Labor Market Risk and the Private Value of Social Security

Sylvain Catherine *

November 6, 2019

Abstract

Social Security provides insurance against idiosyncratic income risk but exposes workers to systematic risk because benefits are indexed to the evolution of aggregate earnings. I calibrate a life-cycle model to compare workers' certainty equivalent valuation of Social Security to its net present value discounted at the risk-free rate. I show that, overall, labor market risk reduces current workers' private value of Social Security by 46%. This adjustment sums up to \$11.4 trillions on the national scale and the equity premium is its main determinant. For workers under 35, the certainty equivalent of Social Security is negative. Exposure to systematic risk through Social Security peaks relatively late in the life-cycle.

Keywords: Household finance, Social Security, Public liabilities, Portfolio choices, Human capital, Labor Income Risk

JEL codes: G11, G18, D91, H55, H06

*The Wharton School, scath@wharton.upenn.edu.

I am thankful to Olivia Mitchell, João Cocco, Joseph Frédéric Bonnans and David Thesmar for their comments as well as conference participants at the London Business School Summer Finance Symposium, the HEC Paris brownbag seminar and the INSEAD PhD workshop. All errors are mines.

1 Introduction

In most developed countries, Social Security pension entitlements represent one of the largest assets owned by households. In 2018, the US Social Security Administration (SSA) estimates that, net of expected contributions, the present value of accrued benefits to current workers and retirees was \$35.5 trillion.[1] By comparison, households' owner equity in real estate was \$15.5 trillion.[2] Unlike real estate, Social Security claims are not traded and their valuation depends on modeling assumptions and, in particular, the choice of an appropriate discount rate.

Understanding households' valuation of Social Security and its risk profile are important for several reasons. First, the sum of private valuations constitutes an upper bound on the cost of transition to a funded system. The size of this cost is essential to determine, for example, whether a privatized system could offer higher rate returns (Geanakoplos et al. (2000)). Second, rising wealth inequalities have become a subject of public concern. Yet, despite its considerable size and redistributive nature, Social Security has generally been left aside in studies of the wealth distribution (e.g. Saez and Zucman (2016)). Finally, in view of its size, Social Security risk is important to understand households' portfolio and consumption choices.

In this paper, I calibrate a life-cycle model to study how labor market risk affects the way households' should value Social Security and how their valuation should evolve over the life-cycle, in response to stock market fluctuations and after labor income shocks. While the SSA assumes that Social Security cash flows should be discounted at the risk-free rate, workers worried by labor market risk may want to

[1]This includes \$32.4 trillion of unfunded obligations for past and current participants reported in the SSA's *Actuarial Note Number 2019.1* and \$2.9 trillion in the Social Security Trust Fund.

[2]According to the Federal Reserve of Saint Louis, for Q42018 (https://fred.stlouisfed.org/graph/?g=nTQ5)

use a different discount rate for two reasons. First, the value of Social Security benefits is indexed to the nationwide evolution of earnings, which is an implicit feature of all pay-as-you-go systems (Samuelson (1958)). In a world where the labor share is relatively stable over long periods, wage-indexation also exposes Social Security participants to the long-run growth rate of corporate profits. In that case, workers can hedge against wage-indexation by short-selling the stock market. The cost of doing so, the equity premium, should therefore be reflected in the discount rate of Social Security. Second, workers with lower than average earnings receive a higher rate of return on their Social Security contributions, which means that negative idiosyncratic income shocks increase the value of past contributions. This feature of Social Security increases its value when idiosyncratic income risk is uninsurable (Merton (1983)). Overall, these relationships between Social Security and income risk suggest that workers should not discount Social Security cash flows at the risk-free rate, even if we ignore political risk. Moreover, the effect of income risk on their valuation of Social Security is ambiguous.

To study workers' private valuation of Social Security, I set up a continuous time life-cycle model in which individuals contribute to a pay-as-you-go system similar to the United States Social Security and compute the certainty equivalent of expected benefits net of future contributions. In the model, an agent with constant relative risk aversion (CRRA) receives a stochastic labor income until retirement and chooses how much to consume or save. Savings can be invested in a risk-free asset or in the stock market portfolio. This model captures important features of US Social Security. First, workers pay a 10.4% payroll tax to current retirees. Second, benefits are computed on the basis of their historical earnings, and low earners enjoy relatively higher rates of return on their contributions than high earners. Third, the value of past contributions is indexed to the nationwide evolution of labor earnings. Finally, as in Benzoni et al. (2007), aggregate dividends and earnings are cointegrated and households can therefore offset their exposure to the evolution of aggregate earnings by investing less in stocks.

My main results are as follows. First, for current workers, the certainty equivalent of Social Security represents only half the sum of future cash flows discounted at the risk-free rate. Second, unlike the risk-free rate valuation, the certainty equivalent is negative for new workers currently entering the labor market and remains negative until around age 35. In the model, new workers would be willing to pay more than a year of earnings to opt out of Social Security. Third, on the national scale, and including retirees, the sum of certainty equivalents is 37% below the risk-free rate valuation. The cost of transition to a funded system with the consent of current participants is therefore significantly lower than implied by SSA estimates. The model suggests that the sum of certainty equivalents was \$19.6 trillion as of 2013, whereas the risk-free valuation was \$31 trillion. Finally, the model implies that workers in the middle of their careers apply a discount rate of 4.5%, which includes a 2.5% premium relative to the risk-free rate. Importantly, all these results assume away political risk.

To calibrate risk aversion and validate the model, I match the evolution of the average equity share over the life-cycle in the Survey of Consumer Finances (SCF). I find that for a relative risk aversion of $\gamma = 5$, the model tracks the life-cycle pattern of the share of financial wealth invested in equity relatively well.

The discount rate used by households to value Social Security depends on their age, and is the highest around 40 years old. As they get closer to retirement, the risk-adjusted discount rate converges to the risk-free rate, because the long run correlation between stock market returns and the growth of the wage index becomes less of a concern. For workers who just entered the labor market, cointegration does not greatly affect the value of Social Security either. From their point of view, positive news regarding the growth rate of wages means higher pensions but also higher payroll taxes. Nonetheless, young households value Social Security negatively because, for the next few decades, it will force them to invest in an asset with low returns given its market beta. After about age 40, households have already contributed for roughly 20 years and are still young enough for the final value of their past contributions to be exposed to cointegration between the labor and the stock markets. Exposure to aggregate risk through Social Security therefore peaks late in the life-cycle, around 10 years before retirement.

Related literature My paper builds on several different literatures on Social Security risk. A first strand of literature on Social Security risk focuses on policy uncertainty. Shoven and Slavov (2006) show that Social Security expected internal rates of return vary significantly across cohorts and over the life-cycle, as a consequence of policy adjustments. More recently, Luttmer and Samwick (2018) infer from survey data that households would be willing to forgo 6 percent of their expected benefits to eliminate policy uncertainty.

Other survey studies have reported substantial uncertainty regarding future benefits. For example, Dominitz and Manski (2006) finds that, from the point of view of the median 40 year old worker, the subjective distribution of yearly benefits has an interquartile range of $10,000. Furthermore, Delavande and Rohwedder (2011) document that US households reporting higher levels of uncertainty have lower equity share, which is consistent with Social Security having a positive market beta.

Relative to this literature, the main contribution of my study is to use a dynamic model to value Social Security and my focus on labor market risk. Despite the ubiquitous use of dynamic programming to price securities, previous papers have mostly used life-cycle models to study retirement choices (e.g. Bodie et al. (1992)) but, to the best of my knowledge, not to compute the value of Social Security. Geanakoplos and Zeldes (2010) and Blocker et al. (2018) use risk-neutral pricing techniques to compute the market value of accrued benefits. In particular, Geanakoplos and Zeldes (2010) estimate that cointegration between the labor and stock markets reduces the market value of Social Security accrued benefits by 19%. Beside the methodology, there are two other important differences between my paper and Geanakoplos and Zeldes (2010) or Blocker et al. (2018). First, these papers provide market values for fully diversified investors, which is not the case of typical Social Security participants who bear a lot of labor market risk. Second, these papers compute the value of accrued benefits, that is the value of future claims based on past contributions. This value is always positive. By contrast, I compute the value of all future benefits, net of future contributions. This turns out to be negative for young workers.

My paper is also related to Benzoni et al. (2007)'s study on portfolio choices. Though they ignore Social Security, these authors document that the labor and stock market are cointegrated and argue that cointegration increases the market beta of human capital and can explain why young workers do not invest in stocks. Huggett and Kaplan (2016) use a life-cycle model to show that cointegration reduces the present value of human capital.

Finally, Geanakoplos et al. (2000) argue that the value of accrued benefits to current Social Security participants affects the comparison of rates of returns between pension systems. For example, if Social Security were privatized, the government would have to issue bonds to honor its past promises. The cost to future taxpayers must be deducted from the gains of moving to a private system. Arguably, the sum of private valuations is a better (and lower) measure of what would need to be paid to current generations than the value of accrued benefits.

2 Social Security and Labor Income Risk

In this section, I present key aspects of the United States Old-Age program and discuss three elements of its risk-profile that are of particular importance to understand households' private valuation: the progressivity of benefits, their indexation to the wage index, and the cointegration between the labor and stock markets.

2.1 Computation of benefits

The Social Security system is a mostly unfunded program in which workers' contributions are directly distributed as pension benefits to current retirees.[3] Since Samuelson (1958), economists know that the performance of an unfunded pension system depends on the nationwide evolution of earnings. In a simple OLG model, the return obtained by one generation is defined as the percentage difference between the payroll taxes it receives from its children and the taxes it pays for its own parents. Holding tax rates constant, this difference is the sum of the population and earnings growth rates. In the US, wage indexation is explicit in the computation of benefits, which is essentially done in two steps.

In the first step, the SSA computes the retiree's average indexed yearly earning (AIYE). This refers to the individual's salary history adjusted for growth in the nominal wage index: hence the series is indexed both for inflation and for growth in real wages. Yearly earnings are capped by an upper limit that has been more than twice the average labor earnings since 1980. The AIYE is then defined as the mean of the best 35 years.

In a second step, the SSA computes the percentage of the AIYE that the agent is eligible for. This percentage is determined by a bend point formula through which

[3] In 2013, the Old-Age and Survivor Insurance (OASI) program paid \$ 672 bn in benefits while its total reserves were 2,674 bn, representing only four years of expenditures.

retirees with higher earnings records get a lower share of their AIYE in benefits. Bend points are limits above which retirees get a smaller percentage of their AIYE and which are themselves wage-indexed. If the agent retires at the full retirement age, the first year benefits are the sum of:

1. 90% of the share of the AIYE below the first bend point;
2. 32% of the share of the AIYE between the first and second bend points;
3. and 15% of the share of the AIYE above the second bend point.

The two bend points being respectively close to 20% and 100% of the wage index, the first year pension benefits B are:

$$B = \begin{cases} 0.9 \times AIYE & \text{if } AIYE/L_1(T) < 0.2 \\ 0.116 \times L_1(T) + 0.32 \times AIYE & \text{if } 0.2 \leq AIYE/L_1(T) < 1 \\ 0.286 \times L_1(T) + 0.15 \times AIYE & \text{if } 1 \leq AIYE/L_1(T), \end{cases} \tag{1}$$

where L_1 is the nationwide average labor earnings, and T the agent retirement year. After retirement, benefits are only adjusted for inflation.

2.2 Cointegration between the stock and labor markets

Because of wage indexation, any statistical relationship that ties the labor and stock markets should tie the discount rate of Social Security to the equity premium. While immediate correlation between labor earnings growth rates and stock market returns is low (Cocco et al. (2005)), Benzoni et al. (2007) find that over the long term, wages and dividends are cointegrated. Adding more recent data, I replicate their analysis and extend it by considering stock market gains instead of dividends. The goal of this analysis is to confirm the possibility of hedging against long-run Social Security risk by short-selling the stock market portfolio.

In economic terms, the cointegration implies that the ratio of wages to dividends tends to revert to an historical mean. I denote $\overline{ld}$ the historical mean of the log difference between the real wage-index (l_1) and the log of S&P500 real dividends ($\hat{d}(t)$), and define y as:

$$y(t) \equiv l_1(t) - \hat{d}(t) - \overline{ld} \tag{2}$$

The variable $y(t)$ measures whether the wage-to-dividend ratio is above or below its historical mean. If this ratio is indeed mean-reverting, then the dynamics of $y(t)$ should be of the form:

$$\Delta y(t) = -\kappa y(t-1) + \epsilon(t) \tag{3}$$

where κ is the speed of mean-reversion and $\epsilon(t)$ an error term. Benzoni et al. (2007) test this model on the 1929-2004 period and find values of κ ranging between 0.046 and 0.264 depending on the specification of the model and the time period.

In the present paper, I compute the average wage as total wages divided by total employees using data from the NIPA tables for the 1929-2011 period.[4] Using stock market data for the S&P500 compiled by Robert Shiller,[5] I run augmented Dickey-Fuller tests on the 1929-2011 period and report the findings in Table 1 (models 1 to 6). Overall, I find values of κ that are consistent with BCG's previous results. In model (2), I allow for a time trend but find it to be insignificant. Models (3) to (6) allow for autocorrelation or errors until the last lag is found insignificant. I find the Dickey-Fuller p-value to be below 10% in all models, except the one with the statistically insignificant time trend.

[4] These data can be retrieved from Emmanuel Saez' website: http://eml.berkeley.edu/~saez/TabFig2012prel.xls - Table B1.

[5] Data available on Robert Shiller's website: http://www.econ.yale.edu/~shiller/data/ie_data.xls

Table 1: Cointegration between the stock and labor markets

Note: I estimate different empirical versions of equation (3) on the 1929-2011 period of the form:

$$\Delta y(t) = -\kappa y(t-1) + b_t t + \sum_{1}^{4} b_i \Delta y(t-i) + c + \epsilon(t)$$

For models (1) to (6), the left hand variable is the variation in the log-difference (y) between S&P500 dividends and the average labor earnings. Model (2) tests the existence of a time trend, and models (3) to (6) allow for autocorrelated errors. For models (7) and (8), the left hand variable is the change in the log-difference between cumulated S&P500 stock returns and the average wage (y_s). Both models assume a time trend. Model (8) tests the existence of autocorrelated errors. ***,**,* indicates p-values below 1%, 5% and 10%. The p-values of $\hat{\kappa}$ are from the Dickey-Fuller distributions and are reported in the lower part of the table.

	Dividends						Stock Gains	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$y(t-1)$	.160**	.161	.233***	.175**	.196***	.180**	.200*	.180
	(.054)	(.054)	(.050)	(.054)	(.055)	(.061)	(.061)	(.065)
t		-.000					-.011***	-.010***
		(.000)					(.003)	(.003)
$\Delta y(t-1)$			.462***	.577***	.732***	.737***		.018
			(.094)	(.096)	(.104)	(.111)		(.109)
$\Delta y(t-2)$				-.355***	-.525***	-.573***		
				(.106)	(.114)	(.139)		
$\Delta y(t-3)$					.340***	.382***		
					(.108)	(.130)		
$\Delta y(t-4)$						-.071		
						(.116)		
DF p-value	.038	.139	.000	.017	.006	.039	.067	.211
$R^2_{Adj.}$	.09	.09	.31	.40	.44	.41	.10	.06
N	82	82	81	80	79	79	82	81

As the main concern of investors are returns rather than dividends, I also test the cointegration of stock and labor markets using stock market gains instead of dividends. To measure stock market gains, I build a series representing the market value of the portfolio of an agent reinvesting dividends in the S&P500 since 1929:

$$S(t+1) = \left(1 + \frac{P(t) - P(t-1) + D(t)}{P(t)}\right) \times S(t), \tag{4}$$

with $S(1929) = 100$. Then, I define y_S similarly to y, using the log of S instead of the log of dividends:

$$y_s(t) \equiv l_1(t) - s(t). \tag{5}$$

If the price-to-dividend ratio P/D is stationary, then y_s should be trend stationary, its trend reflecting the difference between expected stock returns and the growth rate of the wage index. Models (7) and (8) test this hypothesis. In model (7), I allow for a time trend but no autocorrelation and find a speed of mean-reversion of 0.20 and a p-value of 6.7%. Allowing for a lag term increases the p-value to more than 20%, but the autocorrelation term added in model (8) is not statistically significant.

Overall, Table 1 supports the assumption that workers can hedge against the long-run growth rate of earnings by short-selling the stock market portfolio.

3 Model

In the model, an agent receives a stochastic flow of labor earnings until retirement. He also chooses how much to consume and how much to invest in the stock market portfolio. His future labor earnings are subject to permanent idiosyncratic and aggregate shocks.Bad news regarding long run macroeconomic growth affects the stock market immediately because prices reflect all available information. On the other hand, bad news can affect the labor market much later.

3.1 Stock and labor markets

3.1.1 Financial market

The agent can invest in the stock market portfolio or in risk-free asset with interest rate r. The dividend process $D(t)$ of the risky asset is modeled as a geometric Brownian motion:

$$\frac{dD}{D} = g_D dt + \sigma dz_3, \tag{6}$$

where g_D is the expected growth rate of dividends.

For the moment, I assume a time-invariant discount factor. Thus the stock price is proportional to the dividend level and follows the same Brownian motion. The stock return process includes capital gains and dividends and also follows a geometric Brownian motion:

$$\frac{dS}{S} = \frac{dP + Ddt}{P} = \mu dt + \sigma dz_3, \tag{7}$$

where μ is the stock market expected return. In this section, for simplicity, the volatility of stock returns and dividends are the same. I relax this assumption by allowing time-varying discount rates and predictability in returns in section 6.2.

3.1.2 National average wage

To capture the cointegration between stock and labor markets, I use the state-variable y defined in equation (2). Its dynamic is:

$$dy(t) = -\kappa y(t)dt + v_1 dz_1(t) - v_3 dz_3(t), \tag{8}$$

where v_1 represents the standard error of permanent shocks on the average wage and v_3 controls the contemporaneous correlation between stock returns and wages: $\sigma = v_3$ implying no correlation. By combining (6), (2), and (8) we get the following dynamics for the wage index:

$$\frac{dL_1}{L_1} = \left(-\kappa y(t) + g_D - \frac{\sigma^2}{2} + \frac{v_1^2}{2} + \frac{(\sigma - v_3)^2}{2}\right) dt + v_1 dz_1(t) + (\sigma - v_3)\, dz_3(t). \tag{9}$$

The state variable y indicates whether the labor income process is delayed relative to dividends. As such, it determines the expected growth rate of the wage index. A negative shock on dividends increases y and implies a lower expected growth rate of the wage index. This consequently affects the value of human capital and pension entitlements. The parameter v_3 controls the correlation between stock returns and the wage index, with $v_3 = \sigma$ implying no correlation.

3.2 Agent

3.2.1 Labor income

The Social Security wealth of a given individual is affected differently by the nationwide average wage and his own position on the wage scale. For this reason, I decompose his wage as the product of two components: L_1 is the average wage in the economy and $L_{2,i}$ the agent's wage relative to L_1. Hence his labor income is:

$$L(t) = L_1(t)L_2(t) \tag{10}$$

The idiosyncratic component of labor income follows an geometric Brownian motion:

$$\frac{dL_2}{L_2} = \alpha(t)dt + v_2 dz_2(t), \tag{11}$$

where $\alpha(t)$ captures the quadratic effect of experience on wages:

$$\alpha(t) = \alpha_0 + \alpha_1 t. \tag{12}$$

3.2.2 Social Security Benefits

Introducing all the richness of the Social Security formula into a dynamic programming problem is difficult. Therefore, I simplify the problem by assuming that benefits are computed on the basis of a complete career. Because full-career average earnings are lower than the mean of the 35 best years, I multiply the result by a coefficient a chosen such that its expectation equals that of the best-35-years AIYE. Therefore, my proxy for the AIYE is:

$$AIYE(T) = \frac{a}{45}\int_{T-45}^{T} \frac{L_1(T)}{L_1(u)} L_1(u)L_2(u)du, \tag{13}$$

where $L_1(u)L_2(u)$ was the agent's wage at time u and $L_1(T)/L_1(u)$ is the indexation coefficient. This equation assumes a career of 45 years. Within the integral, the $L_1(u)$ terms cancel each others such that the AIYE is the product of the average wage at retirement date $L_1(T)$ and an historical average of $L_{2,i}$. To keep track of historical values of $L_{2,i}$, I define a new state variable H as:

$$H(t) = \frac{a}{45}\int_{T-45}^{t} L_2(u)du, \tag{14}$$

which has the following dynamic:

$$dH = a\frac{L_2(t)}{45}dt. \tag{15}$$

The yearly retirement benefit B depends on the final values of L_1 and H:

$$B = \begin{cases} 0.9 \times AIYE & \text{if } AIYE/L_1(T) < 0.2 \\ 0.116 \times L_1(T) + 0.32 \times AIYE & \text{if } 0.2 \leq AIYE/L_1(T) < 1 \\ 0.286 \times L_1(T) + 0.15 \times AIYE & \text{if } 1 \leq AIYE/L_1(T) \leq 2.5 \\ 0.286 \times L_1(T) + 0.15 \times 2.5 \times L_1(T) & \text{if } AIYE/L_1(T) > 2.5 \end{cases} \tag{16}$$

where

$$AIYE = H(T) \times L_1(T) \tag{17}$$

and where an AIYE above 2.5 times the wage index does not provide any additional benefits to reflect the fact that historical earnings are capped by the maximum earning threshold.

3.2.3 Objective function

The agent chooses his optimal consumption (C) and the share of his financial wealth invested in the market portfolio (π). I explicitly model the working years, between 20 and 65 years old and assume that once retired, the agent face the simpler problem described by Merton (1971): *i.e.* the agent is endowed with some financial wealth and receives a safe stream of cash flows. Assuming a CRRA utility, his objective function is:

$$J(W(t), L_1(t), L_2(t), y(t), H(t), t) \equiv \max_{[C,\pi]} E_t \left[\int_t^T e^{-\psi u} \frac{(C(u))^{1-\gamma}}{1-\gamma} du + J(T) \right], \tag{18}$$

where $J(T)$ is his expected utility on the day of his retirement, γ his coefficient of relative risk aversion, and ψ his preference for the present.

Importantly, we assume a deterministic time of death $T + R$, where R is the number of retirement years. Stochastic mortality would increase the private value of

Social Security because it provides insurance against longevity risk. A deterministic time of death is therefore important if we want to interpret the difference between the certainty equivalent and the risk-free rate valuation as a consequence of labor market risk. Moreover, if we include stochastic mortality, the private value of the longevity insurance provided by Social Security is unlikely to be correct. First, this insurance is only valuable in the absence of an efficient market for annuities, which we do not want to model here. Second, standard life-cycle models with longevity risk generate a high demand for annuity that is not observed in the data. In my model, because there is no other source of risk in retirement and no desire to leave a bequest, the agent would hold all his wealth in annuities if he could (Yaari (1965)).

3.3 Solution method

Optimal consumption, portfolio decisions, and utility levels are computed by solving the Hamilton-Jacobi-Bellman equation backward from the retirement date. This section focuses on the terminal conditions of the problem, i.e. the agent's behavior after retirement, and the computation of the certainty equivalent of Social Security. Further details on the numerical methodology are provided in Appendix A.

3.3.1 Retirement years

To solve the HJB equation backward, I assume that, once retired, the agent owns financial wealth W and receives a safe stream of benefits from Social Security. Thanks to Merton (1971), we know that the problem of the retired agent has an analytical solution. Specifically, the optimal consumption and equity share are:

$$C^*(T) = b(T)^{\frac{-1}{\gamma}}\overline{W}(T) \tag{19}$$

$$\pi^*(T) = \frac{\mu - r}{\gamma\sigma^2}\frac{\overline{W}(T)}{W(T)} \tag{20}$$

where W_T is *total* wealth at retirement, including the present value of Social Security and where:

$$b(t) \equiv \left(\frac{1 - e^{-v(R+T-t)}}{v}\right)^{\gamma}, \tag{21}$$

$$v(t) \equiv \frac{\varphi - (1-\gamma)\left(\frac{(\mu-r)^2}{2\gamma\sigma^2} + r\right)}{\gamma}, \tag{22}$$

and R is the number of retirement years. After retirement, benefits are safe and can be discounted at the risk-free rate. Hence, total wealth at retirement is:

$$\overline{W}(T) = W(T) + \int_0^R e^{-ru} B du. \tag{23}$$

3.3.2 Certainty equivalents

Before retirement, total wealth ($\overline{W}$) is the sum of three components: financial wealth (W) and the certainty equivalents of human capital (HC) and Social Security (SS):

$$\overline{W} = W + HC + SS \tag{24}$$

Total wealth can be computed as the financial wealth that would make the agent indifferent to a drop of L_1 to zero, a scenario in which he would have no more human capital or Social Security wealth. Thus, $\overline{W}$ is solution to the equation:

$$J^{\overline{ss}}(\overline{W}, 0, L_2, y, 0, t) = J^{ss}(W, L_1, L_2, y, H, t), \tag{25}$$

where J^{ss} and $J^{\overline{ss}}$ denote the expected utility functions of the agent when Social Security respectively exists and does not exist. Conveniently, the left-hand-side of (25) does not depend on L_2, y, or H anymore and it has an analytical solution provided by Merton (1969). Indeed, for a CRRA utility function, we know that:

$$J^{ss}(\overline{W}, 0, L_2, y, 0, t) = \frac{e^{-\varphi t}}{1-\gamma} b(t) \overline{W}^{1-\gamma}, \tag{26}$$

where b is defined by (22) and (21). By combining (25) and (26), we can compute total wealth as:

$$\overline{W} = \left[\frac{(1-\gamma)e^{\varphi t}}{b(t)} J^{ss}(W, L_1, L_2, y, H, t) \right]^{\frac{1}{1-\gamma}}. \tag{27}$$

Similarly, the value of human capital can be defined as the total amount of cash that a worker would need to make him indifferent to a drop of L_1 to zero in a world without Social Security. If Social Security did exist, then the agent would not pay payroll taxes so that labor earnings would be higher. Thus, HC is solution to the equation:

$$J^{\overline{ss}}(W + HC, 0, L_2, y, t) = J^{\overline{ss}} \left(W, \frac{L_1}{1-\tau}, L_2, y, t \right). \tag{28}$$

Following the same steps as for the computation of $\overline{W}$, we get:

$$HC = \left[\frac{(1-\gamma)e^{\varphi t}}{b} J^{\overline{ss}} \left(W, \frac{L_1}{1-\tau}, L_2, y, t \right) \right]^{\frac{1}{1-\gamma}} - W. \tag{29}$$

Finally, we can combine equations (24), (27), and (29) to compute the wealth equivalent of Social Security.

4 Calibration

In this section, I detail the model's baseline calibration, which is summarized in Table 2.

4.1 Macroeconomy

Financial markets Historical estimates of the risk-free rate ranges from 0.64% to 3.02%.[6] In its valuation, the SSA assumes a relatively high risk-free rate at 2.9%. I set $r = 0.02$ in the baseline calibration and also report results for $r = 0.029$ and $r = 0.01$. The equity premium is $\mu - r = 0.06$ and stock market volatility is $\sigma = 0.16$.

Wage index In the model, g_D determines both the expected real growth rate of dividends and that of the average wage index. L_1 grew by 1.1% per year between 1947 and 2011, whereas the growth rate of D has been 1.8%. The growth rate of earnings being the most important parameter in the perspective of this paper, g_D is set at 1.2%. The key parameter of the model is κ, the speed of mean-reversion of y. In Table 1, estimates of κ range from 0.16 to 0.23. Following Benzoni et al. (2007), I make a conservative choice of $\kappa = 0.15$ in the baseline calibration. The standard deviation of permanent shocks to the wage-index is $v_1 = 0.025$ and the correlation of these shocks with the stock market is assumed to be zero ($v_3 = \sigma$).

Initial values To capture the diversity of macroeconomic perspectives that an agent may face when he enters the labor market, I initiate $y = 0$ at $t = -1000$ to generate random starting values at $t = 0$. I initialize $L_1(0)$ at $40,000, the level of the average wage around 1970 in 2013 dollars, and set $W(0) = L_1$.

[6] Historical estimates reported by Mehra (2007) using different data sets include 0.64% (Ibbotson, 1926-2004), 1.31% (Mehra-Prescott, 1889-2005), 2.68% (Shiller, 1871-2005) and 3.02% (Siegel, 1802-2004).

Table 2: Baseline calibration

Financial markets		
r	risk-free rate	0.02
g_D	dividends growth rate	0.012
μ	expected stock returns	0.08
σ	stock market volatility	0.16
Labor income		
v_1	SD of permanent labor market shocks	0.025
v_2	SD of permanent idiosyncratic shocks	0.15
v_3	see text	σ
κ	speed of mean-reversion of y	0.15
α_0	quadratic effect of experience	-0.0024
α_1		0.0581
a	see equation (14)	1.14
Preferences		
γ	relative risk aversion	5
ψ	discount rate	0.03
T	working years (from 20 to 65 years old)	45
R	retirement years	20
Initial conditions		
$L_1(0)$	average wage index (in thousands)	40
$L_2(0)$	wage in percentage of L_1	0.60
$W(0)$	wealth (in thousands)	50
$y(0)$	abnormal log-diff between L_1 and D	see text

4.2 Agent

Labor income α_0 and α_1 are calibrated as in BCG, producing a hump-shaped life-cycle profile for L_2. I set $L_2(0) = 0.6$, such that, given α_0 and α_1, L_2 averages

1 over the life-cycle. The agent faces relatively high permanent income shocks ($v_2 = 0.15$ as in BCG).

Preferences Estimates of the psychological discount rate generally lie between 3% and 4%. I set $\varphi = 0.03$. There is little consensus on the degree of relative risk aversion. Mehra and Prescott argue that reasonable values of γ lie between 1 and 10. Following several recent papers in the portfolio choice literature (Gomes and Michaelides (2005), Benzoni et al. (2007), Chai et al. (2011), Lynch and Tan (2011)), I choose $\gamma = 5$ in the baseline scenario because this value allows the model to match the life-cycle profile of the equity share.

Retirement Another important parameter is the number of retirement years. Life-expectancy at age 65 in the United-States is currently estimated between 19 and 20. I thus set $R = 20$.

Financial constraints I assume that the agent cannot short-sell the stock market portfolio nor can he invest borrowed money in equity. Hence $0 \leq \pi \leq 1$.

4.3 Social Security

Computation of benefits Simulations show that the mean of the best 35 years of L_2 is on average 14% higher than the mean L_2 over the whole career. Thus, I set $a = 1.14$ such that I do not penalize retirees by computing their AIYE using their entire career instead of their best 35 years of earnings.

Payroll taxes In order to value Social Security, I also build a benchmark scenario where benefits do not exist and wages are higher. I note τ the percentage difference in wages between the two scenarios. The value τ depends on the expected payroll tax rate. In 2013, the SSA estimated that the present value of future deficits of the OASI program over the next 75 years represents 2.40 % of taxable payroll.[7] On the other hand, 17.3 % of the 2013 OAS benefits (or 2.3 payroll tax percentage points) were allocated to survivors' benefits that are not explicitly modeled in this paper. Overall, assuming that those two corrections cancel out, the current tax rate of 10.6 %[8] of the OAS program may represent the cost of old-age benefits for current workers.

4.4 Model validation

The key ingredients of the model generate predictions regarding the life-cycle pattern of the share of financial wealth that would optimally be invested in the stock market portfolio. In this section, I check that these predictions are consistent with data from the Survey of Consumer Finances.

Figure 1 plots the average equity share simulated in the baseline scenario and empirical life-cycle profiles estimated using the SCF. Overall, the simulated equity share is very low among young households, and then slowly increasing until age 55. A strong rise of the equity share is predicted when the agent is a few years away from retirement because shocks on the stock markets become less likely to affect the wage-index before he retires. Consequently, Social Security wealth rapidly becomes a bond-like asset after 60 years old, inducing a greater willingness to invest financial wealth in stocks.

[7]Source: *The 2013 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds* (p.66)

[8]Although the Social Security tax rate is 12.4%, 1.8% are dedicated to the disability program.

Figure 1: Equity share by age group (SCF data)

Note: This graph represents predicted equity shares by age groups for two OLS models. I use triennial data from the SCF between 1989 and 2013. Individuals that do not participate to the labor force or have non-positive financial wealth are excluded. For the first model, I run an OLS regression with age group and year dummies and plot predicted equity share by age for the year 2013. For the second model, I use cohort dummies and plot the predicted equity share for the cohort born between 1949 and 1951.

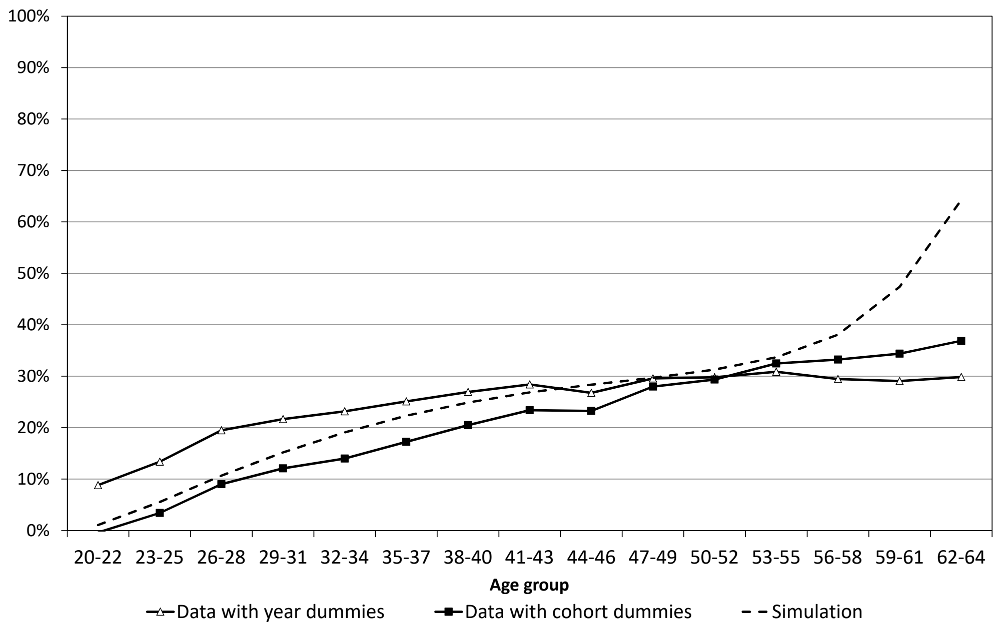

Overall, the model's prediction contrasts with many life-cycle models which, starting with Merton (1971), predict that the equity share should be very high among young households and then fall with age (Viceira (2001), Campbell et al. (2001), Cocco et al. (2005), Chai et al. (2011)). Finance professionals also advice a decreasing equity share. For example, Vanguard Target Retirement Funds have an equity share close to 90% for households under 40 years of age, and the share then drops to reach 50% at age 65.

In contrast, data from the SCF show a very low equity share among young households. As explained by Ameriks and Zeldes (2004), the exact shape of the life-cycle pattern is unclear. Indeed, empirical studies on the life-cycle profile of portfolio choices cannot include cohort and year dummies simultaneously because of perfect multicollinearity (age = year - cohort). As a consequence, the empirical conclusions may depend on the econometrician's choices. In Figure 1, I replicate Ameriks and Zeldes (2004)'s analysis using SCF surveys from 1989 to 2013.[9] I find the equity share to be increasing with age whether I include cohort or year dummies, though the life-cycle profile is flatter.

Overall, the model matches the life-cycle profile of the equity share pretty well for $\gamma = 5$. As illustrated in Appendix Figure A.1, a relative risk aversion of $\gamma = 5$ matches the data better than 4 or 6. However, the data do not support the predicted rise of the equity share in the few years preceding retirement. Portfolio choices after retirement are outside the scope of this paper but possible explanations for the lack of a rise in equity holdings include new sources background risk such as health expenditures (e.g. Rosen and Wu (2004), Goldman and Maestas (2013)), a bequest motive, a decline in cognitive abilities Christelis et al. (2010), or the lost ability to increase labor supply in response to financial losses (Bodie et al. (1992)). In the model, the rise of the average equity share is driven by households with high Social Security-to-financial wealth ratios. In the data, these households have low financial wealth and tend not to participate in the stock market which could be explained by fixed participation costs.

[9]Their own paper uses surveys from 1989 to 1998

5 Results

5.1 Valuation at the household level

Figure 2 plots the average Social Security certainty equivalent by age obtained from 50,000 simulations of the baseline model as well as the NPV of Social Security wealth computed by discounting expected cash flows (including taxes) at the risk-free rate. The latter is estimated for each individual at each point in time by simulating 1,000 paths.

Figure 2: Average Certainty Equivalent of Social Security

Note: This graph plots the average simulated certainty equivalent of Social Security in the baseline scenario as well as the expected value of future benefits net contributions discounted at a risk-free rate of 2%. Labor earnings at 20 are calibrated based on the cohort entering the labor market in 1970 and retiring in 2015.

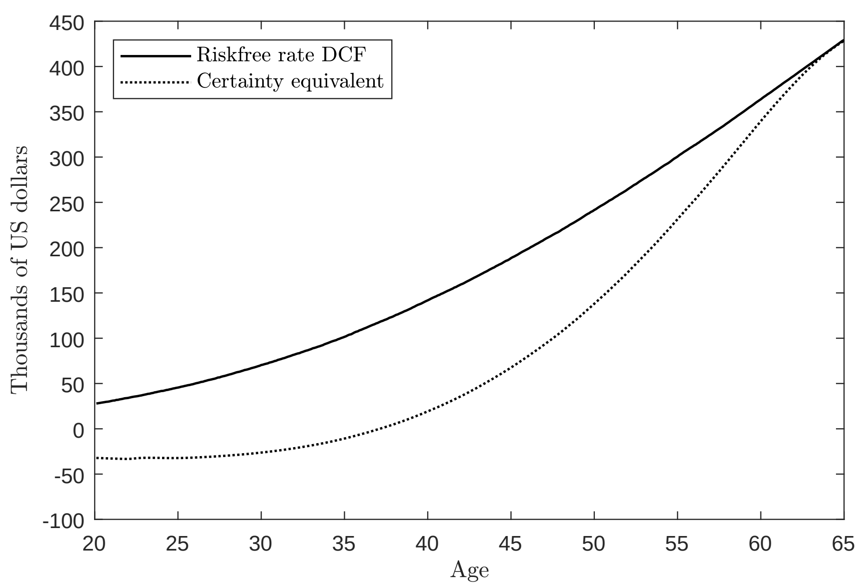

In the baseline scenario, Social Security has a negative certainty equivalent for workers below age 37 and the certainty equivalent is generally below the net present value of future Social Security cash flows (benefits minus contributions) when discounted at the risk free rate. At the time the certainty equivalent reaches zero, the unadjusted present value is already worth nearly 2.5 average earnings.

5.2 Risk-adjusted discount rates

One can also use the certainty equivalent to infer the discount rate that households should use to value Social Security. This can be done by solving for the discount rate for which the net present value equals the certainty equivalent. I conduct this exercise for each simulated worker at each point in time and report the average implied discount rate by age in Figure 3.

Figure 3: Discount rate implied by the certainty equivalent

Note: This graph plots the discount rate for which the present value of expected future benefits net of contributions is equal to the certainty equivalent of Social Security.

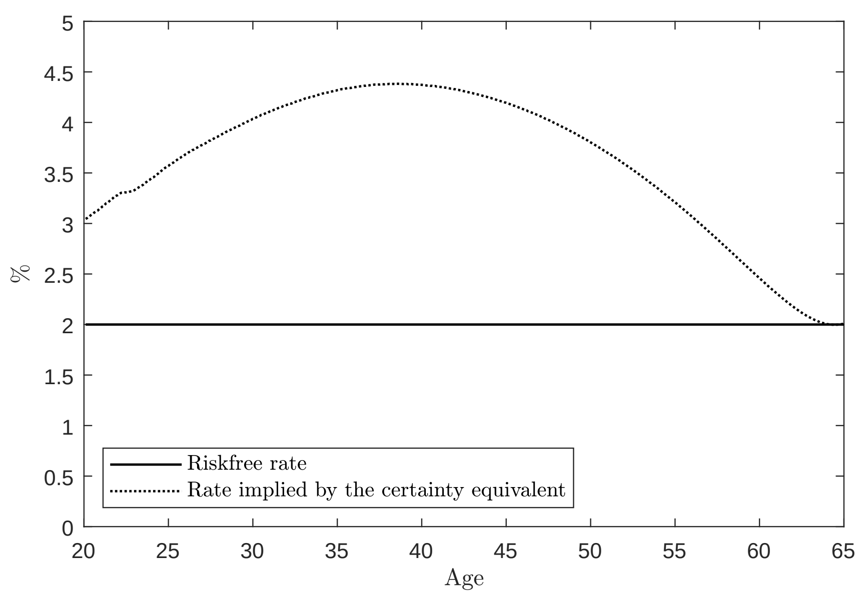

For new labor force entrants, the implied discount rate is about 3%. As detailed in the next section, Social Security does not expose new entrants to any aggregate risk and provides them with insurance against idiosyncratic income shocks. But workers still discount it at a high rate because they know that Social Security will force them to contribute to a system with unattractive returns in the future. At retirement age, and political risk aside, the implied discount rate converges with the risk-free rate because Social Security entitlements are safe at this point. The implied discount rate peaks in the middle of the life-cycle. At this point, workers have already accrued a sizable quantity of benefits which are exposed to aggregate risk through wage indexation. This risk is no longer offset by future contributions for two reasons. First, the amount of contributions still to be paid is smaller. Second, the average maturity of these contributions is lower, which means that they are less exposed to long-run macroeconomic risk.

5.3 Risk profile of Social Security

I now turn my attention to how Social Security value reacts to stock market returns and idiosyncratic shocks. To do so, I run the following regression at each point in time:

$$\frac{dSS_{it} - \tau L_{it} dt}{\overline{W}_{it}} = \beta_{s,t} \times \frac{dS_t}{S_t} + \beta_{l2,t} \times \frac{dL_{2,t}}{L_{2,t}} + u_t + \epsilon_{it} \tag{30}$$

where the left-hand side is the change in the certainty equivalent, net of contribution payments $\tau L dt$, relative to the agent's total wealth $\overline{W}$. The left-hand side can be interpreted as the unexpected change in Social Security value relative to the agent's expected lifetime consumption. On the right-hand side, $\frac{dS_t}{S_t}$ and $\frac{dL_{2,t}}{L_{2,t}}$ are stock market returns and the growth of the idiosyncratic component of the agent's wage. Figure 4 reports the evolution of β_s and β_{l2} over the life-cycle.

The negative value of β_{l2} shows that Social Security provides insurance against idiosyncratic income risk over the entire life-cycle, but much less at the end. Assuming no financial wealth and in the absence of Social Security, a 1% drop in L_2 at the beginning of life-cycle would translate into a 1% drop in expected lifetime consumption. In the presence of Social Security, that drop would be only 0.93%.

The negative value of β_s before 30 years old shows that Social Security also provides some insurance against aggregate shocks early on. For new entrants, a shock to L_1 reduces future contributions and future benefits by the same percentage. However, as illustrated by Figure 2, for young workers, future contributions have higher present value than future benefits. As a consequence, improvements in macroeconomic conditions make Social Security look worse.

Figure 4: Exposure to aggregate and idiosyncratic risks

Note: This graph plots the regression coefficients of equation (30) estimated in simulated data. β_s represents the effect of stock returns on unexpected changes in the certainty equivalent of Social Security, relative to the agent's total endowment and lifetime expected consumption. Hence, at 55, a +1% shock to stock prices translates, through Social Security, into a +.15% increase in expected lifetime consumption . β_{l_2} represents the same thing for idiosyncratic labor income shocks.

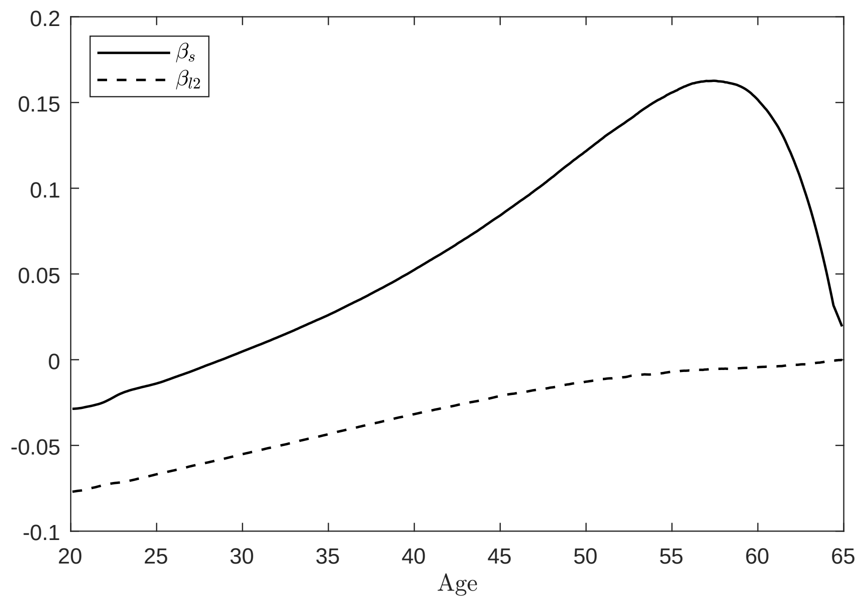

The exposure of the agent to aggregate risk through Social Security peaks in his 50's. At that point, the value of future benefits largely exceeds that of future contributions. Yet workers are still far enough from retirement to be exposed to wage indexation. As they get closer to retirement, wage indexation becomes less of a concern and the market beta of Social Security converges to zero.

5.4 Valuation at the national level

In this section, I estimate the magnitude of risk adjustment at a national scale. I proceed in three steps. First, I use simulated data to compute the average certainty equivalent-to-earnings ratio and the average risk-free valuation-to-earnings ratio at each age. This gives me valuation multiples over the life-cycle that I report in Panel A of Figure 5. In a second step, I compute the total labor earnings received in 2013 by each cohort born between 1949 and 1992 using the American Community Survey (ACS). I report the result in Panel B of Figure 5. Finally, I obtain aggregate Social Security values by age by multiplying the results of step 1 and 2. Panel C reports the aggregate adjusted and unadjusted valuations by age age.

Figure 5: Nationwide valuation of Social Security for current workers

Note: This figure illustrates the nationwide valuation of Social Security. Panel A reports the average certainty equivalent-to-earnings ratio and the average risk-free valuation-to-earnings ratio in simulated data. Panel B reports total labor earnings by age in the 2013 ACS. Panel C reports total Social Security valuations by age, defined as the product of Panel A and B.

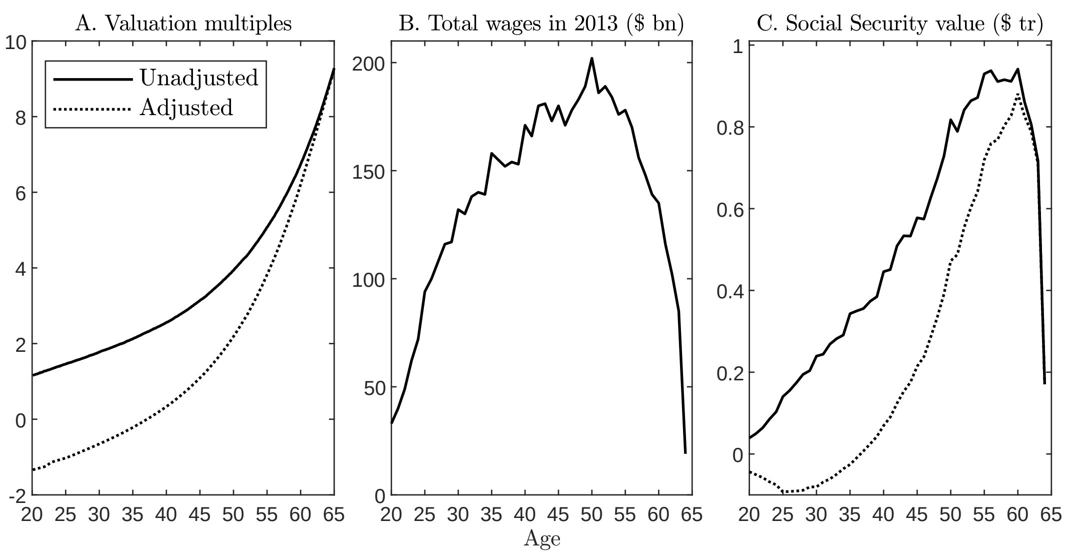

Because the model does not take into account survivors' benefits, I reintroduce their value on a *pro rata* basis, assuming that they represent 17.3% of the total present value. For retirees, I value entitlements as an ordinary immediate annuity ending at 85 years old and discounted at the risk-free rate. The coupon received by each cohort is the sum of Social Security benefits reported in the ACS.

Summing over all cohorts gives a valuation of closed-group obligations (*i.e.* restricted to cohorts that have already participated to the labor force). Nationwide valuations are reported in Table 3. In the baseline model, the unadjusted present value of Social Security entitlements was \$31.00 tr in 2013. This result cannot be compared to the SSA valuation of closed-group obligations since the SSA assumes a higher risk-free rate. When using the same rate (2.9%), I find a value of \$22.29 tr, closer to the SSA estimate of \$23.7 tr for 2013.

The baseline adjusted present value is \$19.6 tr, implying a risk adjustment of 36.8%. For current workers, the unadjusted value of entitlements is \$24.8 tr, but the adjusted one was \$13.4 tr. Hence, labor income risk reduces the value of Social Security by 46

6 Sensitivity analysis

In this section, I test different calibrations of the model and report the sensitivity of my results. Table 3 summarizes the analysis and presents all nationwide valuations discussed in this section.

Table 3: Nationwide adjusted and unadjusted values of Social Security ($tr)

Note: This table reports the sum of Social Security valuations on the national scale. Unadjusted values refers to the sum of future benefits net of future contributions discounted at the risk-free rate. Adjusted values refers to certainty equivalents. As detailled in Section 5.4, nationwide values are estimated by first computing multiples of labor earnings at each age in simulated data and then applying these multiples to the sum of labor earnings of each cohort in the 2013 ACS. For retirees, the value of Social Security is computed assuming that benefits reported in the ACS are safe and received until age 85.

		Workers			Workers and retirees		
	Retirees	Unadj.	Adjusted	Diff.	Unadj.	Adjusted	Diff.
Baseline							
	6.20	24.80	13.40	-46.0%	31.00	19.60	-36.8%
Equity premium							
$\mu - r = 0$	6.20	24.80	20.41	-21.7%	31.00	26.61	-14.2%
$\mu - r = .03$	6.20	24.80	17.61	-29.0%	31.00	23.81	-23.2%
Risk-free rate							
$r = .010$	6.66	37.41	18.48	-50.6%	44.01	25.14	-43.0%
$r = .029$	5.82	16.47	9.85	-40.2%	22.29	15.68	-29.7%
Cointegration							
$\kappa = .12$	6.20	24.80	14.00	-43.5%	31.00	20.20	-34.8%
$\kappa = .18$	6.20	24.80	12.82	-48.3%	31.00	19.02	-38.6%
Relative risk aversion							
$\gamma = 3$	6.20	24.80	11.70	-52.8%	31.00	17.90	-42.2%
$\gamma = 4$	6.20	24.80	12.66	-48.9%	31.00	18.86	-39.2%
$\gamma = 6$	6.20	24.80	13.44	-45.8%	31.00	19.64	-36.6%
Growth rate of wages							
$g_D = .006$	6.20	21.40	11.61	-45.8%	27.60	17.81	-35.5%
$g_D = .017$	6.20	28.16	14.64	-48.0%	34.36	20.84	-39.4%

6.1 Equity premium

How much of the risk adjustment comes from the equity premium? I answer this question by recomputing the certainty equivalent for different levels of expected stock returns μ. As reported in Figure 6, the certainty equivalent gets closer to the risk-free valuation when the equity premium gets closer to zero. For $\mu = r$, workers can insure a large fraction of their long-run human capital and social security risk for free. Consequently, as reported in Table 3, the risk adjustment drops from -46.0% to -20.41%. This drop shows that the equity premium explains most of the risk adjustment.

Figure 6: Social Security value for different levels of the equity premium

Note: The graph reports the average certainty equivalent of Social Security for different levels of equity premium as well as the mean net present value of future benefits net of contributions discounted at the risk-free rate. In all specifications, the risk-free rate is set to $r = .02$.

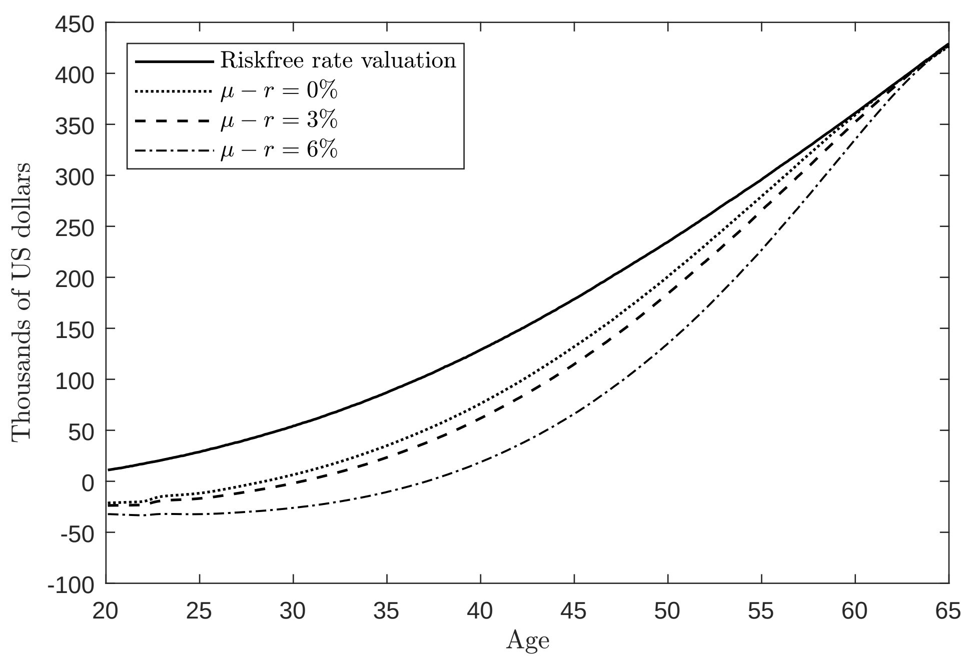

However, even when $\mu = r$, households keep discounting Social Security at a higher rate than the risk-free rate. There are at least two explanations First, their ability to use the equity market as a free insurance is limited by their inability to short-sell stocks. Second, cointegration does not create a perfect hedge against aggregate labor market risk. Therefore, households cannot fully insure and should apply a risk-premium to any asset whose returns are correlated with those of the main component of their portfolio: human capital.

6.2 Predictability in returns

One shortcoming of the model is that stock market volatility comes entirely from the dividends process. Moreover, because of cointegration, the variances of dividend and aggregate earnings are similar in the long run. Consequently, if stock market volatility is partially explained by time-varying discount rates, then the model would not only overestimate the long run volatility of aggregate dividends but also that of aggregate labor earnings.

In principle, this shortcoming should not matter for two reasons. First, because of cointegration, whatever causes variations in the stochastic discount factor is likely to affect the present value of future labor earnings and dividends in a similar fashion. This means that the baseline model should correctly reflect the covariance of stock market and Social Security returns. Second, if stock market volatility is partially explained by time-varying discount rates, then changes in stock prices should partially mean-revert. Hence, the baseline model would overestimate both the long run risk of Social Security and that of investing in stocks. However, a lower long run macroeconomic risk does not imply a higher Social Security certainty equivalent. Because the equity premium remains the same, the long run Sharpe ratio of the stock market portfolio is higher, which means that the cost of hedging against a given quantity of long run macroeconomic risk is higher. Because these two effects offset each other, the certainty equivalent of Social Security should not change much in the presence of time-varying discount rates.

To confirm this intuition, I introduce an alternative modeling choice in which the volatility in stock returns comes from shocks to dividends and but also to the discount factor. The latter tends to mean-revert which causes predictability in returns. Specifically, the stock market gain process introduced in equation (7) becomes:

$$\frac{dS}{S} = (\mu + \phi y(t))\,dt + \sigma dz_3 \tag{31}$$

where ϕ represents the predictability in returns based on the past performance of the stock market index relative to the wage index. The dynamic of y remains described by equation (8) and σ now represents the standard deviation of innovations to the stock market gain process. On the other hand, the dynamics of the wage index becomes:

$$\begin{aligned} \tfrac{DL_1}{L_1} = \quad & \left(-(\kappa - \phi)y(t) + g_D - \tfrac{\sigma^2}{2} + \tfrac{v_1^2}{2} + \tfrac{(\sigma - v_3)^2}{2}\right) dt \\ & + v_1 dz_1(t) + (\sigma - v_3)\, dz_3(t) \end{aligned} \tag{32}$$

Benzoni et al. (2007) estimate empirical versions of equation (31) using US data and find ϕ to be between 0.06 and 0.075. For my robustness check, I assume $\phi = 0.06$ and a risk aversion of $\gamma = 6$,[10] and solve the model again. I also adjust g_D such that $\mathbb{E}[L(T)]$ is the same as in the baseline calibration. Certainty equivalents are computed as in Section 3.3.2, except that equation (26) does not hold anymore, so $J^{ss}(\overline{W}, 0, L_2, y, 0, t)$ is computed numerically.

As reported in Figure 7, the evolution of the certainty equivalent of Social Security over the life-cycle is quantitatively close to the baseline model. The new model

[10] Stocks are more attractive when returns are predictable because stock market losses partially mean-revert and generate investment opportunities. As a results, many households hit the $\pi \leq 1$ constraint when $\gamma = 5$. I set $\gamma = 6$ to avoid that issue. Section 6.3 shows that, if anything, this change in risk aversion should increase the certainty equivalent of Social Security.

has much less aggregate risk. From the point of view of a new entrant, the standard deviation of log aggregate earnings $l_1(T)$ at his retirement date T falls from .98 to 0.59. For $\phi = 0.06$, a one standard deviation increase in $l_1(T)$ corresponds to a 1.3 percentage point increase in the annual growth rate of aggregate earnings over the 45 years of his career. This number seems reasonable: Saez and Zucman (2014) reports that for the bottom 90%, the real annual growth rate of income was 0% between 1917 and 1929, 2.3% between 1929 and 1986, and 0.7% between 1986 and 2012.

Figure 7: Value of Social Security in the presence of time-varying discount rates

Note: This graph reports the average certainty equivalent of Social Security when stock market volatility is partially cause by time variations in the discount factor.

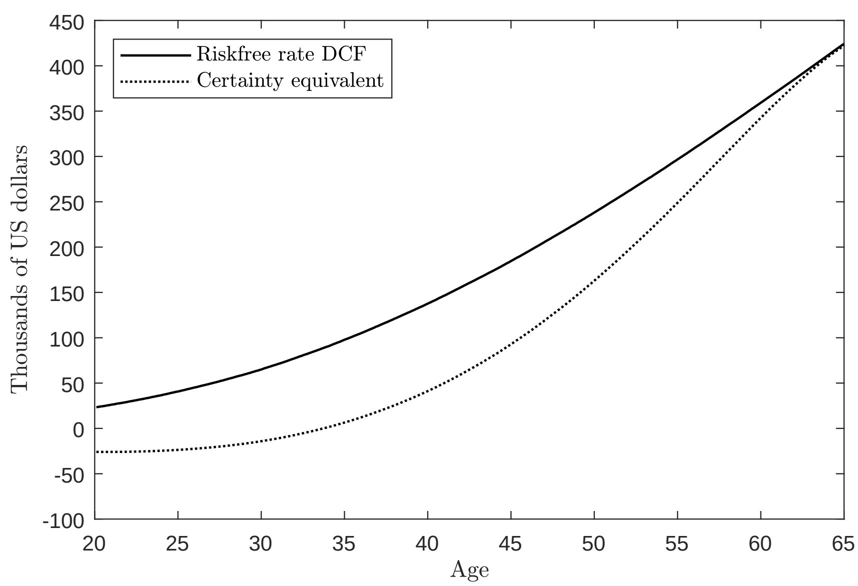

6.3 Risk aversion

Somewhat surprisingly, tests of different coefficients of relative risk aversion ranging from 3 to 6 suggest that smaller coefficients imply greater risk adjustments. For current workers, A coefficient of $\gamma = 3$ implies a risk adjustment of 52.8%, which is reduced to 45.8% when $\gamma = 6$. A natural explanation for this result is that Social Security forces the agent to invest some of his savings at a low rate of return when he would rather invest all of his wealth in stocks. Appendix Figure A.1 shows that for $\gamma = 3$, the vast majority of households want to invest all their wealth in the stock market.

Technically, this is not a risk adjustment, as even the expected utility of a risk-neutral agent would be reduced when forced to invest in an asset with lower expected returns than the stock market. However, this effect is unlikely to be an important driver of my baseline results. Indeed, in the baseline calibration, fewer than one worker out of six ever hits the upper constraint on the equity share ($\pi \leq 1$) during his lifetime. Moreover, relaxing the financial constraint and allowing some leverage ($\pi \leq 1.25$) does not change the size of the risk adjustment for $\gamma = 5$.

A different explanation of my results is that higher risk aversion increases the value of the insurance provided by Social Security against idiosyncratic income risk. While aggregate risk is priced primarily by the equity premium parameter, risk aversion should be the main determinant of the cost of idiosyncratic risk. Therefore, when risk aversion increases, the certainty equivalent of Social Security can increase because its insurance component is more valuable while the price of hedging aggregate risk remains the same.

Figure A.1 presents the life-cycle pattern of the equity share for different levels of risk aversion and suggests that a γ of 5 or 6 is the calibration that fits best the SCF data.

6.4 Cointegration

As expected, the risk adjustment of Social Security is a decreasing function of the speed of mean-reversion of the wage-to-dividends ratio. A conservative assumption such as $\kappa = 0.12$ has only a 2 percentage points effect on the result. From the point of view of a young worker facing a 45 years long career, whether stock market shocks are fully transmitted to the market in 5 or 8 years is not that important. In fact, the model shows no difference in the certainty equivalent of Social Security at age 20 when κ is reduced from 0.18 to 0.12.

When κ is set at 0.18, the nationwide risk adjustment for current workers equals -48.3%. This compares to an adjustment of -43.5% and -46.0% when κ equals respectively 0.15 and 0.12. Overall, the risk adjustment appears to increase by 0.8 percentage point when κ increases by 0.01. Table 1 reports a standard error of 0.55 for κ, which would translate into a standard error of 4.4 percentage points for the magnitude of the risk adjustment.

6.5 Risk-free rate

Figure 8 reports adjusted and unadjusted Social Security values for different levels of the risk-free rate, assuming the same equity premium. Certainty equivalent valuations at the beginning of the life-cycle are barely affected by the level of the risk-free rate. As a result, the size of the risk adjustment is greater when interest rates are low.

Figure 8: Value of Social Security for different risk-free rates

Note: This graph reports, for different levels of the risk-free rate, the average certainty equivalent (CE) of Social Security as well as the value of future benefits net of future contributions discounted at the risk-free rate (DCF). The equity premium is held constant accross specifications.

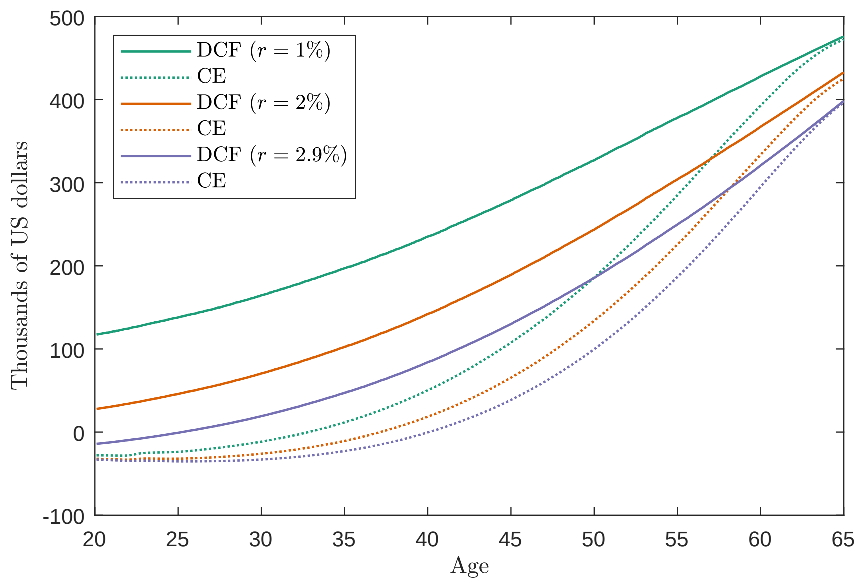

In its *2013 Annual Report*, the SSA assumes a long-term real interest rate of 2.9%, far above current levels. For $r = 0.029$, the certainty equivalent is only 40.2% below the risk-free valuation. However, if we assume that real interest rates will remain low, the required risk adjustment is significantly larger. For $r = 0.01$, the risk-free valuation jumps to \$37.41 tr, whereas the sum of private values only reaches \$18.48 tr, which implies a risk adjustment of -50.6%.

One surprising result is that the certainty equivalent of Social Security for new entrants does not vary much with the risk-free rate. Increasing the risk-free rate reduces the present value of benefits proportionally more than that of contributions. So, if benefits are more valuable than contributions, then the net present value of Social Security is clearly a decreasing function of the risk-free rate. This is what we observe for unadjusted valuations in Figure 8. If contributions are more valuable than benefits, then the relationship between the net present value and the risk-free rate can be reversed. The situation of new entrants seems to be between these two cases.

Interest rates are currently historically low which makes Social Security look extremely attractive to new workers when its cash flows are discounted at the risk-free rate. However, Figure 8 shows that low interest rates do not make Social Security more attractive to new workers when we compute the certainty equivalent.

6.6 Growth rate of wages

Intuitively, the value of Social Security entitlements should be an increasing function of the growth rate of the wage index, which determines the rate of returns on contributions. As a consequence, the unadjusted value is \$6.76 tr smaller when wages are expected to grow at $g_D = 0.006$ than when a higher trend of $g_D = 0.017$ is assumed. The percentage risk adjustment also increases with g_D. When the pension system offers better yearly returns, a larger share of total Social Security wealth goes to young households for which the risk adjustment is more important. Overall, the low and high growth scenarios imply risk adjustments of respectively -35.5% and -39.4%.

7 Conclusion

The valuation of Social Security entitlements is of prime importance in at least two ongoing debates that rank high on the policy agenda: the sustainability of public debts, and likely reforms in old age public benefit systems. The life-cycle model presented in this paper suggests that, from the point of view of households, these entitlements may be greatly overvalued when discounted at the risk-free rate. This means that the cost of transition to a funded system with the consent of current participants may be substantially lower than previously estimated.

A key take-away of this paper is that adjusting for risk is important to determine whether households benefit from Social Security. When discounted at the risk-free rate, future cash flows suggests that Social Security has value for new entrants, in particular in the current context of low interest rates. However, this conclusion is reversed when cash flows are valued using a certainty equivalent approach that takes into account labor income risk.

Because Social Security entitlements are one of households' most important assets, correctly computing its present value and understanding its risk-profile can also be important to understand portfolio choices. Indeed, though I conclude that entitlements are probably worth less than previously thought, I also find that they generate a large exposure to systematic risk, in particular in the second half of a worker's career.

References

Ameriks, John and Stephen P Zeldes, "How Do Household Portfolio Shares Vary with Age?," *Working Paper*, 2004.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Godlstein, "Portfolio Choice over the Life-Cycle when the Stock and Labor Markets Are Cointegrated," *Journal of Finance*, October 2007, *62* (5), 2123–2167.

Blocker, Alexander W., Laurence J. Kotlikoff, Stephen A. Ross, and Sergio Villar Vallenas, "The True Cost of Social Security," in "Tax Policy and the Economy, Volume 33" NBER Chapters, National Bureau of Economic Research, Inc, March 2018.

Bodie, Zvi, Robert Merton, and William F. Samuelson, "Labor supply flexibility and portfolio choice in a life cycle model," *Journal of Economic Dynamics and Control*, 1992, *16* (3-4), 427–449.

Campbell, John Y., João F. Cocco, Francisco J. Gomes, and Pascal J. Maenhout, "Investing Retirement Wealth: A Life-Cycle Model," in "Risk Aspects of Investment-Based Social Security Reform" NBER Chapters, National Bureau of Economic Research, Inc, March 2001, pp. 439–482.

Chai, Jingjing, Wolfram Horneff, Raimond Maurer, and Olivia S. Mitchell, "Optimal Portfolio Choice over the Life Cycle with Flexible Work, Endogenous Retirement, and Lifetime Payouts," *Review of Finance*, 2011, *15* (4), 875–907.

Christelis, Dimitris, Tullio Jappelli, and Mario Padula, "Cognitive abilities and portfolio choice," *European Economic Review*, January 2010, *54* (1), 18–38.

Cocco, Joao F., Francisco J. Gomes, and Pascal J. Maenhout, "Consumption and Portfolio Choice over the Life Cycle," *Review of Financial Studies*, 2005, *18* (2), 491–533.

Delavande, Adeline and Susann Rohwedder, "Individuals' uncertainty about future social security benefits and portfolio choice," *Journal of Applied Econometrics*, April 2011, *26* (3), 498–519.

Dominitz, Jeff and Charles F. Manski, "Measuring Pension-benefit Expectations Probabilistically," *LABOUR*, June 2006, *20* (2), 201–236.

Geanakoplos, John and Stephen P. Zeldes, "Market Valuation of Accrued Social Security Benefits," in "Measuring and Managing Federal Financial Risk" NBER

Chapters, National Bureau of Economic Research, Inc, March 2010, pp. 213–233.

_ , **Olivia S. Mitchell, and Stephen P. Zeldes**, “Would a Privatized Social Security System Really Pay a Higher Rate of Return,” May 2000, (6713).

Goldman, Dana and Nicole Maestas, “Medical Expenditure Risk And Household Portfolio Choice,” *Journal of Applied Econometrics*, June 2013, *28* (4), 527–550.

Gomes, Francisco and Alexander Michaelides, “Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence,” *Journal of Finance*, 04 2005, *60* (2), 869–904.

Huggett, Mark and Greg Kaplan, “How Large is the Stock Component of Human Capital?,” *Review of Economic Dynamics*, 2016, *22*, 21–51.

Luttmer, Erzo F. P. and Andrew A. Samwick, “The Welfare Cost of Perceived Policy Uncertainty: Evidence from Social Security,” *American Economic Review*, February 2018, *108* (2), 275–307.

Lynch, Anthony W. and Sinan Tan, “Labor income dynamics at business-cycle frequencies: Implications for portfolio choice,” *Journal of Financial Economics*, 2011, *101* (2), 333–359.

Mehra, Rajnish, “The Equity Premium Puzzle: A Review,” *Foundations and Trends(R) in Finance*, September 2007, *2* (1), 1–81.

Merton, Robert, “On the Role of Social Security as a Means for Efficient Risk Sharing in an Economy Where Human Capital Is Not Tradable,” in “Financial Aspects of the United States Pension System,” National Bureau of Economic Research, Inc, 1983, pp. 325–358.

Merton, Robert C., “Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case,” *The Review of Economics and Statistics*, August 1969, *51* (3), 247–57.

_ , “Optimum consumption and portfolio rules in a continuous-time model,” *Journal of Economic Theory*, December 1971, *3* (4), 373–413.

Rosen, H.S.Harvey S. and Stephen Wu, “Portfolio choice and health status,” *Journal of Financial Economics*, June 2004, *72* (3), 457–484.

Saez, Emmanuel and Gabriel Zucman, "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," October 2014, (20625).

_ **and** _ , "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," *The Quarterly Journal of Economics*, 2016, *131* (2), 519–578.

Samuelson, Paul A., "An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money," *Journal of Political Economy*, 1958, *66*, 467–467.

Shoven, John B. and Sita N. Slavov, "Political Risk Versus Market Risk in Social Security," April 2006, (12135).

Viceira, Luis M., "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income," *Journal of Finance*, 04 2001, *56* (2), 433–470.

Yaari, Menahem E., "Uncertain Lifetime, Life Insurance, and the Theory of the Consumer," *Review of Economic Studies*, 1965, *32* (2), 137–150.

Online Appendix

A HJB equation and first-order conditions

The financial wealth of the agent has the following dynamics.

$$dW(t) = \left[(r + \pi(t)(\mu - r))\,W(t) + L(t) - C(t)\right]dt + \pi(t)W(t)\sigma dz_3(t) \tag{33}$$

Given the objective function :

$$J(W(t), L_1(t), L_2(t), y(t), H(t), t) \equiv \max_{[C,\pi]} E_t\left[\int_t^T e^{-\psi u}\frac{(C(u))^{1-\gamma}}{1-\gamma}du + J(T)\right] \tag{18}$$

the Hamilton-Jacobi-Bellman equation is

$$-J_t = \max_{[C,\pi]} \theta(C, \pi) \tag{34}$$

where

$$\begin{aligned}
\theta(C,\pi) =& e^{-\psi t}\frac{C^{1-\gamma}}{1-\gamma} + J_W W\left(r + \pi(\mu - r) + \frac{L_1 L_2 - C}{W}\right)\\
&+ J_{L_1}L_1\left(-\kappa y + g_D - \frac{\sigma^2}{2} + \frac{v_1^2}{2} + \frac{(\sigma - v_3)^2}{2}\right)\\
&+ J_{L_2}L_2\alpha(t) + J_H\frac{aL_2}{45} - J_y\kappa y + \frac{1}{2}J_{WW}W^2\pi^2\sigma^2\\
&+ \frac{1}{2}J_{L_1L_1}L_1^2\left(v_1^2 + (\sigma - v_3)^2\right)\\
&+ \frac{1}{2}J_{yy}\left(v_1^2 + v_3^2\right) + \frac{1}{2}J_{L_2L_2}L_2^2v_2^2 + J_{WL_1}WL_1\pi\sigma(\sigma - v_3)\\
&- J_{Wy}W\pi\sigma v_3 + J_{L_1y}L_1\left(v_1^2 - v_3(\sigma - v_3)\right)
\end{aligned}$$

First-order conditions for the controls are

$$0 = e^{-\psi t}C^{-\gamma} - J_W \tag{35}$$

$$0 = J_W W(\mu - r) + J_{WW} W^2 \sigma^2 \pi + J_{WL_1} W L_1 (\sigma - v_3)\sigma - J_{Wy} W \sigma v_3 \tag{36}$$

The resolution of the problem can be simplified by eliminating one state variable. Indeed, since the agent has a CCRA utility function, dividing his labor income, his financial wealth and his pension benefits by a real number z, would also divide his optimal consumption by z and his equity share would remain unchanged. By choosing $z = W$, we can define c as the consumption to financial wealth ratio

$$\begin{aligned} \frac{1}{W} C(W, L_1, L_2, y, H, t) =& C\left(1, \frac{L_1}{W}, L_2, y, H, t\right) \\ \equiv& c\left(X = \frac{L_1}{W}, L_2, y, H, t\right) \end{aligned} \tag{37}$$

Then, we can use the first-order condition for consumption to rewrite the HJB equation in terms of c. To do so, I first differentiate the equation with respect to W using the Envelope Theorem, and then make the appropriate changes of variables.

$$\begin{aligned} 0 =& - c_t - \frac{c}{\gamma}\left[\psi - r - \pi(\mu - r)\right] - \left[r + \pi(\mu - r) + \sigma^2\pi^2 + XL_2 - c\right](c - Xc_X) \\ & - c_X X\left[-\kappa y + g_D - \frac{\sigma^2}{2} + \frac{v_1^2}{2} + \frac{(\sigma - v_3)^2}{2} + \pi\sigma(\sigma - v_3)\right] - c_{L_2} L_2 \alpha(t) \\ & - c_H \frac{aL_2}{45} + c_y\left[\kappa y + \pi\sigma v_3\right] \\ & + \frac{1}{2}\left[(\gamma + 1)c^{-1}(c - Xc_X)^2 - X^2 c_{XX}\right]\sigma^2\pi^2 \\ & + \frac{1}{2}\left[(\gamma + 1)c^{-1}c_X^2 - c_{XX}\right]X^2\left(v_1^2 + (\sigma - v_3)^2\right) \\ & + \frac{1}{2}\left[(\gamma + 1)c^{-1}c_y^2 - c_{yy}\right](v_1^2 + v_3^2) \\ & + \frac{1}{2}\left[(\gamma + 1)c^{-1}c_{L_2}^2 - c_{L_2L_2}\right]L_2^2 v_2^2 \\ & + \left[(\gamma + 1)c^{-1}(c - Xc_X)c_X + c_{XX}X\right]X\pi\sigma(\sigma - v_3) \\ & - \left[(\gamma + 1)c^{-1}(c - Xc_X)c_y - (c_y - c_{Xy}X)\right]\pi\sigma v_3 \\ & + \left[(\gamma + 1)c^{-1}c_X c_y - c_{Xy}\right]X\left(v_1^2 - v_3(\sigma - v_3)\right) \end{aligned} \tag{38}$$

A numerical advantage of this method is that the first-order condition for the equity share is no longer a function of second and mixed derivatives.

$$\pi = \frac{\mu - r}{\gamma\sigma^2} + \left(\frac{\mu - r}{\gamma\sigma^2} + \frac{v_3}{\sigma} - 1\right)\frac{Xc_X}{c - Xc_X} + \frac{v_3}{\sigma}\frac{c_y}{(c - Xc_X)} \tag{39}$$

Finally, the value of J can be derived from c. Let's denote $V(X, L_2, y, H, t) = J(1, X = \frac{L_1}{W}, L_2, y, H, t)$, then replace C and J in equation (35) yields the following differential equation

$$e^{-\varphi t}c^{-\gamma} = (1-\gamma)V - V_X X \tag{40}$$

which can be solved numerically as a system of linear of linear equations. J can then be inferred from V using the properties of the CRRA utility function, that is:

$$J(W, L_1, L_2, y, H, t) = W(t)^{1-\gamma}V(X, L_2, y, H, t) \tag{41}$$

.

B Numerical methodology

For the numerical resolution, I use the logs of X and L_2 as state-variables. This is achieved through a change of variables in the HJB equation and enables to build a grid with high resolution for low levels of X and L_2 while allowing large values. Lower resolution for high values of X also solves important numerical stability issues.

While large values may be reached given the volatility of the different processes, the ratio of wage to financial wealth (X) rapidly converges to small values, while the average L_2 over the life-cycle is 1 by construction.

The equation is solved backward using the alternating direction implicit method for the linear parts of the equation, and explicit finite differences for its nonlinear terms. Since explicit finite differences cause stability issues, an upwind scheme is implemented. I use a time-step of $\Delta t = 0.001$.

The grid is initialized with the parameters presented in Table A.1. I use the fact that large values of L_2 and H cannot be reached at young ages to reduce the size of the grid which improves stability.

Table A.1: Grid parameters at $T = 45$

Variable	max	min	Δ
x	$\ln(0.025)$	$\ln(4)$	0.112
y	-1	1	0.05
l_2	$\ln(0.1)$	$\ln(6)$	0.075
H	0	2.65	0.075

Two optimal control problems are solved: one with Social Security, the other without. The choice of $\Delta x = 0.112$ matches the 10.6 % payroll tax rate and thus enables a convenient computation of the Social Security certainty equivalent. Indeed, as described in Section 3.3.2, this computation requires to compare the value functions of the two problems but assuming a difference in labor income corresponding to payroll taxation.

C Additional figures

Figure A.1: Equity share by age for different relative risk aversions