The Life-Cycle Model Implies that Most Young People Should Not Save for Retirement

Jason S. Scott

John B. Shoven

Stanford University and NBER

Sita N. Slavov

George Mason University and NBER

John G. Watson

Stanford Graduate School of Business

I. Introduction

Retirement policy is often predicated on the assumption that more retirement saving is better, at least at the margin. For example, there is a large literature showing that auto-enrollment dramatically increases participation rates in 401(k) plans, particularly among younger workers (e.g., Madrian and Shea 2001). That outcome is thought to be unambiguously positive, and it has resulted in policy changes and proposals that are designed to increase the use of auto-enrollment (including the Pension Protection Act of 2006 and the proposed SECURE Act 2.0). Moreover, withdrawals from 401(k) plans that occur when someone changes jobs – and are more common among young workers (Munnell and Webb 2015) – are generally viewed in a negative light (Tergesen 2017; VanDerhei 2019).

Concerns over 401(k) participation rates and leakages from retirement saving are based on the claim that less-than-fully rational individuals, who are prone to inertia and myopia, are not saving optimally for retirement. Making that claim, of course, requires determining what optimal retirement saving looks like. A reasonable benchmark is the life-cycle model, which is a foundational tool in economics. The life-cycle model is often used for positive analysis – that is, to predict how individuals may alter their consumption and saving behavior in response to shocks. However, it has also been used for normative analysis – that is, to guide optimal retirement saving (e.g., Bodie, Treussard, and Willen 2007; Scholz and Seshadri 2009; Kotlikoff 2007). Indeed, many economists and financial planners view the life cycle model as a reasonable prescription for behavior, allowing individuals to maintain a steady standard of living despite changes in income. From a normative perspective, the model addresses the key question, “How should I allocate my resources across my lifetime?” In this paper, we argue that if the life cycle model is taken seriously as a benchmark for optimal consumption and saving behavior, then many people in their 20s and 30s should not be saving for retirement at all, a finding that runs counter to policies that aim to boost 401(k) participation rates and discourage leakages uniformly at all ages.

Consider a basic life-cycle model that has two features. First, utility derived from consumption each year is invariant to age[1], and second, aggregate lifetime consumption cannot exceed a given amount. With these assumptions, optimal consumption is constant with the level set to just exhaust lifetime resources. Furthermore, if individuals work while young and earn a constant real income, and if they intend to retire and live off their savings when old, then saving commences immediately upon starting work. In this simple version of the life-cycle model, young people should save for retirement. However, some extensions of this model highlight important conditions under which this result does not apply, and under which saving for retirement at younger ages is not optimal. Moreover, these conditions are realistic and apply to a large fraction of workers.

First, individuals value current consumption more highly than future consumption due to both subjective time discounting and mortality risk. Interest pushes in the opposite direction by lowering the cost of future consumption relative to current consumption and increasing the incentive to consume later in life. However, as of this writing, real interest rates have been zero or negative for over a decade. Moreover, 30-year real interest rates currently hover around zero, with after-tax real rates slightly negative due to the taxation of inflation. With time and mortality discounting pushing consumption earlier, and interest rates doing nothing to favor late-life consumption, the life-cycle model implies a declining path for consumption. Higher optimal consumption earlier in life pushes against starting to save at younger ages.

Second, for a large fraction of the population, wages tend to grow during their 20s, 30s and 40s and then stabilize. This pattern of earnings is especially common among individuals with higher lifetime earnings. Even if time and mortality discounting are offset by asset returns, implying that spending in each year should be constant and equal to lifetime average income, individuals with this pattern of

---

[1] In addition, we need a well-behaved utility function in the sense that first derivative of utility with respect to consumption is positive, and the second derivative is negative.

earnings should not save when they are young because their income at younger ages is below their lifetime average. Young people who anticipate an upward-sloping wage profile are relatively poor, and ideally would like to borrow from their future, wealthier selves.

Finally, for lower-income individuals, who tend to face flatter wage profiles, Social Security replaces a relatively large share of pre-retirement wages. Given the relative generosity of Social Security, optimal savings rates may be low even when real interest rates offset mortality and time discounting. If real interest rates are near zero and do not offset mortality and time discounting, then the desired consumption profile is declining, and optimal savings rates are even lower. Moreover, saving does not begin until middle age or even later.

From a theoretical standpoint, there is nothing surprising about these results, which derive from a straightforward life-cycle model in the vein of Yaari (1965) and Davies (1981). For example, Hurd (1989) finds that mortality risk causes optimal consumption to decline over time, and Hammermsh (1984) documents empirically that spending declines by about 5 percent per year, arguing that the most plausible explanation for this finding is that the subjective rate of time preference exceeds the real interest rate. Leung (1994, 2001, 2007) derives the conditions under which individuals will optimally exhaust their wealth before the end of their lives. Milevsky and Huang (2011) show that once wealth is exhausted, borrowing constraints force individuals into constant consumption equal to their Social Security or pension income. In the absence of borrowing constraints, individuals would prefer to borrow against that income to support a declining consumption profile. In our previous work, we show that under plausible parameter values, including the low real interest rates that have persisted over the last decade, lower-income individuals, who anticipate a relatively high Social Security replacement rate, would optimally accumulate very little retirement saving and exhaust their wealth early in retirement (Scott et al. 2021a).

Young people who anticipate an increasing wage profile are in a similar situation, preferring, if possible, to borrow against future income to support a declining or even flat consumption profile. The presence of a borrowing constraint implies that saving is zero early in an individual's working life. Along those lines, Lachance (2012) finds that it is often not optimal for young people to save for retirement. In previous work, we consider wage profiles that are typical of college graduates and show that access to a defined contribution plan with tax preferences and standard employer matching is far from sufficient to overcome that result (Scott et al. 2021b). A mortality discount and an interest rate that is persistently lower than the rate of time preference – implying that optimal consumption is declining

However, these theoretical findings do not appear to inform retirement policy. Policies designed to encourage retirement saving, like automatic enrollment, are based on the concern that individuals appear to undersave for retirement relative to the predictions of a life cycle model. For example, a large literature documents that many individuals experience a drop in consumption upon retirement (an early and heavily cited study is Bernheim, Skinner, and Weinberg 2001), a finding that is arguably inconsistent with the consumption smoothing implied by the life cycle model. Akerlof (2002) argues that "the best evidence of undersaving is probably [that] … this finding is difficult to explain with the standard life cycle, exponential discounting model." These arguments are not without controversy, and the question of retirement savings adequacy overall is beyond the scope of this paper. Instead, we focus on the age profile of optimal saving, showing that in a standard life-cycle model with realistic parameter values, optimal saving at younger ages is often zero. If a goal of retirement policy is to bring individuals' behavior in line with the predictions of a life cycle model, then encouraging young people to save for retirement may not be desirable.

The contribution of this paper is to generalize the results of our previous work (Scott et al. 2021a,b) and show that, in the life cycle model, it is generally suboptimal for young people to save for retirement even with tax preferences and standard employer matching. For higher-income workers, the argument made in Scott et al (2021b) applies. Even when the real interest rate is 3 percent and equal to the rate of time preference, a modestly upward-sloping wage profile implies that retirement saving is zero until the early 30s. With a more realistic zero real interest rate and the implied downward-sloping consumption profile, saving does not begin until the late 30s or early 40s. Lower-income workers, however, are less likely to face a wage profile that is substantially upward sloping. When the real interest rate is 3 percent and equal to the rate of time preference, saving begins immediately; however, it remains relatively low throughout the working life, reaching between 1 and 4 times the amount of annual, after-tax wages. When the real interest rate is zero percent, saving does not begin until the late 40s at least, and it reaches less than 1.5 times the amount of annual, after-tax wages.

II. Model

We investigate a standard lifecycle model for an individual with a planning horizon of T years. Our individual maximizes her total utility assuming that it is separable with an annual utility contribution of $\alpha_t U(C_t)$ from spending and consuming an amount C_t in year t. Here, $U(\cdot)$ is her felicity function, and $\alpha_t > 0$ is her time and mortality discount. We assume that the felicity function is an increasing, strictly concave function of its argument. Hence, if V_* is the individual's total optimal utility, then

$$V_* = \max \sum_{t=0}^{T} \alpha_t U(C_t) \tag{1}$$

Note that there are $T + 1$ spending events that span T years; the first occurs now at $t = 0$, and the last T years in the future at $t = T$.

The maximization is subject to budget constraints. At time t, let B_t equal an individual's real annual income and benefits[2], and let Θ_t^P be the payroll tax on any wages. Further, let A_t^B equal the

---

[2] For simplicity, we assume income is equal to wages ($B_t = W_t$) during an individual's working career and is equal to Social Security benefits during retirement.

current balance in taxable savings, e.g., money in a brokerage account, and A_t^D equal the current balance in tax-deferred investments, e.g., money in a 401(k) plan or traditional IRA. Then at each time $t \in \{0, \dots, T\}$, these variables are related to real spending C_t and the income tax on wages, benefits, and brokerage returns Θ_t^I by the equalities

$$C_t = B_t - (\Delta_t^B + \Delta_T^D) - (\Theta_t^P + \Theta_t^I) \tag{2a}$$

$$A_{t+1}^B = R_{t+1}^B(A_t^B + \Delta_t^B) \tag{2b}$$

$$A_{t+1}^D = R_{t+1}^D(A_t^D + \Delta_t^D + m_t) \tag{2c}$$

Here, $R_{t+1} \geq 1$ equals the real *gross* return on assets over the interval $[t, t+1]$, Δ_t is the savings (or withdrawal) to an account, and m_t is an employer's match to the tax-deferred account.

We assume that the annual benefit stream, payroll taxes, and account returns are given. The remaining variables are decision variables and subject to bounds. In particular, we assume C_t, Θ_t^I, and A_t^D are non-negative for all t. On the other hand, we allow limited borrowing in the brokerage account, i.e., $A_t^B \geq -l_t$. For example, l_t may equal half an individual's wages W_t during her working years. Further, we impose the following time-dependent upper and lower bounds on tax-deferred contributions

$$\epsilon_t \leq \Delta_t^D \leq \gamma_t \tag{3}$$

These bounds are used to enforce contribution limits and Required Minimum Distributions (RMDs). This topic is explored below.

a. Borrowing Costs

In a taxable account, the interest rates for borrowing and saving are often different. We model this in our optimization problem by splitting an account's value into positive and negative parts and then applying the correct rate to each part. For example, let u_t be the positive part and v_t be the negative part, then

$$A_t^B + \Delta_t^B = u_t - v_t \tag{4a}$$

$$A_{t+1}^B = R_{t+1}^{B,save} u_t - R_{t+1}^{B,borrow} v_t \tag{4b}$$

The non-negative parts u_t and v_t must obey the complementarity condition $u_t v_t = 0$, i.e., only one of the parts may be positive. This prevents an investor from simultaneously saving and borrowing. For example, consider a \$100 investment. If the interest rate for borrowing is less than that for savings, then an investor prefers to borrow another \$100 and invest \$200. The complementarity condition prohibits this arbitrage.

Often, complementarity is a natural consequence of an optimization. If we assume that borrowing rates are greater than savings rates, we can reasonably expect that complementarity is implicitly enforced. In practice, we optimize without an explicit complementarity condition and subsequently verify that it holds. In summary, we model borrowing costs by replacing equation (2b) by equations (4) and introducing pairs of *split variables* $u_t \geq 0$ and $v_t \geq 0$ for each time slice.

b. *Tax-Deferred Contributions, Withdrawals, and Employer Matching*

We assume an individual may only contribute funds ($\Delta_t^D \geq 0$) and receive matching contributions ($m_t \geq 0$) while working ($B_t = W_t > 0$), and she may only withdraw funds ($\Delta_t^D \geq 0$) during retirement. Further, we assume an employer matches an employee's contribution at a rate k_1 up to a maximum wage fraction k_2. For example, $k_1 = 100\%$ and $k_2 = 4\%$ correspond to a full match up to 4% of salary. These parameters are employer specific, but generally $0 < k_1, k_1 \leq 1$. Finally, we assume employers match both regular and catchup contributions.

i. Contributions and Limits

There are three kinds of 401(k) contributions: regular employee contributions (402(g) Elective Deferrals), employee 50+ catchup contributions (414(v) Catch-up Contributions), and employer matching contributions. The IRS limits annual contributions by capping an employee's regular contribution, her 50+ catchup contribution, and the *total* of her regular contribution and her employer's match. A 50+ employee who maximizes her regular contribution and receives the maximum employer match may still make an additional 50+ catchup contribution, however, in no case may the sum of all contributions exceed an employee's annual wage W_t.

For 2021, an employee's regular contribution cap is \$19,500, her catch-up cap is \$6,500, and the cap on the total of her regular contributions and her employer's match is \$58,000. The first two limits are annually adjusted for inflation (CPI-U) and the result, if not a multiple of \$500, is rounded to the next lowest multiple. Similarly, the total limit is adjusted and rounded to a multiple of \$1000. Let K_t equal the sum of an individual's regular contribution limit and her catch-up limit at time *t* adjusted for inflation (CPI-W). If she is under age 50 at time t, then the catch-up limit is zero. Similarly, let Σ_t equal the sum of the total contribution limit and the catch-up limit.

ii. Working Years: $\Delta_t^D \geq 0$

If an employee's contribution $\Delta_t^D \geq 0$ is fully matched by her employer, then $m_t = k_1 \Delta_t^D$. Since the sum of these contributions must not exceed the minimum of an employee's wage W_t and Σ_t, contributions satisfy the inequality $\Delta_t^D \leq \min\,(B_t, \Sigma_t)/(1 + k_1)$.[3] If we combine this bound with the IRS's contribution limit $\Delta_t^D \leq K_t$, we get $0 \leq \Delta_t^D \leq U_t$ where

$$U_t \equiv \min\left(K_t, \frac{B_t}{1+k_1}, \frac{\Sigma_t}{1+k_1}\right) \qquad \text{(5a)}$$

The employer's match is subject to the plan's limit for each employee, i.e., $m_t \leq k_1 k_2 W_t$. Here, contributions are matched at rate k_1 but only a fraction $k_2 W_t$ of wages may be matched. In addition, m_t

---

[3] If an individual contributes more than this limit, then for every extra dollar she contributes, she must give back a dollar of match money. This behavior is sub-optimal and is prevented by this bound.

is bounded by the match rate times the upper bound on contributions U_t. Together we have $0 \leq m_t \leq k_1 M_t$ where

$$\mathrm{M_t} \equiv \min(k_2 W_t, U_t) = \min\left(k_2 W_t, K_t, \frac{B_t}{1+k_1}, \frac{\Sigma_t}{1+k_1}\right) \tag{5b}$$

Note that $M_t \equiv \min(k_2 W_t, U_t) \leq U_t$. In terms of U_t and M_t, the formula for an employer's match m_t m_t as a function of the employee's contribution $\Delta_t^D \in [0, U_t]$ is

$$m_t = k_1 \begin{cases} \Delta_t^D, & 0 \leq \Delta_t^D \leq M_t, \\ M_t, & 0 \leq \Delta_t^D \leq U_t \end{cases}$$

or simply $m_t = k_1 \min(\max(\Delta_t^D, 0), M_t)$. In short, only contributions are matched, and they are matched at rate k_1 until they reach the maximum value $k_1 M_t$.

The employer's match is a piecewise-linear function of an employee's contribution. For our lifecycle optimization, we do not explicitly use the match formula, rather, we get the same result by introducing the linear inequality constraints

$$0 \leq \frac{m_t}{k_1} \leq M_t \tag{6a}$$

and

$$\frac{m_t}{k_1} \leq \Delta_t^D \leq U_t \tag{6b}$$

The first pair of inequalities limits the employer's match to the plan's maximum, and the second pair guarantees that the employee contributes enough to get the employer's match, but no more than allowed by the IRS and plan rules. Note that constraints (6) imply $\Delta_t^D \geq 0$.

Generally, these inequalities are not sufficient to replicate the matching formula; they may require additional logic to guarantee that the match value lies on one of the linear segments of the match formula, and not somewhere below it. However, since matching dollars are "free", and an optimization seeks as many as possible, additional conditions are typically redundant. In practice, we optimize without them and subsequently verify that the result we obtain satisfies the matching formula.

iii. Retirement Years: $\Delta_t^D \leq 0$

Finally, there are Required Minimum Distributions (RMDs) that a retiree must take when she reaches age 72.[4] A distribution is a specified fraction f_t of the value A_t^D of a retiree's tax-deferred account. Thus, during retirement years, we have the constraints

$$L_t \leq -\Delta_t^D \leq A_t^D \tag{7a}$$

where the lower limit L_t is

$$L_t = \begin{cases} 0, & age(t) < 72 \\ f_t A_t^D, & age(t) \geq 72 \end{cases} \tag{7b}$$

In any year, a retiree may withdraw more than the RMD, but no more than the account's value. Note that constraints (7) imply $\Delta_t^D \leq 0$.

c. Taxes

All wages are subject to payroll taxes, e.g., Social Security (FICA), Medicare (MEDFICA), and Additional Medicare taxes. For Social Security, an employees and her employers each pay 6.2% of wages up to the Social Security maximum wage limit, which is \$142,800 for 2021. For Medicare, each pay 1.45% of wages. Since 2014, the ACA mandates that an employee pay an additional 0.9% of wages that exceed \$200,000. Given a wage stream W_t, the corresponding payroll tax stream Θ_t^P is easily calculated.

Next, we turn to income taxes. We explicitly assume that contributions to the brokerage account are made with after-tax dollars, and taxes on gains in the brokerage account are paid annually along with any taxes owed on income. Further, contributions to the tax-deferred account are made with pre-tax dollars, and withdrawals are taxed as income. Under these assumptions, the total taxable income χ_t at time t is

---

[4] See https://www.irs.gov/retirement-plans/plan-participant-employee/retirement-topics-required-minimum-distributions-rmds.

$$\chi_t = B_t - \Delta_t^D + \left(R_t^{B,save} - 1 + \frac{I_t}{1+I_t}\right) u_{t-1} \tag{8a}$$

When an individual is working, all income is from wages ($B_t = W_t$), and contributions to the tax-deferred account ($\Delta_t^D \geq 0$) reduce taxable income. Since contributions may not exceed wages ($\Delta_t^D \leq W_t$), taxable wage income is never negative ($W_t - \Delta_t^D \geq 0$). During retirement, withdrawals ($\Delta_t^D \leq 0$) and benefits add to taxable income. The term in parentheses is the rate at which taxable gains accrue and is multiplied by u_{t-1} , savings at the beginning of the period $[t-1, t]$. The net real gain on brokerage savings over this interval is $R_t^{B,save} - 1$; however, taxes are levied on nominal returns, and this gain must be adjusted for inflation. The adjustment for inflation is $\frac{I_t}{1+I_t}$, where I_t is the inflation rate (CPI-U) over the interval. Basically, inflation is taxed (the numerator I_t), but taxes are paid with inflated dollars (the denominator $1 + I_t$). There is no adjustment to the borrowing rate; for this model, we assume that borrowing has no tax implications.

As an approximation, consider a simple single-bracket tax model[5]

$$\Theta_t^I = \tau \max[0, \chi_t - \mathcal{T}_t]$$

where $\mathcal{T}_t$ is the bracket's lower threshold, τ is the tax-rate on income over the threshold, and χ_t is total taxable income and gains at τ. Here, τ is an individual's marginal tax-rate, and $\mathcal{T}_t$ is a "deduction", which can be chosen to match an individual's average tax-rate. The threshold $\mathcal{T}_t$ is measured in real dollars, and may vary from year to year. For example, we index ordinary income to CPI-U, the index used to adjust contribution limits, whereas the IRS indexes tax-brackets to chained CPI (C-CPI-U), and in this case, the threshold $\mathcal{T}_{tt}$ shrinks annually by approximately 0.25%.

Taxes are a piecewise-linear function of an individual's taxable income. For our lifecycle optimization, we do not explicitly use the tax formula; rather, we get the same result by introducing the linear inequality constraints

---

[5] The single-bracket model is easily generalized to a multi-bracket model with multiple thresholds and tax-rates.

$$\Theta_t^I \geq 0 \tag{8b}$$

$$\Theta_t^I \geq \tau(\chi_t - \mathcal{T}_t). \tag{8c}$$

Since less tax means more spending, an optimization attempts to minimize them, and additional conditions are usually redundant. Again, we optimize without additional conditions and subsequently verify that the result we obtain satisfies the tax formula.

III. Calibration

a. Wage Profiles

We use data from the Social Security 2006 Earnings Public Use File, which includes the full covered earnings histories of a 1 percent random sample of all individuals with Social Security numbers as of January 1, 2007. As the dataset is derived from the Social Security Administration's administrative data, the only demographic information we have is gender. We drop all individuals with unspecified gender. We restrict the sample to individuals whose earnings and quarters of coverage are observed between ages 25 and 50. Earnings prior to 1951 are aggregated, and quarters of coverage are not available for 1951 and 1952. Thus, we restrict the sample to those born between 1928 (who turn 25 in 1953) and 1956 (who turn 50 in 2006). We further restrict the sample to individuals who have four quarters of coverage at age 25. Earnings are converted to constant 2020 dollars using the CPI-U.

For each individual in our sample, we calculate the ratio of earnings at age 45 to earnings at age 25. The distribution of these ratios is summarized in Table 1. Significant fractions of both men and women have zero earnings at age 45. The median male's earnings at age 45 are 1.38 times his earnings at age 25. The median female's age 45 earnings are 1.05 times her age 25 earnings. Large fractions of both men and women experience substantial earnings growth. About 30 percent of men have age 45 earnings that are at least 1.85 times their age 25 earnings. About 30 percent of women have age 45 earnings that are at least 1.62 times their age 25 earnings.

Table 1: Age 45 Earnings / Age 25 Earnings

Percentile	Male	Female
10	0.00	0.00
20	0.16	0.00
30	0.73	0.27
40	1.09	0.70
50	1.38	1.05
60	1.64	1.33
70	1.85	1.62
80	2.19	2.08
90	2.77	2.77

We construct stylized earnings profiles for workers whose earnings at age 45 are a multiple, q, or their earnings at age 25. We assume that earnings increase linearly between age 25 and age 45 and remain constant thereafter. This parameterization allows us to solve for the values of q, given a particular level of initial earnings, for which it is optimal to save during ones 20s or 30s. We can then use the earnings data to estimate the fraction of the population that falls into this category.

These stylized profiles are a reasonable approximation to the earnings patterns of higher earners. To illustrate this point, we define lifetime earnings for individuals in our sample as average real earnings over ages 25-50. We then divide the sample into the top and bottom halves for each gender by lifetime earnings. Let $e_{a,i}$ denote individual i's earnings at age a. For males in the top half of the lifetime earnings distribution, the square markers in Figure 1 show the median, across individuals i, of $e_{a,i}/e_{25,i}$. The line shows our stylized earnings profile, calibrated such that the stylized ratio is equal to the observed median at age 45. The triangle markers show the ratio, at each a, of the median of $e_{a,i}$ to the median of $e_{25,i}$. These ratios of medians turn out to be fairly similar to the medians of the ratios

$e_{a,i}/e_{25,i}$. Figure 2 presents the same data for women in the top half of the lifetime earnings distribution. We obtain a similarly reasonable fit for individuals who have four quarters of coverage in every year between ages 25 and 50 (results are available upon request).

Figure 1: Current Earnings / Age 25 Earnings, Males in Top Half of Lifetime Income Distribution

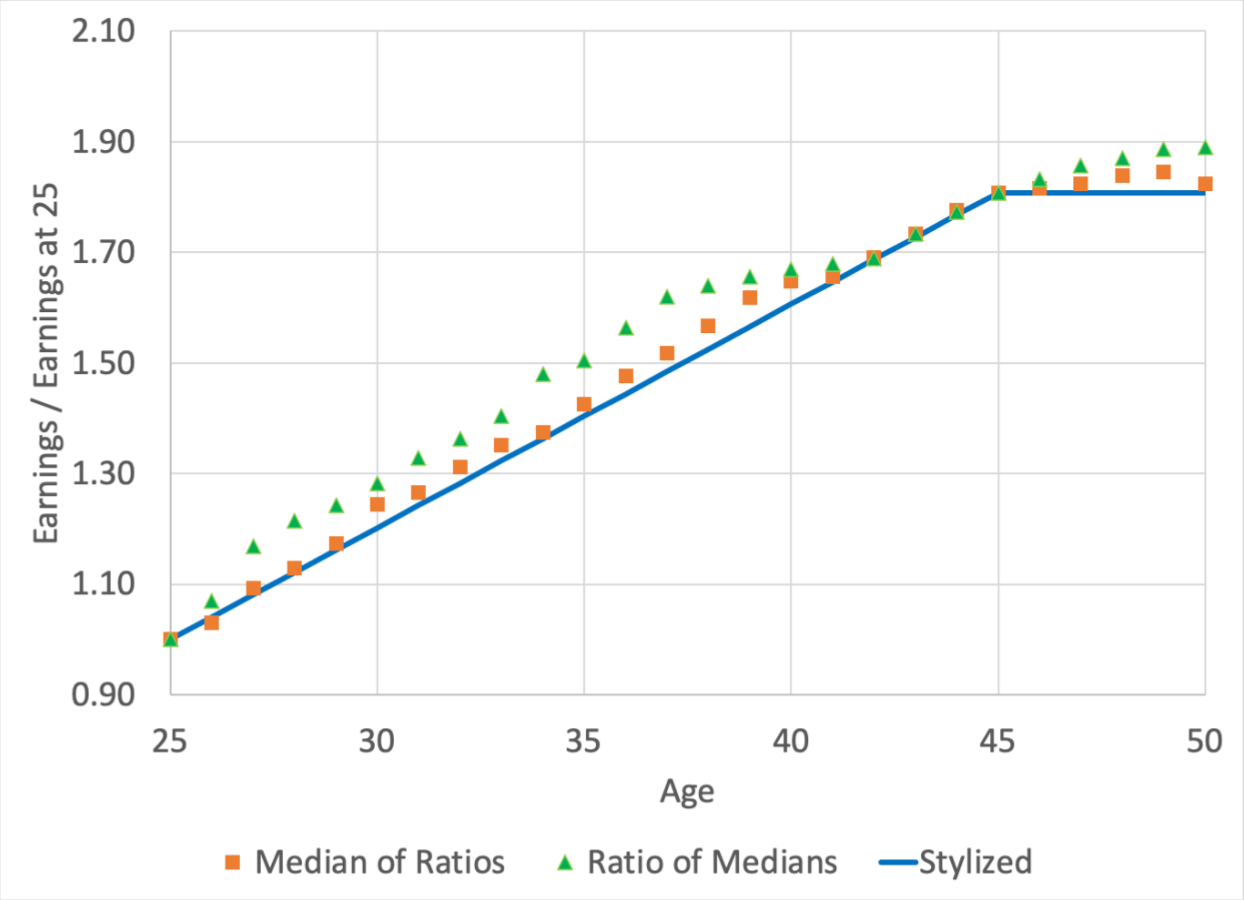

Figure 2: Current Earnings / Age 25 Earnings, Females in Top Half of Lifetime Income Distribution

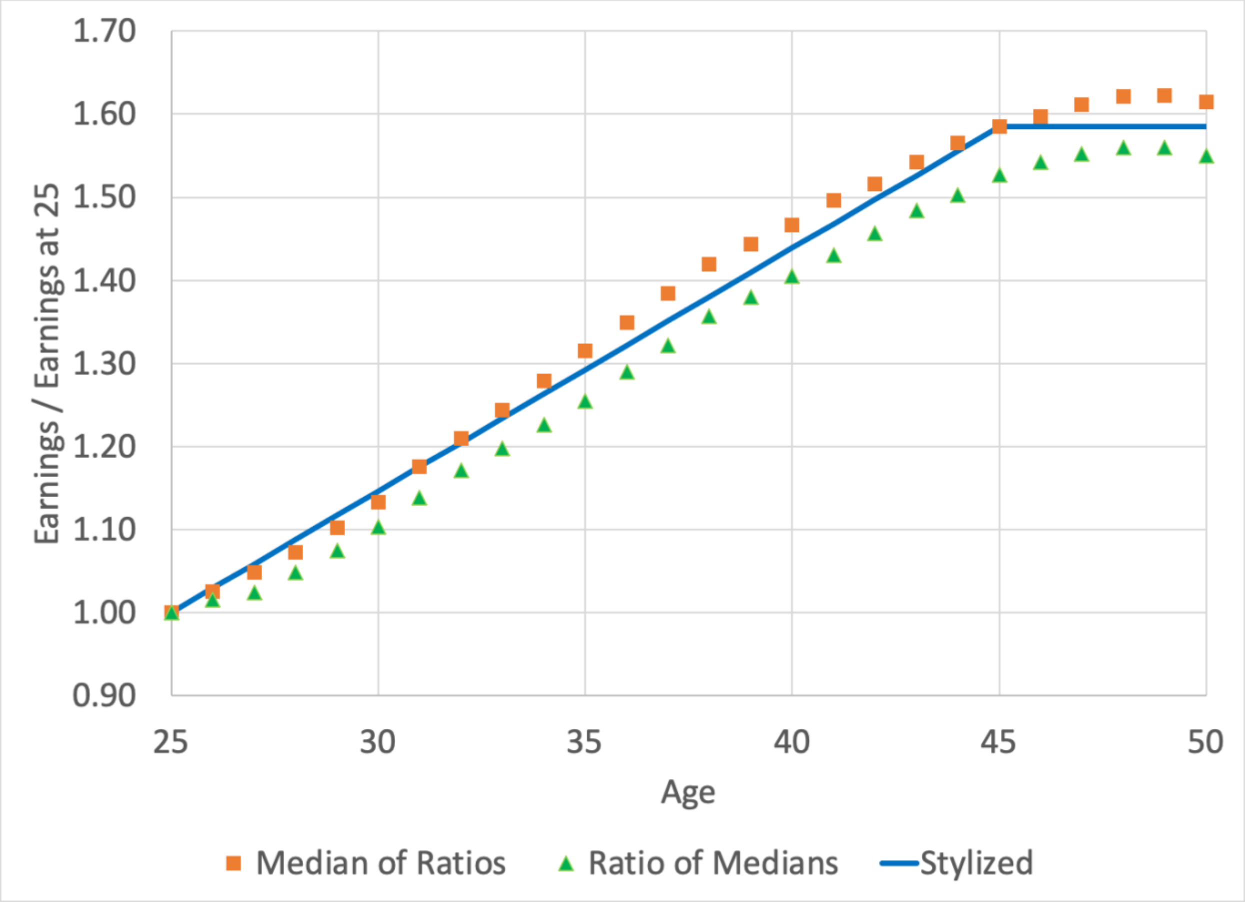

The median level of earnings at age 25 is $41,488 for men in the top half of the lifetime income distribution, and $33,358 for women in the top half of the lifetime income distribution.[6] We set initial earnings equal to these amounts in our simulations.

b. *Mortality and Preference Parameters*

Mortality rates come from the cohort life tables underlying the 2013 Social Security Trustees Report. We use tables for men and women born in 1990.

IV. Results

a. Higher-Income Individuals

---

[6] These medians are similar for men and women with four quarters of earnings in all years.

We assume higher-income individuals face a wage profile that increases linearly between age 25 and age 45, then becomes flat, in line with Figures 1 and 2. We allow the ratio of wages at age 45 to wages at age 25 to vary between 1 and 2, simulating wage profiles of different steepness. We also consider starting wages of $50,000 and $100,000 for men, and $40,000 and $80,000 for women.

Table 2 shows the age at which saving commences for different values of the age-45 to age-25 wage ratio (first column) when the real interest rate is 3 percent and equal to the rate of time preference. In this case, only mortality discounting pushes towards a declining consumption profile, primarily at older ages. Table 2 suggests that when the wage profile is relatively flat (with the age-45 wage 1 or 1.1 times the age-25 wage), retirement saving begins immediately or almost immediately. However, when the wage profile has even a modest slope, with the age-45 wage equal to 1.4 times the age-25 wage, saving does not begin until the early 30s. Table 1 suggests that roughly half of male workers aged 25 have a ratio at or above this level. Thus, even with a relatively flat consumption profile, higher-income individuals in their 20s should not save for retirement.

Table 2: Age of Saving Onset, r = 3 percent

Wage at 45 / Wage at 25	Female $40k Starting Salary	Male $50k Starting Salary	Female $80k Starting Salary	Male $100k Starting Salary
1	25	25	25	25
1.1	28	29	25	25
1.2	32	32	27	26
1.3	33	33	29	29
1.4	34	34	31	31
1.5	34	35	32	32
1.6	35	35	33	33
1.7	36	36	34	34
1.8	36	36	34	34
1.9	36	37	35	35
2	37	37	35	35

Table 3 assumes a more realistic zero percent real interest rate. In this case, saving does not begin until the late 30s (for a starting salary of $100,000 for men) or early-to-mid 40s (for the three other starting salaries). Moreover, the age at which saving begins is relatively insensitive to the slope of the wage profile. Even with a flat wage profile, individuals in their 20s and early 30s should not save.

Table 3: Age of Saving Onset, r = 0 percent

Wage at 45 / Wage at 25	Female $40k Starting Salary	Male $50k Starting Salary	Female $80k Starting Salary	Male $100k Starting Salary
1	43	45	41	39
1.1	43	45	41	39
1.2	43	44	41	39
1.3	43	44	40	39
1.4	43	44	40	39
1.5	43	44	40	39
1.6	43	43	40	39
1.7	43	43	40	39
1.8	43	43	40	39
1.9	43	43	40	39
2	43	43	40	39

b. Lower-Income Individuals

Our simulation of lower-income individuals differs in several important ways. First, we assume a flat wage profile. Second, we assume lower starting salaries. We consider women with constant real wages of $16,000 and men with constant real wages of $27,000. (Gender affects the results only via mortality rates.) These assumptions result in a higher Social Security replacement rate and a lower likelihood that Social Security benefits will be taxed. While we assume that the standard employer match applies, that is less realistic for lower-income workers who often do not have access to defined contribution plans with matching. Thus, we consider scenarios with and without employer matching.

The fourth third column of Table 3 shows the 401(k) balance at retirement with employer matching under different assumptions for wages and interest rates. With an interest rate of 3 percent, saving begins immediately but remains modest. Women accumulate a 401(k) balance that is well below their annual wage. Men accumulate a 401(k) balance that is only 3.3 times their annual wage. When interest rates are zero, implying that declining consumption is optimal, saving does not begin until age 58 for men earning $27,000; women earning $16,000 never begin to save. Moreover, balances at retirement are even lower – less than half the annual wage for men and close to zero for women. As lower-income workers often do not have access to employer-sponsored retirement plans with matching contributions, we perform simulations for these workers without the employer match. The results are shown in the last column of Table 3. Without employer matching, the maximum 401(k) balance falls in all cases and is near zero for female workers with annual earnings of $16,000. With zero real interest rates, any saving commences quite late in life if at all.

Table 3: 401(k) Balance at Retirement for Lower-Income Workers

Interest Rate	Salary	Gender	401k Balance at retirement with Match	401k Balance at retirement without Match
0%	$16,000	F	$1,260	$0
3%	$16,000	F	$6,075	$0
0%	$27,000	M	$28,169	$11,746
3%	$27,000	M	$87,973	$39,666

The bottom line is that under realistic interest rate assumptions, retirement saving commences much later, if at all, for those in the bottom half of the income distribution even if wage profiles are flat. Total optimal retirement savings is low, and exhausting retirement savings early in retirement may be optimal. Shifting money from work life to retirement is not very attractive with a relatively high Social Security replacement rate and zero real interest rate. Our simulations assume that low-income and high-income workers face average mortality rates for their gender. Allowing lower-income workers to (realistically) face higher mortality rates would likely strengthen our findings.

The relatively generous Social Security replacement rate plays a key role in our finding that lower-income workers find it optimal to save very little for retirement. Indeed, we find that a worker with a low, constant real wage receives a replacement rate that exceeds the real wage. Social Security benefits are based on an average of the highest 35 years of earnings, indexed to economy-wide wage growth. A progressive benefit formula gives low-wage workers a benefit that can be as high as 90 percent of this average. (This percentage declines as average indexed wages rises.) A low-income worker with constant real wages has nominal wages that grow in line with the economy-wide average price level. Because wages grow faster than prices, when wages are indexed to economy-wide wage growth to compute the Social Security benefit, the resulting benefit may be greater than the real wage. That is, low-income individuals whose wages grow slower than the economy-wide average can have replacement rates greater than 100 percent. This outcome is related to a point we made in earlier work (Scott et al. 2020): by indexing individual wages to economy-wide wage growth, the benefit formula provides insurance against stagnant individual wages. Figure 1 shows replacement rates for different levels of constant real income (in 2020 dollars). Our simulated workers therefore have replacement rates of either 79 percent (with a constant income of $27,000) or 105 percent (with a constant income of $16,000).

Figure 1: Social Security Replacement Rates for Different (Constant) Real Incomes

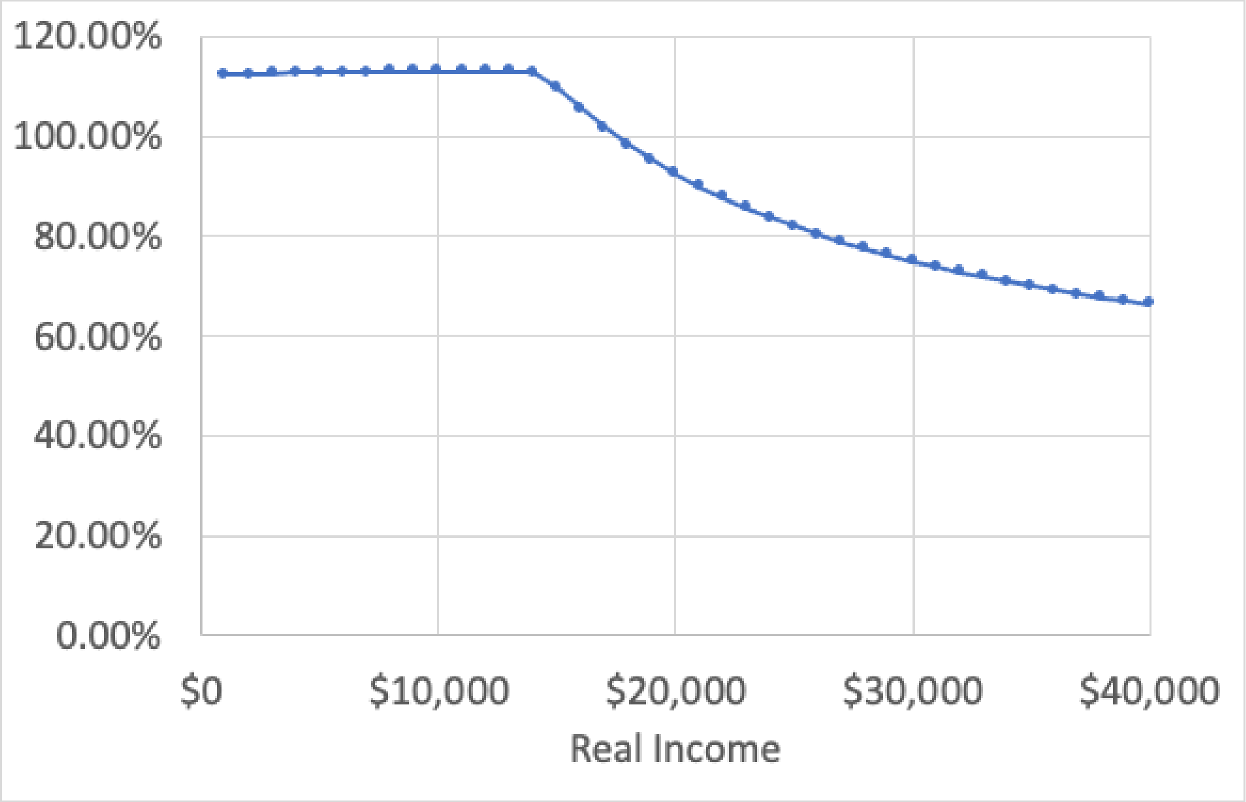

V. Discussion

a. Standard Arguments for Retirement Saving

A standard argument for participating in an employer-sponsored retirement plan is that the employer match is essentially "free money." But this argument doesn't apply to younger workers who are liquidity constrained. One implication of our findings is that the age of retirement saving commencement is not terribly sensitive to the 401(k) match rate. Our results suggest that standard employer matching does not change the optimal (zero) saving of younger, lower-income workers. In our earlier work, we show that inducing higher-income workers in their 20s to save would require match rates of 500 percent or more (Scott et al. 2021b). Intuitively, younger workers are at a corner solution due to borrowing constraints. Ideally, they would like to borrow against future income; however, they are constrained to hold nonnegative assets. While a higher match rate increases the attractiveness of saving for retirement, younger workers are not on the margin of making that decision and are therefore insensitive to the terms of the match.

Another standard argument for starting to save for retirement early in life is the power of compound interest. However, there is no power of compound interest when real interest rates are zero. While individuals could invest in risky assets with higher expected returns, those higher returns are merely compensation for taking on the additional risk. What gets compounded over the years is only the risk premium.

b. Leakages and Non-Retirement Saving

The literature indicates that the participation of young workers increases markedly when automatic enrollment is introduced (Madrian and Shea 2001). However, the life cycle model suggests that these high participation rates may not be optimal. Viewed from this perspective, leakages from 401(k) balances for young workers might be interpreted as correcting a mistake rather than a major problem in need of further government policy.

Our simple life cycle model doesn't include any uncertainty beyond mortality risk. In a richer model, establishing a precautionary fund (a rainy-day fund) would almost certainly be optimal. Similarly, another saving goal for young people may be accumulating a down payment on a house. In the presence of a 50 percent or better match and a 10 percent withdrawal penalty, using a 401(k) as a rainy-day fund or for another savings goal may make sense. Our point is that many, probably most, young people should not be saving for retirement. However, using these plans for other purposes may be optimal.

c. Nondiscrimination Test

Employers that offer defined contribution plans face nondiscrimination tests, which impose restrictions on the contributions of highly compensated employees relative to other employees. In 2020 and 2021, a highly compensated employee is defined as an individual earning more than $130,000. Nondiscrimination tests are intended to ensure that the benefits of tax-preferred retirement saving are not restricted to high-income individuals and are instead enjoyed more broadly. However, there are likely to be firms (such as law firms) where most employees have steep earnings profiles that put them in the highly compensated group at older ages but not at younger ages. At such firms, nondiscrimination tests would force both younger and older workers to participate in the defined contribution plan. The life cycle model suggests that this outcome is suboptimal.

d. Additional Factors

Our model is clearly simplified. Some factors that we abstract away from would push in the direction of lower saving for retirement. For example, our model also only includes discounting for mortality and pure time preference. Work by Goda et al. (2015) suggests that individuals also heavily discount consumption in future states in which their health is poor. Adding health discounting to the model would push against saving for the future strengthen our conclusions. Our model further ignores home production. Aguilar and Hurst (2005) document that while food expenditure declines upon retirement, food consumption does not, as home production is substituted for market purchases. Adding home production to the model would further reduce the need to save for retirement. The additional time available during retirement could allow individuals to spend less during retirement while maintaining consumption through home production. However, other factors that we abstract away from may push in the direction of higher saving for retirement at younger ages. For example, our model does not account for uncertainty about future wages. If the wage profile is uncertain, individuals may wish to begin saving for retirement earlier in life in case future wages do not grow as expected.

Our model also ignores marriage and childbearing decisions. All our stylized individuals are assumed to be single. However, we would anticipate obtaining similar results for couples. Adding children to the model may change saving needs (e.g., saving for college), but it should not alter our conclusions regarding saving at younger ages. The average age for having a first child is in the mid-20s for women and late-20s for men (Stahl 2020). If individuals have children in their 20s, increased consumption needs during those years should push against saving for retirement. On the other hand, it may be optimal for individuals to save for retirement before having children.

While a richer model may alter the level and age profile of saving, the basic point remains. Arguing that individuals undersave for retirement requires providing a benchmark for optimal saving, such as the life cycle model. However, retirement policy that is premised on the assumption that more saving is always better may in fact result in an age profile of saving that is inconsistent with a life cycle model. Policy makers designing interventions to boost retirement saving, such as auto enrollment or nondiscrimination rules, need to carefully consider the optimal level and age profile of spending that they are implicitly using as a benchmark.

References

Aguilar, Mark and Erik Hurst, “Consumption versus Expenditure,” Journal of Political Economy, 113(5), October 2005, pp. 919-948.

Bodie, Zvi, Jonathan Treussard, and Paul Willen, “The Theory of Life-Cycle Saving and Investing,” Federal Reserve Bank of Boston Research Department Public Policy Discussion Paper No. 07-3, 2007.

Goda, Gopi Shah, Matthew R. Levy, Colleen Flaherty Manchester, Aaron Sojourner, and Joshua Tasoff, “The Role of Time Preferences and Exponential-Growth Bias in Retirement Savings.” NBER WP 21482, August 2015.

Gomes, Francisco, Kenton Hoyem, Wei Hu and Enrichetta Ravina, “Retirement Savings Adequacy in United States Defined Contribution Plans,” Kellogg School of Management Working Paper, November 2018.

Hamermesh, Daniel S., “Consumption during Retirement: The Missing Link in the Life Cycle,” *Review of Economics and Statistics*, Volume 66, No. 1, 1984, pp. 1-7.

Hurd, Michael D., “Mortality Risk and Bequests.” Econometrica, 57(4), July 1989, pp. 779–813.

Kotlikoff, Laurence J., “Economists’ Approach to Financial Planning,” November 2007, retrieved from: http://people.bu.edu/kotlikof/New%20Kotlikoff%20Web%20Page/Economics%20Approach%20to%20Fin%20Planning,%20JFP11-08-07%20posted%20Jan%203,%202008.pdf.

Lachance, Marie-Eve, “Optimal Onset and Exhaustion of Retirement Savings in a Life-Cycle Model,” Journal of Pension Economics and Finance, 11(1), 2012, pp. 21.

Leung, Siu Fai, “Uncertain Lifetime, the Theory of the Consumer, and the Life Cycle Hypothesis,” Econometrica, 62(5), 1994, 1233–1239.

Leung, Siu Fai, “The Life-Cycle Model of Saving with Uncertain Lifetime and Borrowing Con- straint: Characterization and Sensitivity Analysis,” Mathematical Social Sciences, 42(2), 2001, pp. 179–201.

Leung, Siu Fai, “The Existence, Eniqueness, and Optimality of the Terminal Wealth Depletion Time in Life-Cycle Models of Saving under Uncertain Lifetime and Borrowing Constraint,” Journal of Economic Theory, 134(1), 2007, pp. 470–493.

Madrian Bridgette C., and Dennish F. Shea, “The Power of Suggestion: Intertia in 401(k) Participation and Savings Behavior,” Quarterly Journal of Economics, 116(4), November 2001, pp. 1149-1187.

Milevsky, Moshe A. and Huaxiong Huang, “Spending Retirement on Planet Vulcan: The Impact of Longevity Risk Aversion on Optimal Withdrawal Rates.” Financial Analysts Journal, 67(2), 2011, pp. 45–58, 2011.

Munnell, Alicia H. And Anthony Webb, "The Impact of Leakages from 401(k)s and IRAs," Center for Retirement Research Working Paper 2015-2, February 2015.

Scholz, John Karl and Ananth Seshadri, "What Replacement Rates Should Households Use?" Michigan Retirement Research Center Research Paper No. 2009-214, November 2009.

Scott, Jason, John B. Shoven, Sita N. Slavov, and John G. Watson, "Retirement Implications of a Low Wage Growth, Low Real Interest Rate Economy" *Social Security Bulletin*, 80(3), pp. 31-45, August 2020.

Scott, Jason, John B. Shoven, Sita N. Slavov, and John G. Watson, "Can Low Retirement Savings Be Rationalized?" *Journal of Retirement*, 8(4), pp. 7-25, Spring 2021.

Scott, Jason, John B. Shoven, Sita N. Slavov, and John G. Watson, "Is Automatic Enrollment Consistent with a Life Cycle Model?" *Economic Inquiry*, forthcoming.

Tergesen, Anne, "The Rising Retirement Perils of 401(k) 'Leakage'," Wall Street Journal, April 2, 2017.

VanDerhei, Jack, "The Impact of Auto Portability on Preserving Retirement Savings Currently Lost to 401(k) Cashout Leakage," Employee Benefit Research Institute Issue Brief No. 489, August 15, 2019.

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