becomes $(r_t^d + \tilde{r}^d(\bar{d}_{t-1}))d_{t-1}$, that is, it is no longer a sum, but identical to the equation in the original model. See [LeGrand and Ragot \(2022\)](#) for a more formal analysis on convergence as N increases.

¹¹⁸Similarly, our model does not capture the transition from a new born to a typical young agent. There the transition occurs in one step as the newborns are allocated across islands with young ones who were already around.

how the middle aged prepare for retirement.

C The Reiter approach and time-varying precautionary savings

The LR algorithm implements the idea proposed in [Reiter \(2009\)](#) to adopt a projection method to deal with idiosyncratic risk and to deal with aggregate uncertainty using perturbation. Specifically, we first solve the model without aggregate risk. This is a set of equations that we solve with a nonlinear-equation solver.¹¹⁹ It is a large system as we have a separate set of first-order conditions for each island and the model consists of several blocks. The benchmark version of the model has 830 equations. Next, we use Dynare to obtain a first-order perturbation solution around this steady state.

The algorithm allows for the possibility of precautionary savings since a projection method is used to deal with idiosyncratic risk. Specifically, consider the environment described in the previous section of this appendix. The right-hand side of the Euler equation contains a weighted sum of possible outcomes reflecting different realizations for the employment status and aging. The corresponding consumption values are different, even in the no-aggregate-uncertainty steady state. And since our utility function implies convex marginal utility, agents will accumulate savings for precautionary reasons.

In this section, we will present a simple example to make clear that the algorithm also allows for the possibility of *time-varying* precautionary savings even though we rely on a first-order perturbation solution to deal with aggregate risk.

Environment. We consider a simple two-period partial equilibrium environment and the interest and discount rate are equal to each other. Income in the first period is equal to y_1 . In the second period, income is equal to y_2 , where $y_2 \in \{1 - \varepsilon, 1, 1 + \varepsilon\}$. The variable ε captures idiosyncratic risk. The corresponding probabilities are equal to $0.25 - 0.5z$, $0.25 - z$, and $0.25 + 0.5z$, where z is a random variable capturing aggregate risk with mean zero. Note that the probabilities always add up to 1.¹²⁰ The idea is that aggregate risk determines the variance of the amount of idiosyncratic risk, but not the expected value. And risk increases with z .

The first-order condition is standard and given by

$$\frac{1}{1-s} = \sum_i p_i \frac{1}{c_i'} = \frac{0.25 + 0.5z}{1 - \varepsilon + s} + \frac{0.5 - z}{1 + s} + \frac{0.25 + 0.5z}{1 + \varepsilon + s} \quad (73)$$

¹¹⁹Some parts can be solved analytically.

¹²⁰The distribution of z must be such that probabilities are always between 0 and 1, that is, the support for z must be between -0.5 and 0.5.

Steady state When $z = 0$, then savings are solved from

$$\frac{1}{1-s} = \frac{0.25}{1-\varepsilon+s} + \frac{0.5}{1+s} + \frac{0.25}{1+\varepsilon+s}. \quad (74)$$

For example, the solution for s is equal to 0.659 when ε is equal to 0.5. The reason for positive savings is the presence of a precautionary motive, since the interest rate is equal to the discount rate and the expected value of next period's income is equal income in the current period.¹²¹

First-order perturbation solution. A key element of the Reiter (and thus the LR) method is that the perturbation is with respect to aggregate risk, i.e., z , not idiosyncratic risk, i.e., ε . Implementing this idea for the first-order condition given above gives¹²²

$$\begin{aligned} (s-\bar{s})\frac{1}{(1-\bar{s})^2} &= -(s-\bar{s})\frac{0.25}{(1-\varepsilon+\bar{s})^2} + (s-\bar{s})\frac{0.5}{(1+\bar{s})^2} - (s-\bar{s})\frac{0.25}{(1+\varepsilon+\bar{s})^2} \\ &+ z\frac{0.5}{1-\varepsilon+\bar{s}} - z\frac{1}{1+\bar{s}} + z\frac{0.5}{1+\varepsilon+\bar{s}}. \end{aligned} \quad (75)$$

When ε is again equal to 0.5, then we get

$$s = \bar{s} + 0.0121z. \quad (76)$$

That is, savings in/decrease with in/decreased uncertainty. And recall that z does not affect expected income growth, it only affects volatility.¹²³ It is key that we do not consider a perturbation solution with respect to ε , which would be around $\varepsilon = 0$. Using the actual three realizations for idiosyncratic risk, an increase in z would push possible consumption values further apart along the convex marginal-utility curve. This increases the expected value of the right-hand side of the Euler equation and thus the value for s . If the perturbation would be with respect to z as well as ε , then the last three terms in equation (75) would cancel out as they would be evaluated around the steady-state of ε , which is equal to 0. Consequently, s would no longer depend on z .

¹²¹Savings would be equal to zero when ε is equal to zero.

¹²²Note that the four constants drop out.

¹²³In our model, aggregate shocks are monetary-policy shocks that affect expected values. But they also affect the extent that agents are subject to idiosyncratic risk as they affect conditions in the matching model. That is, monetary-policy shocks are a combination of a first-order shock (expected income changes) and a second-order shock (individual uncertainty changes). The recent literature is interested in the role of this second-order shock because changes in precautionary savings would strengthen the impact of monetary-policy shocks. Specifically, a monetary tightening would increase idiosyncratic unemployment risk which would lead to a further increase in savings and, thus, a further reduction in demand.

Acknowledgements

We would like to thank for comments and suggestions Roberto Motto, Stefano Neri, Xavier Ragot, Alessandro Secchi, Gianluca Violante, and seminar participants at the European Central bank, Banca d'Italia, Dutch National Bank, Swiss National Bank, Norges Bank, and conference participants at the 4th Dynare Workshop, 19th Dynare conference, 29th International Conference Computing in Economics and Finance, ECB's Research Task Force on Heterogeneity in Macroeconomics. We thank Maria Eskelinen for research assistance with an earlier version.

The opinions expressed here are the sole responsibility of the authors and should not be interpreted to reflect the views of the European Central Bank, or the Eurosystem.

Wouter J. Den Haan

London School of Economics and Political Science, London, United Kingdom; Centre for Economic Policy Research, London, United Kingdom; Centre For Macroeconomics, London, United Kingdom; email: W.Denhaan@lse.ac.uk

Alessandro Ferrari

European Central Bank, Frankfurt am Main, Germany; email: Alessandro.Ferrari@ecb.europa.eu

Falk Mazelis

European Central Bank, Frankfurt am Main, Germany; email: Falk.Mazelis@ecb.europa.eu

Annukka Ristiniemi

European Central Bank, Frankfurt am Main, Germany; email: Annukka.Ristiniemi@ecb.europa.eu

© European Central Bank, 2025

Postal address 60640 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network electronic library](#) or from [RePEc: Research Papers in Economics](#). Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB's website](#).

PDF

ISBN 978-92-899-7443-1

ISSN 1725-2806

doi: 10.2866/5421356

QB-01-25-214-EN-N