Retirement Savings Adequacy in U.S. Defined Contribution Plans*

Francisco Gomes† Kenton Hoyem‡ Wei-Yin Hu§

Enrichetta Ravina¶

February 2020

Abstract

We evaluate retirement savings adequacy using a large panel of U.S. workers with a 401(k) account. We model medical expenditures, longevity, investment risk, and the likelihood of withdrawals due to hardship, job separation, and reaching age 59 1/2. Based on their current account balances, income, saving, and investment behavior, three in four workers in our sample are not saving enough for retirement. The dispersion is related to plan features, account balances, but also worker saving behavior. A bequest motive, lower housing equity, higher risk aversion, lower discount rates or lower future returns worsen the shortfall. We examine various counterfactual policy interventions.

*We thank John Y. Campbell, Joao Cocco, Markku Kaustia, Bob McDonald, Olivia S. Mitchell, Kim Peijnenburg, Giorgio Primiceri, Kathrin Schlafmann, seminar participants at Northwestern University, the 5th Cherry Blossom Financial Literacy Institute, the 4th European Workshop on Household Finance, and the 9th Helsinki Finance Summit on Investor Behavior for their valuable comments and suggestions, and Jingxiong Hu for very helpful research assistance. The views expressed in this paper are not necessarily those of Edelman Financial Engines.

†London Business School, fgomes@london.edu

‡khoyem@gmail.com

§Edelman Financial Engines, whu@edelmanfinancialengines.com

¶Kellogg School of Management, enrichetta.ravina@kellogg.northwestern.edu

1 Introduction

Defined contribution (DC) schemes are gradually replacing traditional defined benefit (DB) pensions in several countries and, given the structural funding problems associated with the latter, this phenomenon is likely to accelerate in the years to come. DC plans give workers the freedom to choose the level and allocation of their retirement savings. However, in a world where many individuals have limited financial literacy (e.g. Lusardi and Mitchell (2014), and Clark, Lusardi and Mitchell (2015)), and several might suffer from time-inconsistent preferences (e.g. Laibson (1997), and Harris and Laibson (2001)) or other behavioral biases, these choices might leave them financially vulnerable at retirement.[1]

The National Retirement Risk Index computed by the Center for Retirement Research at Boston College (Munnell, Webb, and Delorme (2006), and Munnell, Hou, Sanzenbacher (2018)) suggests that a large fraction of the U.S. population is not saving enough for retirement. Yet, several other studies conclude that the vast majority of U.S. workers is actually saving adequately (e.g. Engen, Gale, and Uccello (2005), Scholz, Seshadri, and Khitatrakun (2006) and Hurd and Rohwedder (2012)).

In this paper we revisit this question using data on more than 350,000 U.S. workers enrolled in defined contribution plans and evaluate whether, given their actual savings and investment decisions, they are likely to accumulate enough wealth to maintain their standard of living during retirement. Our data include worker's age, current account balance, contribution rate, salary, portfolio allocations, and tenure at the company. Unlike the majority of the previous literature, we do not assume workers save optimally over the remaining of their working life. Rather, we estimate the evolution of their portfolio shares and contribution rates as functions of their past behavior and observable characteristics, and use these evolution equations as inputs in our simulations. This approach has two main advantages.

[1]Choi, Laibson and Madrian (2011) document the prevalence of highly suboptimal contribution rates in 401(k) plans. Ahmed, Barber and Odean (2016) show that suboptimal asset allocation decisions by individuals are likely to generate lower and more volatile retirement wealth compared to private accounts without choice and to currently promised Social Security benefits. Using HRS data, Lusardi and Mitchell (2018) find that fewer than 1/3 of individuals over the age of 50 have ever tried to devise a retirement plan, and that this lack of planning is associated with low financial literacy. Fisch and Lusardi (20 20) find that many individuals in 401(k) plans only invest through the workplace and have higher degrees of financial illiteracy.

First, to evaluate whether workers are saving optimally for retirement, we need a framework that allows for the possibility they are making mistakes and capturing individual-specific features such as the degree of time-inconsistency or inertia would make the modeling exercise very challenging, especially for features for which there is no modeling consensus (e.g. limited financial literacy). Second, our approach allows us to match individual behavior as accurately as possible without estimating and calibrating a separate version of the model for each worker in the sample.

Our simulations of age-65 retirement wealth incorporate income shocks, the possibility of early withdrawals due to job separations, hardship, and reaching age 59 1/2, employer contributions and plan features, progressive taxation, and IRS limits on the amount the worker and the employer can contribute. We measure the total resources available at retirement as the sum of simulated age-65 retirement wealth, Social Security income, non-DC financial wealth, and net housing equity.[2] We compute Social Security income based on the simulated income profiles and the Social Security Administration formulas. We estimate non-retirement financial wealth and net housing wealth using fitted coefficients from regressions of these two sources of wealth on retirement wealth and individual characteristics in the Health and Retirement Study (HRS).

We check the predictive ability of our contribution and asset allocation evolution equations by estimating them on the first half of the sample and comparing the predicted and actual values on the last sample date.[3] We also use calculate unemployment rates, average probabilities of withdrawal and fraction of the aggregate assets being withdrawn following various life events in our sample, and find that they are comparable with the aggregate statistics from the Bureau of Labor Statistics and Vanguard's *How America Saves.* Finally, in Section 2.1 we show that our sample is representative of the workers in the Current Population Survey (CPS) in terms of age and tenure at their firm, but comprises workers with higher salary; and that it is similar to the Vanguard sample in terms of contribution rates

[2]Since neither reverse mortgages not downsizing are commonly observed in the data (Caplin (2002), Venti and Wise (2004), Poterba, Venti, and Wise (2011), Davidoff (2015)), we consider different scenarios regarding the fraction of housing equity available/used to finance consumption during retirement.

[3]We find that they are not statistically different from each other in 96.40% of the cases for the contribution rates, in 93.04% of the cases for the bond share, and in 95.75% of the cases for the equity share.

and asset allocations.

For each worker in our sample we run 10,000 simulations and compute the level of retirement consumption she can finance with her projected age-65 wealth and Social Security Income. Retirement consumption is estimated from a model of optimal consumption and savings decisions that incorporates medical expenditures, longevity, and investment risks, and takes Medicare and Medicaid into account. Unlike in the wealth accumulation phase of the worker's life where we don't take a stand on whether she is optimizing, at this stage we focus on *optimal* consumption because we want to determine the best a worker can do with the wealth she has accumulated.

We propose two measures of retirement adequacy. The first, which we label the consumption retirement replacement ratio (CRRR), is the ratio of retirement to adjusted current consumption for each simulation. The second, which we label the certainty equivalent ratio (CEQR), is the ratio of the certainty equivalent of future consumption across all simulations to adjusted current consumption, and thus takes the worker's risk aversion into account. We estimate current consumption from the Current Expenditure Survey (CEX) based on individual observable characteristics. We apply a 0.8 adjustment factor to current consumption because several expenditures such as housing, education, and children-related ones take place early in life, but are enjoyed throughout. Moreover, upon retirement, work-related expenses vanish, and individuals might turn to more home production.[4]

In our baseline specification, we find that close to three-quarters of the workers in our sample are not saving enough for retirement. The median individual has more than a 40% probability of having to decrease her consumption after age 65, even after taking the expenditure adjustment into account. Those at the 25th and 10th percentiles of the distribution face a 50% probability of having to cut their standard of living by almost 9% and 22%, respectively. Even those in the top 25th percentile of the distribution face approximately a 25% probability of having to scale down their adjusted consumption upon retiring. The under-saving problem is significantly worse if we consider a bequest motive, decrease the fraction of home equity available to finance retirement consumption, or make more conser-

[4]Academic studies and practitioners propose adjustment factors between 0.7 and 1. Our CRRR results can easily be re-scaled to apply a different adjustment factor.

vative assumptions about future average asset returns. In addition, when we back out the relative risk aversion and discount factor parameters that would make all workers appear adequately prepared, we find that they have in most cases the opposite relationship with the worker demographics and allocation decisions than the ones predicted by theory and the empirical literature.

We also explore how accumulated retirement wealth and preparedness relate to the initial characteristics of the workers in our sample. The results show the importance of the initial contribution rates in explaining the variation in outcomes. A one percentage point higher initial contribution rate is associated to a 2.67 ppt increase in the median CRRR, corresponding to a consumption level 2.67 ppt higher in all retirement years. The size of the account and the equity share have the expected positive and statistically significant coefficients, although their economic magnitudes is smaller: a $10,000 increase in initial account balance corresponds to a 69 basis points increase in CRRR, while a 10 ppt increase in the equity share corresponds, all else equal, to a 58 bps increase in CRRR. A more generous employer match is associated with higher retirement consumption: each 1 ppt higher employer match generates a 29.5 bps higher annual retirement consumption. All else equal, workers employed at companies that are older, privately held, invest more, and have higher net income, tend to have more wealth by the time they reach retirement.The same is true for those living in areas with higher financial literacy and a higher fraction of college educated residents. Our results also indicate that once we take company features into account, the dispersion of outcomes across income levels for workers of the same age increases, especially for the young. Further, we observe a striking difference between those age 35 and below, half of whom have median CRRRs of 1.25 or higher, and the older groups who have median CRRRs of around 1. This result is partially explained by the young being enrolled in plans with more generous employer contributions. In addition, the dispersion of outcomes rises quite noticeably with age, primarily due to an increase in the left tail of the CRRR distribution, a particularly concerning result since those close to retirement have fewer years left to benefit from possible changes in behavior or specific policy measures.

In the final section of the paper, we perform counter-factual experiments to assess the impact of different retirement preparedness policies in the context of our framework. An important caveat is thatin our setting the parameters governing worker behaviorare backward-looking and do not change in response to the policy. While this assumption might appear very restrictive, it is worth noting that the empirical evidence overwhelmingly shows that individuals often respond very passively to changes in the features of their 401(k) plan, both in terms of contribution rates and investment decisions (e.g. Agnew, Balduzzi, and Sunden (2003), Choi, Laibson and Madrian (2009), Choi, Laibson, Madrian, and Metrick (2003, 2004), Madrian and Shea (2001), Chetty et al. (2014)). Finally, we don't attempt to measure the potential crowding out of outside savings induced by these policies, and thus our estimates should be construed as an optmistic assessment of the effect of these policies.

We find that removing penalty-free early withdrawals for workers age 59 1/2 and older would increase retirement consumption by at least 5% for those at the bottom of the distribution. By contrast, setting a minimal contribution rate of either 2% to 5% has negligible effects, increasing the CRRRs by 1% or less. Only an age-dependent contribution rate that starts from a low level of 4.5% and increases gradually to 15.5% before retirement, when the individuals can presumably afford to save more, would generate sizable increases in retirement consumption for all age groups, and particularly for the workers at lower end of the distribution. However, this policy corresponds to a 10% average contribution rate, way above the 6.3% average and the 5.2% median in our sample, further evidence of the magnitude of the under-saving problem we find. Similarly, increasing every worker contribution rate by 5 ppts would generate substantial increases in retirement consumption, but more so for the workers who are already better off, leaving two thirds of the workers still falling short. Finally, introducing automatic rollover of the account balances upon a job switch, while effective, if implemented alone would only increase retirement consumption by 1 to 7 ppts, and more so for the workers who are already better off.

Related literature. Our paper is related to others who have also evaluated retirement preparedness. Poterba (2015) and Skinner (2007) provide comprehensive surveys of the literature and a discussion of the difficulties in evaluating retirement savings adequacy. The paper more closely related to ours are Engen, Gale, and Uccello (2005) and Scholz, Seshadri, and Khitatrakun (2006). Both solve an optimal life-cycle model of consumption and savings decisions and compare the wealth accumulation implied by the model with that of HRS respondents born between 1931 and 1941. Scholz, Seshadri, and Khitatrakun (2006) find that 84% of them are saving optimall y for retirement or close to it, while Engen, Gale, and Uccello (2005) estimate that fraction to be between 56% and 65%. In addition to our more comprehensive data, our study differs from theirs along several important dimensions. First, we do not ask whether individuals behave optimally, but rather model future behavior by projecting forward the estimates from the current data. Second, we study a very different period and more recent age cohorts. The papers above focus on the late 1990s, and estimate the private savings workers close to retirement need in order to supplement the Social Security benefits and DB pensions they are entitled to. They estimate that it is optimal that the workers at the bottom of the wealth distribution finance retirement almost exclusively with Social Security, and everyone else except the very top deciles mostly relies on Social Security and defined benefit pensions. By contrast, in the more recent period we study, the bulk of retirement savings is in the hands of the workers as opposed to the government and the companies they work for, and DB plans are quickly being phased out in favor of DC ones (Munnell and Chen, 2017). Such plans leave more latitiude to individual decisions and might not provide the same level of retirement security. Moreover, we study the age cohorts who are currently working, many of whom are decades away from retirement. Several studies point to the fact that such cohorts might be more financially vulnerable than older ones.[5] Finally, we exploit more granular information about the workers' savings rates, investment choices, plan menus, fees, employer contributions, and leakages, and are able to study the effect of these features on retirement wealth accumulation.

Hurd and Rohwedder (2012) take consumption levels of new retirees in the HRS and project them forward using evolution equations based on demographics and supplemented with shocks. They find that 71% of the respondents can finance the present value of this estimated consumption stream with their retirement wealth, as long as they can make full use of their housing equity. While this approach sidesteps the need of estimating working-life consumption, it has the drawback that the consumption level at the start of the estimation might already reflect a decrease in standards of living that the retirees have put in place upon

[5]For example, Lusardi, Mitchell, and Oggero (2017) document higher late-in-life financial vulnerability of recent cohorts because they have taken on more debt early in life.

reaching/getting closer to retirement. In addition, they focus only on retirees and study an earlier period when retirees might have been better prepared.

Another strand of the literature measures retirement adequacy by estimating income replacement ratios. Based on this approach and data from the Survey of Consumer Finances, Munnell, Webb, and Delorme (2006) conclude that Poterba, Venti, and Wise (2012) take a novel approach and study older retirees in the HRS. They find that almost half of them died with no financial assets, living fully out of their social security income only 43% of retirees are saving enough for retirement. By contrast, Purcell (2012) estimates the median income retirement replacement ratio in the HRS at around 62%, which he views as not too far below the recommended level. While easy to compute, income replacement ratios suffer from several drawbacks, the main one being that they focus on the current period, ignoring changes in income and expenses that might occur later in life. Finally, Poterba, Venti, and Wise (2012) take a novel approach and study older retirees in the HRS. They find that almost half of them died with no financial assets, living fully out of their social security income.

The remainder of the paper is organized as follows. In Section 2 we describe the data and summarize our approach. In Sections 3 and 4 we describe the pre-retirement period simulations, while in Section 5 we estimate optimal retirement consumption. In Section 6, we report our baseline results and sensitivity analysis, while in Section 7 we examine the role of the initial conditions in determining retirement adequacy of the workers in our sample. In Section 8 we report a series of counter-factual experiments, and in Section 9 we conclude.

2 Data and methodology overview

2.1 Data

Our primary data is a proprietary dataset provided by Edelman Financial Engines, the largest independent registered investment advisor in the U.S., which provides advice, and investment management to participants in 401(k) plans. The dataset includes information on worker 401(k) balances and contributions, salary, tenure at the firm, asset allocation over five aggregated asset classes and company stock, demographic characteristics, and zip code over the period between 2005 and 2010. It also includes information on the returns, balance sheets, and income statements of the firms the individuals work at, appended using information from CRSP, Capital IQ, and Compustat; detailed information on plan characteristics, investment options, and employer contributions, appended using information from DOL Forms 5500 that firms submit yearly to the Department of Labor; and fees for a more recent sub-period.

Our sample includes 1.6 million workers age 20 to 64 who have valid tenure data, earn at least the minimum wage, and make their own saving and allocation decisions rather than being enrolled in "managed accounts".[6] To avoid overfitting, we estimate the evolution equations for the contribution rates, asset allocations, unemployment, job separation probabilities, and all the other parameters used in the paper based on this sample, regardless of which subsample of workers we are focusing the analysis on. Panel A of Table 1 shows that, compared to the workers in the 2010 Current Population Survey (CPS), the average worker in our sample has similar age, 42.4 years, but has been working at her current employer for longer, 9.9 vs 7.7 years, and earns a higher salary, \$56,738 vs. \$45,437. She also lives in areas with similar house prices than the rest of the U.S. population (\$256,413 on average in our sample and \$288,172 in the 2010 U.S. Census).

We compare the contribution rates and equity allocations of the workers in our sample to those in Vanguard's *How America Saves* and find that they are similar. The average contribution rate is 6.9% in both datasets, while the medians are 5.6% and 6%, respectively. The average equity allocation is 67% in our sample and 68% in Vanguard's. By contrast, the average account balance is somewhat lower in our sample, \$56,592 vs. \$79,077, and so is the median, \$14,562 vs. \$26,926, consistent with the fact, highlighted by Munnell and Chen (2017), that the Vanguard sample tends to have a disproportionate number of large plans with higher earners (the average salary is \$65,000) and older workers (46 years old on average). Finally, we calculate plan-specific fees for each worker, based on their individual exposure to equity, bonds, and cash, either directly or through target date funds, and the fees charged by the mutual funds in their plan. We find that the average fees are 26.4 bps

[6]The sample excludes individuals who have "managed accounts", whereby Edelman Financial Engines manages the portfolio on behalf of the client and charges a fee on the assets under management. The sample does include workers employed at companies with automatic enrollment and who are auto-enrolled.

for bond funds, and 33.4 bps for the equity ones.

Bekaert, Hoyem, Hu and Ravina (2017) study a larger sample of approximately 3.8 million individuals and 296 firms. They show that the firms our workers are employed at are, on average, larger than those in Compustat, with a median number of employees of 4,600, compared to only 475 in the Compustat sample, have higher ROA, have an average age of 65 years, and are approximately half publicly listed and half privately held. They also display more variety, with very large public firms next to smaller private ones. The same is true for our estimation sample. Panel B of Table 1 shows that about 44% of the firms are public, and they employ 62% of the workers. The median number of employees is 6,948, and the average firm age is 74 years. Firms in both the full sample and our sample tend to be more similar to the S&P500 fir ms: they have similar leverage (23.5% vs. 23.93%) and investment intensity (4.70% vs. 4.17%), defined as capital expenditures over assets, but slightly lower profitability (3.43% vs. 5.84%).

While we use the sample above to estimate our parameters, in the current analysis we focus on the retirement preparedness of the subsample of workers employed at firms that offer only defined contributions plans, rather than both defined benefit and defined contribution ones. This sample includes a total of 350,859 workers. It is more straightforward to analyze and reflects the increasingly common situation of workers having access only to DC plans and being responsible for saving enough for retirement themselves. Indeed, based on data from the Survey of Consumer Finances (SCF), Munnell and Chen (2017) find that between 1983 and 2016 the proportion of workers who only have a 401(k) has risen from 12% to 73%, the proportion of workers who only have a defined benefit plan has decreased from 62% to 17%, and the proportion of workers with both has dropped from 26% to 10%.

Panel C of Table 1 contains the summary statistics for this subsample and shows that the average worker is very similar to the one in the CPS: she is 41.3 years old, and has worked at her firm for 7.9 years. She has however a higher salary ($56,464), like in the larger estimation sample. Compared to the estimation sample, she contributes slightly less to the plan (6.3% of salary vs. 6.9%), invests more conservatively, with a risky share of 62.44%, and pays similar fees. Finally, Panel D of Table 1 presents summary statistics for the firms at which these workers are employed. It shows that the DC-only firms are on average 59 years old, have 13,674 employees, have similar leverage and investment intensity as the estimation sample, but have slightly higher profitability, and are slightly more likely to be private (62.79%). These similarities are comforting, given the further constraints we have imposed on the sample.

2.2 Methodology overview

Our analysis consists of the three main steps which we outline below and describe in detail in Sections 3 to 5.

First, we use simulations to compute, for each worker, the joint expected distribution of total wealth at retirement, W_{i65}^{T}, and social security income, Y_{Si65}, assuming a retirement age of 65.[7]

We decompose total wealth at retirement into four components:

- Retirement wealth associated to the current job and all future jobs (W_{i65})
- Wealth in the retirement accounts associated to previous jobs (W_{i65}^{other})
- Wealth in non-retirement accounts available at retirement age (W_{i65}^{FW})
- Net housing wealth at retirement age (W_{i65}^{HW})[8]

For each worker, we obtain the distribution of W_{i65} by starting from her current account balance and simulating forward until age 65 based on the saving and investment behavior we observe in the data. We run 10,000 simulations for each worker in the sample. The evolution of employee contribution rates is estimated as a flexible function of the employee's own lagged contribution rate, age, salary, tenure at the firm, and their interactions, while the evolution of employer contributions is based on the plan-specific rules reported by each firm in the Form 5500 filed yearly with the Department of Labor (DOL) and on the worker's characteristics. Our simulations take into account the IRS limits on employee and employer contributions, and the fact that, upon turning 50, workers can elect to contribute more. The evolution of the equity, bond, and cash portfolio allocations are estimated based on the worker's own lagged allocation, age, salary, tenure, and their interactions. Consistent with Madrian

[7]Medicare and Medicaid are also taken into account and discussed below.

[8]We consider different scenarios regarding the fraction of net housing wealth accessed during retirement and provide additional details in Section 4.3.

and Shea (2001), Agnew, Balduzzi, and Sunden (2003), and the subsequent literature on 401(k) accounts, we find significant inertia in both employee contribution rates and asset allocations. Finally, based on the worker's age, salary, industry, and other characteristics, we also estimate the probability that in any period she becomes unemployed for some time, switches jobs, or faces economic hardship, and that, as a result, with some probability, she withdraws funds from her retirement account.

We follow a similar procedure to estimate the retirement wealth the worker has accumulated at previous jobs (W_{i65}^{other}). Assuming she started working at age 20 and using information on her tenure in the current job, we calculate how long she has worked for previous employers and compute W_{i65}^{other} by simulating backward based on the same evolution rules for contribution rates, portfolio allocations, and probabilities of withdrawing due to unemployment, job changes, or hardship that we use to compute W_{i65}. Finally, we obtain measures of wealth accumulation in non-retirement accounts (W_{i65}^{FW}) and of net housing wealth (W_{i65}^{HW}) from estimates based on the Health and Retirement Study and the median house price in the area where the worker lives. This approach allows us to assign a different value of W_{i65}^{FW} and W_{i65}^{HW} to each worker in each simulation, instead of assigning the average value to all the individuals in our sample and across all simulations.

In the second step of the analysis we compute the optimal level of retirement-age consumption that can be sustained by each combination of total wealth and social security income generated by our simulations. We use a consumption and savings model that incorporates investment, longevity, and out-of-pocket medical expenditure risk.[9]

Finally, in the third step, we evaluate each worker's degree of retirement preparedness by comparing the retirement consumption computed in the second step with an estimate of the consumption during her working life. We estimate working-life consumption for each worker based on the Consumer Expenditure Survey (CEX) and we account for the fact that upon retirement individuals decrease work-related and other types of expenses and also turn to

[9]While we don't take a stand on whether the workers are optimizing during the asset accumulation phase of their life, but rather we simulate the contribution rates and portfolio allocations based on the behavior observed in the sample, when we estimate the consumption stream that can be financed by the wealth accumulated at age 65, we compute the optimal consumption to determine the best the workers can do with that wealth.

home production for a larger fraction of their consumption needs.

3 Wealth accumulation in current and future DC accounts

We start the simulations at the oldest age for which we observe each worker in the sample, take her current account balance as the starting point, and simulate forward until retirement age.

3.1 Wealth evolution without leakages

To facilitate the exposition, we first describe the evolution of retirement wealth in the absence of pre-retirement withdrawals ("leakages").

3.1.1 Asset returns and fees

In our baseline analysis, retirement wealth can be invested in three asset classes: stocks, bonds, and cash, with gross returns R_t^S, R_t^B, and R^f, respectively.[10] The real return on cash (R_f) is assumed to be constant and calibrated to 0.5%, based on the historical mean real return of 30-day T-Bills from 1926 to 2016. The returns on bonds and stocks are assumed to be normally distributed and i.i.d. over time:

$$R_t^S \sim N(\mu^S, \sigma^S) \tag{1}$$

$$R_t^B \sim N(\mu^B, \sigma^B) \tag{2}$$

The equity return is set equal to the historical real return on the CRSP value-weighted index, with an annual standard deviation of 20%, and an equity premium of 6%. The bond portfolio is a combination of five different types of bonds, each matched to a specific index: the Barclays Capital Intermediate Government Bond Index, the Barclays Capital Long Term

[10] We aggregate the investments in small and mid-cap funds, large cap funds, international equity funds, and company stock into a general equity asset class; investments in bonds into a bond asset class; and short-term treasury bills and cash-like investments in a cash asset class.

Government Bond Index, the Salomon Brothers Non-US Government Bond Index, the Barclays Capital Corporate Bond Index, and the Barclays Capital Mortgage Backed Securities Index. Each index is weighted based on the average monthly holdings by American investors as reported in the ICI Fact Books. The historical return of such bond portfolio is 3.85. We set the standard deviation to 0.08, since the weighted average standard deviation is 8.5%, and the average correlation between the different indices is close to 1 (0.85).

Finally, our simulations also include the fees each worker pays on her stock and bond portfolios, τ_i^S and τ_i^B, respectively. Such fees are calculated based on the fees charged by the mutual funds available in worker's plan menu, and on her exposure to equity, bonds, and cash, either directly or through target date funds. We find that fees vary significantly across plans and investment vehicles, and across workers and firms (Panels A and C of Table 1).

3.1.2 Asset allocations and employee contributions

In the absence of pre-retirement withdrawals, retirement wealth W_{it} evolves based on the net-of-fee returns on previous account balances, and employee and employer contributions, according to the following equation:

$$W_{it} = [\alpha_{i,t-1}^S R_t^S (1-\tau_i^S) + \alpha_{i,t-1}^B R_t^B (1-\tau_i^B) + (1-\alpha_{i,t-1}^S - \alpha_{i,t-1}^B) R^f] W_{i,t-1} + k_{it} Y_{it} + K_{it}^e \quad (3)$$

where k_i denotes the employee contribution rate, K_{it}^e denotes the employer contribution, and α_{it}^S and α_{it}^B denote the shares of her portfolio invested in stocks and bonds, respectively.

In the first year of the simulations, each worker's contribution rate (k_{it}) and portfolio shares (α_{it}^S and α_{it}^B) are set at their most recent sample values, and in the following years they evolve according to the worker's characteristics and the evolution equations estimated from our panel. The employer contribution rates are calculated based on the rules reported by each firm in the DOL Form 5500, and are described in more detail in the next subsection.

Panel A of Table 2 reports the estimation of the worker's contribution rate as a function of its own lag, the worker's demographic characteristics, and their interactions. The specifications in Columns (1) and (3) include age and age squared, the worker's annual salary, and her tenure. The coefficients indicate that workers of higher age, salary, and tenure have on average higher contribution rates. The low R^2 of these regressions points to the fact that demographic characteristics explain a very small amount of the variation in contribution rates across workers, similar to the results in the consumption and investment literature (Browning and Lusardi, 1996; Giglio, Maggiori, Stroebel, and Utkus, 2019). To capture the effect of firm characteristics, plan quality, fees, employer contribution rules, and any time invariant firm characteristics that might affect workers' contribution rates, Columns (2) and (4) include firm fixed effects in the regressions. The coefficients are similar to the ones in Columns (1) and (3), and the R^2 only increases to about 10%. By contrast, Columns (5) to (8) show that including the lagged contribution rate in the regressions increases the R^2 dramatically, to about 70%, and that adding firm fixed effects to such regressions does not further increase their explanatory power (Columns (6) and (8)). We find that contribution rates are a concave and increasing function of age, and that they are very persistent, consistent with prior evidence that individuals don't change their saving behavior by much from one year to the next. Adding zip code fixed effects to better capture the socio-economic background of the worker and the general economic conditions of the area she lives in, winsorizing the dependent variable at the 1st and 99th percentiles, or restricting the sample to cases with positive contribution rates yield qualitatively and quantitatively similar results (available upon request). We choose the evolution equation in Column (5) as an input to our simulations, because of the high R^2 and the parsimonious specification. We thus model the evolution of the worker's contribution rate as

$$k_{it} = -0.0096 + 0.851 * k_{i,t-1} + 0.000692 * a_{it} - 0.00000628 * a_{it}^2 \tag{4}$$

where a_{it} represents the worker's age.[11] In addition, every year the IRS sets the maximum dollar amount that can be contributed to tax deferred retirement accounts, caps the catch-up contributions for workers older than 50, and sets limits on the base salary used to calculate the employer's contributions. Our simulations take these limits into account, and assume they increase with inflation.[12]

[11] In the simulations, we constrain the contribution rates to be non-negative and below 30%, which is higher than the 99th percentile in our sample.

[12] Details about the IRS limits on contributions can be found at https://www.irs.gov/retirement-

In Columns (9) and (10) we model worker contribution rates as a more complex function of their own lag, worker's salary, tenure at the firm, account value, a cubic polynomial in age, and their interactions. The objective is to check whether, as the worker ages, she plays catch up and contributes more if she realizes her account balance is low. Columns (9) and (10) show this is not the case: as the worker ages she contributes more than before if, all else equal, her account balance and past contribution rates are higher, not lower. In Section 6.2.4 we show that our findings are unchanged if use this more complex evolution equation as input for our simulations.

Panel B of Table 2 reports the estimation of the evolution of the portfolio allocations as a function of their own lag, the worker's characteristics, and their interactions. For brevity, we only report the specifications including lagged portfolio shares, as both the bond and equity share display significant persistence. Columns (1) and (3) report the coefficients on age and annual salary, their squared values, and their interactions, while Columns (2) and (4) report the results of running the same specifications with firm fixed effects. Column (1) shows that the most important determinant of the equity share is its own lag, with a coefficient on the lagged share of 0.92, and that the other coefficients are not statistically significant. Column (2) shows that when we add firm fixed effects more coefficients become significant, but that the explanatory power of the regressions does not increase. Column (3) shows that the most important determinant of the bond share is also its own lag, but that age and salary also matter. The coefficient on the lagged bond share is 0.91. All else equal, the bond share in the worker's portfolio increases with age and more so at higher levels of salary, and it decreases with salary, although this effect is economically small. Similar to the case of equity, adding firm fixed effects doesn't increase the explanatory power of the regressions. Thus, we choose the specifications in Columns (1) and (3) as the input for our simulations:

$$\alpha_{it}^{S} = 0.081880053 + 0.91617255\alpha_{i,t-1}^{S} \tag{5}$$

$$\alpha_{it}^{B} = 0.91410212\alpha_{i,t-1}^{B} + (3.467069E-4)a_{it} - (4.088E-8)Y_{it} + (1.199E-11)a_{it}^{2}Y_{it} \tag{6}$$

---

plans/cola-increases-for-dollar-limitations-on-benefits-and-contributions. We verify in the data that such limits do indeed closely track inflation.

where a_{it} represents the worker's age, and Y_{it} her salary.

Finally, we check the predictive ability of our evolution equations by repeating the estimation using only the first half of the sample, and using the estimated coefficients to predict, for each worker, the contribution rate and the asset allocations on the last date she appears in the sample. We find that the predicted and actual values at the last sample date are not statistically different from each other in 96.40% of the cases for the contribution rates, in 93.04% of the cases for the bond share, and in 95.75% of the cases for the equity share.

3.1.3 Employer contributions

We capture the employer contribution schemes with the following flexible specification

$$K_{it}^{e}(Y_{it}) \; = Min\{Min\{k_i^{e0} Y_{it}, \overline{K}_i^0\} + K_i^{match}, \overline{K}_i^{Tot}\} \tag{7}$$

where k_i^{e0} is the portion of the employer contribution independent of the employee's own contribution, expressed as a percentage of her current salary, and capped at $\overline{K}_i^0$, while K_i^{match} is the employer's matching contribution, described below. The total employer contribution is capped at $\overline{K}_i^{Tot}$.

The matching contribution rules usually have multiple tiers. For example, the company might match 100% of the employee's contribution up to 3% of her salary, and 50% of the contribution up to an additional 2% of her salary. Therefore, we specify a fairly general formulation below:

$$K_i^{match} = Y_{it} * \begin{cases} \kappa_i^{e1} * k_{it} & \text{if } k_{it} \leq \overline{k}_i^{e1} \\ \kappa_i^{e1} * \overline{k}_i^{e1} + \kappa_i^{e2} * (k_{it} - \overline{k}_i^{e1}) & \text{if } \overline{k}_i^{e1} < k_{it} \leq \overline{k}_i^{e2} + \overline{k}_i^{e1} \\ \kappa_i^{e1} * \overline{k}_i^{e1} + \kappa_i^{e2} * \overline{k}_i^{e2} + \kappa_i^{e3} * Min\{k_{it} - (\overline{k}_i^{e2} + \overline{k}_i^{e1}), \overline{k}_i^{e3}\} & \text{if } k_{it} > \overline{k}_i^{e2} + \overline{k}_i^{e1} \end{cases} \tag{8}$$

where $\mathrm{k_{it}}$ is the employee contribution rate, defined in Section 3.1.2. In the example above we would have

$$\kappa_i^{e1} = 100\%, \overline{k}_i^{e1} = 3\%, \kappa_i^{e2} = 50\%, \overline{k}_i^{e2} = 2\%, \kappa_i^{e3} = 0\%$$

The features of each pension plan are then represented in our simulations by the following vector of parameters:

$$\{k_i^{e0}, \overline{K}_i^0, \kappa_i^{e1}, \overline{k}_i^{e1}, \kappa_i^{e2}, \overline{k}_i^{e2}, \kappa_i^{e3}, \overline{k}_i^{e3}, \overline{K}_i^{Tot}\}$$

Panel C of Table 2 provides descriptive statistics for these parameters. For comparison, we report values for both the estimation and DC-only samples. The average non-matching employer contribution (k_i^{e0}) is 1.83% of salary for the full sample, and 1.57% for the DC-only sample. The median is 0% in both samples, as only one-third of firms provide non-matching contributions. The median (average) firm matching contribution is 60% (72.08%) of the employee's contributions up to a median (average) limit of 4% (3.82%) for the estimation sample. The median (average) firm matching contribution is 100% (62.77%) of the employee's contributions up to a median (average) limit of 4.5% (3.20%) for the DC-only sample. The averages show that overall the employer contributions in the DC-only sample tend to be less generous than those in the full sample, indicating that firms that are substituting generous defined benefit plans with defined contribution ones also tend to provide better terms in their 401(k) plans. However, the medians tell a different story, making it unclear which set of plans is overall better.

Finally, more than two-thirds of the firms don't make further contributions (i.e. $\kappa_i^{e2} = \kappa_i^{e3} = 0$), and only 1% of the workers have a third tier ($\kappa_i^{e3} > 0$). For this reason, we do not consider additional tiers in our simulations.

3.2 Wealth evolution with leakages

Workers can withdraw funds from their retirement accounts because of a job separation, either due to unemployment or a voluntary job switch, in case of hardship, and upon reaching age 59 1/2.[13] Munnell and Webb (2015) estimate that these "leakages" reduce "aggregate 401(k)/IRA wealth at retirement by about 25 percent."

In the next four subsections we describe how we model such withdrawals in our simulations, and in subsection 3.2.5 we provide details on the empirical estimation of the withdrawal probabilities and amounts for different types of worker.

[13] Hardship includes health care expenditures larger than 10% of the worker's income, permanent and total disability, outlays incurred to prevent foreclosure, costs related to purchasing the principal residence, excluding mortgage payments, and the cost of post-secondary education.

3.2.1 Withdrawal following a voluntary job switch

With probability $\pi^s(.)$ a worker voluntarily switches jobs, and, upon that happening, with probability $\pi_W^s(..)$, she withdraws a fraction $l^s(.)$ of her account balance. The corresponding retirement wealth accumulation equation is given by

$$W_{it} = \begin{cases} R_{it}^W W_{i,t-1} + k_{it} Y_{it} + \mathrm{K}_{it}^e & \text{with probability } 1 - \pi_W^s(.) \\ R_{it}^W W_{i,t-1} + k_{it} Y_{it} + K_{it}^e - l^s(.) W_{i,t-1} & \text{with probability } \pi_W^s(.) \end{cases} \quad (9)$$

where

$$R_{it}^W = \alpha_{i,t-1} R_t^S (1 - \tau_i^S) + \alpha_{i,t-1}^B R_t^B (1 - \tau_i^B) + (1 - \alpha_{i,t-1}^S - \alpha_{i,t-1}^B) R^f$$

and

$\pi_W^s(.)$: probability of withdrawal, conditional on a job switch

$l^s(.)$: fraction of the account balance withdrawn, conditional on a job switch

3.2.2 Withdrawal following a job loss

Workers face a probability $\pi^u(.)$ of becoming unemployed, which is reflected in the evolution of retirement wealth by

$$W_{it} = \begin{cases} R_{it}^W W_{i,t-1} + (1 - u(.))(k_{it} Y_i + K_{it}^e) & \text{with probability } 1 - \pi_W^u(.) \\ R_{it}^W W_{i,t-1} + (1 - u(.))(k_{it} Y_{it} + K_{it}^e) - l^u(.) W_{i,t-1} & \text{with probability } \pi_W^u(.) \end{cases} \quad (10)$$

where

$u(.)$: duration of the unemployment spell (in a fraction of a year)

$\pi_W^u(.)$: probability of withdrawal, conditional on unemployment

$l^u(.)$: fraction of the account balance withdrawn, conditional on unemployment

Notice that in this case the contributions only take place while the individual is employed, for a fraction $(1 - u(.))$ of the year.

Following a job loss, we consider three different scenarios regarding the worker's new pension plan. In the first, we assume that her new DC plan is identical to the current one. In the second and third, we assign her to a new DC plan randomly drawn from the plans of workers in the same income decile, or from all the plans in the sample. Appendix 1 shows that the results are almost identical across these scenarios.[14]

3.2.3 Withdrawal due to hardship

A hardship withdrawal occurs with probability π^h and, in such case, no contribution is made for that period. Wealth accumulation is given by

$$W_{it} = R^W_{it} W_{i,t-1} - l^h(.) W_{i,t-1} \tag{11}$$

where $l^h(.)$ is the fraction of the worker's account balance that is withdrawn.[15]

3.2.4 Withdrawal upon reaching age 59 1/2

Starting at age 59 and a half, workers can withdraw funds from their retirement account without penalty. In our simulations this occurs with probability π^{60}_W. In such an event, no contribution is made to the retirement account, and wealth accumulation is given by

$$W_{it} = R^W_{it} W_{i,t-1} - l^{60}(.) W_{i,t-1} \tag{12}$$

where $l^{60}(.)$ is the fraction of the account balance that is withdrawn.

[14] We are implicitly assuming that if a worker at some points joins a firm with a DB plan she is immediately vested and that the wealth accumulated through this plan is the same as what she accumulates in the DC plan.

[15] In this case, by definition, funds are withdrawn from the retirement account with certainty, $\pi^h_W = 1$.

3.2.5 Estimation of the leakages parameters

Job Separations. Unemployment rates and durations are estimated as a function of industry, age, salary, and education based on the Current Population Survey Merged Outgoing Rotation Groups (MORG), a monthly survey conducted by the Bureau of Labor Statistics to measure labor force participation and employment. In these estimations, we assume the workers enter the labor force at age 23, get no additional education once they start working, and, when changing jobs, stay in the same industry. The coefficients we obtain for the probability of unemployment (π^u) and its expected duration (u) are given by

$$\pi_i^u = constantprob_i - 0.0248848 * a_i + 0.00024636 * a_i^2 \tag{13}$$

$$u_i = constantdur_i + 1.0450717 * a_i - 0.0078187 * a_i^2 \tag{14}$$

where the constant term is worker-specific and includes industry fixed effects, and education effects estimated for each individual based on her industry, age, and salary. As a consistency check, we find that applying the estimates above to our own sample of workers yields an average unemployment rate of 4.73%, and a duration of 28 weeks, in line with the recent aggregate statistics.

We obtain industry-level separation ratios due to voluntary job changes from the Job Openings and Labor Turnover Survey (JOLTS), conducted monthly by the Bureau of Labor Statistics.

$$\pi_i^s = \frac{UnemploymentRate}{SeparationRatio} - UnemploymentRate \tag{15}$$

The fraction of plan participants cashing out upon a voluntary or involuntary job separation and the average fraction of funds withdrawn by such workers in each age bracket are estimated based on the equations above, data from Munnell and Webb (2015), tabulations from the 2008 and 2013 editions of Vanguard's *How America Saves*, and our own data. According to Vanguard data, 9% of workers changed or lost their jobs in 2008. Of these, 28.8% took money out of their retirement accounts, although the fraction of the account balance being withdrawn varied significantly by age and account size, with young and low-balance account holders being more likely to withdraw and withdrawing larger fractions (see Table

2 in Munnell and Webb, 2015). Vanguard estimates that 0.5% of all its 401(k) assets were withdrawn because of this reason.

Panel A of Table 3 reports our calculations. The first two columns reproduce Table 2 from Munnell and Webb (2015). For each age group, the first column reports the percent of plan participants cashing out upon a job separation, either voluntary or involuntary, while the second column reports the fraction of that cohort's aggregate account balance that is being cashed out. Columns (3) and (4) report the probability of losing one's job or switching voluntarily, based on equations (13) and (15) above. Columns (5) and (6) report the fraction of participants cashing out upon unemployment and a voluntary job switch, respectively. They are based on estimates from Engelhardt (2003), who finds that, in case of unemployment, the probability of cashing out is 46.7% higher than for voluntary job switches. For example, a worker in her 20s who leaves her firm has a probability of cashing out equal to 43.84%, if she lost her job (Column (5)), and equal to 30.03%, if she switched jobs voluntarily (Column (6)). These probabilities are calculated to make the overall probability of cashing out equal to 35%, as reported in Column (1).

The other columns in Panel A outline in detail our computations of the average fraction of the worker's account being withdrawn, by age and reason for the separation, based on the total amount of funds available to each age group in our data, the average account value, and the probabilities calculated above. Column (7) reports the total dollar balance available to each age cohort in our dataset, Column (8) the number of people in each age group, and Column (9) their average account value. Columns (10) and (11) calculate the number of workers in our dataset becoming unemployed or switching to another firm, based on the probabilities in Columns (3) and (4). Column (12) reports the aggregate account balance available to such workers, based on the average account value calculated in Column (9). Column (13) reports the total funds withdrawn by those who leave the firm, based on the aggregate balances in Column (12) and the withdrawal fraction estimated by Munnell and Webb (2015) and reported in Column (2). Based on this information, Columns (14) and (15) report the number of workers withdrawing funds because of unemployment and voluntary job switches, respectively. Column (16) report the average amount withdrawn by each worker, as a fraction of the total asset withdrawn (Colum (13)) and the number of people leaving the firm (the sum of Columns (14) and (15)). Finally, Column (17) reports the fraction of the worker's account balance that is being withdrawn, as the ratio of the average amount withdrawn reported in Column (16) and the average account value in that age cohort (Column (9)). The inputs for our simulations are the probability of cashing out due to unemployment (Column (3)) or a voluntary job switch (Column (4)), and the fraction of the worker's own account that she withdraws (Column (17)) given by

$$\pi_W^u = 0.4384 I_{a\in[20,29]} + 0.3997 I_{a\in[30,39]} + 0.3983 I_{a\in[40,49]} + 0.2982 I_{a\in[50,59]} + 02361 I_{a\in[60,69]} \tag{16}$$

$$\pi_W^s = 0.3003 I_{a\in[20,29]} + 0.2738 I_{a\in[30,39]} + 0.2728 I_{a\in[40,49]} + 0.2042 I_{a\in[50,59]} + 0.1617 I_{a\in[60,69]} \tag{17}$$

$$l^s = l^u = 0.4286 I_{a\in[20,29]} + 0.3438 I_{a\in[30,39]} + 0.3125 I_{a\in[40,49]} + 0.2917 I_{a\in[50,59]} + 0.2105 I_{a\in[60,69]} \tag{18}$$

As a consistency check, we use the estimates above to calculate the average fraction withdrawn across all cohorts, and find that it is 30.3%, compared to 28.8% reported by Munnell and Webb (2015).[16] We also calculate the percentage of the total assets being withdrawn and find that it is about 7%, both in our sample and in Munnell and Webb's.

Hardship withdrawals. Based on data from the Survey of Consumer Finances and from Vanguard's *How America Saves,* Munnell and Webb (2015) estimate that people taking hardship withdrawals are 1.2% of the total (see Fig. 3 in their paper) and withdraw 0.3% of total aggregate assets. Based on this information, the fact that in our sample workers age 59 or less are 90.17% of our sample, own 81.67% of the assets, and have an average account value of $55,462, we estimate that 1.33% of them will make a hardship withdrawal and they will withdraw 0.37% of the total assets belonging to their age group. This corresponds to withdrawing on average 27.6% of the funds in their individual accounts. As a consistency check, Vanguard estimates that, on average, people who make hardship withdrawals withdraw 0.28% of the funds available to those younger than 60, and 25% of the funds in their

[16] Taking a weighted average yields a withdrawal fraction equal to 24.31%, as reported at the bottom of Column (17) of Panel A of Table 3.

individual accounts.

Age 59 1/2 withdrawals. Munnell and Webb (2015) estimate that 0.93% of the workers make withdrawals upon reaching 59 1/2 years of age, and cash out 0.2% of the total assets in the system. Since in our sample workers age 60 and higher are 9.83% of the total, own 18.33% of the assets, and have an average account value of $114,081, we estimate that the probability of making a withdrawal conditional on reaching age 60 is 9.49%, the average fraction of the total assets withdrawn is 1.09%, and the fraction of the worker's balance being withdrawn is 11.49%. Both the probability of withdrawing and the fraction withdrawn are estimated as a (negative) function of the worker's salary decile. The estimates are reported in Panel B of Table 3.

It is worth noting that loans are excluded from the analysis, despite being an increasingly important phenomenon, because based on the estimates in Munnell and Webb (2015) they don't affect asset accumulation by much, since the great majority of them is repaid.

3.3 Labor income process

Following the life-cycle consumption and savings literature, we model labor income as

$$Ln(Y_{it}) = \begin{cases} f^Y(t, Z_{it}) + P_{it} + U_{it} & \text{with probability } 1 - \pi_i^u(.) \\ (1-\varphi)(f(t, Z_{it}) + P_{it} + U_{it}) & \text{with probability } \pi_i^u(.) \end{cases} \tag{19}$$

where $f^Y(t, Z_{it})$ is a deterministic polynomial of age and individual characteristics, φ is the fraction of income lost during an unemployment spell, $\pi_i^u(.)$ is the probability of unemployment, defined above, and

$$U_{it} \sim N(0, \sigma_U^2) \tag{20}$$

$$P_{it} = P_{it-1} + N_{it}, N_{it} \sim N(0, \sigma_N^2) \tag{21}$$

We set both σ_N and σ_U to 0.1, as it is standard in the literature (Gourinchas and Parker (2002), Carroll (1997), Cocco, Gomes and Maenhout (2005)). The probability of unemployment is worker-specific, and discussed in Section 3.2.5, while the fraction of yearly income lost during an unemployment spell is taken from Brown, Fang, and Gomes (2012). We initialize our simulations with the actual observed wage income for each worker, and use the stochastic process above to simulate the evolution of income going forward.

4 Other sources of wealth

Most individuals in our sample have worked in at least one previous job before joining their current employer. It is therefore important to estimate the retirement wealth accumulated during any previous employment spells (W_{i65}^{other}). In addition, they might also have accumulated financial wealth outside of their retirement accounts (W_{i65}^{FW}), and they might have access to home equity (W_{i65}^{HW}) they can use to finance consumption in retirement either by downsizing to a smaller house or by entering into a reverse mortgage. Combining all these sources, wealth accumulation at retirement is given by

$$W_{i65}^{T} = W_{i65} + W_{i65}^{other} + W_{i65}^{FW} + \theta W_{i65}^{HW} \tag{22}$$

where θ is the fraction of housing wealth spent during retirement.

As a preview of our results, we find that, except for scenarios with very high θ, total wealth at retirement is largely determined by the worker's retirement account balances ($W_{i65} + W_{i65}^{other}$), which constitute the primary output of our simulations. In our baseline case, where $\theta = 0.5$, the median value of $(W_{i65} + W_{i65}^{other})/W_{i65}^{T}$ is 0.66, i.e. two-thirds of the total wealth available to finance retirement is coming from the retirement accounts. Moreover, even at the 25th percentile of the wealth distribution, this ratio is still larger than 0.50.

4.1 Retirement wealth from previous employment spells

In order to estimate retirement wealth accumulation at previous employers, we first compute the total length of those employment spells by combining information on the worker's age at the start of our simulations (a_0) and her tenure on the current job ($\bar{t}$). In our baseline analysis, we assume that everyone starts working at age 20 (t_0). We then make the same assumptions as above regarding the evolution of wage income, asset returns and fees, the probability of unemployment, voluntary job switches, and leakages. We also assume that during any previous employment spell(s), each worker had access to a retirement plan identical to the current one, or to one drawn at random either from the full sample, or from the plans available to the workers in her income decile. Under these assumptions, the evolution of retirement wealth from earlier employment spells is given by

$$W_{it}^{other} = \begin{cases} R_{i,t-1}^{P} W_{i,t-1}^{other} + K_{it}^{e} + k_{it} Y_{it} & t \in [t_0, a_0 - \bar{t}] \\ R_{i,t-1}^{P} W_{i,t-1}^{other} & t \geq a_0 - \bar{t} \end{cases} \tag{23}$$

where

$$R_{i,t-1}^{P} \equiv \alpha_{i,t-1}^{S} R_t^S (1 - \tau_i^S) + \alpha_{i,t-1}^{B} R_t^B (1 - \tau_i^B) - (1 - \alpha_{i,t-1}^{S} - \alpha_{i,t-1}^{B}) R^f \tag{24}$$

The first branch of the formula captures the wealth evolution while the worker was enrolled in the plan, and includes both worker and employer annual contributions. Labor income is simulated backwards for these years based on the stochastic process in (19). The second branch captures the evolution of the account balances from the time the worker started her current job, at which point no additional contributions were made to these plans. The contribution parameters k_{it} and K_{it}^{e} are derived in Sections 3.1.2 and 3.1.3, respectively, while the returns, the fees, τ_i^S and τ_i^B, and the asset allocations, $\alpha_{i,t}^{S}$ and $\alpha_{i,t}^{B}$, are described in Sections 3.1.1 and 3.1.2.

4.2 Financial wealth outside retirement accounts

Throughout their lives, workers might also save outside their retirement accounts, to finance both pre-retirement and retirement consumption. To measure the wealth saved for retirement purposes outside of retirement accounts, we turn to the Health and Retirement Study, and estimate the relationship between retirement and non-retirement financial wealth at age 65. More precisely, we fit the following regression and use the resulting mapping from retirement to non-retirement wealth in our simulations:

$$W_{i65}^{FW} = \alpha^{FW} + \beta^{FW} * W_{i65} + \varepsilon_i^{FW} \tag{25}$$

Focusing on the sub-sample of married individuals between age 62 and 67 in the 2010 Wave of the HRS, we estimate the relationship between household net wealth excluding retirement and housing wealth, and the balance in the respondent's DC plans, controlling for her salary in the current job or the last job prior to retirement. This wealth measure is computed as the sum of the value of stocks, mutual funds, investment trusts, checking, savings and money market accounts, certificates of deposit, government savings bonds, T-bills, bonds and bond funds, and all other savings, less non-real estate debt. We restrict the estimation to married people, who constitute 67% of the respondents in the 62-67 age bracket, to avoid underestimating households' outside wealth by including single individuals in the regression. Nevertheless, we confirm that in our sample the net wealth of single individuals as defined above is about half of that of married ones.

As a consistency check, Panel A of Table 4 shows that the DC account balance of married respondents between 62 and 67 in the HRS, both overall, from their last job, or measured as total retirement wealth is very similar to the account balances for the workers between age 62 and 67 in our 401 (k) sample, both in terms of their mean and median, and of their overall distribution. The exception is the unconditional total retirement wealth, which is zero for the bottom quartile of the HRS sample, due to non-participation and possibly to issues with this survey question (see Gustman et al. (2014) for a detailed analysis of the issues with the pension wealth data files in the HRS). Because of the findings in Gustman et al. (2014), we estimate outside financial wealth as a function of total DC account balance, rather than total retirement wealth. Finally, Panel A of Table 4 also shows that the balance from the main account, or the account connected to last job, is very similar to the total balance, suggesting that most people close to retirement, or already retired, have consolidated their balances into one account.

The estimates from the regression of outside financial wealth on DC account balance are reported in Panel B of Table 4, both including and excluding salary. We pick specification (1) as an input for our simulations because of its simplicity, and note that the coefficients do not depend much on conditioning on the respondents having positive retirement wealth.

4.3 Housing wealth

During retirement, individuals can use part or all of their home equity by downsizing to a smaller house or entering into a reverse mortgage.

Since worker-level information about housing equity is not available in our sample, we compute it using a three-step procedure. First, we estimate the probability that each worker is a homeowner at age 65 (p_{i65}^{h}). Second, we estimate her house value at age 65 by projecting forward the Zillow median house value H_i in the zip code where she currently lives, using an expected housing price appreciation rate (r^H) of 1%, taken from Cocco (2005).[17] We then combine these values with an estimate of the loan-to-value ratio at age 65 from the Health and Retirement Study. Our estimate of housing wealth is thus given by:

$$W_{i65}^{HW} = p_{i65}^{h}(1 + r^H)^{65-a} H_i (1 - LTV_{i65})$$

The probability of being a homeowner is estimated with a probit model based on the subsample of married individuals between the ages of 62 and 67 in the HRS.[18] Panel A of Table 5 presents the summary statistics. The first four rows of the Panel show that most people in this age group don't own additional real estate beyond their first residence, and that the house values from the HRS and those from Zillow are quite similar over our sample period. The last two rows show that for this sample the average (median) loan to value ratio on the first residence is 27.4% (5.7%), indicating that most people have paid down their mortgages almost completely by the time they have reached retirement age. The values including other real estate are similar.

Panel B of Table 5 reports the regression coefficients. Columns (1) and (2) show the probability of being a homeowner as a function of retirement wealth and salary, overall and conditional on positive retirement wealth, while Columns (3) and (4) show the LTV on the first residence and overall as a function of the same variables. We base our simulations on the

[17] Yao and Zhang (2005) use a 0% appreciation rate. This lower value would decrease estimated housing wealth and exacerbate the retirement shortfall we find in the study.

[18] The data indicate that single individuals in this age bracket are more likely to be renters, and, when they are homeowners, they tend to own cheaper houses and have about half the housing wealth of married couples.

estimates in Columns (1) because of the higher R^2, and, for consistency, on the estimates in Column (3), noting that the LTV coefficients are quite similar across the two specifications.

4.4 Social Security income

During retirement, the workers' accumulated wealth is supplemented by Social Security income. For each simulation, we apply the Social Security formulas to the worker's lifetime income profile estimated based on the stochastic process given by (19). The amount of Social Security income a retiree receives is calculated based on a piece-wise linear function of the average indexed monthly earnings. The formula implies a replacement ratio of 90% up to a certain amount, of 32% for the portion of earnings between that amount and the next kink point, and of 15% for the rest.[19] We let the kink points grow over time at the inflation rate, after verifying the reasonableness of this approach in the data.

5 Consumption during retirement and retirement preparedness

For each simulation, we compute the consumption level that can be financed by the worker's wealth at age 65 and her Social Security benefits. Our approach is to calculate the *optimal* consumption the worker can sustain with the simulated resources she has at the time she retires. While actual consumption patterns during retirement might differ from the optimal path, they would make the worker worse off in terms of expected utility maximization and thus imply worse expected outcomes than the ones we analyze.

We compute optimal consumption during retirement as the solution to an intertemporal consumption and savings problem starting at age 66. Our model includes longevity risk, investment return risk, uncertain out-of-pocket medical expenditures, federal and local taxes, and realistic Social Security rules. By explicitly including out-of-pocket medical expenditures, the model also takes the financial support provided by the Medicaid and Medicare

[19] Details of the Social Security benefit formula can be found at https://www.ssa.gov/policy/docs/chartbooks/fast_facts/2010/fast_facts10.html. For 2010 the kink points were $761 and $4,586.

programs into account.

5.1 Preferences and budget constraint

Individual preferences are given by a time-separable power utility function:

$$U = E\sum_{t=65}^{100}\beta^{t-65}\left(\prod_{j=0}^{t-1}p_j\right)\frac{(C_{it})^{1-\gamma}}{1-\gamma} \tag{26}$$

where p_t denotes the probability that the individual survives to age $t+1$, conditional on being alive at age t. As we explain in more detail below, these probabilities are stochastic, allowing us to incorporate longevity risk in the model. To avoid carrying around an additional preference parameter, we address bequest motives when evaluating retirement adequacy, rather than including them directly in the optimization problem.[20]

We assume retirees have access to a risky asset, stocks, and a riskless asset, T-bills. Based on the extensive evidence that individuals rebalance their 401(k) portfolios only very infrequently (Section 3.1.2 of this paper, but also Agnew, Balduzzi, and Sunden (2003), Brunnermeier and Nagel (2008), and Bilias, Georgarakos, and Haliassos (2010) among others), we assume each individual keeps constant portfolio shares throughout retirement, and present results for different hypothetical allocations. By contrast, a model of optimal portfolio allocation would imply counterfactually high changes in individual allocations over the life cycle for most combinations of the preference parameters. We also note that given that we assume an exogenous portfolio rule and vary the weight in the risky asset to generate different values of the expected return and standard deviation of the retiree's portfolio, our two-asset specification is equivalent to one with a larger set of assets, as long as all risky assets have the same Sharpe ratio.[21] Further, the sensitivity analysis reported in Section 6.2.2 confirms that the main conclusions about retirement preparedness are unchanged under different assumptions about the retiree's asset allocation. Finally, we do not include annuity

[20] Love, Palumbo, and Smith (2008) and DeNardi, French, and Jones (2010), among others, nicely illustrate the role of longevity risk, mortality risk, medical expenditure risk, and bequest motives in explaining wealth evolution during retirement.

[21] Section 3.1.1 shows that the historical Sharpe ratio on stocks is 0.30, and the one on bonds is similar (0.28).

products in our model, as they are not widely used despite having been shown to potentially generate substantial welfare gains at retirement if individuals were to invest (more) in them (e.g. Horneff, Maurer, Mitchell and Stamos (2010)).

The dynamic budget constraint is:

$$W_{i,t+1}^{T} = R_{t+1}W_{it}^{T} - C_{it} - M_{it} + \overline{Y_S} \tag{27}$$

where M_{it} denotes the medical expenditure shocks, R_{t+1} is the return on the fixed retirement portfolio, and $\overline{Y}_S$ is social security income.

5.2 Longevity risk and medical expenditures

We model longevity risk following Lee and Carter (1992), who specify death rates for age t and calendar time x ($d_{t,x} = 1 - p_{t,x}$) as:

$$ln(d_{t,x}) = a_t + b_t \times \phi_x \tag{28}$$

The a_t coefficients capture the average shape of $ln(d_{t,x})$ over the life-cycle, while the b_t coefficients reflect how mortality rates at different ages respond to mortality shocks over time, ϕ_x.[22] The random variable ϕ_x is given by:

$$\phi_x = \mu^{\phi} + \phi_{x-1} + \varepsilon_x^{\phi} \tag{29}$$

where μ^{ϕ} is the drift parameter, and ε_x^{ϕ} is normally distributed with mean zero and standard deviation σ^{ϕ}. We take the values for a_t, b_t, μ^{ϕ} and σ^{ϕ} from Cocco and Gomes (2012).

We estimate the process for out-of-pocket medical expenditures using data from the HRS on the subsample of retirees older than 65. These expenditures reflect what the retirees have to pay beyond any Medicare and Medicaid benefits, thus incorporating the social insurance provided by these programs into our analysis. Since in the data medical expenditures are

[22]The b_t coefficients are a relative measure, normalized to sum to one.

highly correlated with income, we model them as a ratio of disposable income at age 65

$$\frac{M_{it}}{Y_{65}} = f^M(t) + V_{it} \tag{30}$$

where t denotes age and

$$V_{it} = \rho V_{it} + \varepsilon_{it}, V_{it} \sim N(0, \sigma_V^2) \tag{31}$$

We experiment with estimating equation (30) both in logs and in levels and find that the latter better fits the data.[23] We fit $f^M(t)$ as a third order polynomial and estimate $\rho = 0.4377$ and $\sigma = 0.2842$.[24]

5.3 Optimization problem

The full optimization problem is given by:

$$\underset{\{C_t\}_{65}^{100}}{Max}\, E \sum_{t=65}^{100} \beta^{t-65} \left(\prod_{j=0}^{t-1} p_j \right) \frac{(C_{it})^{1-\gamma}}{1-\gamma} \tag{32}$$

subject to the budget constraint in equation (27), to the stochastic process for the survival probabilities in equations (28) and (29), and the process for medical expenditures in equations (30) and (31). We scale wealth by $\overline{Y}_S$, and we are left with four state variables: scaled wealth ($W_{i,t}^T/\overline{Y}_S$), age (t), a persistent medical expenditure shock ($V_{i,t}$), and the current survival probability (p_t). The model is initialized at the wealth level available at the start of retirement, W_{i65}^T. Thus, the optimal consumption path for a given value of $W_{i65}^T/\overline{Y}_S$ is:

$$\{C_{i65}(W_{i65}^T/\overline{Y}_S), C_{i,66}(W_{i65}^T/\overline{Y}_S), \ldots\}$$

We solve the model for a grid of potential values of $W_{i65}^T/\overline{Y}_S$, and, using interpolation, we obtain the implied optimal consumption sequence for each of the 10,000*350,859 paths we

[23]Both in the model and in the simulations we restrict M_{it} to be positive, and we cap M_{it}/Y_{65} at 2.0. Since there are no observations in the HRS for which this ratio exceeded 200%, this constraint is strongly motivated by the data and does not affect the validity of the estimation.

[24]The coefficients of the age polynomial are 0.1463164, 0.0025844, −0.0001708, and 0.0000174, for the constant, linear, quadratic, and cubic term, respectively.

simulate.

5.4 Evaluating retirement adequacy

In the final step of our analysis, we evaluate to what extent, and for which workers, wealth accumulation at age 65 can sustain an optimal consumption level during retirement that is comparable, after making some adjustments, to the pre-retirement one. Below we discuss the methodology used to perform this evaluation.

5.4.1 Our approach

From a utility maximization perspective, the discounted marginal utility of consuming one more dollar today should be equal to the marginal utility of saving that dollar for the future, i.e. the optimal consumption path should satisfy

$$U'(C_{ia_0}) = \beta^{(65-a_0)} \left(\prod_{j=a_0}^{64} p_j \right) E[U'(C_{i65})(R^p_{i,a_0,65})^{(65-a_0)}] \tag{33}$$

where C_{ia_0} denotes current consumption, C_{i65} denotes consumption at retirement age, and $R^p_{i,a_0,65}$ is the return on a feasible investment portfolio from age a_0 till retirement. While this equality holds for all ages, both pre- and post-retirement, we focus on ages a_0 and 65 because we observe a_0 in our data, and because if the optimal consumption at the beginning of retirement, C_{i65}, is too low, the same is true for the optimal consumption in all future years.

Equation (33) implies that optimal consumption is a smooth, although not necessarily constant, function of age. Indeed, in most life-cycle models consumption is a mildly hump-shaped function of age. The hump-shaped pattern is concentrated very early in life, when the precautionary savings motive and liquidity constraints are most prominent, and, absent a strong bequest motive, very late in life, when the survival probabilities drop significantly (e.g. Gourinchas and Parker (2002), and Cocco, Gomes and Maenhout (2005)). For this reason, we assume that optimal consumption smoothing implies

$$U'(C_{ia_0}) = E[U'(C_{i65})] \tag{34}$$

i.e. the level of retirement savings is adequate if, in expectation, the worker will be able to maintain her current standard-of-living when she retires.[25]

It is important to emphasize that several expenditures, e.g. housing, education, and children-related expenses, are incurred early in life, but generate consumption over the entire life cycle. Moreover, during retirement, individuals partially substitute marketplace goods for home production, which has lower financial costs (Aguiar and Hurst, 2005 and 2013).[26] Therefore, we expect that expenditures during retirement will be lower than before even if the individual saved enough and she is actually enjoying the same standard of living. More explicitly, we can rewrite Equation (34) as

$$U'(C_{ia_0}) = U'(\varphi Expenditures_{ia_0}) = E[U'(C_{i65})] \tag{35}$$

where φ is an adjustment factor that based on both academic studies and practitioners' guidelines ranges from 0.7 to close to 1. Since lower values of φ are associated to lower thresholds for retirement preparedness, we start with a conservative baseline value of $\varphi = 0.8$.[27]

Finally, our approach focuses on direct comparisons of consumption rather than income levels. The limitations of evaluating retirement adequacy solely based on income measures are discussed in detail in Poterba (2015) and Biggs (2016), and acknowledged by most of the studies using them (e.g. Munnell, Webb and Delorme (2006), VanDerhei (2006), Brady (2010), and Pang and Schieber (2014), among others). The main one is that, unlike (35),

[25]Since at young ages consumption might actually be lower than in (34) because of precautionary savings and borrowing constraints, our approach understates young workers' retirement preparedness. It is worth noting however that our finding that 75% of the workers are likely to fall short does not depend on the younger cohorts. To the contrary, based on our simulations, younger workers appear to be the best prepared.

[26]Medical expenditures increase substantially after-retirement, and indeed they are directly incorporated in our model of optimal retirement consumption.

[27]One of the measures of retirement adequacy we report below, the consumption replacement ratio, can be easily recomputed for different values of φ by multiplying it by the ratio of 0.8 to the new adjustment factor. For example, results for $\varphi = 0.9$ can be computed by multiplying the consumption replacement ratios reported in Sections 6 and 8 by 0.8/0.9.

which is the result of an optimization problem over the entire retirement period, focusing on income replacement ratios at the date of retirement ignores changes in income and expenses that may occur later in life. Also, target income replacement ratios do not consider the heterogeneity in individual circumstances. Finally, the way both the pre- and post- retirement income is measured can sometimes substantially affect the fraction of households deemed adequately prepared for retirement.

5.4.2 Certainty equivalent ratio and consumption replacement ratios

Following the approach outlined above, we calculate two measures of retirement preparedness: the certainty equivalent ratio (CEQR), and the distribution of the consumption retirement replacement ratios (CRRRs).

The certainty equivalent ratio (CEQR) is obtained by computing the certainty equivalent of consumption from the right-hand-side of (35) after imposing a specific functional form (power utility in our case):

$$CEQR \equiv \frac{\overline{C}_{i65}}{\varphi Expenditures_{ia_0}} = \frac{\overline{C}_{i65}}{C_{ia_0}} \tag{36}$$

where $\overline{C}_{i65}$ is obtained from:

$$E[U'(C_{i65})] = U'(\overline{C}_{i65})$$

By contrast, for each worker, the consumption retirement replacement ratios (CRRRs) are the percentiles of the distribution of the ratio of post- to pre-retirement consumption across the 10,000 simulation paths:

$$CRRR(\omega) \equiv \frac{C_{i65}(\omega)}{\varphi Expenditures_{ia_0}} = \frac{C_{i65}(\omega)}{C_{ia_0}}, \, \omega = 1, \ldots, 10,000 \tag{37}$$

A risk-neutral worker will aim for an average $CRRR = C_{i65}/C_{ia_0}$ of 1. For such an individual, a value less than 1 represents a shortfall in retirement savings, with the actual distance between 1 and her CRRR measuring the shortfall percentage expressed in consumption units. A risk averse worker would instead aim for a CRRR greater than 1 and a CEQR of 1.

Retirement consumption, C_{i65}, is computed, for each simulation path, by inputting the total wealth at retirement, W_{i65}^T, and social security income, $\overline{Y}_S$, into the model of retirement consumption described in Section 5.3, after choosing the parameters of the retirement model $Z \equiv \{\alpha^R, \gamma, \beta\}$. Consumption during working years, C_{ia_0}, is estimated based on Consumer Expenditure Survey (CEX) data, following the procedure described in the next section.

5.4.3 Estimating current expenditures

We use the CEX to estimate expenditures during the individual's working life as a function of the workers' observable characteristics in our dataset, just as we used the HRS to estimate other forms of wealth. Total expenditure is defined as the sum of food, alcohol, and tobacco, apparel and services, entertainment, personal care, housing and shelter, health, reading and education, transportation, and miscellaneous. We use the methodology proposed by Deaton and Zaidi (2002) to calculate adult equivalents for each household and convert household-level expenditures into individual-level ones.

Table 6 reports the coefficients of the regression of total annual expenditures on age and salary, for respondents between age 20 and 65 who were interviewed by the CEX between 2006 and 2011, and whose ratio of total expenditures over salary falls in the interquartile range.[28] We explore various specifications in which age enters the regressions both linearly and as a set of dummy variables, and both as a stand-alone variable and interacted with salary and salary squared. The regression in Column (1) estimates total expenditure as 50.4% of pre-tax salary, plus a term dependent on the respondent's age. For example, a 42 year old worker earning the median salary, $50,000, would have annual total expenditures of $30,045, equal to 60% of her pre-tax salary. The more flexible specifications in Columns (2) to (5) yield similar results. We pick the specification in Column (4) as input for our analysis because of the high R^2 and the flexible specification.[29] Based on these coefficients, we calculate total expenditure in our 401(k) dataset as a function of worker age and salary.

[28] This corresponds to a ratio of expenditures over salary ranging between 0.4 and 1. We impose this restriction to avoid the effect of outliers and unusual circumstances.

[29] According to this specification, a 42 year old worker with $50,000 salary would spend 61% of her pre-tax salary.

5.4.4 Adjusting for taxes

Since expenditures in the CEX are financed by after-tax income, we convert the retirement consumption obtained from each simulation into its after-tax equivalent by taking both federal and local taxes into account.

We also account for differences in taxes during working life and retirement by taking tax brackets into account. In addition, for each individual, we calculate state-specific taxes based on the zip code where she currently lives, and, to economize on computations, her median simulated income.[30] For simplicity, we assume that after retirement both Social Security payments ($\overline{Y}_S$) and dis-saving from existing financial wealth ($R_{t+1}W_{it}^T - W_{i,t+1}^T$) are taxed at the same rate, and apply the appropriate tax rates directly to each simulated value of consumption:[31]

$$C_{it} = R_{t+1}W_{it}^T - W_{i,t+1}^T - M_{it} + \overline{Y}$$

6 Baseline results and sensitivity analysis

6.1 Baseline results

Table 7 shows the distribution of the consumption replacement ratios (CRRRs) the certainty equivalent ratio (CEQR), and the distribution of age-65 total simulated wealth for a baseline case where we assume power utility, risk aversion of 5, a discount factor of 0.95, and a 50% equity allocation at retirement, i.e.

$$Z \equiv \{\alpha^R, \gamma, \beta\} = \{50\%, 5, 0.95\}$$

We pick these parameter values because they are commonly used in the literature on portfolio choice over the life-cycle (see for example Brown (2001), and Horneff, Maurer,

[30] This approach assumes, for lack of a better alternative, that the individual remains in the same state until retirement. If he moves to a state with lower taxes upon retirement, our approach conservatively overestimates the degree of retirement preparedness.

[31] This approach will slightly overstate the individual's tax burden since it assumes $M_{it} = 0$. Note however that at age 65, the time of our analysis, the expected value of medical expenditures is very low.

and Mitchell (2018)). In Section 6.2.2 we illustrate how changing these parameters affects the results and also re-examine our findings in light of the well documented cross-sectional heterogeneity in risk aversion and discount factors (Harrison, Lau and Rutstrom (2007) and Paravisini, Rappoport, and Ravina (2016)).

Columns (1) through (6) of Table 7 present the consumption replacement ratio (CRRR) results.[32] The rows refer to workers ranking at different percentiles in the worker population, and display, for each individual, the CRRRs corresponding to different percentiles of the distribution across her 10,000 simulations. For example, the row labelled "$50th$ percentile" refers to the worker with the median CRRR in the population. Based on our simulations, she has a 50% probability of obtaining a CRRR of 1.11, i.e. of being well prepared and having retirement consumption equal 1.11 times her adjusted pre-retirement consumption, where the adjustment factor is $\varphi = 0.8$. However, with 40% probability her CRRR at retirement will be 1, i.e. her accumulated wealth will be just enough to maintain her standard of living. With 30% probability her CRRR will only be 0.90, and she will be falling short by 10%, i.e. each retirement year her consumption will be 10% lower than the one preserving her standard of living; and with 10% probability her CRRR will be as low as 0.69, a 31% shortfall. The prospects are worse for the worker in the second row, the 25th percentile of the distribution of CRRRs in the population. His probability of falling short is more than 50%: the CRRRs are only 0.91, 0.84, 0.76, and 0.60, respectively. Column (6) reports the mean CRRR for the same workers and shows that, if all workers were risk neutral, the worker at the 25th percentile and above would be saving enough for retirement. However, the results worsen if we take risk aversion into account.

Column (7) reports the certainty equivalent ratio ($CEQR$), which integrates over the full distribution of outcomes using the baseline utility function, and thus measures the shortfall taking risk aversion into account. It shows that for the baseline risk aversion coefficient of 5, the certainty equivalent ratio is only 0.76 for the 25th percentile worker and 0.86 for the median one, indicating that, in risk-adjusted terms, their wealth accumulation falls short by

[32] We report the results for the case in which, following unemployment, the worker joins a company with the same DC plan as her current firm. Appendix 1 shows the results are quantitatively indistinguishable when we assign her to a new plan drawn randomly from our sample, or from the set of plans of workers in her income decile.

24% and 14%, respectively. Indeed, Column (7) shows that more than 75% of the workers are not saving enough for retirement, with the CEQR at the 75th percentile falling just below 1 (0.99). The picture is much worse for the worker at the 10th percentile of the distribution, with a CEQR of 0.68, but even the worker at the $90th$ percentile, despite saving adequately in risk-adjusted terms (CEQR greater than 1), still faces close to 10% probability of not having accumulated enough wealth by age 65, with a CRRR of 0.92 at the 10th percentile of her CRRR distribution.

Finally, Columns (8) to (12) report the simulated age-65 total wealth in 2010 constant dollar terms corresponding to each of the CRRRs in the first five columns.

6.2 Sensitivity analysis

6.2.1 Bequest motive and housing wealth availability

The baseline results in the previous section assume no bequest motive, apart from the remaining housing equity.[33] If the workers in our sample wished to leave an additional bequest, then their retirement savings would be even less adequate. Panel A of Table 8 reports the results from introducing a target bequest equal to 10% of age-65 wealth, and assuming that only savings in excess of this amount can be used to finance consumption during retirement. By contrast, Panel B of Table 8 reports the results when we assume that all of housing equity is left unused.[34]

Panel A of Table 8 shows that, relative to the baseline results reported in Table 7, a bequest motive affects the CRRRs of a given individual more in the good scenarios than in the bad ones. The reason is that, regardless of her ranking in the population, when a worker gets a bad draw and ends up in the bottom 10th or 20th percentile of her CRRR distribution, a larger fraction of her consumption will be financed by Social Security, and thus setting 10% of her remaining wealth aside for a bequest has a smaller impact on her

[33] Since the baseline results only allow for 50% of home equity to be used to finance retirement consumption, the remaining 50% is available to bequeath.

[34] This could happen because of bequest motives, but also because of lack of access or information about reverse mortgages and other financial products that allow for home equity release. The costs of accessing housing equity can be both direct financial costs and indirect ones like obtaining and processing the necessary information.

consumption. Similarly, a bequest proportional to the total wealth accumulated at age 65 affects the CRRRs of wealthier workers disproportionately more, as they rely less on Social Security to finance consumption during retirement. Compared to the baseline specification, the reduction in CRRRs ranges from 3 percentage points for the worker at the 10th percentile of the distribution and facing a worst case scenario, to 12 percentage points for the 90th percentile worker getting the median draw from her CRRR distribution. Similarly, the CEQRs are 4 percentage points lower for those at the $10th$ percentile of the worker distribution, and 7 percentage points lower for those in the $90th$.

Similarly, Panel B of Table 8 shows that restricting access to housing equity hits the right tail of the distribution more strongly, since both wealthier workers, and, for any given worker, better draws from their own wealth distribution, tend to be associated with more housing wealth. This pattern reverses for workers in the 75th percentile and above for whom housing equity is a larger fraction of their total accumulated wealth when they get a bad draw compared to cases when they get a good one and have proportionally more financial wealth. Yet, unlike in the bequest analysis above, compared to the baseline results in Table 7, the drop in consumption replacement ratios and in CEQRs is substantial regardless of the worker's ranking, or the draw from her CRRR distribution. The workers at the 10th and 25th percentiles suffer drops in their CEQRs of 10 and 13 percentage points, respectively. This result is quite worrisome, since poorer workers are more likely to lack access to home equity release products or to be deterred by their high costs, and thus to end up in this scenario. By contrast, while individuals in the 90th percentile face a larger drop in their CEQR, 27 percentage points, and end up with a CEQR below 1, they are also substantially more likely to have access to housing equity and not end up in this situation.

6.2.2 Alternative preference parameters and asset allocations during retirement

As discussed in Section 5, retirement consumption is obtained from a structural model which computes the optimal consumption that can be sustained by the worker's simulated wealth and Social Security income. For this reason, the level of retirement consumption depends on the choice of risk aversion γ, subjective discount factor β, and retirement-period portfolio allocation (α^R).

Table 9 shows the median worker's CEQR and her median CRRR for different values of these parameters.[35] Panel A reports the CEQR for the baseline equity allocation of 50%, and risk aversion parameters of 2, 5, and 8, respectively. Panel B reports the median CRRR for the same combination of parameters. Panel C reports the CEQRs for the baseline risk aversion coefficient of 5, and equity allocations during retirement ranging from 0 to 100%. The rows refer to values of the discount factor, ranging from 0.925 to 0.975. Panel A and B show that as risk aversion falls both CEQRs and CRRRs increase, as individuals with lower risk aversion care less about medical expenditures, longevity risk, and investment returns risk, and thus need to save less for retirement. Based on our simulations, a risk aversion parameter of 2 generates enough retirement savings for the median worker regardless of her discount factor, and will likely allow her to leave a bequest or avoid accessing home equity. By contrast, higher risk aversion than our baseline value generates an even more severe retirement under-saving problem. Notice that since the distribution of age-65 wealth is independent of risk aversion, the variation in the median CRRR across values of risk aversion reportedl in Panel B reflects how strongly the worker cares about retirement period risks, i.e. medical expenditure, longevity, and investment risk during retirement. Panel B shows that such variation is substantial, highlighting that our finding that in the low risk aversion scenario most workers are saving enough for retirement relies heavily on them not caring much about post-retirement risks.

Panel C shows that the median CEQR doesn't vary much with the retirement portfolio equity share. Our baseline assumption of α^R=50% yields the highest CEQR, as a no equity exposure investment strategy forgoes the equity premium and results in low wealth accumulation, and a 100% equity exposure generates excessive risk-taking. For each discount factor, the different equity allocations we consider generate a difference in average yearly consumption of only a few percentage points, indirectly confirming that our approach of keeping retirees' equity shares constant in the optimization problem in Section 5 is reasonable given the limited effect of retirement-period equity shares on replacement ratios.

Finally, an increase in the discount rate corresponds to higher CRRRs and CEQRs,

[35] The full set of results, including other CRRR percentiles, is available upon request.

although the effect is smaller than that of changing risk aversion. Indeed, even with a discount factor of 0.925, the median CEQR is still below 1, indicating that more than half of the workers is not saving enough for retirement.

In Fig. 1 we further illustrate how the fraction of workers adequately saving for retirement varies with the risk aversion and the discount factor parameters. Panel A shows that, when we hold the discount factor constant at 0.95, for values of risk aversion above 3.6 more than half of the workers in our sample have a CEQR lower than 1, i.e. are not be adequately prepared for retirement. For values of 2.8 and above, about 1/3 of the workers fall short. Similarly, Panel B of Fig 1 shows the fraction of workers adequately saving for retirement as a function of the discount factor. In this case, the curve is much flatter than the risk aversion one and, even for a discount factor of 0.90, more than half of the workers are not saving adequately.

Finally, an alternative way to look at our results is to ask what level of risk aversion (discount factor) would make each worker appear adequately prepared, i.e. have a CEQR equal to 1. Panel A (B) of Fig. 2 shows the cross-sectional distribution of risk aversion (discount factor) parameters we obtain by conducting this exercise on 10,000 randomly drawn accounts. While both distributions look very sensible, the analysis below shows that the risk aversion and discount factor parameters consistent with all workers saving enough have in most cases the opposite relationship with the worker demographics and allocation decisions than the ones predicted by theory and the empirical literature.

Panel A shows that the distribution of relative risk aversion we obtain from the exercise above has a mean of 4.17, a median of 4.12, and displays substantial heterogeneity. For a small number of workers there is no value of the risk aversion parameter that make their CEQR equal to one, but overall this distribution is consistent not only with those assumed in the portfolio choice literature, but also with the findings from surveys, lab, and field studies both in the U.S. and in developing countries. For comparison, Paravisini, Rappoport, and Ravina (2016) find a mean RRA of 2.85, a median of 1.62, and significant heterogeneity, when estimating the risk aversion of actual investors on a peer-to-peer lending platform. Choi et al. (2007) find an implied risk aversion of 1.8 in their study of rainfall in Vietnam. Holt and Laury (2002, 2005) estimate relative risk aversion below 1 when presenting lab participants with a series of lotteries designed to elicit their risk attitudes.[36] Panel B shows that the discount factor distribution is more tilted toward impatience and under-weighting of the future than the findings in the literature. The average discount factor is 0.93 and the median is 0.94, with values ranging from 0.87 at the 10th percentile to 0.996 at the 90th.[37]

In Panel D of Table 9 we examine how the values of risk aversion and discount factors derived above relate to the workers' demographic characteristics and asset allocation decisions. Columns (1) to (3) show that our implied risk aversion measure is either unrelated or negatively related to age, counter to the vast empirical evidence finding that the elderly are more risk averse than the young (Harrison et. al., 2007, among others). Our measure of implied risk aversion is also positively related to account value, income, and the equity share, counter to other evidence and to theory. Columns (4) to (6) show that the discount factor implied by our results is unrelated to age, contrary to the findings in Green, Fry, and Myerson (1994) and the subsequent literature. It is however positively related to the account value, salary and the equity share, as we would expect. In Columns (7) to (12) we confirm that these results are not due to a peculiarity of our dataset, as the equity share and the account value are indeed related to these demographics and have highly statistically significant coefficients of the expected sign. The equity share decreases with age, is higher for workers with larger account values and higher salaries, and it is positively related to contribution rates and negatively related to tenure. Similarly, the size of the retirement account is positively related to age, the equity share, salary, contribution rates, and tenure.

Most importantly, in Column (1) of Panel E, we show the strongly positive and extremely large relationship between risk aversion and simulated age-65 wealth: to make our results consistent with retirement preparedness, a $10,000 higher retirement-age wealth requires the worker's risk aversion to be 2.74 higher. In the following three columns we control for the same demographics as in Column (1) to (3) of Panel D, and find that the coefficients have the same sign and similar or larger magnitude. In Column (4) we show that the relationship between simulated age-65 wealth and the discount factor is also strongly positive

[36] They also show that risk aversion estimates increase with the stakes, which partially explains their lower estimates.

[37] See Angeletos et al. (2001), Shapiro (2005), Ashraf, Karlan, and Yin (2006) for estimates based on consumption, savings, asset allocation, and voluntary adoption of forced savings technologies.

and extremely large: a $10,000 increase in wealth is associated to a 0.9 higher discount factor, which corresponds to a 1.5 standard deviation increase. In the next three columns we control for the same demographics as in Column (4) to (6) of Panel D and find similar results.

While a vast literature in both psychology and economics has established risk aversion is context-dependent and not necessarily consistent across domains (Weber and Milliman (1997), Barseghyan et al. (2011)), the numerous studies of its relationships with demographics and asset allocations have yielded quite consistent findings (Harrison et. al. (2007), Dohmen et al. (2011)). The fact that the risk aversion and discount factor parameters that can explain our results have in most cases the opposite relationship with the worker demographics than the one overwhelmingly predicted by the literature, and specifically the fact that the poor, or those who get a bad draw, are predicted to be very risk loving and yet don't have portfolio allocations reflecting these preferences cast doubt that the derived risk aversion and discount factor parameters are a realistic depiction of workers' preferences.

6.2.3 Lower average returns

In this Section we explore the sensitivity of our results to the possibility that future returns on risky assets might be lower than in the past. Specifically, we repeat our simulations under a more conservative scenario in which the expected returns on both stocks and bonds are 1 percentage point lower going forward.[38]

Table 10 reports the CRRRs and CEQRs under this scenario, and includes below each entry the difference relative to the baseline case. The results show that those most affected by lower future returns are the well-off, as a larger fraction of their wealth is in financial assets, rather than housing or Social Security claims. Workers at the 75th percentile of the distribution have a 30% probability of having a CRRR lower than 1, and thus being forced to lower their standards of living during retirement. The losses diminish as we move down the distribution. The CRRR of the median worker is 8% lower for the median simulation, and the CRRRs of those at the 10th and 25th percentiles are only modestly affected. The CEQRs fall by an amount between 4 and 2 percentage points, implying a retirement consumption

[38] Horneff, Maurer and Mitchell (2018) study the impact of low future returns on 401(k) wealth accumulation in the context of a life-cycle model.

between 2 and 4 ppt lower every year.

6.2.4 Catching up on contributions when account balances are low

While our analysis takes into account that after age 50 the cap on employee contributions increases, so that workers who have not saved enough can contribute more if they want to, we have not yet examined which type of worker does so, and what happens when we allow the contribution rate to depend directly on the account balance accumulated up to that point.

Fig. 3 shows the average contribution rate by age bracket and account decile. It shows that account balances and contribution rates are positively correlated, and that older workers tend to have higher contribution rates, but more so if they have high account balances and, most likely, high salary. The slope of the relationship between contribution rates and account balances is lower for workers in the 60-65 age bracket though, indicating that for workers close to retirement, the differences in contribution rates between wealthier and poorer workers while substantial are not as large as for the younger cohorts.

The last two columns of Panel A of Table 2 report the results of including the lagged account balance in the contribution rate regressions. Column (9) shows that the main determinant of future contribution rates is the worker's past contribution rate and that the coefficient is stable compared to the specification we have used in our baseline simulations (Column (5)). In addition, over the relevant values of the age parameter, the cubic polynomial in age has the same concave shape as the quadratic age function in Column (5). Besides salary, tenure, and the lagged account balance, the regression also includes various interaction terms of the lagged account balance, age, and the lagged contribution rate. The coefficients point to a positive relationship between the account balance and future contribution rates. This is true for all workers, even those close to retirement.[39] Column (10) shows that these coefficients are robust to the inclusion of firm fixed effects in the regressions. Table

[39]If we compare a 55 year old worker with the average contribution rate, salary, tenure, and account balance of his cohort with a 55 year old with the same characteristics but in the 10th percentile of account balances, a \$10,000 higher account balance would translate in a 4.89 ppts higher contribution rate for the worker with the higher balance, and a substantial, although smaller, 3.69 ppts higher contribution rate for the worker with the lower one.

11 hows that the CEQR s do not change by much, and if anything they are a few ppts lower than the baseline. Similarly, most of the CRRRs stay the same, except for the workers at the 75th and 90th percentiles for whom this specification results in CRRRs between 2 and 6 ppts lower. It is worth noting that this more complex specification generates a lower R^2 than the one in Column (5) we picked for the baseline analysis.

7 Cross-sectional heterogeneity

In this section we regress the median wealth accumulated at retirement, the median CRRR, and the CEQR on the *initial* characteristics of the workers in our sample. The objective is to study how the outcomes and our measures of retirement preparedness relate to the initial heterogeneity in our sample.

7.1 Regression Analysis

Panel A of Table 12 reports the coefficients from cross-sectional regressions of the median wealth accumulated by age 65 on worker's characteristics, while Panel B reports the coefficients from similar regressions of the median CRRR and the CEQR.

Column (1) of Panel A shows that median wealth at age 65 is a convex function of age at the starting point of the simulations, and a concave function of salary. All else equal, a worker who is 41 years old today and earns the median salary in our sample, \$54,522, can expect to have accumulated \$584,270 at retirement, 3.2 times the wealth of a worker of the same age who earns the 10th percentile of salary, \$21,925, and slightly more than half the wealth of a worker with the same age who earns the 90th percentile of salary, \$96,794. These effects are qualitatively robust when we control for account balance, contribution rates, tenure at the firm, and the percent invested in equity. Column (4) of Panel A shows the importance of the initial contribution rates and fraction invested in equity in explaining the variation in age-65 wealth. A one percentage point higher contribution rate is associated to a \$30,580 higher age-65 retirement wealth, while a ten percentage point higher equity allocation is associated to a \$7,120 higher retirement wealth, on average. Adding these account features to the regressions increases the R^2 substantially, from 58.6% to 78.1%.

Column (5) of Panel A shows the coefficients from including plan and firm characteristics in the regressions. We find that, all else equal, workers employed at companies with more generous employer contributions, companies that are older, privately held, invest more, and have higher net income, tend to have more wealth by the time they reach retirement. Workers employed at companies that all else equal are larger in terms of assets and number of employees tend to have lower wealth, on average. Adding firm and plan characteristics increases the R^2 from 78.1% to 83.4%. Column (6) shows that the coefficients are robust to including company fixed effects instead of controlling for the available firm-level characteristics. Finally, Column (7) shows that the coefficients on the state-level financial literacy scores and the zip code-level education dummies have the expected sign, and are statistically significant at the 5% or 10% level, although adding them doesn't increase the R^2.[40]

We next turn our attention to the regressions of the median CRRR and the CEQR reported in Panel B of Table 12.[41] The coefficients show that the consumption replacement ratios are a convex function of age, and that, unlike for retirement wealth, initial salary is not statistically significant once we add regressors to the analysis. Column (4) shows the important role of contribution rates in determining retirement adequacy. A one percentage point higher contribution rate is associated to a 2.67 ppt increase in the median CRRR, corresponding to a consumption level 2.67 ppt higher in all retirement years. The size of the account and the equity share have the expected positive and statistically significant coefficients, although their economic magnitudes is smaller: a $10,000 increase in account balance corresponds to a 69 basis points increase in CRRR, while a 10 ppt increase in the equity share corresponds, all else equal, to a 58 bps increase in CRRR. Similar to the findings

[40]When we estimate the evolution equations for our simulations, we include the features and level of employer contributions among the explanatory variables, but exclude some of the firm and local characteristics such as employer size and profitability, financial literacy, and education measures for the sake of tractability and because of their small additional contribution to the explanatory power of the regressions. We include them in the analysis here, as they are part of the retirement preparedness discussion and therefore it is interesting to see how they are related to the simulation outcomes. The results in this Section confirm that these variables are significantly related to retirement wealth and preparedness at age 65, but that including them in the regressions doesn't significantly increase the R^2 once all the other variables are taken into account.

[41]To save space we only report the coefficients of specifications (1), (4), (5), and (7). The full table is available upon request.

on median retirement wealth, workers with longer tenures at the firm end up with lower CRRRs: a 5 year higher tenure corresponds to a 4.87 ppts lower CRRR. Column (5) of Panel B shows that a more generous employer match, proxied by a higher tier one match rate, is associated with higher retirement consumption: each 1 ppt higher employer match generates a 29.5 bps higher annual retirement consumption.[42] The same column shows that, all else equal, workers employed at firms that are private, older, and have higher capital expenditures and net income tend to have higher CRRRs on average. On the contrary, workers at larger firms and firms with more employees have on average lower CRRRs. Finally, Column (7) shows that workers living in areas with higher financial literacy and with a higher fraction of college educated people tend to have higher CRRRs. The results are similar for the CEQR, except that salary is significant, while the equity portfolio share is not.

7.2 Worker Profiles

In this section we further illustrate the retirement preparedness outcomes by looking at how workers with different initial profiles fare in our simulations. Panel A shows the outcomes by age bracket, while Panel B shows the outcome for workers with the median characteristics of their age bracket.

The last three rows of Panel A show that for workers ranking at the 50th percentile and higher in the outcome distribution, the median CRRR and the CEQR are a U-shaped function of age. By contrast, the CRRRs fall from 1.08 to 0.83 for the worker at the 25th percentile of the distribution, and from 0.96 to 0.67 for the one at the 10th. The Panel also shows that, as a result of this pattern, the dispersion in outcomes increases significantly with age: the difference in CRRRs (CEQRs) between the 10th and 90th percentiles increases from 0.67 (0.41) for the younger cohort to 1.01 (0.61) for the oldest one. This increase in dispersion is primarily driven by the left tail of the distribution: more than 85% of those in the 20 to 35 years old age bracket have a median CRRR higher than one, while only about 50% of those in the older groups have this level of retirement preparedness. This finding is particularly

[42] Regressions available upon request indicate that the results are robust to more complex specifications of the employer matching contribution schemes, and that the effect of higher fees is negative but not statistically significant.

concerning since older workers have less time to adjust their wealth accumulation paths. One possible reason for these findings is that younger cohorts are enrolled in plans with more generous matching contributions: their employers match, on average, 66.24% of the worker's contributions up to a cap, compared to an average of 57.22%for the older cohorts.[43]

In Panel B of Table 13 we take the coefficients from Panels A and B of Table 12 and use them to compute the expected age-65 wealth, median CRRR, and CEQR for selected worker profiles progressively controlling for more worker characteristics. The first three columns illustrate the results for a worker who is about 41 years old, the median age in our sample, Columns (4) to (6) report the results for a 26 year old, the 10th percentile of age in our sample, and Columns (7) to (9) for a 57 year old, the 90th percentile.

In the first section of the Panel, we report fitted values for these three worker profiles based on the coefficients in Column (1) of Panels A and B of Table 12, i.e. controlling for their age and salary, but no other characteristics. The first three columns show that a 41 years old worker would have accumulated an estimated \$508,770 at retirement age if, on the last day we observe her in the sample, she had the median income in her age group (40-42 years old), \$48,323, instead of the median for the general population. This is 2.5 times the wealth she would accumulate if at the start of the simulations she earned a salary at the 10th percentile of the distribution for her age group, \$23,492, and 44% of the wealth she would have accumulated if she earned the 90th percentile for her age, \$103,316. Columns (7) to (9) show that the variance in median age-65 wealth is slightly higher for the 57 years old. By contrast, Column (4) to (6) show that the variance in wealth i s much smaller for younger workers: a 26 year old is predicted to have median wealth at retirement equal to \$524,520 if today she earns the median salary in her age group, \$34,292. This is only 1.56 times the wealth she would accumulate if she earned the 10th percentile salary, i.e. \$19,195, and 58% of the wealth if she earned the 90th percentile salary, \$65,625.

The second section of Panel B shows that controlling for account balance, contribution rates, tenure at the firm, and fraction of the portfolio invested in equity at the start of the simulations accentuates the variance across salary levels for the younger worker, and decreases it for the middle and older ones. The predicted wealth at retirement for a 26 year

[43]The subsequent tiers and caps of the employer contribution rules are also better for the younger cohort.

old earning the median salary is now 1.61 times the wealth she would have if her salary was at the 10th percentile for people her age, and 53% of the wealth she would have if she earned the 90th percentile salary. By contrast, the ratio of the median to 10th percentile worker drops to 2.1 for the 41 year old and to 2.3 for the 57 years old, and the ratios of the median to 90th percentile worker are now 44% and 39%, respectively. The values for the 57 and the 26 year old workers illustrate how those in the 10th and 90th percentiles in these age groups are on different trajectories in the labor market. A 57 year old at the 10th percentile of salary has typically been at his current firm for only 3.4 years, has a very small account balance, $3,057, low contribution rates for his age, 4.3%, and has only 60% of his portfolio invested in equities. By contrast, a 57 year old worker at the 90th percentile of her age-adjusted salary distribution has typically been at her current firm for more than 16 years, has a very large account balance, $171,567 on average, takes advantage of the additional contributions allowed after age 50 and contributes a large fraction of her salary, 16.1%, and on average invests 64% of her retirement portfolio in stocks. By contrast, relative to her counterpart at the 90th percentile, a 26 years old earning the 10th percentile salary for her age has been at her current company longer (2.33 vs. 1.94 years), contributes a smaller fraction of his salary (2.29% vs. 6.68%), and allocates a significant lower fraction of her account to equities (59% vs. 78%). These profiles illustrate the interplay of labor market trajectories, saving behavior, and investment choices in determining variation in wealth levels at retirement.

The results in the third section of Panel B, based on the coefficients in Column (5) of Table 12, indicate that workers sort into companies that exacerbate the cross sectional variance in retirement wealth highlighted above, and that this is especially the case for younger ones. Once we take company features into account, the ratios for a 26 years old become 2.44 and 47%. This spread is significantly larger than those calculated based on just age and salary (1.56 times and 58%) or age, salary, and account-level variables (1.61 times and 53%). The Panel shows that the spread also increases for the 41 and 57 year old workers, albeit to a lesser extent. As one might expect, wealth heterogeneity is even larger if we consider other values of the account- and plan-level variables, instead of the mean for the specific age and corresponding salary intervals. For example, a 41 year old worker with salary, account balance, contribution rates, and equity allocations at the 10th percentile of her age group, would have retirement wealth of $116,390, while a 41 year old worker with salary, account balance, contribution rates, and equity allocations at the 90th percentile of her age group, would have retirement wealth of $1.4 million.

Finally, the patterns for the median CRRRs 's sake By contrast, the variation in CEQR across workers with different initial salaries is much smaller, especially when we only use initial age and salary to predict retirement preparedness. Panel B shows that the cross-sectional variation increases when we use other worker and plan characteristics to compute the fitted values, and that the workers with of median age are the ones with overall a lower degree of preparedness as measured by the CEQR, while the older workers are the ones displaying the biggest amount of variation within their age group.

8 Counter-factual experiments

In this Section we perform a series of counterfactual experiments in an attempt to quantify, within our framwork, the impact of alternative policy interventions aimed at improving retirement savings adequacy.

A limitation of our approach is that the results are subject to the Lucas critique, since we are not using a structural model with optimizing agents for the pre-retirement period, and assume instead that the stochastic processes for contributions and portfolio allocations remain unchanged following the policy change. In addition, in our counterfactual experiments we will keep non-retirement financial wealth and housing wealth ($W_{65}^{FW} + \theta W_{65}^{HW}$) constant at the level estimated in the baseline scenario. In reality, when forced to increase their contributions to DC accounts, some workers might respond by saving less in their non-retirement accounts. For these reasons, we view our results as a best-case scenario for these policies. It is worth noting that the empirical evidence overwhelmingly shows that individuals often respond very passively to changes in the features of their 401(k) plan, both in terms of contribution rates and investment decisions (e.g. Choi, Laibson and Madrian (2009), Choi, Laibson, Madrian and Metrick (2003, 2004), Madrian and Shea (2001), and Chetty et al. (2014)).

The policies we consider are removing penalty-free withdrawals upon reaching age 59 1/2, except in case of hardship and unemployment; imposing a minimum contribution rate for all workers; increasing all workers' contribution rate by a fixed percentage; and automatic rollovers following a job switch.

8.1 Removing penalty-free age 59 1/2 withdrawals

Munnell and Webb (2015) show that individuals withdraw significant amounts from their retirement accounts before age 65. Restricting such withdrawals could therefore potentially decrease the large shortfalls in wealth accumulation that we have documented.

We evaluate the effect of removing the ability to withdraw funds from age 59 1/2 onwards without a justification, while still allowing those due to hardship and unemployment at any age. Table 14 shows that this measure would have the largest impact on the wealth accumulation of workers in the bottom half of the distribution, causing increases in retirement wealth between 15% and 21%. The reason is that both the probability of a withdrawal and the percentage of the account being withdrawn decrease with wealth (see Section 3.2). Yet, the increases in CRRRs and CEQRs are higher for workers who were already better-off in the baseline scenario, since the fraction of retirement consumption financed by DC wealth, as opposed to Social Security benefits, is increasing in total wealth. The results in Table 14 show that preventing individuals from withdrawing funds from their DC accounts without justification after age 59 1/2 can increase annual retirement consumption levels by 5% for poorer workers to 9% for wealthier ones. Yet, despite these substantial improvements, about 2/3 of the workers would still end up with a CEQR below 1, and close to 25% of them would still have a median CRRR below 1, due to the low CRRRs and CEQRs in our baseline results.

8.2 Minimum contribution rate

In this Section we explore the effect of a mandatory minimum contribution rate of either 2.5% or 5%. While this kind of measures might be more effective for workers who are not saving in a 401(k) altogether, we explore them here to see if they could also help increase the saving rates of the workers who already have a 401(k) but are saving very little.

For comparison's sake, the average contribution rate in our sample is 6.3%, the median is 5.2%, and the 25th percentile is 2.2%. Table 15 shows that imposing a minimum contribution rate of 2.5% (Panel A), or even 5% (Panel B), has negligible effects on the retirement savings of the workers in our sample. The average increases in CRRRs never exceed 2% and, as discussed above, such increases might be lower or inexistent if we allowed for crowding out of non-retirement wealth.

Panel C reports the results across age groups for the case of a 5% minimum contribution rate, and shows that, while the improvements are negligible for the overall sample, workers younger than 35 experience increases in median CRRRs of about 4%. The reason is that the young have the lowest contribution rates in the sample, due to their lower income at this stage of their life cycle.[44] In addition, most of their savings are precautionary and should not be invested in an illiquid retirement account. Carroll (1997), Gourinchas and Parker (2002), and Gomes, Michaelides and Polkovnichenko (2009) indeed show that the optimal retirement saving rate of the young can be much lower than 5%. Forcing them to contribute more may be suboptimal and could lead several of them to opt out of the pension plan altogether. Thus, in Table 16 we explore the possibility of setting the following age-dependent minimum contribution rate:

$$k_{\min} = 4.5\% + (age - 21) * 0.25\% \tag{38}$$

where the minimum contribution age averages 10% between ages 21 to 65, but starts at a low level of 4.5% and increases gradually to 15.5% just before retirement, when the worker can presumably afford to save more. An average contribution rate of 10% is certainly substantial, since the average contribution rate in our sample is 6.3%, and might not be feasible for the older workers with low income. Yet, given the magnitude of the under-savings problem we have documented, this is the change needed to get significant improvements. Panel A of Table 16 shows that those in the left tail of the distribution would experience increases in their median CRRR of about 5 ppts, and increases in their CEQR of 3 ppts. More importantly, Panel B shows that the gains would be similar for all age groups within a given decile.

[44]In our sample, 58% of workers younger than 35 have contribution rates below 5%, while only 43% of older workers do.

8.3 Increase in contributions rates for all workers

In this section we study the impact of an increase in every worker's contribution rate by either 2 or 5 percentage points, regardless of their current saving rate. Again, our goal is to evaluate the additional savings effort required to address the shortfalls that emerged from our baseline results, rather than being prescriptive.

The results reported in Table 17 are promising. A 2 ppts increase in contribution rates improves yearly retirement consumption levels by 2% to 9% (Panel A). A 5 ppts increase raises the standard of living at retirement by between 4% to 20% (Panel B).[45] It is worth noting that a 5 ppt increase in retirement contributions represents a quite significant savings effort. Also, the improvements are highest for those who are already better off, as these workers tend to have higher salaries and end up saving larger amounts as the result of this policy. Compared to the baseline results in Table 7, the percentage of workers not saving enough for retirement, i.e. with a CEQR less than 1, falls from 3/4 to about 2/3. While this is quite a significant improvement, these figures are still worrisome and highlight the magnitude of the current under-saving problem.

Panel C shows the results across age groups for an increase in contribution rates of 2 ppt. The largest effect is for younger workers who would have more time to benefit from the increased saving rate. Their CRRR would be between 9% to 15% higher. Those in 35 to 49 age group would increase annual retirement consumption by about 4% to 5%, while those currently in the 50 to 64 age bracket would enjoy a more modest boost between 1 and 3 ppts, as they have less years left before retirement.

8.4 Automatic rollover following a job switch

One of the recently proposed policy measures to decrease leakages and improve retirement savings is to make the plans portable so that they are automatically transferred to the new employer in case of a job separation. In this Section we estimate the potential effect of such policy by completely removing the possibility of withdrawals following a job switch.

[45]Notice that our simulations take into account IRS limits on total employee annual contributions and age 50 catch-up additional contributions.

Table 18 shows that both the CRRR and CEQR increase across the worker distribution, and, within each worker, across the outcome distribution. The size of the consumption increase ranges between 1 and 7 ppts and is significantly larger for workers who are already better off. While effective, this policy alone, in the context of our model and assumptions, would not address the retirement savings crisis. Additional changes related to increasing savings for everyone, but especially for the poorer workers, or a combination of various policy measures are necessary to address the shortfall.

9 Conclusions

Are Americans adequately prepared for retirement? In this paper we have explored retirement savings adequacy of a large and representative sample of U.S. workers saving in their company 401(k) by simulating their wealth accumulation forward till age 65 and examining if, after considering longevity, medical, and investment risks, it is sufficient to maintain their standard of living in retirement. We find that, based on their current account balances, income, saving, and investment patterns, on the probabilities of withdrawing funds following unemployment, job switches, hardship, and reaching age 59 1/2, about 3/4 of the workers in our sample are not saving enough for retirement. Several factors contribute to the dispersion in outcomes across individuals. The most significant ones are the heterogeneity in the generosity of employer contributions, individual saving rates, and asset allocations.

Our results are robust to various alternative calibrations. Only if we assume both low risk aversion and very high discount rates, do we conclude that the median worker is saving enough.Yet, the risk aversion and discount factor parameters that can explain our results have in most cases the opposite relationship with the worker demographics and allocation decisions than the ones predicted by theory and the empirical literature.

While the picture of retirement preparedness emerging from our analysis is somber, there are various reasons to believe it is an understatement of the paucity of retirement savings. Data from the Bureau of Labor Statistics indicate that only 65% of private sector workers have access to a retirement plan, and only 48% participate in them. Furthermore, we have not included in our analysis risks such as potential reductions in social security benefits, or increases in medical costs. On the other hand, we don't study an increasingly small fraction of workers with access to a defined benefit plan, which might, if its promises are fulfilled, provide better retirement adequacy.

Finally, the analysis we have conducted here is in many ways exploratory, and many open questions remain. We have only analyzed a few policies aimed at increasing retirement savings, while others, such as postponing retirement, mandatory automatic enrollment for all workers, and financial education could also be beneficial. In future work we also plan to explore a wider range of preference parameters and return environments, to allow for more asset classes, and for the possibility that some workers end up in firms with no retirement plan and fail to make up the lack of savings on their own.

Appendix: Results with different assumptions regarding new pension plan following a job loss

As discussed in Section 3, we consider three different scenarios regarding the worker's new pension plan following a job loss:

- The new DC plan is identical to the current one

- The new DC plan is randomly drawn from our sample of workers within the same (initial) income decile.

- The new DC plan is randomly drawn from our full sample.

Table $A1$ shows that the baseline results are quantitatively almost identical across all three scenarios. Panel A reports the results when the new DC plan is randomly drawn from our sample of workers within the same income decile, and Panel B reports the results for a random plan. The CEQs are all exactly identical to the baseline ones reporter in Table 8 in the paper up to the 2 decimal points. The vast majority of CRRRs are almost identical, and the few differences are in the order of 0.01.

Agnew, J., P. Balduzzi, and A. Sunden, 2003, “Portfolio Choice and Trading in a Large 401(k) Plan”, American Economic Review, 93, 1, 193-215.

Aguiar, M., and E. Hurst, 2005, “Consumption vs. Expenditure,” Journal of Political Economy 113, 4, 919-948.

Aguiar, M., and E. Hurst, 2013, “Deconstructing life cycle expenditure,” Journal of Political Economy 121, 3, 437-492.

Ahmed, J., Barber B., and T. Odean, 2016, “Made poorer by choice: Worker outcomes in social security vs. private retirement accounts,” Journal of Banking and Finance, 23, 1-12.

Angeletos, G., Laibson, D., Repetto, A., Tobacman, J. and Weinberg, S., 2001, “The hyperbolic consumption model: calibration, simulation, and empirical evaluation,” Journal of Economic Perspectives, 15, 3, 47-68.

Ashraf, N., Karlan, D., and W. Yin, 2006, “Tying Odysseus to the mast: evidence from a commitment savings product in the Philippines,” Quarterly Journal of Economics, 121, 673-697.

Barseghyan, L., Prince, J., and Teitelbaum, J., 2011, “Are Risk Preferences Stable Across Contexts? Evidence from Insurance Data,” American Economic Review, 101, 2, 591-631.

Bekaert, G., K. Hoyem, W. Hu and E. Ravina, 2017, “Who is internationally diversified? Evidence from the 401(k) Plans of 296 Firms”, The Journal of Financial Economics, 124, 1, 86-112.

Biggs, A. 2016, “Retirement Replacement Rates: What and How”, in Retirement System Risk Management: Implications of the New Regulatory Order, edited by Olivia Mitchell, Raimond Maurer and J. Michael Orszag.

Bilias, Y., Georgarakos, D., and M. Haliassos, 2010, “Portfolio Inertia and Stock Market Fluctuations,” Journal of Money, Credit, and Banking, 42, 4, 715-742.

Brady, P., 2010, “Measuring Retirement Resource Adequacy,” Pensions Economics and Finance, 9, 2, 235-262.

Brown, J., C. Fang and F. Gomes, 2012, “Risk and Returns to Education Over Time ”, NBER Working paper.

Browning, M., and A. Lusardi, 1996, “Household Saving: Micro Theories and Micro Facts”, Journal of Economic Literature, 34, 1797-1855.

Brunnermeier, M., and S. Nagel, 2008, “Do wealth fluctuations generate time-varying risk aversion? Micro-evidence on individuals’ asset allocations,” American Economic Review, 98, 3, 713-736.

Caplin, A., 2002, “The Reverse Mortgage Market: Problems and Prospects,” in Innovations in Retirement Financing. Eds Olivia S. Mitchell, Zvi Bodie, Brett Hammond, and Steve Zeldes. Philadelphia, PA: University of Pennsylvania Press: 234-253.

Carroll, C. D., 1997, “Buffer-Stock Saving and the Life-Cycle/Permanent Income Hypothesis,” Quarterly Journal of Economics, 112, 1–55.

Chetty, R., J. Friedman, S. Leth-Petersen, T. Nielsen, and T. Olsen, 2014, “Active vs. Passive Decisions and Crowd-out in Retirement Savings Accounts: Evidence from Denmark”, The Quarterly Journal of Economics 129(3): 1141-1219.

Choi, J., D. Laibson and B. Madrian, 2009, “Mental Accounting in Portfolio Choice: Evidence from a Flypaper Effect,” American Economic Review, 99, 2085-2095.

Choi, J., D. Laibson and B. Madrian, 2011, “$100 Bills on the Sidewalk: Suboptimal Investment in 401(k) Plans” Review of Economics and Statistics, 93, 3, 748-763.

Choi, J., D. Laibson, B. Madrian and A. Metrick, 2003, “Optimal Defaults,” American Economic Review: Papers and Proceedings, 93, 180-185.

Choi, J., D. Laibson, B. Madrian and A. Metrick, 2004, “For Better or For Worse: Default Effects and 401(k) Savings Behavior.” Perspectives in the Economics of Aging, edited by David A. Wise. Chicago, IL: University of Chicago Press.

Choi, S., Fisamn, R., Gale, D., and S. Kariv, 2007, “Consistency and heterogeneity of individual behavior under uncertainty,” American Economic Review, 97, 5, 1921-1938.

Clark, R., A. Lusardi and O. Mitchell, 2015, ““Financial Knowledge and 401(k) Investment Performance,” Journal of Pension Economics and Finance, 1–24.

Cocco J., 2005, “Portfolio Choice in the Presence of Housing,” Review of Financial Studies, 18, 2, 535-567.

Cocco J. and F. Gomes, 2012, “Longevity Risk, Retirement Savings and Financial Innovation,” Journal of Financial Economics, 103, 3, 507-529.

Cocco, J., F. Gomes and P. Maenhout, 2005, “Consumption and Portfolio Choice over the Life-Cycle,” The Review of Financial Studies, 18, 2), 491-533.

Davidoff, T., 2005, “Maintenance and the Home Equity of the Elderly,” manuscript, University of California at Berkeley.

De Nardi, M., E. French and J. B. Jones, 2010, “Why Do the Elderly Save? The Role of Medical Expenses,” Journal of Political Economy, 118, 1, 39-75 .

Deaton, A., and S. Zaidi, 2002, “Guidelines for Constructing Consumption Aggregates,” LSMS Working Paper 135, The World Bank.

Dohmen, t., Falk, A., Huffman, D., Sunde, U., Schupp, J., Wagner, G., 2011, “Individual Risk Attitudes: Measurement, Determinants and Behavioral Consequences,” Journal of the European Economic Association, 9, 3, 522-550.

Engelhardt, G., 2003, “Reasons for job change and the disposition of pre-retirement lump-sum pension distributions,” Economic Letters, 81, 333-339.

Engen, E., W. Gale, and C. Uccello, 2005, “Lifetime Earnings, Social Security Benefits, and the Adequacy of Retirement Wealth Accumulation,” Social Security Bulletin, 66, 1, 38-57.

Fisch, J., and A. Lusardi, 2020, “Defined Contribution Plans and the Challenge of Financial Illiteracy”, Cornell Law Review, Forthcoming.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus, 2019, “Five Facts about Beliefs and Portfolios”, Working Paper.

Gomes, J., F. Michaelides and V. Polkovnichenko, 2009, “Optimal Savings with Taxable and Tax-Deferred Accounts,” Review of Economic Dynamics, 12, 718-735.

Gourinchas, P.-O., and J. Parker, 2002, “Consumption Over the Life Cycle,” Econometrica, 70, 47–91.

Green, L., Fry, A., and J Myerson, 1994, “Discounting of delayed rewards: a life-span comparison,” Psychological Science, 5, 1, 33-36.

Gustman, A. , T. Steinmeier, and N. Tabatabai, “Updated Pension Wealth Data Files in the HRS Panel: 1992 to 2010 Part III”, 2014.

Harris C, and D. Laibson, 2001, “Dynamic Choices of Hyperbolic Consumers,” Econometrica, 69, 4, 935-957.

Harrison, G., Lau, M., and E. Rutstrom, 2007, “Estimating risk attitudes in Denmark: a field experiment,” Scandinavian Journal of Economics 109, 2, 341-368.

Holt, C., and S. Laury, 2002, “Risk aversion and incentive effects,” The American Economic Review, 92, 5, 1644-1655.

Horneff, W., R. Maurer, and O. Mitchell, 2018, “How will persistent low expected returns shape household economic behavior,” Journal of Pension Economics and Finance, 1-11.

Horneff, W., R. Maurer, O. Mitchell, and M. Stamos, 2010, “Variable Payout Annuities and Dynamic Portfolio Choice in Retirement,” Journal of Pension Economics and Finance. 9(2): 163-183.

Hurd, Michael D. and Susan Rohwedder, 2012, “Economic Preparation for Retirement,” in Investigations in the Economics of Aging, 77–113. University of Chicago Press, David Wise (ed.).

Laibson D. 1997, “Golden Eggs and Hyperbolic Discounting,” Quarterly Journal of Economics. 112, 2, 443-477.

Lee, R. D. and L. R. Carter, 1992, “Modelling and Forecasting U.S. Mortality,”Journal of the American Statistical Association, 87, 419, 659-671.

Love, D. M. Palumbo and P. Smith, 2009, “The Trajectory of Wealth in Retirement,” Journal of Public Economics, 93, 191-208.

Lusardi, A. and O. Mitchell, 2014, “The Economic Importance of Financial Literacy: Theory and Evidence.”Journal of Economic Literature, 52, 5-44.

Lusardi, A., O. Mitchell, and N. Oggero, 2017, “Debt and Financial Vulnerability on the Verge of Retirement,” NBER Working paper.

Lusardi, Annamaria and Olivia S. Mitchell, 2018, “Financial Literacy and Planning: Implications for Retirement Wellbeing,” Working Paper.

Madrian, B. and D. Shea, 2001, “The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior,” The Quarterly Journal of Economics, 116, 4, 1149-1187.

Munnell, A., and A. Chen, 2017, “Who contributes to individual retirement accounts?” Issues in brief, Center for Retirement Research.

Munnell, A., Hou, W., and G. Sanzenbacher, 2018, “National Retirement Risk Index shows modest improvement in 2016,” Issues in brief, Center for Retirement Research.

Munnell, A., A. Webb, and L. Delorme, 2006, “A New National Retirement Risk Index,” Working Paper, Center for Retirement Research at Boston College.

Munnell, A., A. Webb, 2015, "The impact of leakages on 401(k)/IRA assets," Working Paper, Center for Retirement Research at Boston College.

Paing, G. and S. Schieber, 2014, "Why American Workers' Retirement Social Income Security Prospects Look so Bleak: A Review of Recent Assessments," Journal of Retirement, 2, 35-54.

Paravisini, D., Rappoport, V., and E. Ravina, 2016, "Risk Aversion and Wealth: Evidence from Person-to-Person Lending Portfolios," Management Science, 63, 2, 279-297.

Poterba, J., 2015, "Saver Heterogeneity and the Challenge of Assessing Retirement Saving Adequacy," National Tax Journal 68, 2, 377-388.

Poterba, J., S. Venti, and D. Wise, 2011. "The Composition and Drawdown of Wealth in Retirement." Journal of Economic Perspectives, 25(4):95-118.

Poterba, J., S. Venti, and D. Wise, 2012. "Were They Prepared for Retirement? Financial Status at Advanced Ages in the HRS and the AHEAD Cohorts," Investigations in the Economics of Aging, 21–69. University of Chicago Press, Wise, David (ed.).

Purcell, P., 2012,"Income Replacement Ratios in the Health and Retirement Study", Social Security Bulletin, vol. 72, no.3.

Scholz, J.K., A. Seshadri and S. Khitatrakun, 2006, "Are Americans Saving 'Optimally' for Retirement," Journal of Political Economy, 114, 4 607-643.

Shapiro, J., 2005, "Is there a daily discount rate? Evidence from the food stamp nutrition cycle," Journal of Public Economics, 89, 303-325.

Skinner, J., 2007, "Are You Sure You're Saving Enough for Retirement?," Journal of Economic Perspectives, 21(3), 59-80.

Yao, R. and H. Zhang, 2005, "Optimal Consumption and Portfolio Choice with Risky Housing and Borrowing Constraints," Review of Financial Studies, 18, 1, 197-239.

VanDerhei, J., 2006, "Measuring Retirement Income Adequacy and Calculating Realistic Replacement Rates," EBRI Issue Brief, 297.

Venti, S., and D. Wise, 2004, "Aging and Housing Equity: Another Look," Perspectives on the Economics of Aging, 127–175. University of Chicago Press. David Wise (ed.)

Weber, E.U., and Milliman, R., 1997, "Perceived risk attitudes: Relating risk perception to risky choice," Management Science 43, 122-143.