CPB Netherlands Bureau for
Economic Policy Analysis

The impact of a lower return on wealth on optimal wealth accumulation

The return on safe assets has been declining over the past 40 years. Should people save more for retirement to compensate for the lower return on wealth?

In this document we quantify the effect of a lower return on wealth using a life-cycle model of optimal consumption and savings calibrated with the profile of mean gross income of Dutch households.

The numerical results indicate that optimal wealth accumulation is lower throughout a person's lifetime due to both lower savings in the beginning of the career and the lower rate at which savings is accrued.

CPB Background Document

Nicoleta Ciurila

August 2020

The impact of a lower return on wealth on optimal wealth accumulation

Nicoleta Ciurilă[1]

[1]CPB Netherlands Bureau for Economic Policy Analysis

Background document to the Policy Brief “Lage rente en de toekomst van pensioenen”

Abstract

This paper evaluates the impact of a permanent decrease in the return on wealth on optimal savings, wealth and consumption of a person using a life-cycle model calibrated on Dutch data. For a person that is just entering the labor market, a one percentage point permanent decline in the return on wealth leads to less consumption and more savings over most of the life. Total wealth is 7% lower before retirement due to slightly less savings in the beginning of the person's career and the lower return at which savings accrue.

1 Introduction

What effect would a lower return on wealth have on optimal savings, wealth and consumption over a person's life? This is currently an important question given the steady decline in the safe interest rate over the past 40 years. More specifically, should people save more for retirement when the return on wealth is lower (or, alternatively, should pension funds increase their premiums)? What implications will this have for individual and aggregate consumption?

In this document I describe how the lower return on wealth affects people's choices through different channels and quantify their effects using a life-cycle model of optimal consumption and savings. I calibrate this model using the life-cycle profile of mean gross income of Dutch households. For people's preferences (time preference, elasticity of intertemporal substitution) I consider values from the literature. I set the return on wealth to 3% in the baseline scenario and simulate the optimal wealth accumulation and consumption across the life-cycle. In order to determine the impact of a permanent decline in return, I simulate the same model with a return on wealth of 2% and compare it to the baseline scenario.

The numerical evaluation indicates that, for a person that is just starting her work career, a permanently lower rate of return implies less wealth accumulation at the retirement age and lower consumption for almost the entire life. I find that in the low return scenario optimal wealth accumulation is around 7% smaller at the retirement age. People save more when the return declines for the most part of their life, but total wealth is still lower due to the lower rate at which savings are accrued over a person's life. As a consequence, also consumption is lower over most of a person's life.

There is in general no clear answer to the question whether people optimally save more or less following a permanent decline in the return on wealth (Summers (1981), Elmendorf (1996), Jones (2018)). This is because a permanently lower return impacts the optimal level of savings for retirement through channels that act in opposite directions. The relative strength of these channels depends on the preferences of a person, the age, the income profile and the wealth accumulated at the moment when the return on wealth declines.

A lower return on wealth affects a person's choices via three channels: i) a substitution effect: people are more reluctant to give up a unit of current consumption in order to save it and obtain more consumption in the future because the return on savings is lower; hence savings decrease; ii) an income effect: because the income obtained from savings during retirement is lower, people have to save more in order to prevent a too high decline in consumption after retirement; hence savings increase; iii) a wealth effect: a lower return on wealth makes it possible for a person to borrow more against her human and financial capital, as their value increases; in this sense people are wealthier, they can consume more and save less. For realistic calibrations of the intertemporal rate of substitution[1], the income effect dominates the substitution effect, so a lower return on wealth leads to higher optimal savings. However, it is unclear whether it also dominates the wealth effect. That is why a numerical evaluation of the three different channels is necessary.

The numerical results obtained with the life-cycle model indicate that a permanently lower return on wealth leads to slightly lower optimal savings in the first part of a person's career (the substitution and human wealth effect dominate the income effect). However, as people approach the retirement age, the present value of labor income and, hence the human wealth effect, declines, so a lower return on wealth leads to higher optimal savings later in a person's career. Total wealth is lower throughout a person's lifetime due to both lower savings in the beginning of a person's career and the lower rate at which savings is accrued.

2 A two period model

I begin the analysis with a two period model. The effect of a decline in the return on wealth on savings in the first part of a person's career can easily be determined with this model. However, the effect on savings later in a person's career can differ and cannot be determined in a two period model.

I consider the steady state of an economy in which each generation lives for two periods:

[1] $\frac{1}{\sigma} < 1$, i.e. $\sigma > 1$ in a CRRA utility function.

young age (between 25 and 55 years) and old age (between 55 and 85 years). A person maximizes the present value of utility from consuming c_y when young and c_o in the old age. I consider that individuals discount the utility from future consumption with β and have a probability equal to ψ to reach the old age. People earn an income equal to y when they are young and $y \cdot G\lambda$ when they are old. G represents the gross income growth rate over the person's working life, while λ represents the fraction of the old age period that is spent working. The savings of the young a are invested in capital and earn a gross rate of return R. The wealth accumulated by a person in this two period model is equal to aR. I consider the case of a CRRA utility function:

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma} \tag{1}$$

Formally, the problem that a person solves is given by:

$$\max_{\{c_y, c_o, a\}} u(c_y) + \beta\psi u(c_o) \tag{2}$$

$$c_y + a = y \tag{3}$$

$$c_o = aR + yG\lambda \tag{4}$$

The budget constraint of a person when old (equation 4) shows how the *income effect* works. A lower return on wealth R implies that people have less income during old age. Hence they must increase savings a during the first period to ensure that consumption does not drop too much during old age. Consequently, a lower return on wealth implies higher savings through the income effect.

The budget constraints (3) and (4) can be consolidated in a life-time budget constraint:

$$c_y + \frac{c_o}{R} = y + \frac{yG\lambda}{R} \tag{5}$$

From this we can identify another channel through which the return on wealth impacts on the optimal savings and consumption of a person.

A lower rate of return on wealth increases the present value of the income earned by a person (right hand side of equation (5)). The present value of labor income can be seen as the *human wealth* of a person that is just entering the labor market. All else equal, a higher present value of future labor income implies that the life-time budget of the person expands and she can consume more in both time periods. A higher consumption in the first period implies that people **save less**. This is the *human wealth effect*[2].

In order to get some intuition for the *human wealth effect*, let us assume that human wealth

[2] Summers (1981) first identifies the existence and quantitative importance of the human wealth effect on the elasticity of savings with respect to taxation. Elmendorf (1996) and more recently Ordonez and Piguillem (2020) show that the human wealth effect is quantitatively important savings reacts to changes in the return on wealth.

is tradable. Instead of waiting for the income $yG\lambda$ to be realized during old age, we assume that people can obtain the present value of this income when young (at a discount). They must then repay this when they are old and actually receive the income. More specifically, when young, in addition to the labor income y, they also receive the sum $\frac{yG\lambda}{R}$. They then pay $yG\lambda$ back when they are old [3]. What happens if the return R declines? The discounted value of future income $\frac{yG\lambda}{R}$ will increase. As a result the human wealth of a person will be higher, she will receive a larger sum when young and hence consume more (spend less).

Optimal savings can also be influenced by a decline in R through a *financial wealth effect*. Until now we assumed that a young person starts her life with no wealth. Suppose now that she receives a bequest of a bond with a fixed coupon equal to aR payable when she is old. If R declines today, the value of the bond increases since the coupon it pays is unchanged. Hence, the wealth of the young person increases and she can afford to consume more and save less. I will not consider the *financial wealth* further in the two period model.

The first order condition of the problem summarized in equations (2) - (4) is given by the following Euler equation:

$$c_o = c_y(\beta R\psi)^{\frac{1}{\sigma}} \tag{6}$$

This result is derived in the Appendix. The Euler equation captures the *optimal* consumption (and hence savings) behaviour of a person. It shows that people smooth consumption across periods: any extra income earned by a person when young implies more consumption in both time periods (higher c_y but also higher c_o through the Euler equation). If the return on wealth R falls, consumption when old becomes less attractive compared to consumption when young. As a result, consumption will increase less throughout a person's life and the consumption profile will be flatter, i.e. $\frac{c_o}{c_y}$ will decline. This implies that when the rate of return on wealth R declines, people will **save less**. This is the *substitution effect*.

There are two parameters that determine how strongly savings react to a lower return on wealth: the rate of time preference β and the elasticity of intertemporal substitution $\frac{1}{\sigma}$. The substitution effect will be higher if people are more patient (higher β). Similarly, the substitution effect will be higher if people are more willing to accept high fluctuations in consumption over their life (higher elasticity of intertemporal substitution, lower σ).

The solution of the model is given in the following three equations:

[3] Allowing human wealth to be tradable does not alter optimal consumption. However, people will save more, namely $a + \frac{yG\lambda}{R}$. They then use aR to finance old age consumption and pay back $yG\lambda$.

$$c_y = y \frac{1 + \frac{G\lambda}{R}}{1 + \frac{(R\beta\psi)^{\frac{1}{\sigma}}}{R}} \tag{7}$$

$$c_o = (R\beta\psi)^{\frac{1}{\sigma}} y \frac{1 + \frac{G\lambda}{R}}{1 + \frac{(R\beta\psi)^{\frac{1}{\sigma}}}{R}} \tag{8}$$

$$a = y - y \frac{1 + \frac{G\lambda}{R}}{1 + \frac{(R\beta\psi)^{\frac{1}{\sigma}}}{R}} \tag{9}$$

One can identify the three above mentioned channels in the solution of optimal savings a:

- substitution effect: this corresponds to the R in the term $(R\beta\psi)^{\frac{1}{\sigma}}$. This term comes from the Euler equation. Through this effect a lower R implies a higher c_y and a lower a.
- income effect: this corresponds to the $\frac{1}{R}$ in the term $\frac{(R\beta\psi)^{\frac{1}{\sigma}}}{R}$. Through this effect a lower R implies a lower c_y and a higher a.
- human wealth effect: this corresponds to the R in the term $\frac{G\lambda}{R}$. Through this effect a lower R implies a higher c_y and a lower a. The higher is the difference between the increase in income over a person's career G and the rate of return on wealth R, the more important is the human wealth effect. More specifically, if R is low, a steep labor income profile (high G) implies an important human wealth effect and a more important decline in savings following a lower return on wealth.

A decrease in R increases savings through the income effect but decreases savings through the substitution and wealth effects. In order to determine the overall effect of a lower return on wealth, I take the derivative of savings (as defined by equation (9)) and wealth aR with respect to the return on wealth R. The sign of the derivative is given by the following expression:

$$\begin{aligned} sgn\left(\frac{\partial a}{\partial R}\right) &= sgn\left(\frac{G\lambda}{R^2}\left(1 + (\beta\psi)^{\frac{1}{\sigma}} R^{\frac{1-\sigma}{\sigma}}\right)\right) + sgn\left(\frac{1-\sigma}{\sigma} R^{\frac{1-\sigma}{\sigma}} (\beta\psi)^{\frac{1}{\sigma}} \left(\frac{1}{R} + \frac{G\lambda}{R^2}\right)\right) \qquad (10) \\ &= + \text{ if } \sigma \leq 1 \\ &= ? \text{ if } \sigma > 1 \\ sgn\left(\frac{\partial aR}{\partial R}\right) &= sgn\left(R\frac{\partial a}{\partial R} + a\right) = \qquad (11) \\ &= + \text{ if } \sigma \leq 1 \\ &= ? \text{ if } \sigma > 1 \end{aligned}$$

The above relation shows that if $\sigma \leq 1$, then $\frac{\partial a}{\partial R} > 0$ and $\frac{\partial aR}{\partial R} > 0$: the substitution and human wealth effect dominate the income effect and a lower return on wealth leads to lower optimal savings of households a and lower optimal wealth aR. This implies more consumption when young but less when old.

However, the empirically relevant situation is $\sigma > 1$ (Havranek et al. (2015)). In this case the signs of the derivatives are unclear [4]. The income effect dominates the substitution effect, but it is unclear whether it also dominates the human wealth effect.

In conclusion, the overall impact of a lower R on savings and wealth is unclear in this simple setting. I therefore turn to a model with a more realistic number of periods. This setting enables me to obtain a more realistic numerical quantification of the impact of a lower value of R.

3 A multi-period model

Let us assume now that a person lives for T years. She earns an income equal to y_t at each age t, until reaching the retirement age ra. After the retirement age, the household consumes by decumulating her wealth. The household chooses each period how much to consume c_t and how much wealth to accumulate a_t by maximizing the present value of the utility of consumption. People are born without wealth and consume all wealth by the maximum age of T. The probability to survive until period t is equal to ψ_t.

$$\max_{\{c_t,a_t\}} \sum_{t=1}^{T} \beta^{t-1} \psi_t u(c_t) \tag{12}$$

$$y_t - c_t = a_t - Ra_{t-1}, t \leq ra \tag{13}$$

$$c_t = Ra_{t-1} - a_t, t > ra \tag{14}$$

$$a_0 = a_T = 0 \tag{15}$$

In the above, the term $a_t - Ra_{t-1}$ represents the savings of a person, while the term a_t measures the wealth of a person (accrued savings).

I assume again a standard CRRA function for consumption preferences. I show in the appendix that the solution to the above problem is the following:

[4] One might be tempted to assume that in old age people don't work ($\lambda = 0$) and hence conclude that for $\sigma > 1$ there is a negative relation between the return on wealth and the savings of households. However, in a realistic setting where households live more than 2 periods and hence also receive an income after the first period, the wealth effect will have an impact on optimal household savings.

$$c_1 = \frac{\sum_{t=1}^{ra} \frac{y_t}{R^{t-1}}}{\sum_{t=1}^{T} \frac{\psi_t^{\frac{1}{\sigma}} (\beta R)^{\frac{t-1}{\sigma}}}{R^{t-1}}} \tag{16}$$

$$a_1 = y_1 - c_1 \tag{17}$$

$$a_t - Ra_{t-1} = y_t - c_t \tag{18}$$

$$c_t = c_1 (\beta R)^{\frac{t-1}{\sigma}} \psi_t^{\frac{1}{\sigma}}, \ t = \overline{1, T} \tag{19}$$

We notice in the above relations that the return on wealth affects first period consumption and savings through the same three channels that I highlighted in the two period model:

- substitution effect: a lower R makes it less attractive for people to shift resources to the future. Hence, they prefer to consume relatively more earlier in their life. This is the R corresponding to the term $(\beta R)^{\frac{t-1}{\sigma}} \psi_T^{\frac{1}{\sigma}}$ in equation (16). Consumption in the first period c_1 increases and savings in the first period a_1 decreases.

- income effect: a lower R requires people to save more to prevent a too high reduction in consumption after retirement. This is the R corresponding to the term $\frac{1}{R^{t-1}}$ in equation (16). Consumption in the first period c_1 decreases, savings in the first period a_1 increases.

- human wealth effect: a lower R increases the present value of earnings $y_1 + \frac{y_2}{R} + ... + \frac{y_{ra}}{R^{ra-1}}$. People can afford to consume more throughout their life and hence save less. As people approach retirement age, the present value of future income declines as does the effect of the wealth effect on savings.

In the above analysis we determined how savings $a_t - Ra_{t-1}$ changes when R is lower. However, the impact on wealth a_t is also affected by the fact that wealth accumulation will proceed more slowly when R is lower. Therefore even if savings turns out to be higher after a certain age, total wealth can still end up being lower.

I rewrite the optimal consumption in the first period of a person's life (16) in the following way:

$$c_1 = \frac{\sum_{t=1}^{ra} \frac{y_t}{R^{t-1}}}{\sum_{t=1}^{T} \psi_t^{\frac{1}{\sigma}} \beta^{\frac{t-1}{\sigma}} R^{\frac{(t-1)(1-\sigma)}{\sigma}}} \tag{20}$$

From the formula above, I can draw the following conclusions regarding the impact of R on consumption and savings:

- for $\sigma > 1$, just as in the small model, the income effect dominates the substitution effect, i.e. a lower R implies higher level of savings in the beginning of a household's life a_1. However, the overall effect (including the human wealth effect) is difficult to determine.

- if the household retires after the first period, i.e. $ra = 1$, then there is no human wealth effect. For $\sigma > 1$, a lower R implies higher level of savings in the beginning of a household's life a_1. The further in the future is the retirement age of a household, the more important is the human wealth effect.

4 Results in a calibrated version of the multi-period model

In order to determine the overall impact of a lower return on wealth R on the savings, wealth and consumption of a household, I calibrate the model presented in section 3 using inputs relevant for the Dutch economy. In order to determine the life-cycle income profile relevant for the Dutch economy, I use the Inkomenspanelonderzoek (IPO) data provided by Statistics Netherlands (CBS). The sample period covers the years 2006-2013. The measure of income which I consider is the gross income of households net of financial income. I deflate the income to real values using the Consumer Price Index. In order to correct for the size of the household, I transform the measure of income into income per household member by dividing it by the CBS equivalization scale[5]. This scale assigns a value of 1 to a household consisting of one adult, adds 0.37 for every adult and a value between 0.15 and 0.33 for every under age child. From the gross income I deduct all taxes paid using the tax schedule of the Dutch economy[6]. I then add back the contributions paid to the first and second pension pillar[7].

I compute the mean of this income measure across all households with a head aged t years, with t ranging between 25 years and 65 years. This will be the input for the variable y_t in the model (see equation (16)). The resulting profile of mean income is presented in Figure 1. The profile of income is increasing over the life until the age of 57 years and then declines. As labor force data shows (see OECD Labor Force Statistics), this decline in household income is likely due to a decline in the number of hours worked and to early retirement. In this paper, I determine the optimal savings behavior of a person that works full-time until the retirement age and leave aside the impact of a lower rate of return on the labor supply decision of the household. Hence, I eliminate the decline in the income profile of Dutch households due to labor supply adjustments by setting the income constant after the age of 57 years.

[5] In this analysis income is equivalized. Consequently, the fact that the composition of a household changes over time has no impact on savings per household member. Explicitly modelling the size of the household may provide different quantitative results. There is an ongoing project at the CPB that looks specifically into how family size impacts on the optimal path of consumption and savings in a low interest rate environment.

[6] I use the income tax schedule corresponding to the year 2010. This is consistent with the fact that values are made real by dividing through the CPI with a fixed base in year 2010. For this year, the income tax brackets were: 33.45% for the income up to 18.218 EUR, 41.95% for the part of the income between 18.218 EUR and 32.738 EUR, 42% for the part of the income between 32.738 EUR and 54.367 EUR and 52% for the income that exceeds 54.367 EUR.

[7] In the year 2010, the contribution to the first pension pillar (AOW contribution) was 17.9% of the income up to 32.738 EUR. For the second pillar, I considered that contributions are equal to 23% of the income that exceeds the AOW francize. The value of the AOW francize was 12.800 EUR in the year 2010.

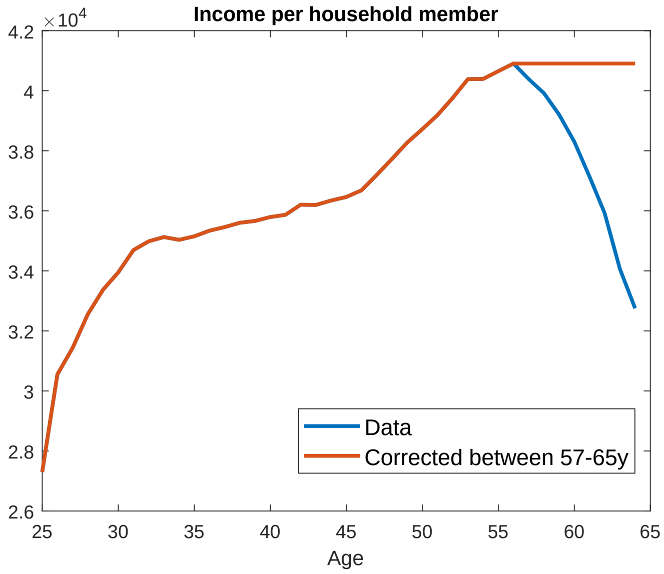

Figure 1: Net income plus pension premiums

In the model, a person starts earning an income at the age of 25, stops receving an income at the retirement age of 65 years, i.e. $ra = 40$ and dies with certainty when she reaches 98 years, i.e. $T = 74$. The survival probability from one year to the other is age-varying and equal to the surviving probability of the Dutch population (source CBS, figure 2). The calibration of all parameters is summarized in table 1.

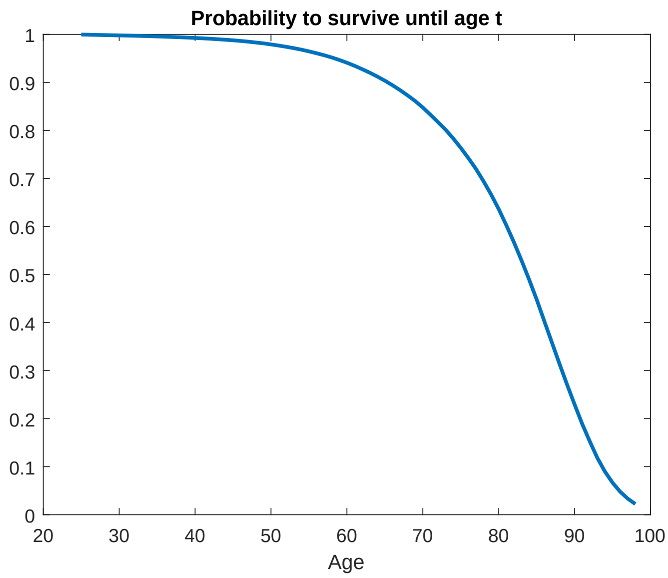

Figure 2: Probability to survive until age t

Table 1: Calibration

Parameter		Value
σ	Inverse of elasticity of intertemporal substitution	2
β	Time preference	0.98
R	Return on wealth	1.03
T+25	Maximum number of years	98
ra+25	Retirement age	65

I consider an elasticity of intertemporal substitution $\frac{1}{\sigma}$ equal to 0.5. This is the mean of the values estimated in the literature (Havranek et al. (2015)). The gross return on wealth is set to 1.03. The estimations for the time preference parameter β range in the literature between 0.93 and 0.99, depending also on the education level (Alan and Browning (2010), French (2005), Gourinchas and Parker (2002), Guvenen and Smith (2014), Cagetti (2003)). For the baseline scenario, I choose a value of 0.98 that corresponds with an annual discount factor of 2%. I present robustness checks of the results for the values of σ and β.

With the assumptions on the parameters presented in Table 1, I compute the profile of optimal consumption, savings and wealth using equations (16)-(19). I will call this *the baseline scenario*. Then I consider a permanent decline in the return on wealth R from 1.03 to 1.02. The permanent decline in R takes place in the moment that a person enters the labor market. I recalculate the profile of optimal consumption, savings and wealth with the lower return on wealth. I will call this *the lower R scenario*.

The profile of the optimal savings rate is upward sloping over the career of a person (figure 3, lower graph). The optimal savings rate starts a bit above the value of 0%, increases over the career of a person and stabilizes around the value of 25% of net income after the age of 50 years.

A lower return on wealth R decreases the optimal savings rate until the age of 35 years (figure 3, lower graph). This implies that the human wealth effect dominates the income effect early in a person's career. The savings rate is slightly negative in the first two years of a person's career, indicating that people would optimally like to borrow at that age. If we take into account that people use student loans to finance their education, this is exactly what happens in reality as well. After the age of 35 years the savings rate becomes higher in the lower R scenario, a consequence of the fact that human wealth declines as people reach the retirement age. Correspondingly, consumption is higher in the first years of a person's career and then becomes lower (figure 3, right panel). Overall, the consumption profile becomes flatter across the life of the individual.

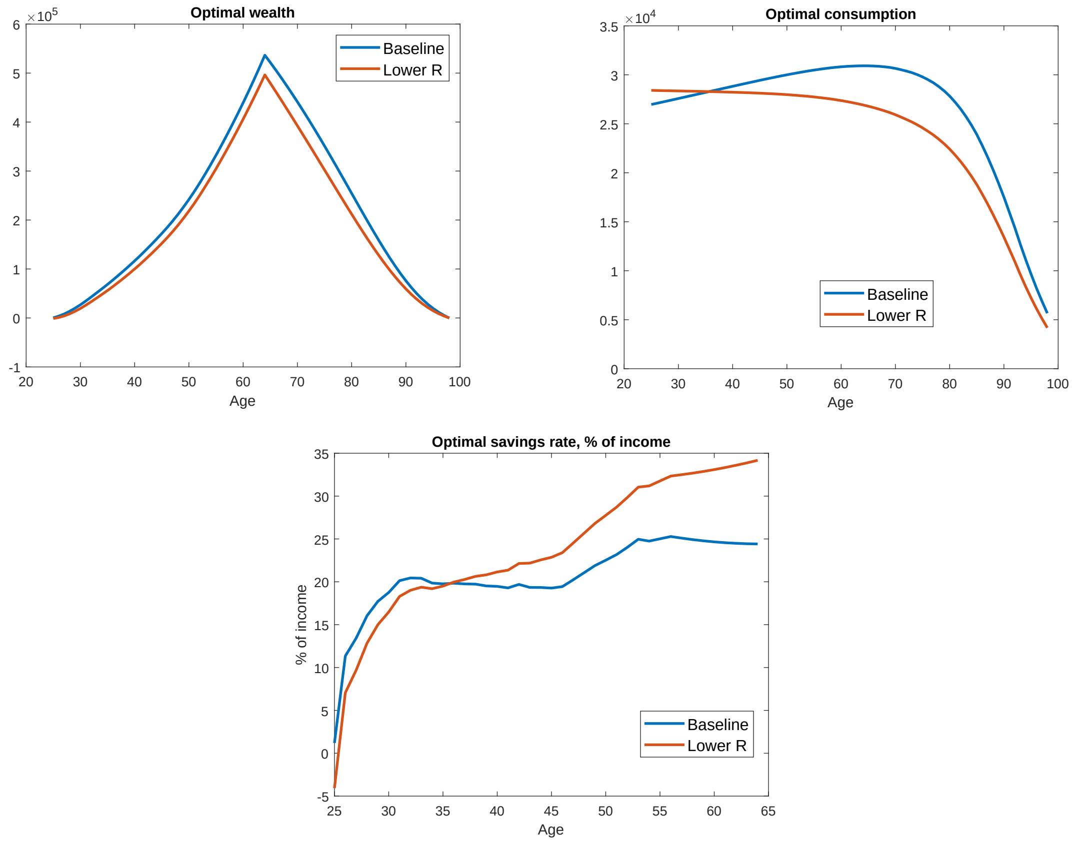

Figure 3: The impact of a lower R

The optimal level of wealth accumulated by a person (the sum of accrued savings) is lower at every age: at the peak, optimal wealth is 7% lower in the low R scenario. This result has two causes: i) in the lower R scenario savings are lower during the first part of a person's career and ii) wealth accumulates at a slower pace due to the lower R.

We next investigate how sensitive our results are to the parameters of the model.

Sensitivity with respect to the retirement age. In the baseline scenario I assumed a retirement age of 65 years. However, this is expected to increase in the future and may not be a relevant retirement age for the cohorts that are currently entering the labor market. Consequently, I consider a variant of the model in which households retire at 70. The differences between the baseline and the lower R scenario are qualitatively the same (figure 4): wealth is lower at every age, consumption is flatter (higher at younger ages and lower at older ages), people have a lower saving rate when young and a higher saving rate when old. Quantitatively, when R declines, the savings rate is lower for a longer period of a person's life (until the age of 40 instead of 35 in the baseline model) because the human capital effect is higher. After the year of 40, the savings rate increases less than in the case that retirement takes place sooner. This is because people obtain a return from income for longer and they have to build up less savings in order to consume during retirement. The optimal wealth accumulation is 8% lower instead of the 7% drop obtained in the baseline model.

Overall, when people work longer, they need to increase saving less (or even decrease it) following a decline in R.

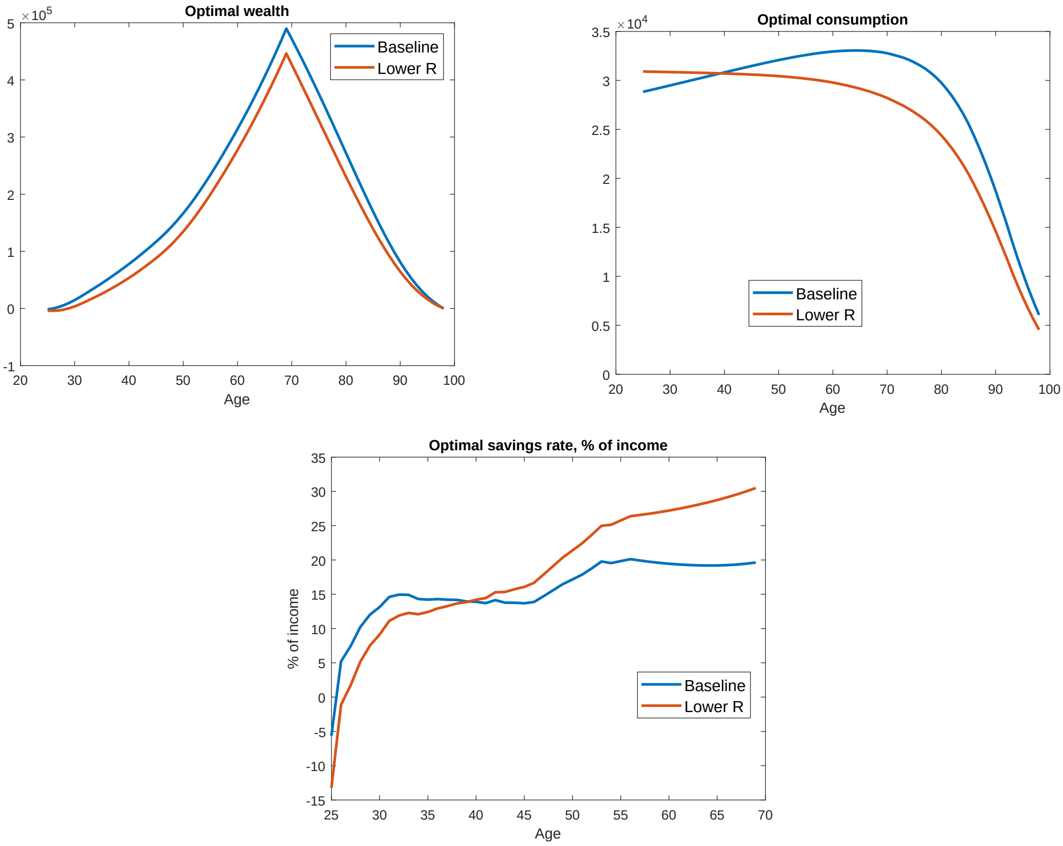

Figure 4: The impact of a lower R - higher retirement age

Sensitivity with respect to time preference. The value of time preference β chosen in the baseline scenario lies at the higher end of the values estimated in the literature (Alan and Browning (2010), French (2005), Guvenen and Smith (2014), Gourinchas and Parker (2002)). I analyse here the implications of people being less patient than assumed in the baseline scenario. I set β to a lower value $\beta = 0.96$ and construct the same scenarios: $R = 1.03$ and then $R = 1.02$. Because people are less patient, they prefer current consumption to future consumption. Hence consumption is decreasing over a person's life (figure 5, top right panel). The comparison between the baseline and the lower R scenario yields the same qualitative results: in the lower R scenario, consumption is lower after retirement, wealth is lower throughout a person's life (figure 5, top left panel) and the savings rate is lower in the first part of a person's career and higher in the second part (figure 5, bottom panel).

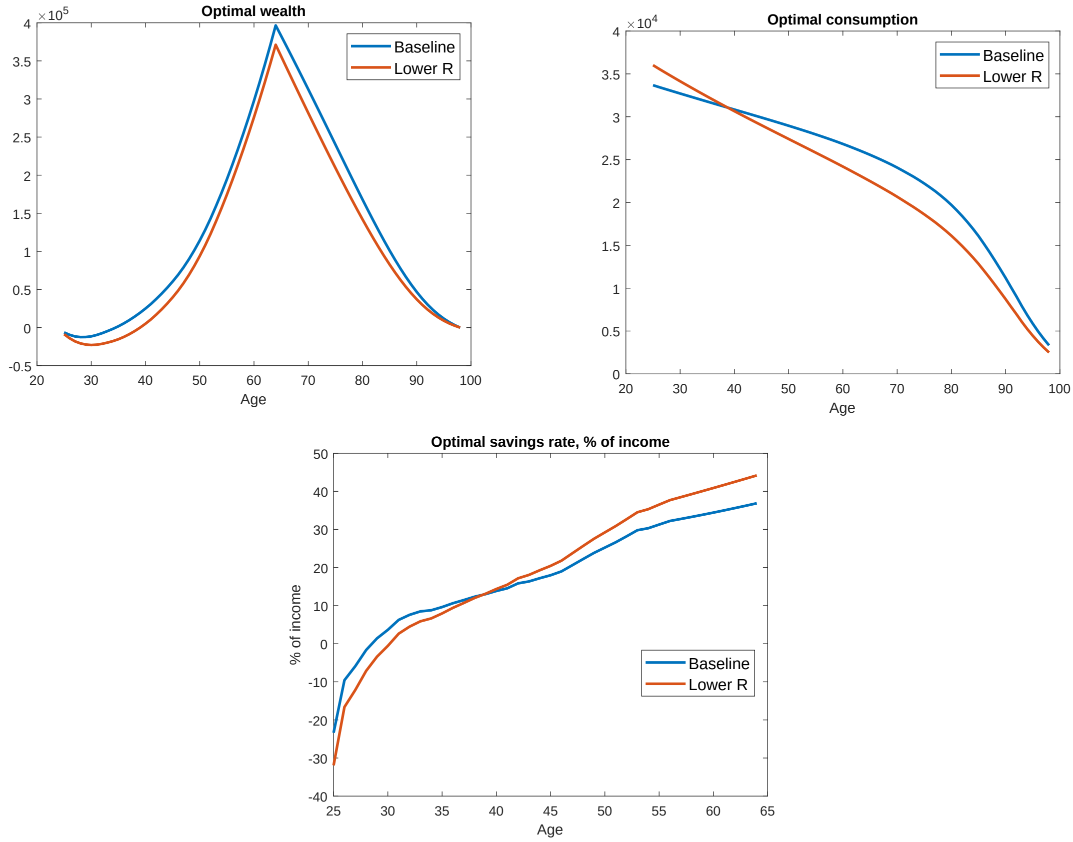

Figure 5: The impact of a lower R - lower patience

Sensitivity with respect to the elasticity of intertemporal substitution. The value of the elasticity of intertemporal substitution chosen in the baseline scenario (0.5, corresponding to $\sigma = 2$) represents the mean of the values estimated in the literature (Havranek et al. (2015)). However, a large part of the macroeconomic literature uses lower values of σ ($\sigma = 1$, implying a high elasticity of intertemporal substitution), while many structural estimations of this parameter find substantially higher values that imply a very low willingness of people to substitute consumption across periods (see for example French (2005)).

I check what different values of σ imply for the results of the analysis. I simulate the baseline and the low R scenario for $\sigma = 1$ and $\sigma = 4$. The results are presented in figure 6 and figure 7, respectively.

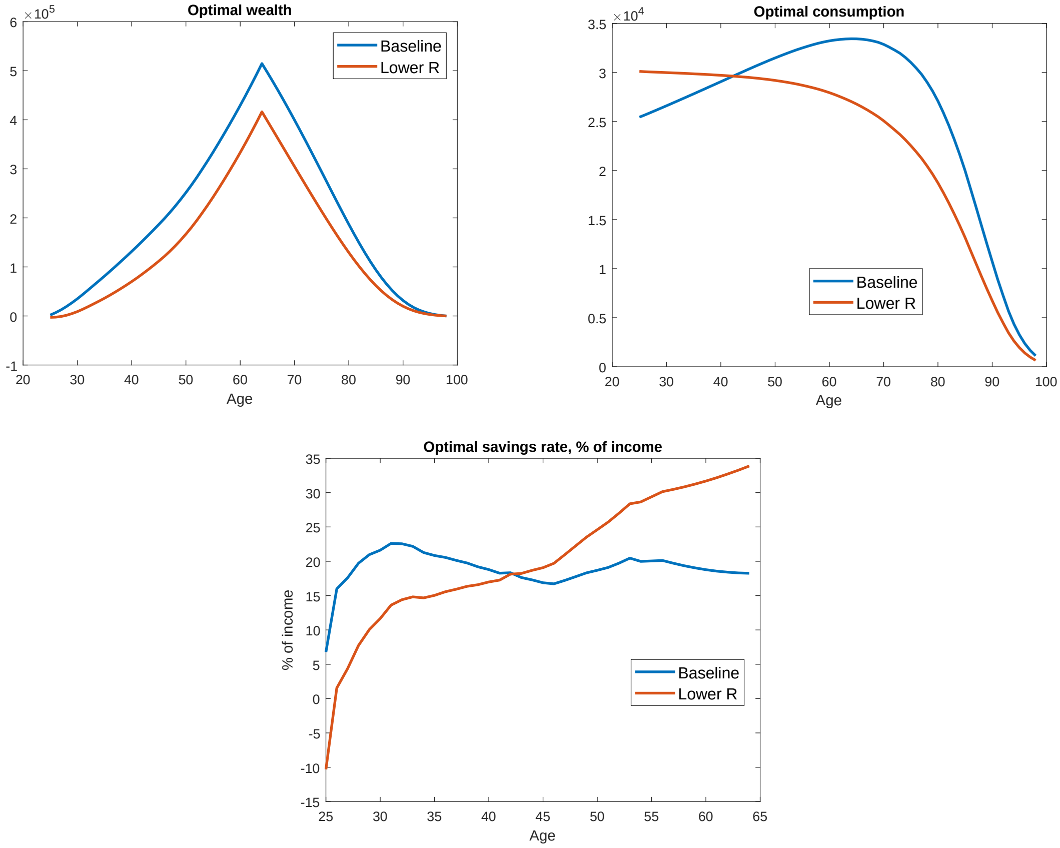

Figure 6: The impact of a lower R - higher elasticity of intertemporal substitution ($\sigma = 1$)

Compared to the results obtained with $\sigma = 2$, in the model with $\sigma = 1$ the substitution effect is quantitatively more important. The sum of the substitution and the human wealth effect is higher than the income effect for a longer part of a person's life when $\sigma = 1$. Consequently, in the lower R scenario, savings are lower for a longer period of time and are also substantially lower than in the baseline scenario (compare figure 6, bottom panel with figure 3, bottom panel). People have 14% less wealth accumulated at retirement age when $\sigma = 1$ than when $\sigma = 2$ (compare figure 6, top left panel with figure 3, top left panel). Intuitively, when σ is low, the elasticity of intertemporal substitution is high. Hence, a change in R has a more important quantitative impact on the profile of consumption and savings.

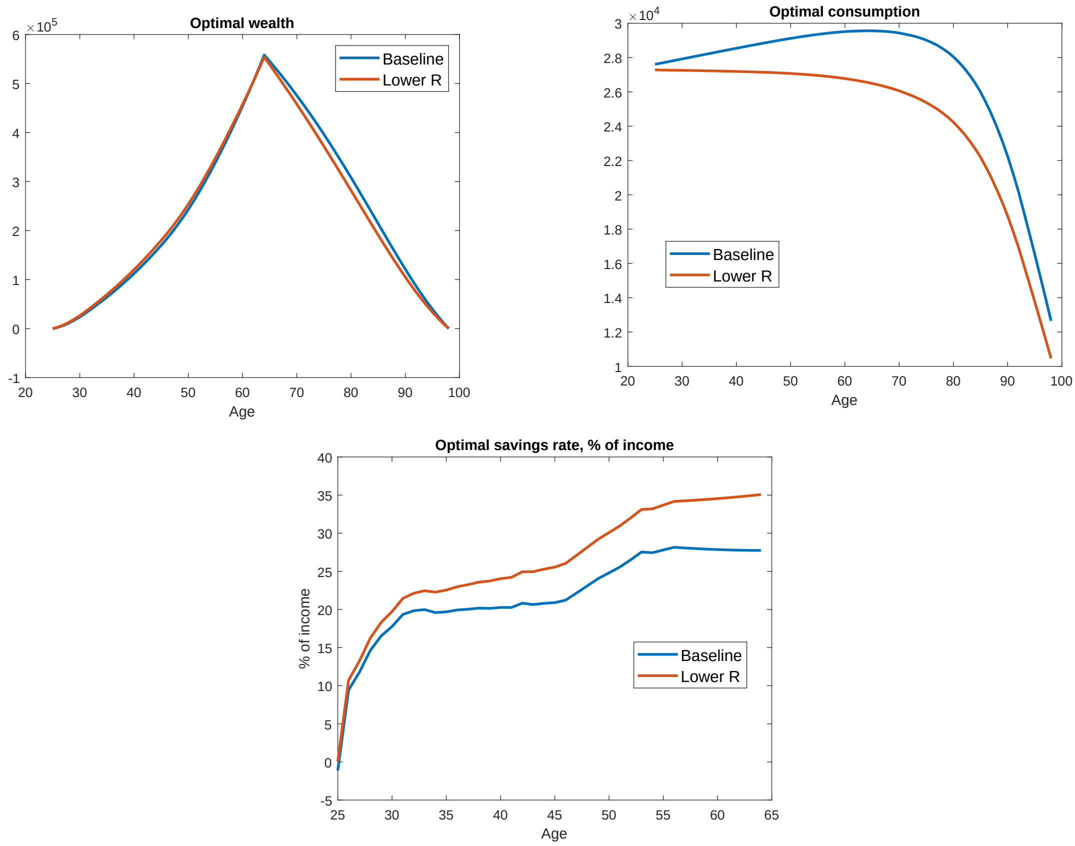

Figure 7: The impact of a lower R - lower elasticity of intertemporal substitution ($\sigma = 4$)

When $\sigma = 4$ the substitution effect is low, so the income effect is higher than the sum of the substitution and human wealth effect at every age. People have a higher savings rate throughout their life (compare figure 7, bottom panel with figure 3, bottom panel). A consequence of the higher amount of savings is that the optimal wealth held by a person is almost unchanged when R is lower (compare figure 7, top left panel with figure 3, top left panel). Intuitively, when σ is high, the elasticity of intertemporal substitution is low. Optimal consumption is lower throughout a person's life (figure 7, top right panel).

Accounting separately for AOW premiums and benefits. In this subsection I take into account the fact that part of the premiums paid throughout an individual's lifetime represents AOW premiums. As AOW benefits are not affected by the lower R, it is useful to see the changes in optimal wealth accumulation net of AOW premiums and benefits. Consequently, I deduct the standard taxes paid by a person and do not add back the part that represents the AOW contributions (17.9% of gross income up until an income of 32.738 EUR in the year 2010). The income received after the retirement age is equal to the AOW benefits net of taxes (on average across Dutch households, this was equal to 8564 euro). Figure 8 presents the results of this scenario.

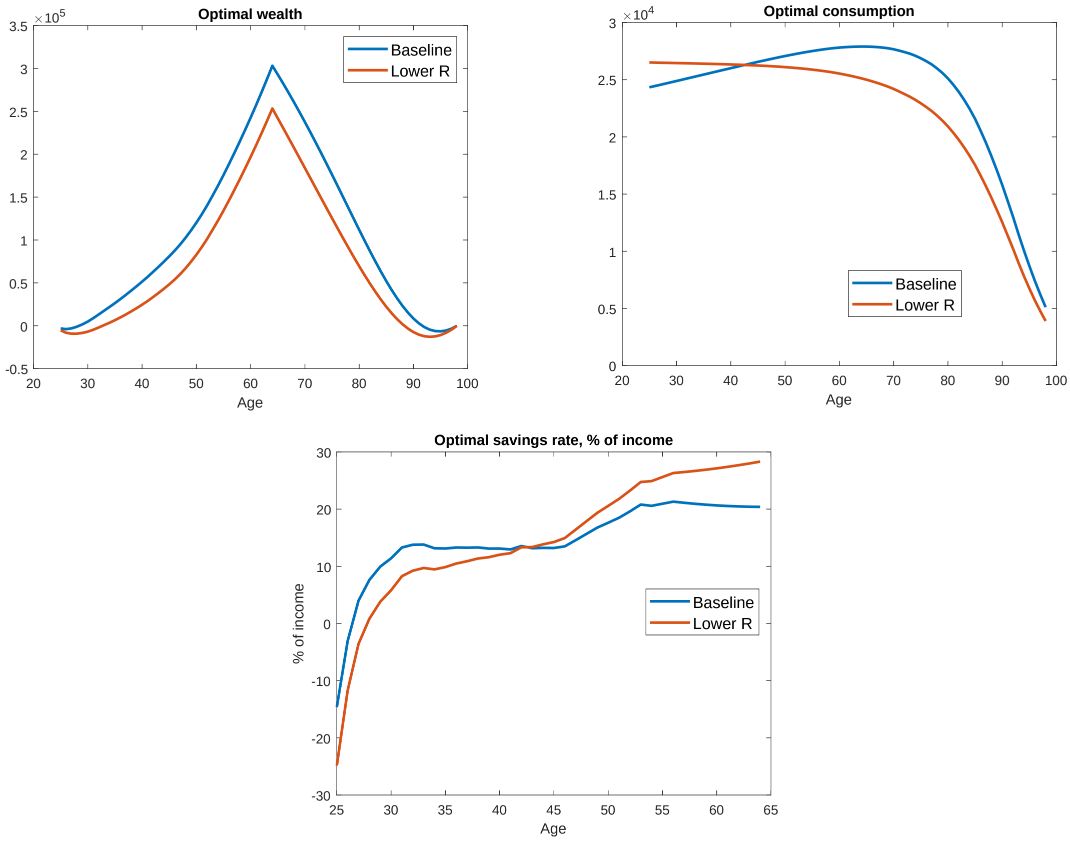

Figure 8: The impact of a lower R when AOW is separately accounted for

The impact of a lower R on optimal savings and consumption choice is qualitatively the same. However, the savings rate is lower for a more extended period of a person's life (until the age of 43, figure 8, bottom). This is because people earn an income even after the retirement age through the AOW. As a result the human wealth effect is more important. The decline in optimal wealth accumulation is more substantial when we take AOW benefits in account. More specifically, optimal wealth accumulation is now 16% lower when R decreases by 1 p.p.

5 Other factors than influence the impact of a lower return on consumption and savings

The analysis performed in the previous sections does not include some factors that may prove important for the decision of households.

First, I did not take into account the impact of a lower R on savings through the financial wealth effect. I analyzed how a permanently lower return on wealth impacts on the consumption and savings decision of a person that just entered the labor market. However, the impact can be different for people who are at a different point of their career when the shock takes place. For example, people that approach the end of their career have already accumulated some wealth in the form of stocks, long term bonds or housing. A lower return on wealth leads to an increase in the value of the wealth already acquired[8]. Hence, these people will not need to save more for retirement as the main analysis shows.

Second, throughout this analysis I assumed that people retire at the AOW retirement age. If the return on wealth becomes permanently lower, people can prevent a too high decline of consumption after retirement by also choosing to retire later. In this case, the impact of the lower R on savings will be lower than quantified in this analysis.

Third, a lower return on wealth makes it worthwhile to borrow and invest in one's human capital. This would help people achieve a higher increase in income over the life. In this case, people could achieve the same level of consumption even with a lower R without increasing their savings much.

Fourth, one result of the analysis is that people would optimally save less (borrow more) in the first part of their career if the return on wealth declines. However, if people are already at their maximum borrowing limit, it will be impossible to further increase borrowing. Hence the impact of a lower R on savings will be small in the beginning of a person's career.

Finally, people have also other reasons for saving than retirement. Saving for a "goal" (for example to afford the education of children) will require an increase in savings when the return on wealth is lower (only the income effect is relevant in this case).

References

Alan, S. and Browning, M. (2010). Estimating intertemporal allocation parameters using synthetic residual estimation. *The Review of Economic Studies*, 77(4):1231–1261.

Cagetti, M. (2003). Wealth Accumulation over the Life Cycle and Precautionary Savings. *Journal of Business & Economic Statistics*, 21(3):339–353.

Elmendorf, D. W. (1996). The effect of interest-rate changes on household saving and consumption: a survey. Finance and Economics Discussion Series 96-27, Board of Governors of the Federal Reserve System (U.S.).

French, E. (2005). The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour. *Review of Economic Studies*, 72(2):395–427.

Gourinchas, P.-O. and Parker, J. A. (2002). Consumption Over the Life Cycle. *Econometrica*, 70(1):47–89.

Guvenen, F. and Smith, A. A. (2014). Inferring Labor Income Risk and Partial Insurance From Economic Choices. *Econometrica*, 82:2085–2129.

[8] We must note that if people hedge against the return risk, then a lower return on wealth has no impact on savings through the financial wealth effect. Pension funds partially hedge against the risk of a lower return.

Havranek, T., Horvath, R., Irsova, Z., and Rusnak, M. (2015). Cross-country heterogeneity in intertemporal substitution. *Journal of International Economics*, 96(1):100–118.

Jones, C. I. (2018). *Macroeconomics*. W.W.Norton and Company.

Ordonez, G. and Piguillem, F. (2020). Savings and saving rates: Up or down? Working Paper 27179, National Bureau of Economic Research.

Summers, L. H. (1981). Capital Taxation and Accumulation in a Life Cycle Growth Model. *American Economic Review*, 71.

Appendix

Proofs from Section 2

The optimization problem that a person has to solve is:

$$\max_{\{c_y, c_o, a\}} u(c_y) + \beta\psi u(c_o) \tag{21}$$

$$c_y + a = y \tag{22}$$

$$c_o = aR + yG\lambda \tag{23}$$

The budget constraints (22) and (23) are consolidated in a life-time budget constraint:

$$c_y + \frac{c_o}{R} = y + \frac{yG\lambda}{R} \tag{24}$$

I construct the Lagrangian:

$$\max_{\{a, c_o, c_y\}} \frac{c_y^{1-\sigma}}{1-\sigma} + \beta\psi\frac{c_o^{1-\sigma}}{1-\sigma} + \lambda\left(y + \frac{yG\lambda}{R} - c_y - \frac{c_o}{R}\right) \tag{25}$$

First order conditions with respect to c_o, c_y, λ:

$$c_y^{-\sigma} = \lambda$$

$$\beta\psi c_o^{-\sigma} = \frac{\lambda}{R}$$

$$c_y + \frac{c_o}{R} = y + \frac{yG\lambda}{R}$$

Combining the first two equations from above I obtain the Euler equation from the text:

$$c_o = c_y\left(R\beta\psi\right)^{\frac{1}{\sigma}} \tag{26}$$

Finally, I substitute the relationship between c_y and c_o from (26) into the life-time budget constraint (24) to obtain the final solution to the household's problem:

$$c_y + \frac{c_y\left(R\beta\psi\right)^{\frac{1}{\sigma}}}{R} = y + \frac{yG\lambda}{R} \Rightarrow c_y = \frac{y + \frac{yG\lambda}{R}}{1 + \frac{(R\beta\psi)^{\frac{1}{\sigma}}}{R}}$$

$$c_o = c_y\left(R\beta\psi\right)^{\frac{1}{\sigma}} = \left(R\beta\psi\right)^{\frac{1}{\sigma}} \frac{y + \frac{yG\lambda}{R}}{1 + \frac{(R\beta\psi)^{\frac{1}{\sigma}}}{R}}$$

$$a = y_1 - c_y = y - \frac{y + \frac{yG\lambda}{R}}{1 + \frac{(R\beta\psi)^{\frac{1}{\sigma}}}{R}}$$

Proofs from Section 3

The problem that a person must solve is:

$$\max_{\{c_t, a_t\}} \sum_{t=1}^{t=T} \beta^{t-1}\psi_t u(c_t) \tag{27}$$

$$c_t + a_t = y_t + Ra_{t-1}, t \leq ra \tag{28}$$

$$c_t + a_t = Ra_{t-1}, t > ra \tag{29}$$

$$a_0 = a_T = 0 \tag{30}$$

The period by period budget constraints from (28)-(29) can be consolidated in a lifetime budget constraint:

$$c_1 + \frac{c_2}{R} + ... + \frac{c_T}{R^{T-1}} = y_1 + \frac{y_2}{R} + ... + \frac{y_{ra}}{R^{ra-1}} \tag{31}$$

The Lagrangian of the problem:

$$\max_{\{c_t\}} \sum_{t=1}^{t=T} \beta^{t-1}\psi_t \frac{c_t^{1-\sigma}}{1-\sigma} + \lambda\left(y_1 + \frac{y_2}{R} + ... + \frac{y_{ra}}{R^{ra-1}} - c_1 - \frac{c_2}{R} - ... - \frac{c_T}{R^{T-1}}\right) \tag{32}$$

The first order conditions of the above maximization problem are:

$$c_1^{-\sigma} = \lambda \tag{33}$$

$$\beta^{t-1}\psi_t c_t^{-\sigma} = \frac{\lambda}{R^{t-1}}, \forall t = \overline{2, T} \tag{34}$$

$$y_1 + \frac{y_2}{R} + ... + \frac{y_{ra}}{R^{ra-1}} - c_1 - \frac{c_2}{R} - ... - \frac{c_T}{R^{T-1}} = 0 \tag{35}$$

I rewrite the above first order conditions for consumption:

$$c_1^{-\sigma} = (\beta R)^{t-1}\psi_t c_t^{-\sigma}, \forall t = \overline{2, T}$$

$$c_1 = (\beta R)^{-\frac{t-1}{\sigma}}\psi_t^{-\frac{1}{\sigma}} c_t \tag{36}$$

It follows that the relationship between first period consumption c_1 and the consumption in period t is given by:

$$c_t = c_1(\beta R)^{\frac{t-1}{\sigma}} \psi_t^{\frac{1}{\sigma}}, t = \overline{1, T}$$

I substitute the above in the life-time budget constraint of the household:

$$c_1 + \frac{c_1 \psi_2^{\frac{1}{\sigma}} (\beta R)^{\frac{1}{\sigma}}}{R} + ... + \frac{c_1 \psi_T^{\frac{1}{\sigma}} (\beta R)^{\frac{T-1}{\sigma}}}{R^{T-1}} = y_1 + \frac{y_2}{R} + ... + \frac{y_{ra}}{R^{ra-1}}$$

I obtain a solution for the first period consumption and savings of the household:

$$c_1 = \frac{y_1 + \frac{y_2}{R} + ... + \frac{y_{ra}}{R^{ra-1}}}{1 + \frac{\psi_2^{\frac{1}{\sigma}} (\beta R)^{\frac{1}{\sigma}}}{R} + ... + \frac{\psi_T^{\frac{1}{\sigma}} (\beta R)^{\frac{T-1}{\sigma}}}{R^{T-1}}}$$

$$a_1 = y_1 - c_1$$