Welfare Effects of Automatic-IRAs

Erin Cottle Hunt*
Lafayette College†

October 21, 2022

Abstract

Several states and municipalities in the United States require employers who do not offer retirement benefits to automatically enroll their employees in state-sponsored individual retirement accounts (IRAs). We quantify the welfare effects of participating in an automatic IRA for individuals who follow a rule-of-thumb to make savings decisions. We find college graduates who saves more than 5% and high school drop-outs/graduates who saves more than 3% would be willing to give up 0-1.4% of lifetime consumption to *avoid* being enrolled in an IRA with a 3% contribution rate. Workers who save at lower rates, or who do not save at all, are made better off by being enrolled in an IRA by up to 4.2% of lifetime consumption. We consider several extensions including a fully rational model with credit market frictions, stochastic income and pre-retirement withdrawals, and costly IRA opt-outs, and find that being enrolled in an IRA is most likely to benefit workers who do not want to borrow during the life-cycle.

JEL Classifications: D14, D15, H75, J26, J32, G51

*cottlee@lafayette.edu

†Special thanks to undergraduate research assistants Pooja Kumar and Joseph Polifonte (Lafayette College), and to Elina Huang (Barnard College). Joe revived research funding through the Lafayette EXCEL student research program.

1 Introduction

Several states and municipalities in the United States require employers who do not offer retirement benefits to automatically enroll their employees in state-sponsored individual retirement accounts (IRAs). These state-run programs establish a means for employees to contribute to retirement accounts that otherwise may not have been easily accessible to them. As of 2020, six states and the city of Seattle have adopted auto-enrollment IRA programs as a state-facilitated effort to assist retirement savings among private sector workers.[1]

Automatic IRAs are often motivated as a way to help workers increase their savings for retirement and improve their lifetime well-being. However, from a theoretical perspective it's not clear that automatic enrollment in an IRA increases total life-cycle saving or improves lifetime well-being. If there are no credit market frictions, workers are fully rational, and the IRA earns the same rate of return as private savings automatic enrollment in an IRA has *no* effect on total retirement saving, and thus no effect on lifetime well-being.[2] For automatic enrollment to change total savings and have any effect on lifetime well-being, one of the assumptions above needs to be relaxed, for example by introducing a behavioral friction that results in under-saving (in the absence of the IRA), or by introducing credit market frictions. We explore both possibilities, as well as a handful of other extensions, in this paper.

As our baseline, we quantify the welfare effects of being automatically enrolled in an IRA in a life-cycle model with credit market frictions for an individual who makes savings decisions

[1] Programs in Oregon, California, Illinois are notable early examples of state-run IRA programs. See https://www.oregonsaves.com, https://www.calsavers.com, https://www.ilsecurechoice.com.

[2] This is the case if IRA contributions are completely offset by decreases in other types of savings and/or increases in debt. In the absence of credit market frictions, this is the prediction of neoclassical model as well as a model with present bias or quasi-hyperbolic discounting (Findley and Cottle Hunt (2019)).

following a simple rule-of-thumb. We model three income profiles based on educational attainment. As a starting point, we assume that workers who are automatically enrolled in the IRA do not opt-out; we also consider an extension in which workers faces a monetary cost to opt-out (which proxies for the cognitive, financial, time, or other costs associated with opting out a default). Our assumption that workers who are enrolled in the IRA continue participating in the retirement account is consistent with the behavior of the majority of workers enrolled in such programs.[3]

We find that being enrolled in the IRA increases the lifetime well-being of workers who do not save very much. As an upper bound, we consider workers who do not save at all privately (their rule-of-thumb savings rate is zero). We find that college graduates who do not save privately would be made better off by the equivalent of 4.2% of lifetime consumption if they are enrolled in an IRA that contributes 3% of their after-tax income. The welfare gain is smaller but still positive for high school graduates (2.2%) and drop-outs (2.2%) who do not save privately because their life-cycle disposal income profiles are flatter (i.e., their after tax income Social Security benefit).

The welfare gain of participating in an IRA is decreasing in the rule-of-thumb savings rate. Above a threshold savings rate, participating in the IRA reduces lifetime well-being because the individual saves so much that their life-cycle consumption is titled heavily towards the retirement years. We find that college workers who save more than 5% privately and high school graduates and drop-outs who save more than 3% privately are made worse off if they are enrolled in an IRA. The average personal savings rate in the U.S. is around 5%, so

[3]According to Oregon's annual report to the legislature in 2018, 72.25% of workers stay enrolled in the program. Similarly Belbase and Sanzenbacher (2018) find that 29% of plan participants opt-out and 4% set the contribution rate to zero in the OregonSaves program. The Center for Retirement Research (2015) estimates that the opt-out rate for the Connecticut Retirement Security Program will likely be around 19%.

this suggests that individuals could be harmed by enrolled in an IRA if they use a simple rule-of-thumb to make their saving-consumption decisions.

Our baseline model assumes that individuals follow a simple rule-of-thumb to make savings decisions. We relax this assumption in an extension and consider instead a fully rational individual who makes optimal saving consumption decisions in the presence of credit market frictions. In this setting, we find that workers of all three education levels are made worse off by participating in the IRA. College graduates are harmed the most and would be willing to give up 0.7% of lifetime consumption to avoid being enrolled in (and remaining in) an IRA that contributes 3% of their after-tax wage to an IRA. The welfare cost for high school graduates and drop-outs are 0.4% and 0.2%, respectively. Participating in the IRA harms workers in this setting because they would prefer to borrow (or save less than 3%) during their early working years. Credit market imperfections make it costly for the individual to offset their IRA contributions to achieve their desired total saving or borrowing rate. Thus, although the IRA offers a small tax advantage, workers would prefer not to be enrolled.

Our baseline analysis assumes that workers who are enrolled in the IRA continue to participate. Of course, in reality, a worker could opt-out of an automatic IRA. This would require some financial knowledge, cognitive effort, and time to fill out the appropriate paperwork (or on-line form) to make the change. We consider a model extension that incorporates these opt-out costs as a one-time monetary cost. Including a monetary opt-out cost does not change the welfare metrics for individuals for whom being enrolled in the IRA improves lifetime well-being. This is because the individuals would not choose to opt-out, even if given the opportunity, so considering a cost to opting out does not change the welfare calculation.

Including this opt-out cost has quantitative impacts on our welfare metrics in cases where individuals are harmed by the IRA, but does not qualitatively change our results. For example, a college graduate who saves more than 5% privately would prefer to pay the opt-out cost to avoid participating in the IRA. Paying this opt-out cost reduces their lifetime well-being by a small amount (around 0.1% of lifetime consumption). In our baseline analysis, we assumed these individuals continued participating in the IRA, which also reduced their lifetime well-being (0.5-1.1% of lifetime consumption).

As a final extension, we consider an individual who faces a risky income stream (stochastic periods of unemployment) who can make pre-retirement withdrawals from the IRA. Early empirical evidence from OregonSaves found 20% of individuals made pre-retirement withdrawals (Quinby et al. (2019)). Having an IRA could potentially make workers better off if they experience a negative income shock (unemployment) and they either do not have sufficient private savings to cover their needs and/or it is costly to borrow. In our dynamic stochastic life-cycle model with risky income, we find that the welfare effect of being enrolled in an IRA and being able to make pre-retirement withdrawals to be positive and equivalent to 0.85% of lifetime consumption. However, the majority of the welfare gain (over 90%) is driven by the higher rate of return of the tax-preferred IRA rather than from the ability to make pre-retirement withdrawals. When the income shocks are turned off, the individual doesn't want to borrow, and the only welfare effect of the IRA is to increase total lifetime resources through a higher rate of return.

Considering our baseline results and our model extensions, we find that being enrolled in an IRA is most likely to benefit individuals who either do not save very much on their own, or who do not want to borrow to unwind their IRA contributions at any stage in the life-cycle. Workers who already save for retirement can be made worse off when they are enrolled in an IRA if it is difficult of unwind their automatic contributions (because of a credit wedge, best opting out is costly, or because the individuals follow a simple rule-of-thumb and do not save less privately when they are enrolled in the IRA). Our paper offers theoretically support for Harris et al. (2018) who find that nation-wide automatic IRAs could affect 24.2 million workers aged 25–64. Many of these workers have high-interest debt. Automatic enrollment in IRAs could negatively impact these workers if they are unable to unwind IRA savings (without increasing their high-interest debt) or if they do not actively opt-out.

Our analysis relates to a growing body of papers exploring the effects of automatic enrollment in IRAs and other defaults in retirement saving accounts. Preliminary data from the auto-enrollment program in Oregon show a majority of employees provided access to the program are participating and the average contribution rate is about 5% (Chalmers et al. (2020)). Empirical evidence from other settings also suggests that many households respond to defaults in making savings decisions including Madrian and Shea (2001), Choi et al. (2003), Carroll et al. (2009), DellaVigna (2009), Chetty et al. (2014), Bernheim et al. (2015), Beshears et al. (2016), and DellaVigna (2018), among many others. These researchers generally find opt-out costs to be in the range of $1,000-$3,000.[4] In contrast to these papers, we do not search for optimal default or for the opt-out cost, rather we take the parameters of the state-run IRA program as given, and calculate the welfare effect of staying in the program.

Our analysis also relates to Hurst and Willen (2007), who study the welfare effects of Social Security contributions. They show that Social Security forces young households with

[4]In a rich dynamic model that allows individuals to change their default contribution rate multiple times over the life cycle Choukhmane (2019) finds opt-out costs to be $250 every time an individual changes their contribution rate.

upward-sloping income profiles to save a sizable portion of their income for retirement, when their optimal consumption plan would likely have them save little or borrow. The negative welfare effects in our paper are conceptually similar: workers who are harmed by the enrollment in the IRA program are worse-off because they are forced to save while young when their income is relatively low and they cannot easily unwind that saving by borrowing because of the credit wedge.

Our results suggest that automatic enrollment in an IRA may make some workers worse off by forcing them to save at stages of the life cycle when their wage is low and the marginal utility of consumption is high. If workers face credit constraints, such a higher interest rate on borrowing than the return on savings, it may be too costly for them to unwind the forced saving of the IRA program, and as a result they may end up saving more than they would prefer when young which reduces their overall lifetime utility. Although well-intentioned, automatic enrollment in state-run IRAs may have unintended negative consequences for workers who are already saving optimally, face credit constraints, or are unable or unsure how to opt out of the program. Workers who are not saving at all, or who follow a simple rule of thumb when making savings decisions, have the greatest potential to be helped by automatic enrollment in an IRA.

2 Model

Time is continuous and indexed by t. An individual is born at $t = 0$, retires at date t_R, and dies at $t = T$. The individual receives disposable income $y_1(t)$ during the working life, and $y_2(t)$ during retirement. The individual earns wage income, pays Social Security taxes, and contributes to an IRA during the working years and receives Social Security benefits and an IRA annuity in retirement. Consumption is denoted $c(t)$ and private asset holdings are denoted $k(t)$, which can be positive (saving) or negative (borrowing). The individual is born with no assets $k(0) = 0$ and also dies with no assets $k(T) = 0$.

We allow for the interest rate to depend on the state of asset holdings. If assets are positive and the individual is saving, the after-tax interest rate is r_S. If assets are negative and the individual is borrowing, the interest rate is r_B. If the individual is enrolled in an IRA, the assets invested in the IRA grow rat rate r_{IRA} which is higher than the private return on saving, reflecting the tax advantages of saving in a Roth IRA. We assume $r_s = (1 - \tau_k) r_{IRA}$ where τ_k represents capital gains taxes. We assume that the individual neither opts out of the IRA nor changes the default contribution rate. In our welfare analysis we assume future utility is discounted at the rate ρ and utility is given by the CRRA function:

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

2.1 Social Security Details

Individuals pay Social Security taxes at rate τ on wage income up to a wage cap $w_c < 1$. After retirement, individuals receive Social Security benefits $b(\bar{w})$ that depend on their wage earnings over the last 35-years of working period, which we will denote $\bar{w} = \int_{t_R-35}^{t_R} w(v)dv$. Social Security benefits (PIA) are a piece-wise linear function of earnings that replace 90% of earnings ($\bar{w}$) up to the first bend point, 32% of earnings ($\bar{w}$) between the first and second bend points, and 15% of earnings ($\bar{w}$) between the second and third bend points. Beyond the third bend point, the function is flat. We use a conventional estimate of the bend points relative to average wages, $0.2\mathbb{E}(w)$, $1.24\mathbb{E}(w)$, and $2.47\mathbb{E}(w)$ (as in, e.g., Alonso-Ortiz (2014)).[5]

2.2 IRA Details

The individual has two types of saving assets: a Roth IRA, $A(t)$, that grows at the tax-free rate r_{IRA}, and a private saving or borrowing account, $k(t)$. If enrolled in an IRA, the individual contributes after tax income $\hat{w}(t) = w(t) - \min[\tau w(t), \tau w_c]$ to the IRA at rate δ, up to a maximum contribution amount m.

During the working-life the IRA evolves according to the law of motion:

$$\dot{A} = \min[\delta \hat{w}(t), m] + r_{IRA}A(t) \qquad \text{for } t \in [t, t_R]. \tag{2}$$

The IRA is annuitized at retirement. Following retirement, the IRA evolves according to the law of motion:

$$\dot{A} = r_{IRA}A(t) - a \qquad \text{for } t \in [t_R, T], \tag{3}$$

where a is the constant annuity that exhausts the IRA balance at the date of death. The equation for a is

$$a(w) = \frac{\min[\delta \hat{w}(t), m] \int_0^{t_R} e^{-r_{IRA}v} dv}{\int_{t_R}^{T} e^{-r_{IRA}v} dv}. \tag{4}$$

[5]Technically, US Social Security benefits depend on the highest 35 years of earnings, rather than the last 35 years. We assume that wages are hump-shaped with peak earnings in middle age and lowest earnings early in life. Thus, the basing benefits on the last 35 years of earnings is close to the true US system.

2.3 Income Details

Given the assumptions above, after-tax, after IRA contribution, take-home pay is given by:

$$y_1(t) = (w(t) - \min[\tau w(t), \tau w_c]) - \min[\delta(w(t) - \min[\tau w(t), \tau w_c]), m] \tag{5}$$

$$y_2(t) = a + b(\bar{w}). \tag{6}$$

In our numerical analysis, we will consider three different wage profiles $w(t)$ to correspond to three different education levels.

2.4 Rule-of-thumb saver

To being our analysis, we assume individuals save a constant fraction s of their disposable income every instant as in Deaton (1992), Browning and Crossley (2001), and Winter et al. (2012).[6]

We find the rule-of-thumb model to be a compelling benchmark for our analysis since it does not assume individuals solve a complex optimization problem to make their saving consumption decision. Rather, the individual follows a simple rule of thumb, saving a constant fraction of their income every period. This simple assumption is consistent with some advice given by financial planners and it is also broadly consistent with the vast empirical literature on financial literacy that suggests that many households do not save optimally for retirement (see Lusardi (2002), Lusardi and Mitchell (2008), and Lusardi et al. (2017) among many others). The simple rule-of-thumb model is also conceptually consistent with

[6] Winter et al. (2012) also considers more complicated rules-of-thumb for saving behavior. Our analysis is closest to their Rule 1 in which the individual consumes all of their current income.

the policy design—an individual following a simple rule-of-thumb will be influenced by a change in default and will save more for retirement.

The individual saves s of their disposal income in every instant during the working years in a private savings account that is annuitized at retirement. The individual consumes:

$$c(t) = (1-s)y_1(t) \quad \text{for } t \in [0, tR) \tag{7}$$

$$c(t) = S(w) + a(w) + b(\bar{w}) \quad \text{for } t \in [tR, T] \tag{8}$$

where $\hat{w}(t) = w(t) - \min[\tau w(t), \tau w_c]$ is after tax income, $a(w)$ is the IRA annuity payment, $b(\bar{w})$ is the Social Security benefit, and $S(w)$ is the annuity that depletes the private savings account.

The private annuity $S(w)$ is found by solving the differential equations:

$$\dot{k} = r_s k(t) + s y_1(t) \quad \text{for } t \in [0, tR] \tag{9}$$

$$\dot{k} = r_s k(t) - S(w) \quad \text{for } t \in [tR, T], \tag{10}$$

which implies

$$S = \frac{\int_0^{t_R} e^{-r_s v} s(\hat{w}(v) - \min[\delta \hat{w}(v), m])dv}{\int_{t_R}^{T} e^{-r_s v} dv}. \tag{11}$$

Here r_s is the return on private savings. In the limit as the savings rate approaches zero, the model collapses to a simple, hand-to-mouth model in which the individual consumes their income in every period and never saves.

2.5 Parameterization

We parameterize the model to match the US economy. The wage profile is hump-shaped over the life-cycle, according to the functions in Cocco et al. (2005), who use the Panel Study on Income Dynamics (PSID) to estimate income as a function of age and education level.[7] The wage profile for a worker who did not graduate high school is $w_{noHS}(t)$, for a high school graduate is $w_{HS}(t)$, and for a college graduate is $w_{college}(t)$ according to:

$$w_{noHS}(t) = \lambda e^{-2.1361+0.1684(t+20)-0.0353\frac{(t+20)^2}{10}+0.0023\frac{(t+20)^3}{100}} \tag{12}$$

$$w_{HS}(t) = \lambda e^{-2.1700+0.1682(t+20)-0.0323\frac{(t+20)^2}{10}+0.0020\frac{(t+20)^3}{100}} \tag{13}$$

$$w_{college}(t) = \lambda e^{-4.3148+0.3194(t+22)-0.0577\frac{(t+22)^2}{10}+0.0033\frac{(t+22)^3}{100}}. \tag{14}$$

We set the fraction of the economy that are high school dropouts, χ_N, and high school graduates, χ_H, equal to 0.22 and 0.56, respectively, which are the corresponding estimates from Cocco et al. (2005) and also in line with Love and Schmidt (2015). Figure 1 plots the income profiles for all three education levels.

Table 1 summarizes the parameters. We assume agents enter the market at age 22. All workers retire at age 65 ($t_R = 43$) and pass away at age 82 (model age $T = 60$), following Cocco et al. (2005). We set the CRRA parameter $\sigma = 1.6$ (Attanasio et al. (1999)), and we set the discount rate $\rho = r_{IRA}$. As our baseline, we set the private interest rate $r_s = (1 - \tau_k)r_{IRA}$ where $\tau_k = 0.15$ represents capital gains taxes, and $r_{IRA} = 0.04$ is the real

[7] They model three separate third-order polynomials corresponding education groups because Attanasio (1994) and Hubbard et al. (1995) find evidence that that wage hump is not the same for different income levels.

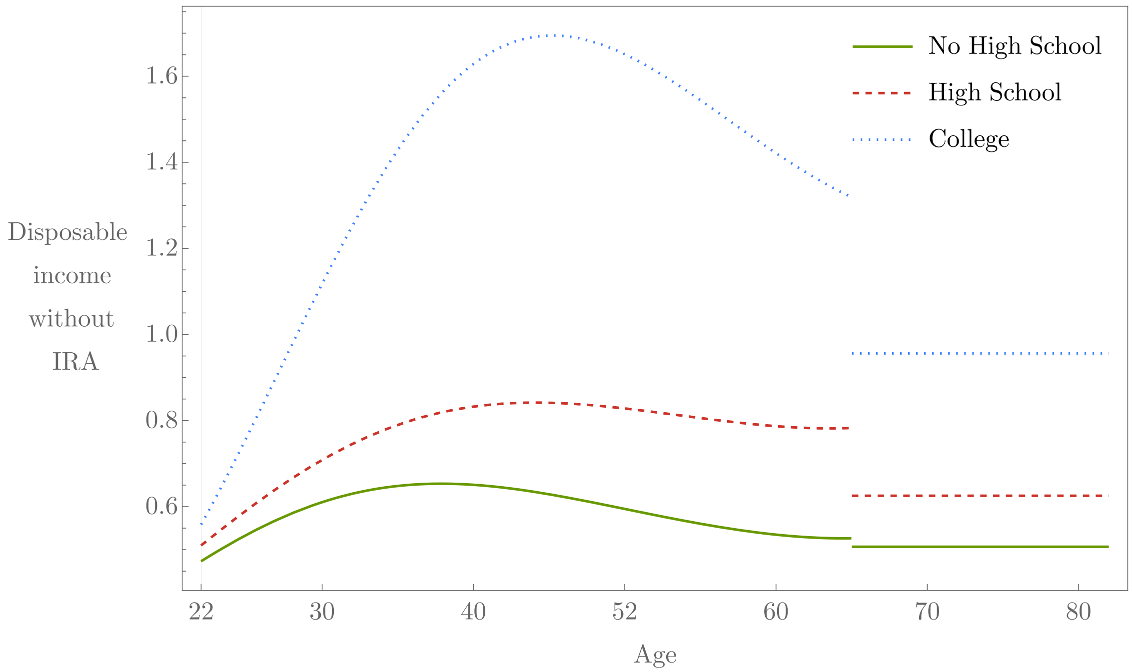

Figure 1: Disposal income without an IRA for the three life-cycle education levels.

interest rate (McGrattan and Prescott (2003)). In our baseline rule-of-thumb model, the household does not borrow. When we relax that assumption, we set the rate on borrowing to be $r_b = r_{IRA} + \tau_b$ where $\tau_b = 0.08$ is the mark-up on borrowing (Livshits et al. (2007), and Davis et al. (2006)).

We parameterize Social Security using the actual benefit earning rule with bends points normalized to median income as in Alonso-Ortiz (2014). The tax rate is set to 10.6% to correspond to the Old Age Survivors portion of Social Security payroll taxes. For most of our analysis, we assume the IRA contribution rate $\delta = 0.03$. Contribution rates vary by state and range between 3% to 5%. We set the contribution max m to the ratio of the the actual contribution limit of \$6,000 and the average wage.[8]

[8]Note that m does not influence our analysis because the contribution rate is relatively small (3%). A worker would have to earn over \$120,000 and contribute 5% of the wage to their IRA in order for the contribution max to be relevant, which is higher than the maximum wage in any of our profiles.

Parameter		Value	Target
λ	Normalization	0.491	Average wage equal to 1
ρ	Discount rate	0.04	equal to interest rate
r_{IRA}	Real interest rate	0.04	McGrattan and Prescott (2003)
τ_b	Mark-up on borrowing	0.08	Livshits et al. (2007)
τ_k	Capital gains tax	0.15	statutory rate
τ_{SS}	Social Security tax rate	0.106	statutory rate
δ	IRA contribution rate	0.03-0.05	rate set by states
σ	Inverse Elasticity of Substitution	1.6	Attanasio et al. (1999)
t_R	Model retirement age (retire at age 65)	43	Cocco et al. (2005)
T	Length of lifetime (die at age 82)	60	Park and Feigenbaum (2018)
χ_N	Fraction of high school dropouts	0.22	Cocco et al. (2005)
χ_H	Fraction of high school graduates	0.56	Cocco et al. (2005)

Table 1: Baseline Parameterization

3 Results

The consumption profiles for the three education groups are plotted in Figure 2 assuming a private savings rate $s = 0.021$, and an IRA contribution rate of $\delta = 0.03$. This implies a combined personal saving rate of 5.1% which is the average personal savings rate in the U.S.[9] Consumption tracks income and is hump-shaped over the working years of the life-cycle. Retirement consumption is higher than working life consumption for high school drop-outs, nearly flat at retirement for high school graduates, and falls at retirement for college graduates. This is partially due to the relative steepness of the college graduate's income profile, and also because the Social Security benefit is progressive.

3.1 Rule of Thumb Welfare Results

We use a consumption equivalent variation technique to measure the welfare effect of being enrolled in an IRA. This welfare metric calculates the fraction of lifetime consumption an

[9]B.E.A. personal savings rate June 2022, accessed at https://www.bea.gov/data/income-saving/personal-saving-rate.

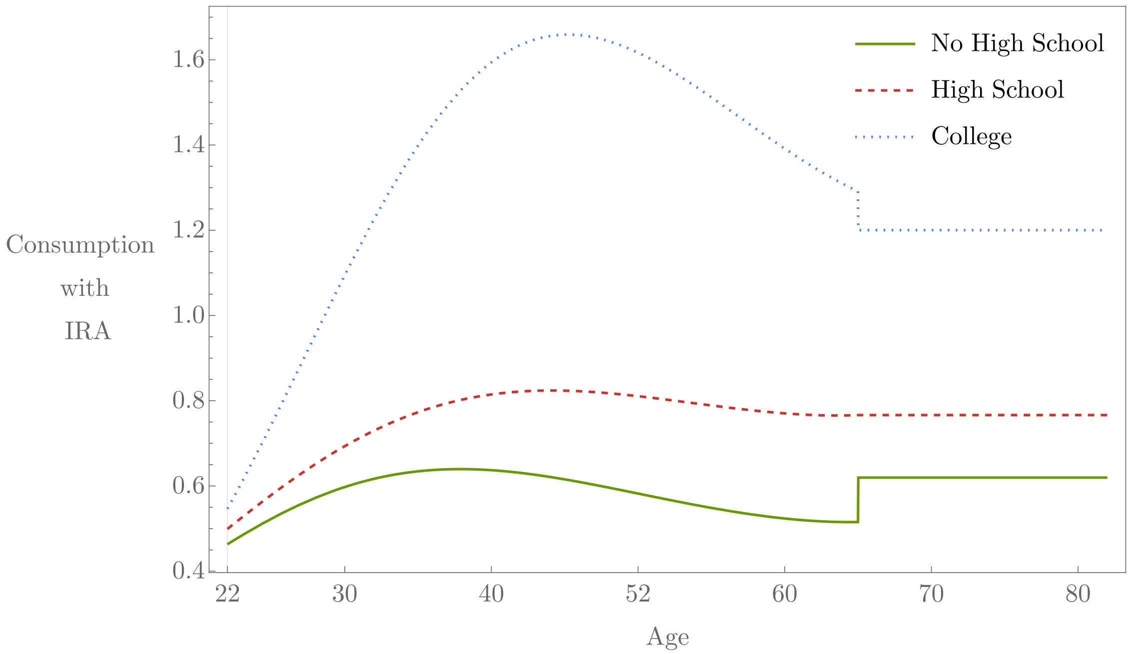

Figure 2: Consumption in a simple rule-of-thumb saving model with $s = 0.021$ given IRA contribution rate $\delta = 0.03$.

individual with an IRA would be willing to give up such that their discounted lifetime utility is equal to their discounted lifetime utility without an IRA. The consumption equivalent variation for a particular education/income group is calculated using the following equation:

$$\int_0^T e^{-\rho t} u(c(t))dt = \int_0^T e^{-\rho t} u(c^{IRA}(t)(1 - \Delta))dt. \tag{15}$$

Here $c(t)$ is consumption without an IRA and $c^{IRA}(t)$ is consumption with an IRA, and the consumption equivalent variation is Δ. Our results are shown in Table 2 and in Figure 3. A positive number indicates that the IRA is welfare enhancing.

The welfare gain of being enrolled in the IRA is decreasing in the individual's private savings rate s. As s approaches zero, the welfare gains increase. For college graduates, the CEV is 0.042 when the IRA contribution rate $\delta = 0.03$ and $s = 0$. College graduates who do

Consumption Equivalent Variation			
s	No High School	High School	College
0	0.021	0.022	0.042
0.02	0.005	0.006	0.017
0.04	-0.004	-0.004	0.003
0.06	-0.010	-0.010	-0.005
0.08	-0.014	-0.014	-0.011

Table 2: Consumption Equivalent Variation in a simple rule of thumb saving model for various savings rates s given IRA contribution rate $\delta = 0.03$.

not save privately, but consume hand-to-mouth, are made better off by the equivalent 4.2% of lifetime consumption. The CEV is smaller for the other two education groups, because their income profiles are flatter and also because Social Security replaces a larger fraction of their income in retirement. The CEV is 0.021 and 0.022 for high school drop-outs and high school graduates who do not save privately (i.e., $s = 0$.) The welfare gain of participating in the IRA falls as the private saving rate s increases. For high savings rates, the welfare effect is negative. This is because when s is high, individuals are already saving a large percentage of their take-home pay. The auto-IRA may force them to save so much that their consumption is higher in retirement that during the working years, which could reduce their lifetime well-being. The welfare effect switches from positive to negative around $s = 0.03$ for high school drop-outs and graduates and around $s = 0.05$ for college graduates. This is notable because the average U.S. personal saving rate of 5.1%. Our results suggest that individuals who follow a simple rule-of-thumb and save the average amount for the U.S. would be harmed by being enrolled in (and remaining in) an IRA.

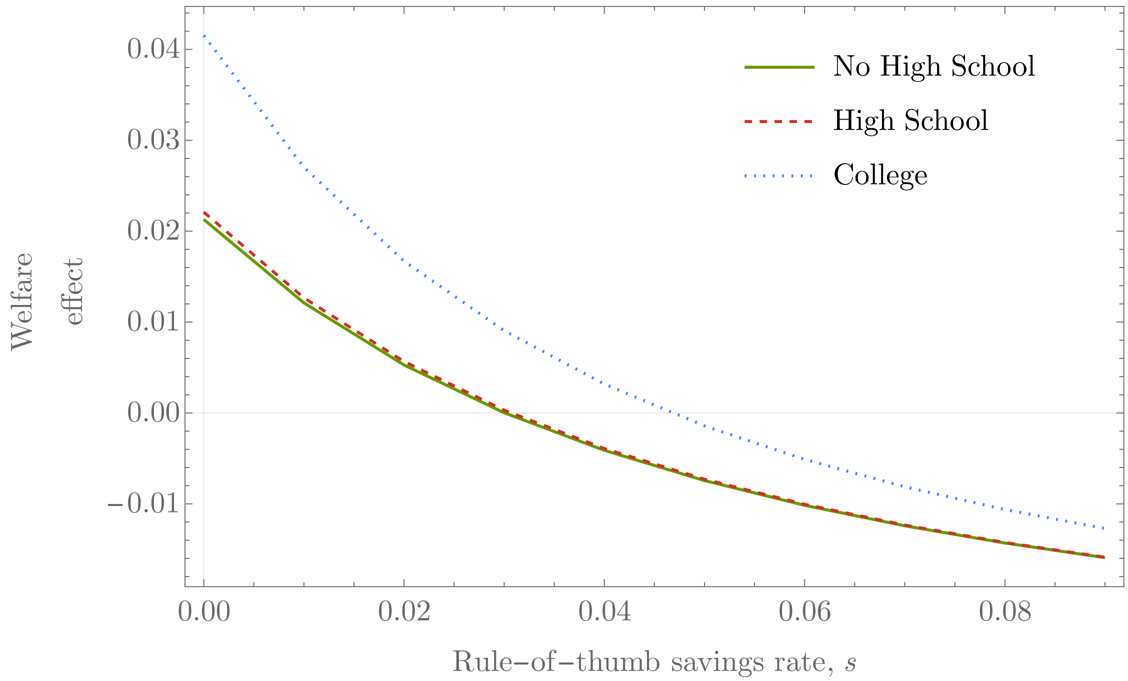

Figure 3: Consumption Equivalent Variation in a simple rule of thumb saving model for various savings rates s given IRA contribution rate $\delta = 0.03$.

3.1.1 Robustness, Lower income Profiles

The income profiles used in this paper represent the average wage earnings of an individual in a particular educational attainment group. One possible concern with this assumption is that workers who are enrolled in state-run automatic IRAs might have below-average earnings for their educational group. Recall, it is only employers who do not offer retirement benefits who are required to opt their workers into state-run IRAs. Thus employees who already receive retirement benefits at work, and likely have higher wages, will not participate in this program.

Oregon, California, and Illinois have released preliminary data about the average monthly contributions and average contribution rates for workers enrolled in automatic IRAs. For example in California, the average monthly contribution is $114 and the average contribution rate is 0.512% which implies an average monthly wage of \$2,230.[10] This corresponds to 44% of the average wage used in the paper.[11]

As a robustness check, we scale down the income profiles such that the average wage is 0.44, rather than 1. This reduces the welfare gain (increases the welfare cost) of the automatic IRA, since lower income households receive a Social Security benefit that replaces a larger fraction of their wage and thus do not need to save as much privately in order to fund the same level of retirement consumption. The welfare effects for this robustness check are presented in Table 3. The CEVs for workers who do not save privately are 0.002, 0.004, and 0.018 for high school drop outs, graduates, and college graduates, which are lower than the baseline CEVs of 0.021, 0.022, and 0.042. Similarly, the welfare cost to individuals who save 8% are −0.018, −0.017, and −0.014 with reduced income profiles compared to −0.014, −0.014, and −0.011 in the baseline. Qualitatively, the same pattern persists with lower income profiles: individuals who save more than about 3% privately are harmed by being enrolled in an automatic IRA.

[10] The implied average monthly wages are \$1,871 and \$2,528 for Illinois and Oregon. Monthly contribution amounts and contribution (or deferral) rates available at https://illinoistreasurer.gov/Individuals/Secure_Choice/Secure_Choice_Performance_Dashboards, https://www.oregon.gov/treasury/financial-empowerment/Documents/ors-board-meeting-minutes/2020/2020-12-Program-Report-OregonSaves-Monthly.pdf, and https://www.treasurer.ca.gov/calsavers/reports/participation/calSavers-participation-and-funding-snapshot-20210131.pdf.

[11] The model has been calibrated such that the economy-wide wage average is equal to one with corresponding Social Security taxes and benefits (i.e., the benefit earning rule flattens out for wages 2.47 times the average wage). The Social Security Administration reports the average wage earnings in 2021 were \$60,575.07 per year, which is \$5,047 per month. The implied average monthly wage in California is \$2,230 which is 44% of \$5,047.

Consumption Equivalent Variation			
s	No High School	High School	College
0.00	0.002	0.004	0.018
0.02	-0.006	-0.005	0.004
0.04	-0.011	-0.010	-0.004
0.06	-0.015	-0.014	-0.010
0.08	-0.018	-0.017	-0.014

Table 3: Consumption Equivalent Variation with lower wages. CEV in a simple rule of thumb saving model for various savings rates s given IRA contribution rate $\delta = 0.03$.

4 Extensions

Our baseline analysis assumes that individuals follow a rule-of-thumb when making their saving consumption decisions. Although the simple model provides a useful benchmark to evaluate the effectiveness of auto-IRAs, the model lacks many of the features that are standard in more traditional life-cycle models such as borrowing. In this section we explore three different extensions that relax the assumptions of our baseline model and allow us to further investigate the welfare effects of being enrolled in a state-run IRA. In the first extension we consider a fully rational life-cycle model. In the second, we model a counter-factual auto-IRA that requires individuals to pay a monetary cost in order to opt-out. In the final extension we add idiosyncratic income shocks and allow the individual to make withdrawals from their IRA before retirement.

4.1 Fully Rational Life-cycle Model

In this section we assume the individual solves a dynamic optimization problem to make their saving consumption decisions over the life-cycle; all other details remain the same as in the baseline model in section 2.

An individual of a particular education level solves the following utility maximization problem:

$$\max \int_0^T e^{-\rho t} u(c(t)) dt \tag{16}$$

subject to:

$$\dot{k} = y(t) - c(t) + \mathbf{1}\{k(t) < 0\} r_b k(t) + \mathbf{1}\{k(t) > 0\} r_s k(t), \tag{17}$$

$$k(0) = 0, \tag{18}$$

$$k(T) = 0, \tag{19}$$

where $\mathbf{1}\{k(t) < 0\}$ is an indicator function that equals 1 if assets are negative and zero otherwise, and $\mathbf{1}\{k(t) > 0\}$ is an indicator function that equals 1 if assets are positive.

This is a multi-stage control problem with state-dependent and time-dependent switches the state equation. The most common approach to solve a model with a credit spread ($r_b > r_s$) is to break asset holdings into two separate variables, one for debt and one for savings, each with corresponding inequality constraints, as in Davis et al. (2006) or Hurst and Willen (2007). We follow Caliendo and Findley (2019), and solve the problem directly with a single asset variable.[12]

The Hamiltonian for this problem is

$$\mathcal{H} : e^{-\rho t} u(c(t)) + \mu(t) \Big[y(t) - c(t) + \mathbf{1}\{k(t) < 0\} r_b k(t) + \mathbf{1}\{k(t) > 0\} r_s k(t) \Big] \tag{20}$$

[12]In an appendix, Caliendo and Findley (2019) prove that the direct approach is equivalent to the two-asset approach of Davis et al. (2006).

with necessary conditions

$$\mathcal{H}_c = e^{-\rho t} u'(c(t)) - \mu(t) = 0 \tag{21}$$

$$\mathcal{H}_k = \mu(t)\Big[\mathbf{1}\{k(t) < 0\} r_b k(t) + \mathbf{1}\{k(t) > 0\} r_s k(t)\Big] = -\frac{d\mu(t)}{dt}. \tag{22}$$

Solving the costate equation for $\mu(t)$ yields:

$$\mu(t) = \mu(0) \exp\big[-\int_0^t \big(\mathbf{1}\{k(t) < 0\} r_b k(t) + \mathbf{1}\{k(t) > 0\} r_s k(t)\big) dv\big]. \tag{23}$$

Combining the equation above with the maximum condition yields an expression for consumption:

$$c(t) = \mu(0)^{-1/\sigma} \exp\left[\frac{1}{\sigma} \int_0^t \big(\mathbf{1}\{k(t) < 0\} r_b k(t) + \mathbf{1}\{k(t) > 0\} r_s k(t)\big) dv - \rho t\right]. \tag{24}$$

From the maximum condition equation, we know that $\mu(0)^{-1/\sigma} = c(0)$. Therefore,

$$c(t) = c(0) \exp\left[\frac{1}{\sigma} \int_0^t (\mathbf{1}\{k(t) < 0\} r_b k(t) + \mathbf{1}\{k(t) > 0\} r_s k(t))\, dv - \rho t\right]. \tag{25}$$

Together equations (25) and (17) implicitly define optimal consumption $c(t)$ as a function of assets $k(t)$ and an unknown constant $c(0)$. Using the terminal condition $k(T) = 0$, we find the initial value of consumption $c(0)$ computationally using the shooting method.

4.1.1 Results: Fully Rational Model

The life-cycle consumption and saving profile for a high school drop-out in the fully rational model are plotted in Figure 4. Were it not for the credit wedge, the individual would perfectly smooth their consumption over the life cycle. However, because the cost of borrowing is higher than the return on saving, the individual chooses not to borrow. As a result, their consumption perfectly tracks their disposal income for the first 7.5 years of their working life (without an IRA) when their income is low. Then the consumption profile flattens out and is smooth through the remainder of the life cycle. The asset path is completely flat for the first 7.5 years when the individual does not save or borrow. The individual begins to save at age 29.5 and accumulates assets until the date of retirement. They perfectly draw down their assets during retirement.

The consumption profile is similar if the individual is enrolled in (and remains enrolled in) an IRA. The IRA reduces take-home pay, so the individual chooses to delay saving until they are 9.4 years into their working life. Following that point, consumption is higher with an IRA than without. The asset profile with an IRA is qualitatively different than without. When the individual has an IRA, they do not begin saving until age 31 (a little less than 2 years later). Their asset accumulation peaks *prior* to retirement. The individual begins to gradually draw down their private assets in their early 60s as their income declines. The asset profile is kinked at the date of retirement when the individual starts to draw down their assets more rapidly. The individual begins to draw down their private assets before retiring because they are still contributing to their IRA during those years.

The consumption and savings paths for high school graduates are similar to that of high school drop-outs and are plotted in Figure 5. The high school graduates choose to consume their take-home pay early in the working life. They begin saving 14.6 years into the work life without an IRA and 12 years into the work life when they are enrolled in the IRA. The consumption profile is flat as soon as the individual begins to save. The asset profiles are monotonically increasing from the date the individual begins to save until the retirement date, after which the individual draws down their assets.

College graduates are the only group that choose to borrow in this framework. The consumption and savings profiles for college graduates are plotted in Figure 6. College graduates borrow when young in order to be able to consume more than their take-home pay. However, because borrowing is costly, they do not perfectly smooth their consumption. Their consumption profile is kinked; it increases early in life when their income profile is very steep. Then, their consumption profile flattens out and is perfectly smooth until retirement. The kink in the consumption path corresponds to moment in time when the individual switches from borrowing to saving. The individual borrows a little when young, repays the debt, and then accumulates assets until the date of retirement.

Participating in an IRA has two competing effects for an individual. Saving in a tax-advantaged Roth IRA increases the lifetime resources of the individual, which increases their lifetime consumption. However, the IRA contributions reduce the individual's take-home pay This harms the individual early in life when they would prefer to borrow rather than save. The individual can borrow privately to offset the forced IRA savings, but this is costly because of the credit market wedge. This reduces lifetime welfare.

We use a consumption equivalent variation technique to measure the net effect. For our welfare analysis, we assume that the individual stays enrolled in the IRA. If our welfare results are negative, they can be interpreted as the cost (in terms of consumption) the household would be willing to incur to opt-out of the IRA. We feel this comparison offers a clear baseline to evaluate the effect of being enrolled in (and remaining in) an IRA.[13] The consumption

[13]We relax this assumption and model an explicit opt-out cost in Section 4.2.

equivalent variation (CEV) is calculated using equation 15 and the consumption paths that solve the individual's optimization problem. The CEV is negative for all three education levels. The CEV for high school drop-outs is -0.0016, for high school graduates is -0.004, and for college graduates is -0.007.[14]

The welfare cost of being enrolled in an IRA is increasing in education. This is because the highest educated group has the steepest income profile and would like to borrow the most. The IRA contributions reduce their income in precisely the years they would want to borrow, which reduces their utility. The welfare cost is between 0.1%-0.7% of consumption, which is similar in magnitude to the welfare cost in the rule-of-thumb model when the savings rate is 6%.

As an additional robustness check we also consider the case where the cost of borrowing is parameterized to match credit card debt, $r_b = 0.18$ (this interest rate is within the range estimated by Davis et al. (2006)). This change does not alter the optimal consumption choices of high school drop-outs or graduates, as the baseline cost of borrowing was high enough to deter them from taking on debt. Raising the interest rate on borrowing changes the optimal consumption path for college graduates and increases the welfare cost of enrolling in and remaining in the IRA. The CEV is -0.009 in this case compared to -0.007 with the baseline parameters. The change in CEV is small because the college graduates do not borrow much, even in the baseline case.

[14] As a robustness check, we scale down the wages such that the average monthly wage corresponds to data released by the state of California (as described in section 3.1.1). This marginally increases the welfare cost of the IRA, since lower income individuals receive a Social Security benefit that replaces a larger share of their wages and thus would prefer to save less. The CEV for high school drop-outs is -0.002, for high school graduates is -0.005, and for college graduates is -0.008, in this case.

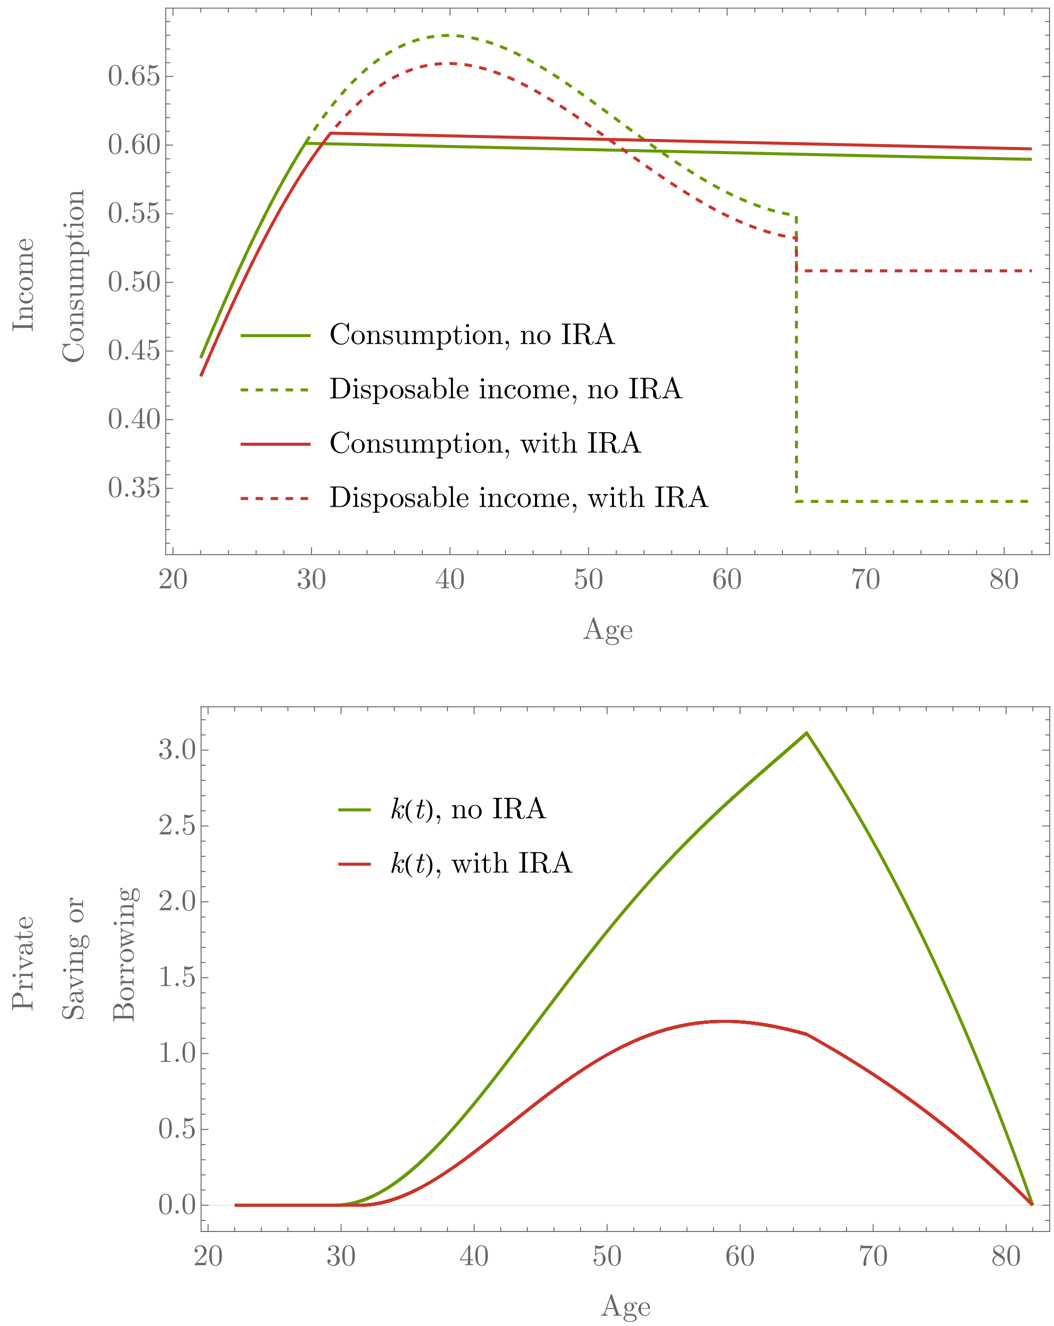

Figure 4: Fully rational consumption, disposable income, and private savings $k(t)$ with and without an IRA for high school drop-outs given IRA contribution rate $\delta = 0.03$.

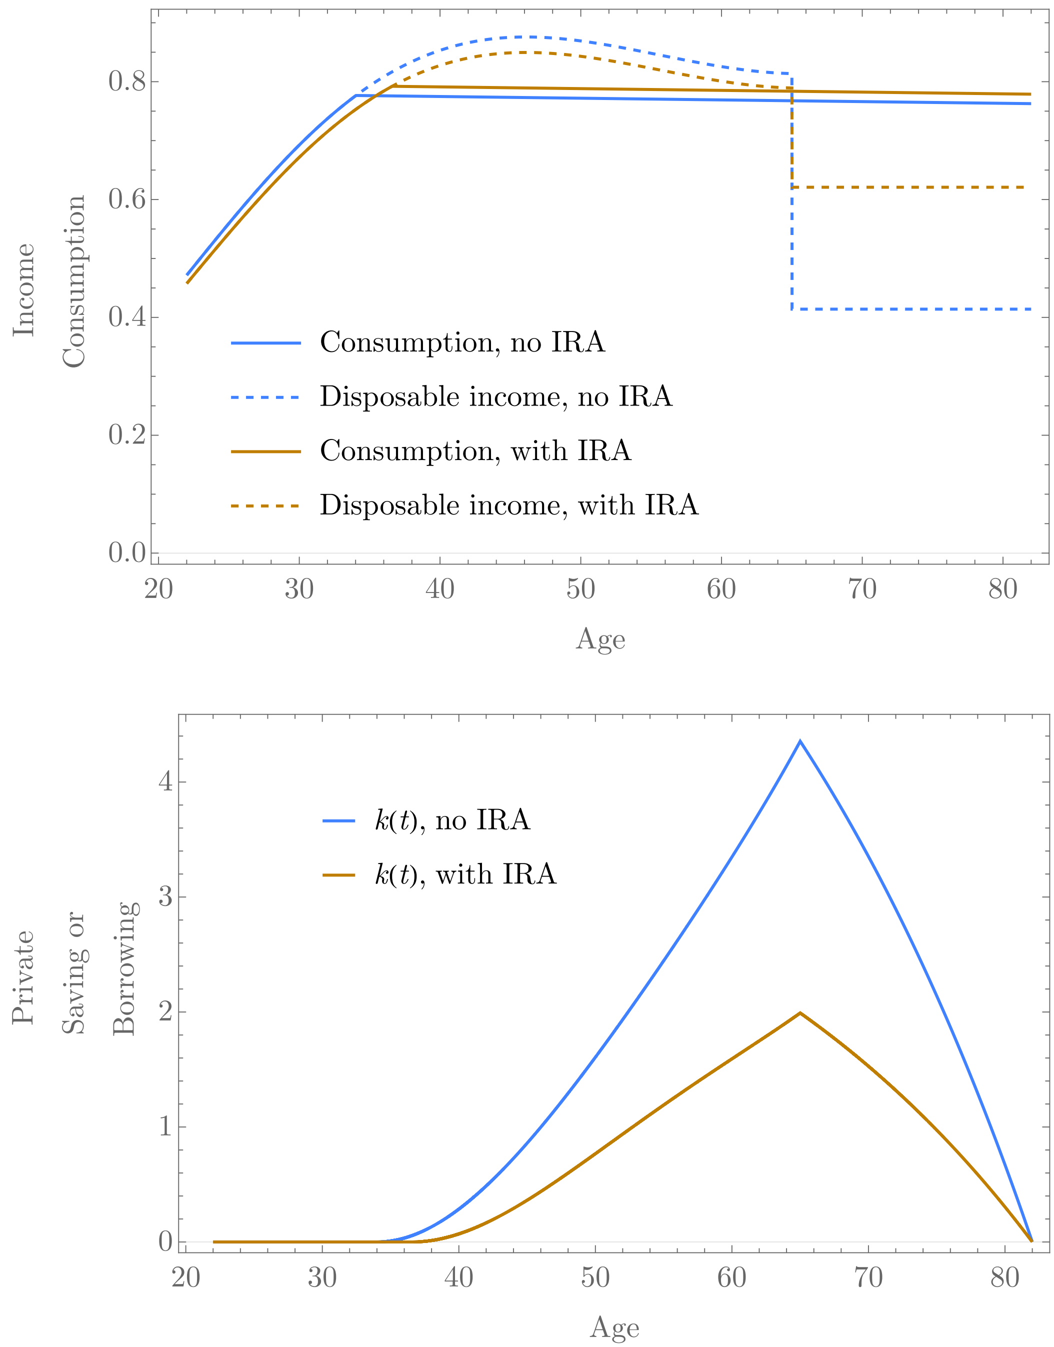

Figure 5: Fully rational consumption, disposable income, and private savings $k(t)$ with and without an IRA for high school graduates given IRA contribution rate $\delta = 0.03$.

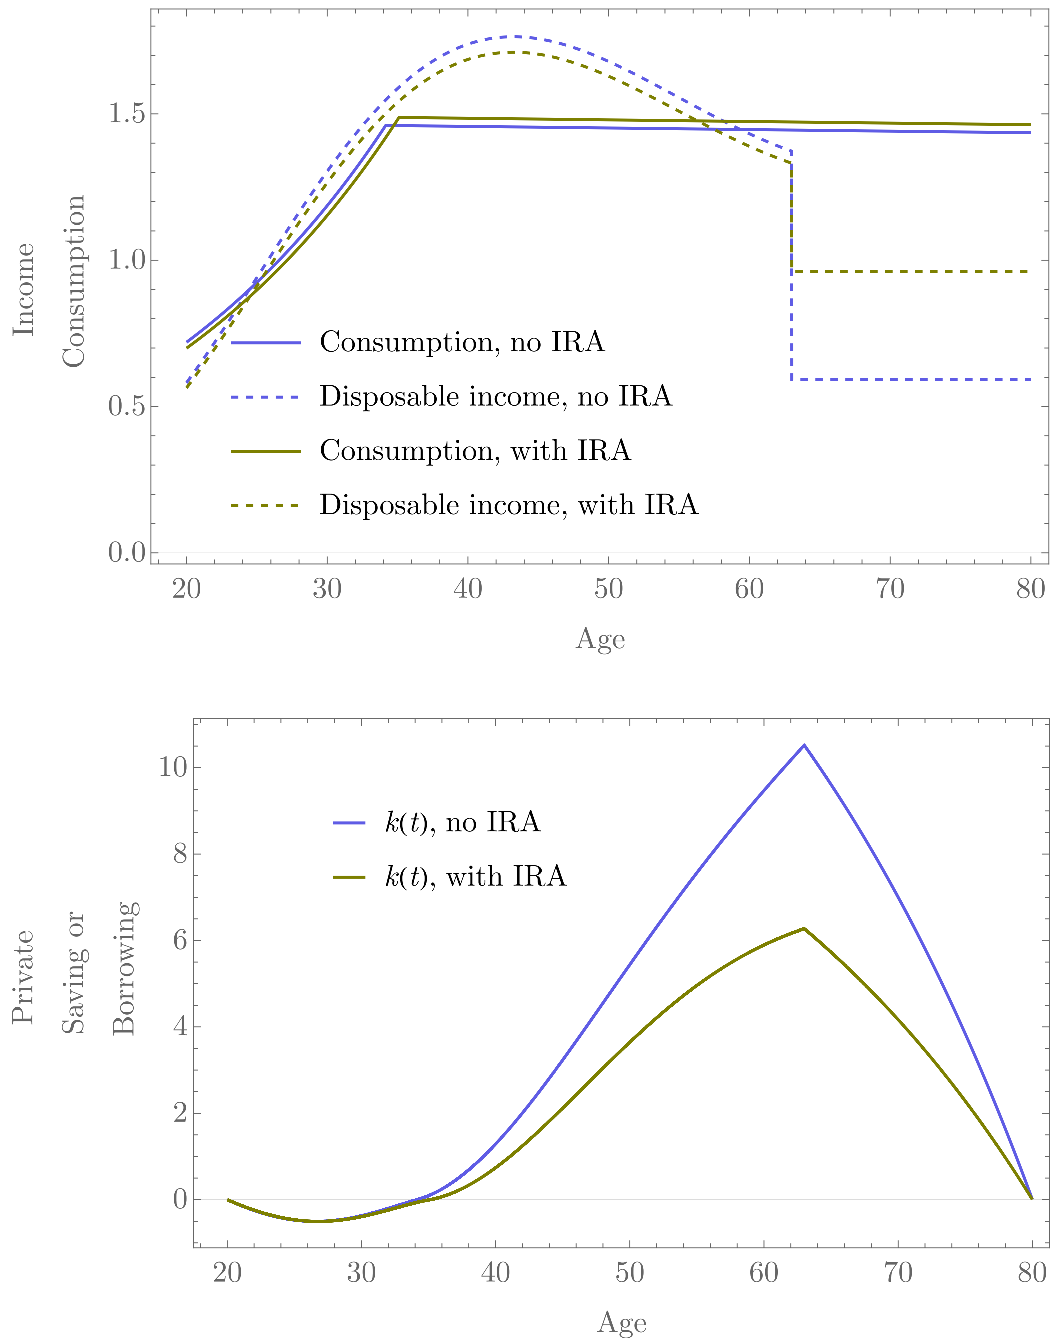

Figure 6: Fully rational consumption, disposable income, and private savings or borrowing $k(t)$ with and without an IRA for college graduates given IRA contribution rate $\delta = 0.03$.

4.2 Monetary Opt-out Cost

Thus far in the paper we have assumed that households who are enrolled in an IRA remain in the IRA and make automatic contributions until they retire. Empirical evidence suggests that the majority of households who are enrolled continue making contributions, so we feel this is a reasonable baseline. However, in reality, a worker could opt out of the IRA. This requires some time or effort, which we proxy with a monetary cost. We augment our baseline model and allow the household to pay a one-time fixed cost at $t = 0$ as they enter the model in order to opt out of the IRA.[15]

We re-run our welfare analysis including this monetary opt-out cost. We compare two alternative states of the world: (i) the individual is opted into an IRA and they can either continue participating or pay a monetary cost to opt out, and (ii) the individual does not have an IRA (nor do they pay an opt-out cost). If the individual is opted into the IRA, we assume they only continue participating in the IRA if gives them higher lifetime utility than paying the opt-out cost and not having an IRA. Our welfare metric calculates the fraction of lifetime consumption an individual who is enrolled in an IRA that they can opt out if they pay a monetary cost would be willing to give up such that their discounted lifetime utility is equal to their discounted lifetime utility if they did not have an IRA. A positive number indicates the individual would prefer being enrolled in an IRA with a costly opt-out option to not participating in an IRA at all.

We parameterize the opt-out cost to be $1,000 to match the opt-out cost estimated in several papers, including DellaVigna (2009), DellaVigna (2018), Bernheim et al. (2015).[16]

[15] Handel (2013) use a similar one-time fixed cost at $t = 0$ to model the opt-out cost of changing health insurance.

[16] Choukhmane (2019) calculates a much smaller opt-out cost in a complex model where households con-

Our results for our baseline rule-of-thumb model with an opt-out cost are presented in Table 4. The welfare gains of being enrolled in an IRA are the same as our baseline for individuals who choose to remain enrolled. This is because the IRA improves the individual's well-being, so they do not exercise the option to opt-out of the IRA and the cost does not impact them. High school drop-outs and graduates with savings rates less 3.5% benefit from the IRA and would not choose to opt-out. College graduates with savings rates less than 5% do not opt-out.

The monetary opt-out cost is only relevant for individuals who would choose to opt-out for free if given the option. For these individuals being enrolled in an IRA with a costly opt-out option makes them worse off either because they have to pay the cost to opt-out, or because the cost is high enough that they choose to stay enrolled in the IRA even though it reduces their lifetime utility. In our numerical analysis, the welfare cost of the later (staying in the IRA when they would otherwise choose to opt-out) is quite small and only relevant for a small window of private savings rates. The monetary opt-out cost is relatively low compared to lifetime income, so most savings rate for which individuals are harmed by the IRA, the individuals choose to opt out.

The negative CEVs in Table 4 all represent the welfare cost for individuals who chose to opt out of the IRA. Note that the CEVs are closer to zero than in the baseline case in Table 2 that assumed individuals did not opt-out. This is because paying to opt-out of the IRA reduces the lifetime utility of an individual by less than remaining enrolled in the IRA. Qualitatively, the welfare results are similar with out without an opt-out cost. Being enrolled in an IRA is only beneficial for a rule-of-thumb saver if their private savings rate is low (less

tinuously re-optimize and understand that their choices today become the defaults for tomorrow.

Consumption Equivalent Variation			
s	No High School	High School	College
0	0.021	0.022	0.042
0.02	0.005	0.006	0.017
0.04	-0.001	-0.001	0.003
0.06	-0.001	-0.001	-0.001
0.08	-0.001	-0.001	-0.000

Table 4: Consumption Equivalent Variation with a monetary opt-out cost in a simple rule of thumb saving model for various savings rates s given IRA contribution rate $\delta = 0.03$.

than 3.5% for high school drop-out and graduates and less than 5% for college graduates).

The results in this section confirm the intuition of our baseline model and our previous extensions. If workers are given the option to opt-out of the IRA, but that option is costly (either in terms of a monetary expense as in this model extension, or in terms of utility due to cognitive, financial, or time costs), they will only opt-out if participating in the IRA reduces their lifetime well-being. However, because opting out is costly, this will also reduce their well-being. Qualitatively, this is similar to our analysis in previous sections. Our welfare metrics would be similar if were to incorporate an opt-out cost into any of our extensions. Workers for whom staying enrolled in the IRA increases lifetime well-being have the same CEV with or without an opt-out cost. The CEVs only change for individuals who would want to opt-out; however, even in those cases the sign of the CEV metric doesn't change, only the magnitude, which gets smaller.

4.3 Income Shocks and Withdrawals

In this final extension, we explore the interaction between risky income and pre-retirement IRA withdrawals. Our baseline model and the extensions of the previous sections assume an individual earns income, contributes to an IRA, borrows or saves privately during their working years, and then retires and draws down the IRA and private saving accounts. These standard modeling assumptions allowed us to focus on the effect of having an additional asset to save for retirement. In reality, many individuals in the U.S. make withdrawals from their IRA accounts *prior* to retirement.[17] This is particularly relevant if the household has experienced a negative income shock. We consider an alternative modeling framework in this section to jointly account for negative income shocks and pre-retirement IRA withdrawals.

We model the dynamic stochastic optimization problem for an individual with a risky income stream (they can experience periods of unemployment). If enrolled in an IRA, the individual can make pre-retirement withdrawals. It is necessary to step away from our baseline line assumption of a rule-of-thumb savings rate if we want to explore the welfare effects of pre-retirement withdrawals. We assume the the individual is fully rational and chooses their consumption, saving, and IRA withdrawals to maximize their lifetime utility.

Time is discrete; a worker enters the model at $t = 0$, retires at $t = t_R$, and dies at $t = T$. Prior to retirement, the net labor income of individual i is given by

$$y_{i,t} = \psi_{i,t} \left(w(1 - \tau_{SS}) - \min[\delta w(1 - \tau_{SS}), m] \right), \quad \text{for } t < t_R \tag{26}$$

where $\psi_{i,t}$ is an indicator function that is equal to one if labor income is positive and zero otherwise (i.e. employed or unemployed). We assume the indicator function $\psi_{i,t}$ evolves according to a two-state Markov process with transition matrix Π, with $\pi_{p,q} = \Pr(\psi_{i,t+1} = q | \psi_{i,t} = p)$. Wage income is denoted w, τ_{SS} is the Social Security tax rate, and the automatic

[17]For example, Biggs et al. (2019) finds that retirement account withdrawals were common even during the economic expansion of the late 2010s. Similarly Quinby et al. (2019) finds that 20% of workers enrolled in Oregon's state-run IRA "OregonSaves" made a pre-retirement withdrawal of at least $1,000.

IRA contribution is either the fraction δ of the after tax wage, or the contribution maximum m.[18] The individual only contributes to the IRA in periods in which they have positive labor income. After retirement, the individual receives a social security benefit b, in every period. We set b such that government runs a balanced budget:

$$y_{i,t} = b = \frac{t_R(1 - \frac{\pi_{1,0}}{\pi_{1,0}+\pi_{0,1}})w\tau_{SS}}{T - t_R}, \quad \text{for } t \in [t_R, T]. \tag{27}$$

The IRA account is denoted $A_{i,t}$ and earns interest at rate r_{IRA}. The individual can choose to make a pre-retirement withdrawal $x_{i,t} \geq= 0$ in any period. After retirement the individual optimally chooses withdrawals to jointly deplete their private savings and their IRA account by the date of death.

The recursive problem of the household can be written using a value function:

$$V_{i,t}(k, A, \psi) = \max_{c,k,A} u(c) + \beta \mathbb{E}_{\psi'}[V_{i,t+1}(k', A', \psi')|\psi] \tag{28}$$

subject to

$$x = (1 + r_{IRA})A + \min[\delta\psi w(1 - \tau_{SS}), m] - A' \tag{29}$$

$$k' = (1 + \mathbf{1}\{k < 0\}r_b + \mathbf{1}\{k > 0\}r_s)k + y - c - x \tag{30}$$

$$\psi' = \begin{cases} 1 & \text{with probability } \pi^{\psi,1} \\ 0 & \text{otherwise} \end{cases} \tag{31}$$

[18] We will only consider contribution rates δ that are low enough such that the maximum IRA contribution does not bind.

where β is the discount factor. As in the previous section, $\mathbf{1}\{k < 0\}$ is an indicator function that equals 1 if assets are negative and zero otherwise, and $\mathbf{1}\{k > 0\}$ is an indicator function that equals 1 if assets are positive. The interest rate on saving r_s is lower than the interest rate on borrowing r_b (and is parameterized as in section 2.5). Period utility $u(c)$ is CRRA as in equation 1.

We parameterize the stochastic model to correspond to an 80 year life span with the work life beginning at age 25 and retirement at age 65. We set $\pi_{1,1} = 0.977$ and $\pi_{0,0} = .564$ as in Ashman and Neumuller (2020). This implies a steady-state unemployment rate of $\frac{\pi_{1,0}}{\pi_{1,0}+\pi_{0,1}} = 0.05$. The interest rates, CRRA parameter σ, social security tax rate τ_{SS}, capital gains tax τ_k, IRA contribution rate δ and maximum contribution m are all parameterized as in Table 1. We set the discount factor $\beta = 1/(1 + r_{IRA})$ and normalize the wage to $w = 1$.

4.3.1 Results: Income Shocks and Withdrawals

The average consumption profile is hump-shaped over the life-cycle. This is because of the precautionary savings motive of the young. Workers face a probability of becoming unemployed in any period during their working life. The prospect of being unemployed (and not having labor income) compels young individuals to save a relatively large fraction of their earnings. As they age and accumulate assets, the precautionary savings motive is not as strong, and consumption increases. The average income and consumption profiles are plotted in Figure 7.

Average wealth over the life cycle is presented in Figure 8. Average total wealth when enrolled in an IRA (private savings plus IRA assets) is similar across the life cycle to average wealth without an IRA. Individual who participate in the IRA save less privately such that

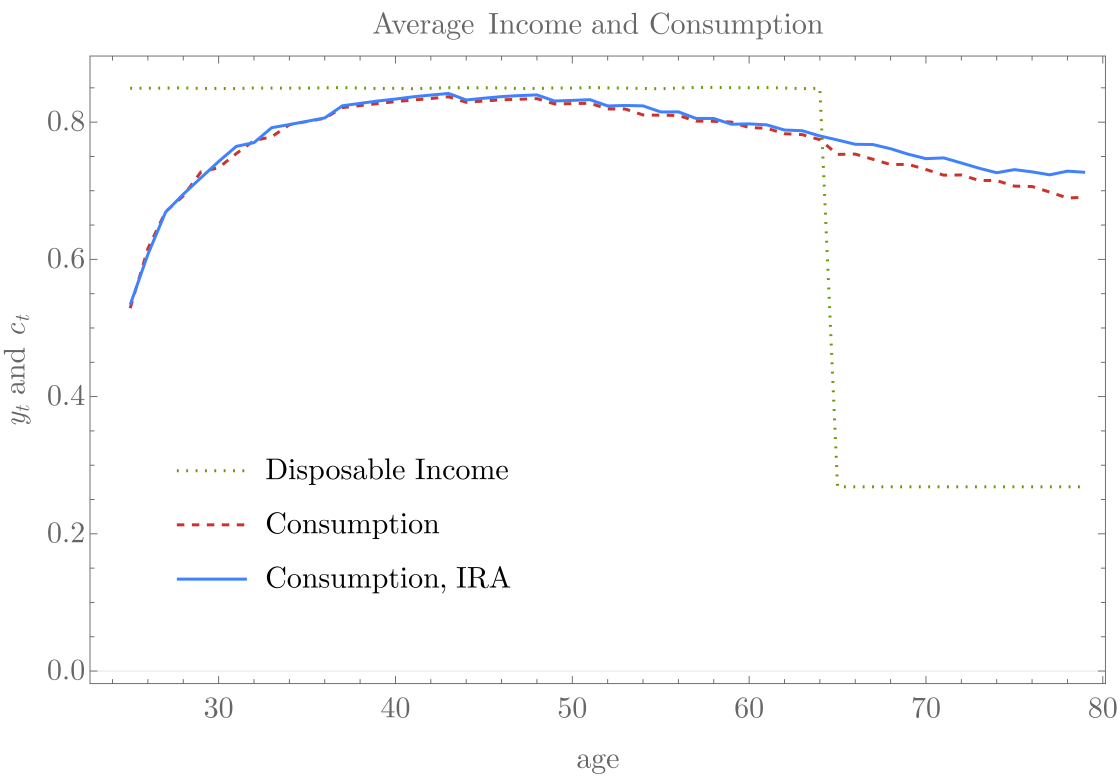

Figure 7: Average disposable income (without an IRA) and average consumption with and without an IRA with stochastic income, given IRA contribution rate $\delta = 0.03$.

their total wealth is similar to what they would have saved without an IRA. Individuals with an IRA choose to draw down their private savings before making IRA withdrawals in retirement because of the different rates of return.

Participating in an IRA gives the worker an additional asset they can draw down if they receive a negative income shock (i.e., when they are unemployed). The IRA earns a higher rate of return than private savings because of capital gains taxes, so the individual will only want to draw down the IRA if they do not have enough positive assets to cover their desired consumption level. The IRA could increase the ex-ante well-being of the individual through two channels: the higher rate of return, and the increased liquidity at the time of an income shock. We conduct two different welfare experiments to disentangle these effects.

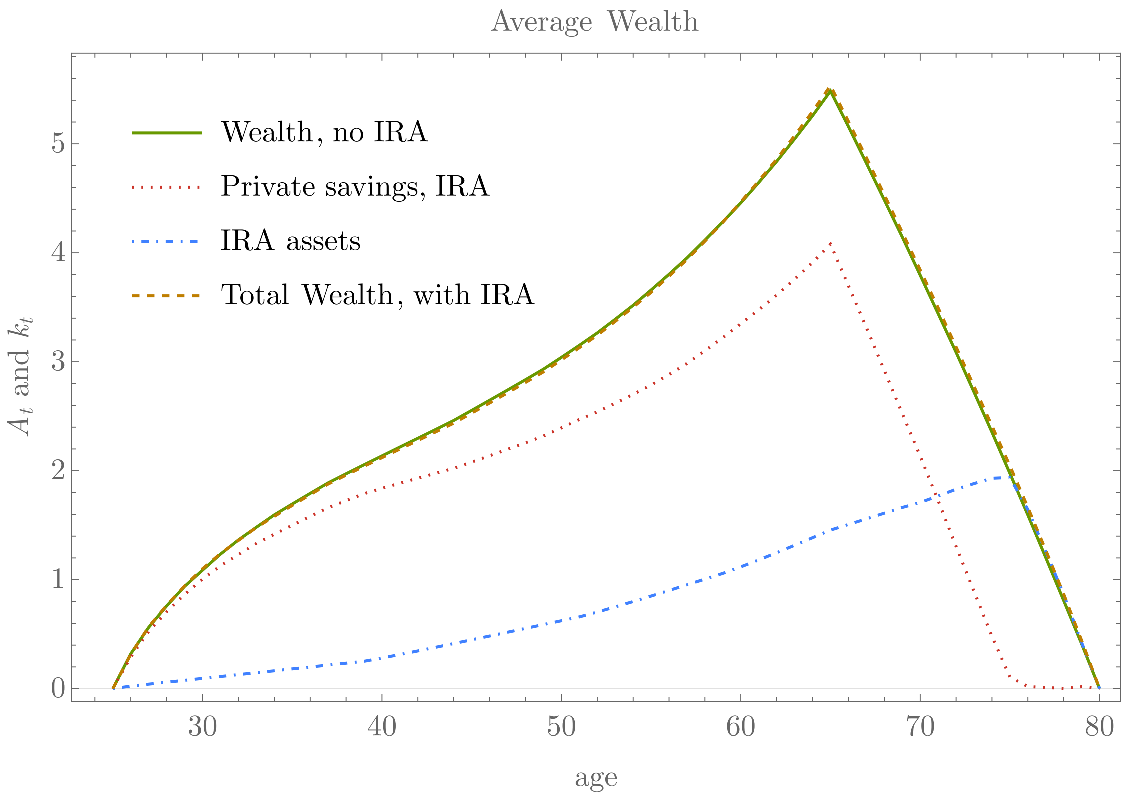

Figure 8: Average wealth with and without an IRA with stochastic income, given IRA contribution rate $\delta = 0.03$.

We calculate the welfare effect of being enrolled and remaining in the IRA using an ex-ante consumption equivalent variation. This is similar to the CEV in the previous sections; the only difference is that now the individual faces a risky income stream, so we compare *expected* lifetime utility with and without an IRA. We find the CEV to be 0.0085. Participating in an IRA increases the individual's lifetime resources and they are able to make a withdrawal from the IRA when a negative income shock hits which improves their ex-ante expected utility.

To isolate the effect of having access to a higher rate from the effect of additional liquidity, we turn off the income shock by setting $\psi_{i,t} = 1 \ \forall i, t$. Then we compute the discounted lifetime utility of the worker with and without an IRA. When the income shocks are turned off, the only effect of the IRA is providing the worker higher rate of return. The wage profile is flat over the life-cycle in this extension, and so in the absence of income shocks, the individual doesn't have an incentive to borrow. The IRA crowds out positive private savings, but has a higher rate of return, so this makes the individual better off. The CEV is 0.0078 when the income shocks are turned off. This CEV is a little smaller than the CEV with income shocks which was 0.0085. Comparing the CEVs with and without income shocks, over 90% the welfare gain from participating in the IRA is coming from the higher rate of return, rather than from being able to strategically withdraw assets from the IRA during periods of low income.

5 Concluding Remarks

Automatic enrollment in an IRA has the potential to increase a worker's lifetime consumption and well-being by directing saving into a tax-advantaged account. However, in the presence of behavioral or credit frictions, automatic enrollment in an IRA could reduce a worker's well-being by increasing saving early in life when the worker would otherwise want to borrow (or save less). Additionally, program fees could reduce the net rate of return of saving in an IRA and could potentially eat away any tax advantage associated with the account. We quantify these competing effects using a life-cycle model with a credit wedge. We consider both a rule-of-thumb model and a fully rational model to put upper and lower bounds on the potential effects of automatic enrollment on lifetime well-being.

Workers do not save optimally but instead follow a simple of rule-of-thumb to make savings decisions, can be made better off if they are automatically enrollment in an IRA. Workers of all education levels would be willing to give up some of their lifetime consumption in order to be enrolled in an IRA if they are following a rule of thumb and save less than 3% of their income. The welfare gains are largest for highest earners, who have the steepest income profile. A college educated worker who does not save at all would be willing to give up 4.2% of their lifetime consumption to be enrolled in an IRA with a 3% contribution rate. Workers with flatter life-cycle income profiles benefit less from the IRA. Similarly, workers who already save privately do not benefit as much. Workers of any education level who save more than 5% privately are harmed by being enrolled in an IRA.

In a fully rational model, workers who face a credit wedge and are already following an optimal saving-consumption path are harmed by automatic enrollment in an IRA. Workers would be willing to give up between 0.1%-7% of lifetime consumption to avoid being enrolled in the IRA. This is because the IRA forces the worker to save early in life when they would prefer to borrow (or save less). The worker can unwind some of this forced saving by borrowing, but it is expensive because the interest rate on borrowing exceeds the interest rate on saving. This is true even though the IRA offers a higher rate of return than private savings because of the tax advantage associated with a Roth IRA.

Our results suggest that automatic enrollment in an IRA is most beneficial for workers who save very little on their own. In contrast, workers who are already saving optimally and face a credit wedge or other borrowing constraint could be harmed by enrollment in the program. Our calculations assume that workers stay enrolled in the IRA, or pay a monetary cost to opt-out. Policymakers face a trade-off of creating barriers to opting-out that are sufficiently strong to keep workers who could benefit from the program enrolled, while simultaneously ensuring that workers who are harmed by the program are able to opt-out. The characteristics of who remains enrolled in state-sponsored IRAs and who opts out is an open question that merits future research.

6 References

Alonso-Ortiz, J. (2014). Social security and retirement across the OECD. *Journal of Economic Dynamics and Control*, 47:300 – 316.

Ashman, H. and Neumuller, S. (2020). Can income differences explain the racial wealth gap? a quantitative analysis. *Review of Economic Dynamics*, 35:220 – 239.

Attanasio, O. P. (1994). The intertemporal allocation of consumption: Theory and evidence. *working paper*.

Attanasio, O. P., Banks, J., Meghir, C., and Weber, G. (1999). Humps and bumps in lifetime consumption. *Journal of Business & Economic Statistics*, 17(1):22 – 35.

Belbase, A. and Sanzenbacher, G. T. (2018). How Have Workers Responded to Oregon's Auto-IRA. Working Paper 18-22, Center for Retirement Research at Boston College.

Bernheim, B. D., Fradkin, A., and Popov, I. (2015). The welfare economics of default options in 401(k) plans. *American Economic Review*, 105(9):2798–2837.

Beshears, J., J.Choi, J., Laibson, D., Madrian, B. C., and Wang, S. Y. (2016). Who Is Easier to Nudge? *working paper*.

Biggs, A. D., Munnell, A. H., and Chen, A. (2019). Why are 401(k)/ira balances substantially below potential?.

Browning, M. and Crossley, T. F. (2001). The life-cycle model of consumption and saving. *The Journal of Economic Perspectives*, 15(3):3 – 22.

Caliendo, F. and Findley, T. S. (2019). Time-inconsistent preferences and saving policies. *Utah State University working paper*.

Carroll, G. D., Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2009). Optimal defaults and active decisions. *The Quarterly Journal of Economics*, 124(4):1639–1674.

Center for Retirement Research (2015). Report on Design of Connecticut's Retirement Security Program. Working paper, Center for Retirement Research at Boston College.

Chalmers, J., Mitchell, O. S., Reuter, J., Zhong, M., and Sanzenbacher, G. T. (2020). Auto-enrollment Retirement Plans for the People: Choices and Outcomes in OregonSaves. Working Paper WP2020-15, Pension Research Council.

Chetty, R., Friedman, J. N., Leth-Petersen, S., Nielsen, T. H., and Olsen, T. (2014). Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from denmark. *Quarterly Journal of Economics*, 129(3):1141 – 1219.

Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2003). Optimal defaults. *The American Economic Review*, 93(2):180–185.

Choukhmane, T. (2019). Default options and retirement saving dynamics. *working paper*.

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

Davis, S. J., Kubler, F., and Willen, P. (2006). Borrowing costs and the demand for equity over the life cycle. *The Review of Economics and Statistics*, 88(2):348–362.

Deaton, A. (1992). *Understanding Consumption*. Oxford University Press.

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2):315 – 372.

DellaVigna, S. (2018). Structural behavioral economics.

Findley, S. and Cottle Hunt, E. (2019). Quasi-hyperbolic discounting and the save more tomorrow program. *working paper*.

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 103(7):2643–82.

Harris, T. F., Troske, K., and Yelowitz, A. (2018). How will state-run auto-iras affect workers? *The Journal of Retirement*, 6(2):27–33.

Hubbard, R. G., Skinner, J., and Zeldes, S. P. (1995). Precautionary saving and social insurance. *Journal of Political Economy*, 103(2):360–399.

Hurst, E. and Willen, P. (2007). Social security and unsecured debt. *Journal of Public Economics*, 91(7):1273 – 1297.

Livshits, I., MacGee, J., and Tertilt, M. (2007). Consumer bankruptcy: A fresh start. *The American Economic Review*, 97(1):402–418.

Love, D. and Schmidt, L. (2015). Comprehensive Wealth of Immigrants and Natives. (University of Michigan Retirement Research Center Working Paper 328).

Lusardi, A. (2002). Preparing for retirement: The importance of planning costs. *Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association*, 95:148–154.

Lusardi, A., Michaud, P.-C., and Mitchell, O. S. (2017). Optimal financial knowledge and wealth inequality. *Journal of Political Economy*, 125(2):431–477.

Lusardi, A. and Mitchell, O. S. (2008). Planning and financial literacy: How do women fare?. *The American Economic Review*, 98(2):413 – 417.

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

McGrattan, E. R. and Prescott, E. C. (2003). Average debt and equity returns: Puzzling? *The American Economic Review*, 93(2):392–397.

Park, H. and Feigenbaum, J. (2018). Bounded rationality, lifecycle consumption, and Social Security. *Journal of Economic Behavior and Organization*, 146:65 – 105.

Quinby, L. D., Hou, W., Belbase, A., and Sanzenbacher, G. T. (2019). Participation and pre-retirement withdrawals in oregon's auto-ira.

Winter, J. K., Schlafmann, K., and Rodepeter, R. (2012). Rules of thumb in life-cycle saving decisions. *The Economic Journal*, 122(560):479 – 501.