

Preliminary and Incomplete
Please Do Not Quote
Comments Welcome

Smooth Retirement Accounts*

Thomas J. Brennan[†]

March 12, 2014

Abstract

I introduce the concept of “smooth retirement accounts” (SRAs) to provide a method for taxing retirement savings evenly over time. I contrast this with the back-loaded taxation of traditional accounts, and I use lifetime utility maximization models to demonstrate that future non-linear and uncertain tax brackets can distort savings incentives and portfolio allocations for traditional account holders. I also contrast SRAs with the front-loaded taxation of Roth accounts, and I argue that SRAs would bring a reasonable portion of retirement account taxes into the current budget window without leading to the extreme result of Roth accounts that leave no tax receipts beyond the year of contribution. Because SRAs can eliminate investment and savings distortions for taxpayers, as well as help set government budgetary incentives correctly, I recommend that they be created by Congress as a replacement for the current choices of Roth and traditional accounts.

*I thank Alan Auerbach, Cathryn Brennan, Charlotte Crane, Mihir Desai, Michael Doran, Dan Halperin, David Kamin, Louis Kaplow, Michael Knoll, Mike Schler, Dan Shaviro, Reed Shuldiner, Ted Sims, Norman Stein, Josh Teitelbaum, Al Warren, David Weisbach, Kathy Zeiler, as well as participants at the NYU Tax Policy Colloquium in February 2014, the Harvard Law School Seminar on Current Research in Taxation in August 2013 and participants at the Georgetown University Law Center Faculty Research Workshop in November 2010 for very useful discussions and comments on prior versions of this paper. In addition, I am grateful to Christopher Lee, Duong Nguyen, and Ben Wallner for providing excellent research assistance over the course of this project. I also thank Pegeen Bassett and Erin Conway for very able assistance in locating data and information. Research support from the Northwestern University School of Law Faculty Research Program is gratefully acknowledged.

[†]Professor of Law, Northwestern University School of Law and Professor of Finance (Courtesy), Northwestern University Kellogg School of Management, Chicago, IL 60611. Phone: 312-503-3233. E-mail: t-brennan@law.northwestern.edu

Contents

1	Introduction	1
2	Choices and Problems for Taxpayers	3
2.1	Traditional vs. Roth: Overview	3
2.2	Utility Maximization: Contribution Choice	5
2.3	Utility Maximization: Allocation Choice	8
3	Incentives for the Government	10
4	Implementing SRAs	11
4.1	Virtualization: The Accounting Approach	13
4.2	Actualization: The Cash-Flow Approach	16
5	Conclusion	17
A	Appendix: Numerical Lifetime Utility Model	18
A.1	Monte Carlo Simulation	20
A.2	Stochastic Tax Rates	24
B	Appendix: Model and Simulation Parameters	24

1 Introduction

The past several decades have witnessed a significant transformation in the nature of retirement planning, with employer managed pensions and defined benefit plans being replaced by individually managed accounts and defined contribution plans.¹ For a time, all of the new defined contribution plans were taxed according to a methodology similar to that of the old defined benefit plans: amounts set aside for the future are deductible from income up front, and amounts are included in income when they are ultimately distributed to account holders. However, in 1997 a new breed of plan was introduced, a so-called Roth account,² under which no deduction is allowed for amounts initially set aside, but also under which ultimate distributions are not included in the income of account holders. Roth accounts are popular because they allow the government to collect revenue up front and give taxpayers an opportunity to escape the effect of future tax rate increases. These same features may have negative aspects, however, to the extent that myopic revenue acceleration by the government comes at the cost of revenue available to future generations.³

Roth accounts are already numerous and growing, but they are not yet nearly as common or as well funded as traditional IRAs.⁴ The landscape continues to change, and both the President and Congress have demonstrated an interest in creating more opportunities to contribute to Roth-type accounts.⁵ This environment makes it an ideal time to review and consider the possibilities for retirement account design and reassess the best path forward. In this paper I take advantage of this opportunity and analyze the options currently available in detail. I employ theoretical arguments based on utility maximization models to provide

¹The traditional individual retirement account (IRA) was created in 1974 by the Employee Retirement Income Security Act (ERISA). For some statistics and details about the transition away from defined benefit plans and to defined contribution plans, see Section 2.1 of Poterba, Venti and Wise, “New Estimates of the Future Path of 401(k) Assets” in “Tax Policy and the Economy, Volume 22” (2008).

²The Roth individual retirement account (IRA) was created by the Taxpayer Relief Act of 1997. It was named for its sponsor, Senator Roth.

³This cost can be offset by increases in future taxes, but there may still be societal costs that make myopic revenue acceleration undesirable. See Section 3 for a more complete discussion.

⁴As of 2012, only \$310 billion in assets were in Roth IRAs, as contrasted with \$4,740 billion in traditional IRAs. Investment Company Institute, “The U.S. Retirement Market, First Quarter 2013” (June 2013), Table 7, available at http://www.ici.org/info/ret_13_q1_data.xls. The figures for 2002, by contrast were \$78 and \$2,322 for Roth and traditional IRAs, respectively. *Id.* Thus, although Roth accounts hold only about 6.5% of the assets of traditional accounts currently, they have grown in asset size by a factor of nearly four over the last decade, while traditional accounts only doubled in asset value during the same period.

⁵See, for example, President Obama’s endorsement of “myRAs” in his January 28, 2014 State of the Union Speech. These accounts are conceptualized as starter accounts marketed toward non-high-income taxpayers who do not have access to employer retirement plans. See also the discussion draft released by Chairman Camp and the House Ways and Means Committee on February 26, 2014, which eliminates the income eligibility limit for Roth IRAs.

insight into the choices facing retirement account holders.⁶ I also analyze the revenue raising tools and the incentives that the structure of the current system gives for the government.

Not surprisingly, from the point of view of taxpayers, I find good and bad features throughout the current system. Traditional accounts may provide more utility to taxpayers than Roth accounts if they give them access to lower tax rates in the future, although this utility will come at the cost of over-encouraging savings relative to the Roth situation. In addition, traditional accounts, but not Roth accounts, can distort investment incentives and savings behavior by levying a non-linear progressive tax on distribution amounts that include investment returns. Finally, the choice that confronts taxpayers between a traditional account and a Roth account is fraught with complexity and uncertainty, particularly because of the risk of future changes in tax rates or brackets.

From the point of view of the government, I also find difficulties with the current system. In particular, I find that the ready availability of Roth accounts as a short-term revenue raising tool may give the government an excessive incentive for myopic acceleration of revenue receipts. Such an acceleration need not be innocuous, as theories of Ricardian equivalence would suggest, but will likely produce a long-term net societal cost.⁷

Having analyzed the current system, I propose a new entity, the smooth retirement account. The key idea underlying my proposal is that payment of taxes in connection with a retirement account should be spread over time, rather than lumped all in the beginning (when contributions are made) or all at the end (when distributions are made).⁸ In addition, much like Roth accounts, taxes payable should not be determined by reference to investment performance in the account.

My proposal preserves the most desirable attributes of the current system while also eliminating the negative ones. The fact that taxes are reckoned without reference to account investment performance means that taxes on the new accounts do not distort investment decisions. In addition, the requirement that taxes be paid over time decreases the likelihood that the government will need to engage in myopic acceleration of revenue, since a regular stream of revenue will be available.

⁶In the appendix I supplement these theoretical models with a Monte Carlo simulation analysis designed to incorporate many realistic features of the current law regarding retirement taxation and social security.

⁷See Section 3 for further discussion.

⁸[NB: The notion of “smoothness” will be explicated in more detail in the final paper. In particular, from the government’s point of view, smoothness may not be a concern if the population is homogeneous over time and as many people are contributing to retirement accounts as are leaving them at any given point. In this case, my attempt to smooth government receipts may not actually make much difference because things would already be very smooth. In reality, however, the population is not homogeneous, and overall contributions to retirement accounts are still increasing. The precise degree to which the demographics are un-smooth is an empirical question that needs to be discussed further and supported with data, and I plan to do so in the final draft.]

The remainder of this paper proceeds as follows. In Section 2, I analyze the two main choices, traditional and Roth accounts, faced by taxpayers in the current system. I develop a simple lifetime utility maximization model for investor choice, and I determine contribution amounts and their sensitivities to parameters analytically.⁹ In Section 3, I analyze the current system from the perspective of the government, and I argue that it provides an undesirable incentive for the government to behave myopically and accelerate revenue now at the cost of future years. In Section 4, I articulate the details of my new proposal, highlighting the ways in which it fixes existing problems, and also addressing potential concerns and problems. In Section 5, I summarize and conclude.

2 Choices and Problems for Taxpayers

In this section, I focus on the choices and problems faced by taxpayers under the current rules. In Section 2.1, I analyze the choice between traditional and Roth accounts under the simple assumption that tax rates remain constant over time. In Sections 2.2 and 2.3, I analyze savings and portfolio allocation decisions for each type of existing accounting using a lifetime utility maximization model.

There are a variety of problems under the current rules that I do not address here because they are already well understood and can be fixed relatively easily. For example, the annual contribution limits for Roth accounts are effectively greater than they are for traditional accounts because of the fact that the numerical limits are the same but the former is funded with after-tax dollars.¹⁰ In addition, the required minimum distribution rules for traditional accounts generally require faster distributions than the corresponding rules for Roth IRAs. I recommend that these and all similar types of inconsistencies between the rules for various retirement accounts be eliminated so as to reduce complexity in decision-making. For purposes of this paper, I assume that these inconsistencies have already been eliminated.

2.1 Traditional vs. Roth: Overview

The decision between a traditional and a Roth type of account appears at first not to be very complex.¹¹ The simplest possible situation to analyze occurs when tax is levied at a constant

⁹In the appendix, I also address the problem numerically using Monte Carlo simulations.

¹⁰Consider, for example, the \$5,000 annual contribution limit applicable to both types of IRAs in 2012. For a taxpayer in the 25% bracket, a pre-tax contribution of \$6,667 to a Roth IRA is possible, since this amount is \$5,000 after taxes. However, for the same taxpayer, only a \$5,000 pre-tax contribution is possible to a traditional IRA account.

¹¹The option to choose one of the two types of accounts may not be available. For example, taxpayers with sufficiently high incomes are generally not permitted to make contributions to Roth IRAs.

rate, τ , on all income and at all times. In this case, if a contribution of C pre-tax dollars is contemplated, and a return of r is what will be earned over the time until distribution on any amount invested, then the final distribution amounts from a traditional account and a Roth account are the same, namely

$$C \underbrace{\underbrace{(1 - \tau)}_{\text{Reduction for Initial Tax}} \underbrace{(1 + r)}_{\text{Asset Growth Factor}}}_{\text{Roth distribution}} = C \underbrace{\underbrace{(1 + r)}_{\text{Asset Growth Factor}} \underbrace{(1 - \tau)}_{\text{Reduction for Final Tax}}}_{\text{traditional distribution}}. \quad (1)$$

This analysis fails to capture reality, however, because tax are not general constant over time. Changes in a taxpayer's income may place him in a different bracket during distribution years than during contributions years. It may also happen because Congress alters the applicable tax rates between the time of the contribution and the time of the distribution. As a result, the assumption that tax rates remain the same may not be appropriate. A more correct comparison to make is

$$C \underbrace{\underbrace{\left(1 - \frac{T_1[C]}{C}\right)}_{\text{Reduction for Initial Tax}} \underbrace{(1 + r)}_{\text{Asset Growth Factor}}}_{\text{Roth distribution}} = C \underbrace{\underbrace{(1 + r)}_{\text{Asset Growth Factor}} \underbrace{\left(1 - \frac{T_2[C(1 + r)]}{C(1 + r)}\right)}_{\text{Reduction for Final Tax}}}_{\text{traditional distribution}}, \quad (2)$$

where T_1 is the function that determines the tax on the incremental amount of income C at the time of contribution and T_2 is the function that determines the tax on an incremental amount of income $C(1 + r)$ at the time of distribution, with the understanding that the exact value of $T_2[C(1 + r)]$ will generally not be known with certainty until the time the distribution is made. It is not unlikely that

$$\frac{T_1[C]}{C} \neq \frac{T_2[C(1 + r)]}{C(1 + r)}, \quad (3)$$

and so the final distribution amounts may well fail to be the same.

It is clear from (2) that the right hand side of the equation will be greater if T_2 represents a sufficiently low tax rate relative T_1 . Using this insight, it is often said that taxpayers who anticipate lower income levels, and consequently lower tax rates, in the future may expect higher average retirement cash flows by choosing traditional accounts instead of Roth accounts. In this way they will be able to take better advantage of the lower effective tax rates they will enjoy in the future. In fact, this common wisdom may not fully reveal the

truth of the incentives taxpayers have to choose traditional accounts. The economic model in the following section, along with the numerical simulations in the appendix, suggest that traditional accounts may be preferable to taxpayers, but only at the cost of significant distortions in savings behavior and portfolio allocation.

2.2 Utility Maximization: Contribution Choice

A taxpayer’s decision about how much to invest in a retirement account can be thought of as a lifetime utility maximization problem. In this section, I analyze this problem with a formal economic model. To keep the analysis manageable, I consider a simple situation in which there are only two periods. The taxpayer has an exogenously given amount of pre-tax income, W during the first period. He chooses an amount, C , to place in a retirement account, he pays any taxes due, and then he consumes the portion of W that remains after the contribution and payment of taxes. In the second period, he withdraws all amounts from his retirement account, pays any taxes due, and consumes all after-tax amounts remaining. This model is stated broadly enough to encompass the possibility that the retirement account is either a traditional or a Roth IRA, and I will consider each possibility in turn.

Traditional IRAs

Suppose that the taxpayer chooses to invest in a traditional IRA. The discounted expected utility he seeks to maximize is

$$V_{\text{trad}}(C) = U((W - C)(1 - \tau_0)) + \delta E[U(CR(1 - \tau_1))], \quad (4)$$

where τ_0 is the tax rate in the first period, $0 < \delta < 1$ is a discount factor, R is the uncertain pre-tax return factor for the retirement account portfolio, τ_1 is the uncertain future tax rate. The function U is a the taxpayer’s utility function, and the operator E represents the expectation, i.e., probability weighted outcome, of unknown future values. I assume that there are two states of the world in the second period. I denote these states with “ d ” and “ u ”, indicating a downside and an upside possibility, respectively, in which the investment portfolio performs relatively worse or better. Table 2.2 indicates the probability of each state of the world and provides further notation regarding the values of the variables in each state.

To determine the optimal level of contribution, C , for a taxpayer, it is necessary to specify a particular utility function. I assume constant relative risk aversion so that

$$U(x) = \frac{x^{1-\alpha} - 1}{1 - \alpha}, \quad (5)$$

	State d	State u
Probability	p	$1 - p$
Portfolio pre-tax return factor (R_1)	R_d	R_u
Tax rate in second period	τ_d	τ_u

Table 1: Notation for variables in two possible future states of the world.

for a suitable $\alpha > 0$. Under this assumption, the second derivative of V_{trad} with respect to C satisfies $\frac{\partial^2 V_{\text{trad}}}{\partial C^2} < 0$, provided that $0 < C < W$, $0 < \tau_0, \tau_1 < 1$, and $R > 0$. The maximum value of V_{trad} is therefore obtained at the choice of C for which $\frac{\partial V_{\text{trad}}}{\partial C} = 0$. Performing the necessary differentiation, one finds that

$$\frac{\partial V_{\text{trad}}}{\partial C} = -(1 - \tau_0)^{1-\alpha} (W - C)^{-\alpha} + \delta C^{-\alpha} \|R(1 - \tau_1)\|_{1-\alpha}^{1-\alpha},$$

where I have used the notation $\|\cdot\|_{1-\alpha}$ to denote the $L^{1-\alpha}$ norm, namely,

$$\|R(1 - \tau_1)\|_{1-\alpha} = (pR_d^{1-\alpha}(1 - \tau_d)^{1-\alpha} + (1 - p)R_u^{1-\alpha}(1 - \tau_u)^{1-\alpha})^{\frac{1}{1-\alpha}},$$

in order to make the formulas easier to read. The value of $\frac{\partial V_{\text{trad}}}{\partial C}$ is zero when

$$C = W \left(1 + \delta^{-\frac{1}{\alpha}} \|R\rho\|_{1-\alpha}^{1-\frac{1}{\alpha}}\right)^{-1}, \quad (6)$$

where $\rho = \frac{1-\tau_1}{1-\tau_0}$ is the ratio of the factor by which returns are reduced for taxes in the second period to the corresponding factor for the first period. The value of ρ will either be ρ_d or ρ_u , according to whether τ_1 is equal to τ_d or τ_u . The value of C given by (6) is the optimal level of retirement contribution for the taxpayer.

Using the formula for C in (6), we can determine the sensitivity of C with respect to changes in parameters, such as τ_d and τ_u . For example,

$$\frac{\partial C}{\partial \tau_d} = \left(\frac{\alpha - 1}{\alpha}\right) \left(\frac{C^2}{W}\right) \delta^{-\frac{1}{\alpha}} \|R\rho\|_{1-\alpha}^{\alpha-\frac{1}{\alpha}} (p(R_d\rho_d)^{1-\alpha}) (1 - \tau_d)^{-1}. \quad (7)$$

The formula for the derivative with respect to τ_u is entirely similar, except that τ_d , ρ_d , p , and R_d are replaced by τ_u , ρ_u , $(1 - p)$, and R_u , respectively.

From (7), one sees that the sign of $\frac{\partial C}{\partial \tau_d}$ depends on whether α is less than or greater than one. When $\alpha > 1$, the taxpayer assigns a very high penalty to a very small amount of consumption in either period, with the penalty being infinite when consumption is equal to zero. The intuition, then, is that consumption should be roughly balanced between the two periods, with neither period receiving much less than the other. Accordingly, as τ_d increases,

so that the tax rate in the second period in the “ d ” state of the world becomes relatively larger compared to that in the first period, the optimal value of C will increase as well, somewhat for the higher tax burden in the second period and leave consumption across the two periods more balanced. This explains why $\frac{\partial C}{\partial \tau_d} > 0$ when $\alpha > 1$, and similarly for the partial derivative with respect to τ_u .

The result of the last paragraph runs contrary to what conventional wisdom might predict. For relatively risk averse investors, with $\alpha > 1$, higher tax rates in the future mean more money invested in retirement plans. As explained above, this makes intuitive sense because the taxpayer desires to avoid excessively low consumption in the second period. If there were a floor to possible consumption, provided perhaps by a system like social security, the need might not be as great to put very large amounts of funds into the IRA to offset the effects of the high future tax rate.

When $0 < \alpha < 1$, so that the investor is not so risk-averse, the sign of $\frac{\partial C}{\partial \tau_d}$ is negative, and similarly for the derivative with respect to τ_u . This reflects the fact that the taxpayer does not associate arbitrarily large penalties with very low levels of consumption. As a result, a relatively high tax in the second period means that the taxpayer will choose to save less in an IRA and consume more in the relatively low-tax environment of the first period instead.

Roth IRAs

The analysis for Roth IRAs is very similar to that for traditional IRAs, but it is simpler because there is only one tax rate to consider, namely the one applicable in the first period. The value the taxpayer seeks to maximize in this case is

$$V_{\text{Roth}}(C) = U((W - C)(1 - \tau_0)) + \delta E[U(CR(1 - \tau_0))]. \quad (8)$$

In this formulation, C represents the pre-tax amount that is set aside to be contributed to the Roth IRA. Because the Roth IRA requires after-tax dollars to be contributed, only the amount $C(1 - \tau_0)$ is actually invested, and this grows to the size $CR(1 - \tau_0)$ in the second period, where R is unknown from the perspective of the first period, but C and τ_0 are fixed.

Because (8) has exactly the same form as (4), with τ_u and τ_d both set equal to τ_0 , the above analysis for traditional IRAs carries over perfectly. If we continue to assume a utility function of the form (5), then the optimal value of C satisfies

$$C = W \left(1 + \delta^{-1/\alpha} \|R\|_{1-\alpha}^{1-1/\alpha} \right)^{-1}, \quad (9)$$

which is analogous to (6), except that all the tax rates are the same and so $\rho \equiv 1$. From (9)

it is clear that the optimal value of C does not depend in any way upon the tax rate. As a result, the savings choice for the Roth is what it would be in the absence of taxes and is not distorted by the taxation.¹²

2.3 Utility Maximization: Allocation Choice

The utility maximization framework can also be used to analyze the choice a taxpayer makes of what fraction of his retirement portfolio he should allocate to risky assets. The idea is to decompose the return factor R into components, namely,

$$R = R_f + \omega(R_m - R_f), \quad (10)$$

where R_f is the return factor for a risk-free investment, R_m is the return factor for an investment in a risky asset, and ω is the fraction of the portfolio invested in the risky asset. As above, these return factors represent the amount by which the amount initially placed in a retirement account should be multiplied in order to determine how much is available for distribution before any tax is imposed on the withdrawal.¹³ The value of $R_m - R_f$ is unknown as of the first time period, and it will have a higher or lower value depending upon whether the state of the world is u or d . It is convenient to write Δ_u and $-\Delta_d$ for the value of the excess return, $R_m - R_f$, in these two respective cases. Thus, the return factor for investments is either $R_u = R_f + \omega\Delta_u$ or $R_d = R_f - \omega\Delta_d$, and it is assumed that $\Delta_u, \Delta_d > 0$, so that it is possible for the risky asset to provide either a better or a worse return than the risk-free asset.

Traditional Accounts

The function V_{trad} from (4) can be written in a way that emphasizes its dependence on ω as well as C , namely,

$$V_{\text{trad}}(C, \omega) = U((W - C)(1 - \tau_0)) + \delta E[U(CR(1 - \tau))]. \quad (11)$$

The dependence on ω is present by virtue of the fact that R depends on ω , as described in (10). The second derivative of V_{trad} with respect to ω satisfies $\frac{\partial^2 V}{\partial \omega^2} < 0$, and so there is a

¹²It is still of course possible that taxes could distort factors other than those considered here. For example, the level of tax may impact labor decisions which could affect the size of W . For purposes of the model in the text, though, W is exogenous, and conditional on this assumption, there is no distortion to savings behavior when a Roth IRA is used.

¹³Because of this notational convention, R_f is a number greater than 1, equal to 1 plus the risk-free rate of return.

unique maximum at the value of ω for which $\frac{\partial V}{\partial \omega} = 0$. This optimal value of ω satisfies

$$\frac{U'(CR_u(1 - \tau_u))}{U'(CR_d(1 - \tau_d))} = \frac{p\Delta_d(1 - \tau_d)}{(1 - p)\Delta_u(1 - \tau_u)}.$$

If we assume that U is of the same form specified in (5), the optimal value of ω is seen to satisfy

$$\frac{R_f + \omega\Delta_u}{R_f - \omega\Delta_d} = \frac{D_u}{D_d},$$

where

$$D_d = p^{\frac{1}{\alpha}}\Delta_d^{\frac{1}{\alpha}}(1 - \tau_d)^{\frac{1}{\alpha}-1} \quad \text{and} \quad D_u = (1 - p)^{\frac{1}{\alpha}}\Delta_u^{\frac{1}{\alpha}}(1 - \tau_u)^{\frac{1}{\alpha}-1}.$$

Solving for ω , one finds that the optimal value of ω is¹⁴

$$\omega = R_f \left(\frac{D_u - D_d}{D_u\Delta_d + D_d\Delta_u} \right). \quad (12)$$

The derivative of ω with respect to τ_u is

$$\frac{\partial \omega}{\partial \tau_u} = R_f \left(1 - \frac{1}{\alpha} \right) \left(\frac{D_u D_d}{1 - \tau_u} \right) \left(\frac{\Delta_u + \Delta_d}{(\Delta_d D_u + \Delta_u D_d)^2} \right). \quad (13)$$

The formula in (13) shows that the sign of $\frac{\partial \omega}{\partial \tau_u}$ is positive or negative according to whether $\alpha > 1$ or $\alpha < 1$, respectively. The two possibilities arise because an increase in τ_u produces somewhat mixed effects, both a lower possible return but also less risk, in the “ u ” state of the world. For a taxpayer with a sufficiently low level of risk aversion, namely $\alpha < 1$, the effect of the lower return dominates and the optimal investment ω in the risky asset decreases as τ_u increases. For a taxpayer with a level of risk-aversion above $\alpha = 1$, the effect of the lower level of risk dominates, and the value of ω increases as τ_u increases.

[More discussion to be included here. The final draft will include further discussion of the foregoing results regarding contribution and allocation choices. The key insights and results have already been identified above. The additional material will include discussion of results from the Monte Carlo simulations in the appendix, as well as an evaluation of the economic size and importance of the various distortions under traditional accounts.]

¹⁴To accomplish the calculation, it is convenient to use the fact that the inverse of the Möbius transformation $y = \frac{ax+b}{cx+d}$ is simply $x = \frac{dy-b}{-cy+a}$.

3 Incentives for the Government

Section 2 examined in detail the choices and problems that the current retirement account rules create for taxpayers. In this section, I consider the ways in which the existing rules create sub-optimal and undesirable incentives from the government's perspective.

A simple but important point is that Roth accounts provide an immediate supply of revenue, since tax is paid now instead of later, but traditional accounts, in contrast, push revenue into the future. From a theoretical perspective, under appropriate assumptions, the timing of a tax should not matter, since the government can simply issue debt to finance future taxes if it has a short-term cash flow problem.¹⁵ From a practical vantage point, however, the timing of governmental tax revenue can make a difference.¹⁶ In particular, one of the justifications for allowing taxpayers to invest in Roth-type accounts has been that this will generate revenue for the government in the near term.¹⁷

The short-term revenue raising mechanism provided by Roth accounts is of concern because it provides a particularly easy path by which funds can be raised. In more typical situations, it is politically costly to raise revenue through taxes. By providing taxpayers with the option (but not the obligation) to invest in Roth accounts, however, the government can avoid political objections from taxpayers. Moreover, many taxpayers will be willing to take the option and pay tax now, particularly if they are in a high bracket and believe that tax rates will increase in the future.¹⁸ From a societal perspective, this choice is likely to be myopic and suboptimal. The revenue collected now will come at the cost of revenue collected in the future, and future revenue needs will likely be greater than the current receipts, assuming investors in Roth accounts are proven to have made the right choice, from a personal perspective.¹⁹ This will place a larger future tax burden on taxpayers who did not opt for Roth accounts, and this shifting of burdens may be socially suboptimal.

In response to the concerns raised by the preceding paragraph, it can be argued that the

¹⁵This is the principle of Ricardian equivalence. See, for example, Barro, R. "Are Governmental Bonds Net Wealth?" *Journal of Political Economy* (1974).

¹⁶The limited 10-year budget window is an example of the principle that the government behaves in ways that focus disproportionately on the near term at the cost of the long term.

¹⁷This has been offered as a partial explanation for the original creation of Roth IRAs in 1997. It has also been cited as a reason for the temporary rule change that allowed wealthy taxpayers to convert from traditional to Roth accounts in 2010. Investment in Roth accounts during 2010 increased nearly 10 times over the investment in prior years. See Bryant, Victoria L. and John Goyer, "Accumulation and Distribution of Individual Retirement Arrangements, 2010," *Statistics of Income Bulletin*, Fall 2013.

¹⁸See Section A.2.

¹⁹The idea here is that many investors in Roth accounts are essentially making a bet that future tax rates will increase to such a degree that, even after adjusting for the time value of money, the payment of tax now at current rates will leave them better off than the payment of a future tax. As a result, if these taxpayers have bet correctly, the payment of taxes on the creation of Roth accounts does not adequately compensate the government for the future taxes that are foregone.

government may change the rules regarding distributions from Roth types of accounts in the future and thereby make up for any revenue shortfall that might be experienced. Indeed, the possibility of such a rule change is a central concern to those deciding whether to invest in Roth accounts. Scholars have examined the question of whether a future rule change of the required type would be permissible, and it appears that it would.²⁰ Nevertheless, the prospect of such a potential future change creates an uneasy environment in which taxpayers split into groups with competing political agendas for how the government should manage future taxes.²¹ In addition, even if legally permissible, the disappointment of investor expectations by the government in this area may erode confidence and trust of government promises more broadly.

In summary, the current rules create a structural environment in which the raising of revenue through Roth accounts is politically easy and desirable by the government in the short-term. This tack, however, is likely to come with long-term societal costs, either because non-Roth holders will be forced to take on a disproportionate share of future tax burdens or because the government will change the rules on distributions in the future, thereby incurring the political costs and loss of trust that come from disappointing taxpayer expectations. As a result, it would be preferable to restructure the rules in a way that does not make it as likely for the government to head down an undesirable and myopic path. The smooth retirement account that I propose in Section 4 accomplishes much of this goal.

4 Implementing SRAs

I propose a new type of retirement account, a smooth retirement account or “SRA”, to address the shortcomings of the retirement vehicles available under current law. There are two fundamental components to the proposal. The first addresses the problem of taxpayer investment distortions by levying tax at a rate that does not depend on investment performance. The second addresses concerns related to the government’s budget window by requiring tax on the account to accrue over the life of the account, rather than either entirely at the time of contribution or the time of withdrawal.

An SRA provides a single solution for both taxpayer and governmental problems. The two issues are analytically distinct, however, and accordingly each can also be addressed separately. The key property of SRAs that avoids distortion of taxpayer decisions is the specification that the applicable tax rate does not depend upon investment returns in the

²⁰See Crane, C., “Honoring Expectations About Taxes: Are Roth IRAs Different?” (November 12, 2009), available at SSRN: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1505120.

²¹Those who own Roth accounts would lobby for no change in the rules, and the rest would likely be in favor of a change, if the alternative was to bear a disproportionate share of future taxes.

account. A Roth account accomplishes this goal by levying all tax due at the time of contribution, before there are any returns to be considered. The same goal could be reached alternatively, however, by applying the same tax rate as a Roth account would levy on the pre-tax contribution dollars, but simply imposing the rate at the time of withdrawal instead. In this case, the rate would need to be applied not just to the original contribution amount but rather to the quantity to which the contribution has grown over time as a result of investment performance. Also, if there are many contributions over the life of the account, then the applicable tax rate would be an appropriately blended rate, taking into consideration the relative sizes of the contributions as compared to the existing account. In addition to levying the tax all up-front, as with a Roth, or all upon withdrawal, as just described, the tax could be levied over time, and the details of this intermediate methodology are laid out in the description of SRAs below. If the goal is simply to avoid taxpayer distortions, any of these methods—front-loading, back-loading, or taxation along the way—will achieve the same goal. The only requirement is that the applicable tax rate be the appropriately blended rate determined as of times of contribution, rather than determined after investment performance has been realized.

The key property of SRAs that addresses governmental concerns is the spreading out of tax over time, rather than all at the time of contribution or all at the time of distribution. If we focus solely on this perspective, the particular tax rate applied does not matter, and it may even be chosen to depend upon investment performance. What is critical is that the government only accrue an appropriate portion of taxes with regard to the account during any given budgetary window. Currently, Roth accounts and traditional accounts either overstate or understate, respectively, the tax ultimately to be collected by the government when the artificial cut-off of the budget window is imposed. The SRA corrects this problem by spreading out the timing of tax evenly between the time of contribution and the end of the life of the account, or withdrawal, whichever occurs sooner. This can be accomplished either in a virtual way, through specification of accounting rules, as described further in Section 4.1, or by adjusting actual cash flows from tax collection, as described further in Section 4.2.

It is important to note that implementation of SRAs would not be a simple task. With respect to determining a correct blended rate for taxpayers, it would likely be necessary for taxpayers to compute their tax due in the year of a contribution in two separate ways. One would include the contribution in income, and the other would not. The difference between the two computations of tax due would indicate the appropriate tax rate to use with respect to that contribution. With respect to adjusting the government's rules, a significant change would need to be made either to governmental budgeting rules or to the timing of collection

of cash flows from retirement accounts, and this would likely be a substantial undertaking. Nevertheless, it is entirely possible to implement SRAs in principle. The decision whether it is advisable to do so should be based on the costs of implementation just mentioned, and how they compare to the savings solving the problems of taxpayer distortions and governmental myopia.

4.1 Virtualization: The Accounting Approach

In this section, I describe a way to implement SRAs by changing the rules the government uses to account for future tax receipts from retirement accounts. This “virtualization” approach allows all actual tax cash flows to occur either at the time of distributions or contributions, as with traditional or Roth accounts, respectively, but it also requires that for budgeting purposes the government must act as though the accounts were SRAs instead of traditional IRAs. In addition, I describe how the tax rate levied on distributions can be adjusted to represent a blended rate reflecting prevailing rates over the course of the contribution years, rather than using the income tax brackets otherwise in force at the time of the distributions. These two steps, changing government accounting rules and changing the rate of tax on distributions, are designed to address the issues under the current system for the government and for taxpayers, respectively, that were discussed in Sections 2 and 3.

It is easiest to start with the case of a single contribution and a single distribution. Suppose that a taxpayer wishes to contribute C pre-tax dollars to create an SRA at the end of year 1 and plans to take a distribution of everything in the SRA at the end of year L . The idea is basically to account for tax on a fraction $f = 1/L$ of account value during every year of its life. In addition, adjustments must be made each year to previously accounted for amounts to reflect ongoing changes in account value. As of year L , the full amount of tax will thus have been accounted for. The actual levying of tax and resultant cash-flow may occur either at the time of contribution or distribution. To avoid too much complexity I will focus on a scheme that taxes at the time of distribution. In addition, keeping actual cash-flows at the time of distribution better serves the purpose of addressing concerns about government spending raised in Section 3.

Let τ be the rate of tax that would be applicable to the last C dollars of the taxpayer’s income in year 1, but for the contribution. That is, if the taxpayer’s income in year 1 before making the contribution is I , the value of the product $C\tau$ is the difference between the tax on an income of I and a tax on an income of $I - C$.²² Let V_t be the account value at the end of year t , without adjustment for any taxes accounted for, so that $V_1 = C$ and V_L is the

²²I define the rate τ in this way to be account for the fact that tax rates applicable to income may not be flat.

final pre-tax distribution amount. For each year $1 \leq t \leq L$, the government should take into account an amount determined in such a way that the cumulative amount accounted for up through that time is $tf\tau V_t$. In the first year, the amount of tax accounted for should be exactly $f\tau V_1 = f\tau C$. The additional amount taken into account by the government in future years is exactly the amount required to bring the total accrual up to the level $tf\tau V_t$ at each time t . The incremental tax accounted for may be negative in certain years, if the value of the account has diminished by a sufficiently large factor.²³ At the time of the distribution at the end of year L , an actual tax should be levied in the amount τV_L . Because $Lf\tau V_L = \tau V_L$, the virtual tax that has been accounted for exactly matches the amount ultimately collected.

It is important to the scheme just described that the value of the account is known each year. To the extent retirement account holdings consist of market securities, this should not generally be a problem. If less liquid assets are held in accounts, however, exact values may be difficult to determine. In this latter case, special rules would be necessary for determining approximate account values or proxies therefore each year. One example would be to use the last known value, increased by the compounded risk-free rate of return since that prior date. At a future time when the value becomes known, that new value can be used and the procedure described above will automatically take into account the adjustment from the previous proxy value to the real new value. It is important to note, however, that the vast bulk of retirement account holdings are liquid, and so special rules of this sort would not be required in most cases.²⁴

To move beyond the case of a single contribution and single distribution, consider a taxpayer who makes multiple contributions to his SRA. With each successive contribution an update to the fractions described above is generally necessary. For example, the remaining life of the account will now be shorter than it was in the past, and so tax on the contribution should be taken into account on an annual basis in fractions corresponding to this shorter length. In addition, the income tax rate applicable at the time of the new contribution may not be the same as it was in the past. In order to deal with these issues, it is convenient to form weighted averages of the previously applicable parameters and new parameters corresponding to the new contribution.

Suppose that contributions are made at the ends of various years t over the life of the account. For each such year t , let C_t denote the amount of the contribution, and let V_t denote the value of the account at the end of year t immediately *prior* to the contribution. As above, the value V_t is pre-tax and does not reflect any taxes accrued by the government.

²³As an extreme example, suppose that the account value falls to zero. In this case, the appropriate amount of tax to reflect on the government's books will also be zero, even though prior amounts were positive.

²⁴[Cite needed here about statistics regarding retirement account holdings.]

For each contribution C_t , let L_t be the remaining years of account life at that time, counting year t as the first year. Thus, the life of the account ends at time $t + L_t - 1$. Also, let τ_t be the rate of tax that would be due on the final C_t dollars of income during year t .

At the time of the initial contribution, C_1 , define $g_1 = \tau_1 f_1$, where $f_1 = 1/L_1$. For $t > 1$, define

$$g_t = \left(\frac{V_t}{V_t + C_t} \right) g_{t-1} + \left(\frac{C_t}{V_t + C_t} \right) \tau_t f_t. \quad (14)$$

In addition, define $G_1 = g_1$, and for $t > 1$, define

$$G_t = g_t + \left(\frac{V_t}{V_t + C_t} \right) G_{t-1}. \quad (15)$$

With this notation, the incremental tax that should be taken into account by the government for accounting purposes in year 1 is $G_1 V_1$. For each year t with $t > 1$, the correct amount is $G_t V_t - G_{t-1} V_{t-1}$. As before, this amount may be negative in years in which the value of the account decreases by a sufficient amount. The total amount of taxes taken into account by the government at each time t is thus $G_t V_t$.

In the case of a single contribution over the life of the account, $G_t = t g_1$, and the situation is exactly the same as described above in the context of one contribution. In the general situation, G_t can be thought of as the cumulative amount that would have been taken into account by the government if each contribution were to a separate retirement account, so that each account was a single-contribution account, with the further assumption that the investment strategy is the same throughout all the accounts. Thus the contributions are all treated as homogeneously co-mingled and invested uniformly across the assets held in the account.

With the ideas and notation just developed, it is straightforward to specify the treatment for distributions prior to the end of the account life. Assume that a distribution D_t is made at the end of year t , immediately *after* tax has been accrued for the entire account by the government for that year. In this case, without any further contributions, the fraction of account value that will be taxed in the aggregate in future years is $g_t(L_t - t + 1)$, that is, the currently applicable annual fraction g_t multiplied by the remaining years in the life of the account. Accordingly, the tax due upon the distribution should be an amount $D_t g_t(L_t - t + 1)$. This cash flow may be collected from the taxpayer, and accounted for, by the government at the time of the distribution, and thereafter the distributed dollars will have been fully taxed.²⁵ The remaining value in the account will continue to be taxed as specified above, with no change either to the g or the G values as a result of the distribution. If there should

²⁵If it is desired to assign a penalty to distributions that occur before a specified time, such a penalty may be assessed in addition to the tax due as described.

be further contributions, then adjustment to the g and G values for application to amounts remaining in the account would be made at the time of such future contributions, but they would have no effect on the previously completed distribution.

Throughout the foregoing discussion I have assumed that there is a fixed life for the account. This is necessary so that the values L_t can be determined with certainty. I contemplate that an SRA will have a target life, which may depend on the life-expectancy of the taxpayer. During the life of the account, there will generally be contributions during working years and distributions during retirement years. This target life can be used as the basis for determining the L_t values. If the account life is short, because of the death of the taxpayer or for other reasons, the remaining tax due may be imposed at that time. This remaining tax would be collected and accounted for in the same way as the tax on a distribution prior to the end of the life of the account. If the target life is exceeded, then no further tax would be due with respect to amounts in the account. I would recommend that distribution requirements be established so that all amounts should be withdrawn by the taxpayer by the end of the target life of the account, or perhaps until the death of the taxpayer, if that comes later.

4.2 Actualization: The Cash-Flow Approach

Instead of taking the accounting approach described above, it may be desirable to the government to undertake actual collections and receive real cash flows as the accounting takes place. This can be accomplished using a scheme entirely similar in spirit to that in the accounting approach. The one wrinkle is that payment of taxes due during the life of the account will decrease the account value.²⁶ The value thus needs to be grossed up in order for the above accounting rules to carry over, because they are structured to work with an account from which no taxes have actually been withdrawn.

In order to accomplish the necessary gross-up, one need only keep track of a fraction by which actual account value, A_t , needs to be multiplied in order to obtain the virtual account value, V_t , which is the same as the value of V_t discussed above. With each payment of tax, this fraction will be multiplied by the ratio the actual account value just prior to payment of the tax bears to the actual account value just after payment of the tax. Multiplication by the cumulative ratio at any point in time converts A_t to V_t , and then the procedure for accounting for taxes above can be used to determine actual taxes due.

There are difficulties with the cash-flow approach. It puts more pressure on the difficulties of valuing assets in the account. Also, it will often require liquidation of certain account assets

²⁶It is important that taxes not be paid using funds outside the account because this would lead to the possibility of implicitly exceeding contribution limits that may be imposed on the account.

to pay the tax. In addition, if the account values fall sufficiently far, taxes may be due from the government to the accounts, and paying such actual cash-flows may be undesirable from the government's point of view.

On the positive side, actual cash-flows may seem less like an accounting "gimmick," and it may be harder for Congress to tweak and adjust a cash-flow based system in ways that may only serve myopic short-term interests of lawmakers currently in power. Nevertheless, Congress has arguably shown a willingness to make just such adjustments in our current system of cash-flow based retirement accounts by making conversions to Roth accounts temporarily easier. Thus a cash-flow system may not offer complete protection from such problems.

5 Conclusion

In conclusion, I have analyzed the current retirement account rules from a theoretical perspective using lifetime utility maximization models to understand optimal savings and investment allocation choices for investors. I have found both good and bad features throughout the current system. Traditional accounts can allow for enhanced utility for many individual investors, but they also distort portfolio allocation and savings decisions. Roth accounts do not cause the same such distortions, but they also do not provide as much utility to investors. In addition, Roth accounts increase the likelihood of myopic revenue acceleration by the government. Finally, the choice between traditional and Roth accounts that faces many taxpayers leads to significant complexity and uncertainty, particularly because of the possibility of future changes in income tax brackets and rates.

In light of the findings of my analysis, I propose creation of smooth retirement accounts. The key idea of the SRA is that payment of taxes on the account are spread evenly over time. In addition, tax on the account is levied at a rate that does not depend on account investment performance. The SRA therefore does not create distortions in portfolio allocation decisions, and if the tax rate is chosen appropriately, the *SRA* can encourage saving at the desired level. In addition, it stabilizes the stream of revenue available to the government over time, and it decreases the likelihood that the government will need or be able to engage in myopic acceleration of revenue. I therefore urge adoption of the SRA as an alternative that can keep the best aspects of tax-advantaged retirement accounts while leading to improvements for taxpayers and the government alike.

A Appendix: Numerical Lifetime Utility Model

[This appendix is a very incomplete and preliminary version of the analysis I intend ultimately to present, but it provides an indication of my methodological approach and presents some initial findings.]

In this appendix I use a Monte Carlo simulation methodology, coupled with realistic tax rules and rates, including social security taxes, to develop further understanding of distortions faced by taxpayers investing in traditional retirement accounts.

To formalize the taxpayer’s decision-making process in choosing between a Roth and a traditional account, assume that he seeks to maximize his aggregate discounted expected utility from consumption over the course of his life. That is, he seeks to maximize the value of

$$\text{Lifetime Utility} = \sum_{t=0}^L \delta_t E[U(C_t)],$$

where the investor makes decisions at time $t = 0$ for a lifetime of length L , where δ_t is a discount factor applicable to utility at time t in the future, and where $U(C_t)$ is the utility of an amount of consumption C_t at time t . If the taxpayer must choose either a Roth or a traditional account for his lifelong savings, then he can determine the maximum lifetime utility he can attain through each choice and make the choice that leads to the higher overall utility.

Constraints on the taxpayer’s characteristics and choices are in order. I assume that he makes his decision at age 32, that he will retire at age 67, and that he will live with certainty until exactly age 87. I choose 67 as the retirement age at which full Social Security benefits can be had. The age of 32 provides 35 years until retirement, and this happens to be the number of years of maximum earnings taken into account in computing Social Security benefits. I take 87 to be the age of death of the taxpayer based on IRS life expectancy tables.²⁷ I do not allow for mortality risk so as to keep the problem manageable and more narrowly focussed.

I assume further that the taxpayer earns a constant annual wage from age 32 to age 67, and the taxpayer cannot change this wage, so that I am not modeling the labor/leisure decision. I also assume that the taxpayer saves money only through his chosen retirement account.²⁸ Thus, all income in retirement years will come from benefits paid from his re-

²⁷See Publication 590 (2012), Table I, p. 94.

²⁸Sophisticated taxpayers can engage in complex planning with respect to joint management of investments inside and outside tax advantaged accounts. See, e.g., Robert M. Dammon, Chester S. Spatt, and Harold H. Zhang, “Optimal Asset Location and Allocation with Taxable and Tax-Deferred Investing,” 59 Journal of Finance 999 (2004). I assume a much simpler situation in which the taxpayer chooses only the option of investing through his chosen type of retirement account.

irement account and from social security.²⁹ The retirement account pays out a fraction $1/(87 - T)$ of its remaining value when the investor is T years old, for T between 67 and 87. This methodology smooths the payment of benefits over time and is in accord with the existing minimum required distribution rules for traditional retirement accounts.

With respect to investment options, I assume that the investor can choose between a risky asset, which I refer to as stock, and a risk-free asset. The stock has a higher expected return, but also has risk, while the risk-free asset provides a definite return.

The taxpayer does have three things he can choose. The first is the amount, C , of his annual contribution in pre-tax dollars to a retirement account. He may contribute any amount from 0 up to the limit of \$17,500 per year or his total salary, whichever is lower.³⁰ His second choice is the fraction ω_W to invest in stocks, with the remainder invested in the risk-free asset, during his working years. His third choice is the fraction ω_R to invest in stocks, with the remainder invested in the risk-free asset, during his retirement years. Each of these fractions must be between 0 and 1.³¹ The taxpayer's choice of C , ω_W , and ω_R must be made at the outset, when the investor is age 32.³²

Finally, the taxpayer is assumed to consume all after-tax un-invested earnings in the year in which they become available to him, either through payment of wages or benefits. The taxpayer thus seeks to maximize

$$V(C, \omega_W, \omega_R) = \sum_{t=32}^{66} \delta_t E[U(W - C - T_w)] + \sum_{t=67}^{86} \delta_t E[U(S + B_t - T_t)], \quad (\text{A.1})$$

where W is the fixed wage, C is the annual pre-tax contribution amount, T_w is the fixed tax on the wage after making the contribution, S is the Social Security benefit payable, B_t is the uncertain benefit payable from the retirement account in year t , and T_t is the uncertain tax due in year t . The tax payments due include Social Security withholding and adjustments for Social Security benefits paid. They also take into account income tax brackets, standard deductions, and personal exemptions.³³ The terms B_t and T_t are uncertain from the perspective of the taxpayer at age 32 because they depend upon investment performance over time.

Because the only stochastic terms in (A.1) are B_t and T_t , it is possible to simplify the

²⁹Details on my modeling of the Social Security benefits paid are included in the Appendix.

³⁰I choose \$17,500 as the limit for contributions to 401(k) accounts in 2013. This limit is not binding except for the highest-income taxpayers I consider.

³¹I thus prohibit borrowing or short selling by the taxpayer.

³²This rigidity is somewhat unrealistic, but it is imposed to make the problem more tractable.

³³See the appendix for more details about the parameters for the model relating to tax and social security benefits.

optimization problem by eliminating the expectation in the terms in the first summation, so that

$$V(C, \omega_W, \omega_R) = \sum_{t=32}^{66} \delta_t U(W - C - T_w) + \sum_{t=67}^{86} \delta_t E[U(S + B_t - T_t)], \quad (\text{A.2})$$

The problem is still difficult to solve analytically, however, because of the significant non-linearity in the income tax rules. To deal gain a better understanding of the solution to this problem, I introduce a numerical approach based upon Monte Carlo simulations in the next section.

A.1 Monte Carlo Simulation

Because the lifetime utility maximization problem is analytically difficult, I use Monte Carlo methods to solve the problem numerically for a particular type of utility function and particular assumptions about investment options. I use a utility function with constant relative risk aversion, namely

$$U_\alpha(C) = \frac{C^{1-\alpha} - 1}{1-\alpha}, \quad (\text{A.3})$$

for a suitable $\alpha > 0$. In what follows, I use $\alpha = 8$.³⁴ For the investment options, I model the stock using sampling of actual historic returns for the value-weighted average of the stock market. I correct returns for inflation so that the analysis can be carried out in real dollars. This is important because it allows me to keep tax brackets, social security rules, etc., all constant in today's dollars. Also, I assume that dividends from the stock are not reinvested in the stock but are instead invested in the risk-free asset. For the risk-free asset, I use the average historic return for T-Bills, adjusted for inflation. I also use this risk-free rate, r_f , to estimate δ_t in (A.1), with $\delta_t = e^{-r_f(t-32)}$. Summary statistics for the historic data I use are presented in Table A.1.

I construct 10,000 Monte Carlo paths, with 45 years represented along each path. A single path is a representation of an outcome that may occur with regard to stock price. For each year, and along each path, I randomly draw historic values from the market-value weighted stock return, adjusting the stock asset by the return exclusive of dividends, and adding the dividend to the investment in the risk-free asset. For any choice of (C, ω_W, ω_R) that the taxpayer makes in attempting to maximize V in (A.2), I can estimate the lifetime

³⁴It is important to choose α large enough to make the level of risk aversion realistic, and I found that values of α at the level of 2 or below tended to make the taxpayer invest exclusively in equity in all situations. This seemed to indicate to little risk aversion. It is also important not to choose α to be too large, or the utility function becomes so extreme that results can be unreliable. See, e.g., Rajnish Mehra and Edward C. Prescott, "The Equity Premium: A Puzzle," 15 *Journal of Monetary Economics* 145 (1985), p. 155, note 5. My results do not change in significant ways for α in the range of at least 6 through 10.

	Unadjusted		Inflation Adj.	
	Mean	SD	Mean	SD
Market Index	5.3%	20.0%	2.3%	20.2%
Dividend Rate	3.8%	1.4%	3.8%	1.4%
90-Day T-Bills	3.5%	3.0%	0.5%	3.6%

Table 1: Market returns, dividend rates, and T-Bill returns for the period from 1926 through 2012. The market index return is computed using the CRSP value-weighted market index, and it is exclusive of dividends. The dividend rate is also based on the CRSP value-weighted index. All returns are stated on a continuously compounded basis, so that a return of r corresponds to a gross change in value by a factor e^r at over the course of a year. Inflation adjustments are made by subtracting the rate of return to the CPI-U index, non-seasonally adjusted. Dividends are not adjusted for inflation.

expected utility. The procedure is simple. With respect to each path, I calculate the cash flows at all times, including all taxes paid, etc., I determine the value consumed in each year, and I calculate the utility of this amount using (A.3) with $\alpha = 8$. I then find the discounted sum of such utilities as prescribed in (A.1), but without taking expectations, because only one path is used. Finally, I average the sums over all paths in order to reflect the necessary expectations, and this average result is the numerical estimate of $V(C, \omega_W, \omega_R)$. This process can be repeated by a computer searching for the optimal value of V , and the optimizing set of choices $(C^*, \omega_W^*, \omega_R^*)$ can be found in this way.³⁵

Once the optimum choice of $(C^*, \omega_W^*, \omega_R^*)$ is known for a particular wage, I calculate the constant annual after-tax consumption amount that would lead to the same level of lifetime utility. That is, I find the unique constant K such that

$$V(C^*, \omega_W^*, \omega_R^*) = \sum_{t=32}^{86} \delta_t E[U(K)],$$

where the expectation is not necessary because K is constant. Thus,

$$K = U^{-1} \left(\frac{V(C^*, \omega_W^*, \omega_R^*)}{\sum_{t=32}^{86} \delta_t} \right).$$

The usefulness of this quantity K is that we can express the optimum utility in the equivalent form of constant lifetime consumption, and this is often an easier value to interpret.

The first notable result of my simulations is that, under the assumptions I have described, traditional accounts are generally preferable to Roth accounts. Figure A.1(a) shows this result graphically. The average difference between the equivalent lifetime constant con-

³⁵I used the Octave software package to do these calculations.

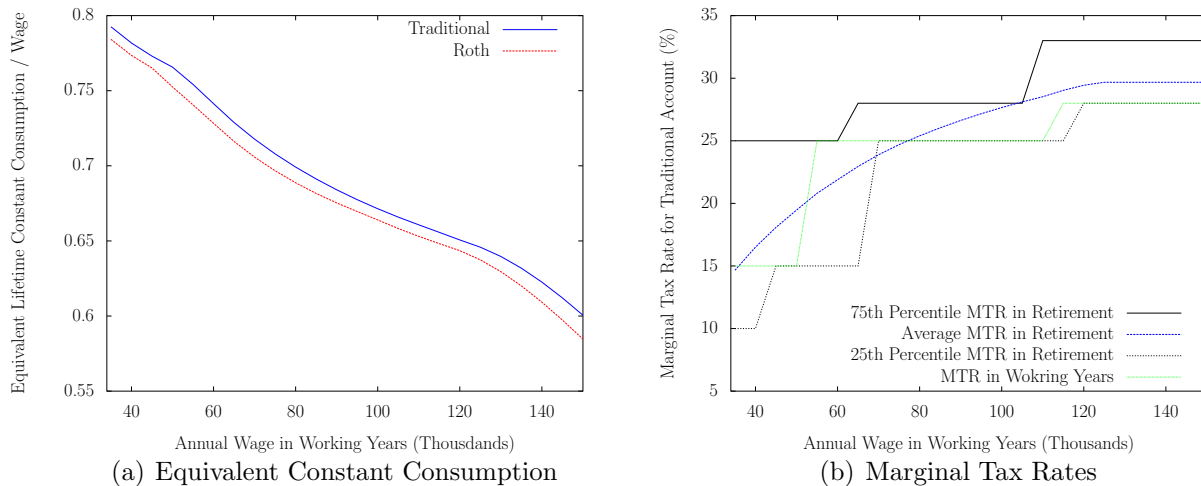


Figure 1: The diagram on the left shows the equivalent constant annual consumption amount for optimal lifetime utility, expressed as a fraction of wage during working years. The red and blue lines represent the optimal values in the case of Roth and traditional accounts, respectively. The diagram on the right shows the marginal effective tax rates for the traditional account holder. The marginal effective tax rate during working years is known with certainty. The marginal effective tax rate in retirement years is stochastic, and the mean value, as well as the 25th-percentile and 75th-percentile levels are shown. Calculations in both graphs are carried out for wages from \$35,000 to \$250,000 in increments of \$5,000.

sumption amounts is around 1.5% of their value. This lifetime utility gain does come with risk. Figure A.1(b) shows marginal tax rates in retirement for the traditional account, and the ends of the inter-quartile band are always in different brackets, meaning that large variability in income outcomes are occurring. This risk is factored into the optimization of lifetime utility, however, and so it is acceptable within the model. It is important to stress the dependence of this result upon the choice of utility function. Use of an alternative choice that penalized risk in retirement years more heavily might yield a different conclusion.

The second major result of my simulations is that traditional accounts lead to distortions in investment behavior, with taxpayers placing more assets in equities in retirement than they would if they were not subject to taxation, as is the case with Roth investors. Figure A.1(a) shows this graphically.³⁶ The average difference is around 4%, and it can be as high as 7%. The intuition for this result is relatively straightforward. The progressive rate structure penalizes risky returns relative to certain ones. When the progressivity is not too great, a taxpayer may seek to “undo” the additional tax burden by shifting more of his assets into the stock. This is particularly true in the case of the utility function in (A.3), because

³⁶Note that Figure A.1(a) only deals with the value of ω_R , which corresponds to investment allocation in retirement years only. The value of ω_W is equal to 100% for all investors in either type of account at all wages I considered.

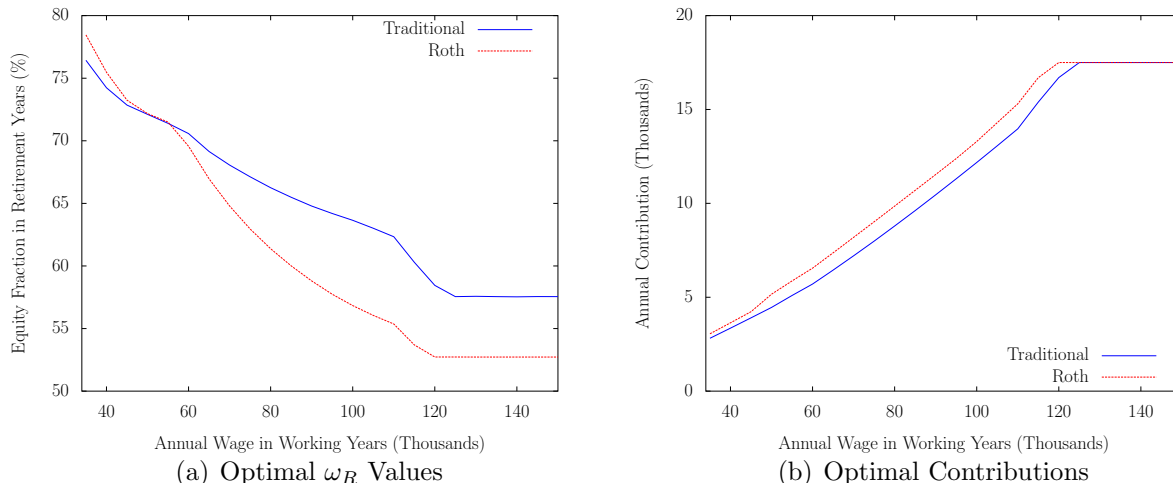


Figure 2: The diagram on the left shows the optimal choice of ω_R , the fraction of investment in equity during retirement. The diagram on the right shows the optimal choice of contribution amount. For both diagrams, the red and blue lines represent the optimal values in the case of Roth and traditional accounts, respectively.

the optimum value of that utility function applied to portfolio returns is scale-invariant, meaning that an investor will try to rebalance after-tax assets to get something equivalent, up to scale, to the pre-tax optimal balance. This logic is only approximately true, but it is compatible with the results for most wages in Figure A.1(a). Interestingly, below a wage of about \$50,000, where the progressivity of the applicable brackets becomes much more pronounced,³⁷ the opposite divergence from the Roth optimum occurs. In this situation, the penalty on the stock's risk is so large that it causes a distortion toward the risk-free asset.

The final important result of my simulations is that, as shown in Figure A.1(b), traditional account holders contribute fewer pre-tax dollars than Roth account holders with the same wage. Putting this result together with the previous ones paints a clearer picture of what is driving the preference for traditional accounts shown in Figure A.1. Traditional account holders generally pay a higher average marginal tax rate in retirement than their marginal rate during working years.³⁸ This means that a marginal dollar kept out of the retirement account is generally taxed less heavily, and so additional utility can be had by reducing savings. Moreover, further utility increase can be had by distorting the retirement portfolio to invest disproportionately heavily in stocks, relative to Roth investors who are tax-insensitive with respect to their retirement portfolio performance and can thus be taken as a benchmark of optimality. In sum, traditional account holders are achieving greater utility than Roth account holders, but not by getting lower tax rates in retirement. Instead, they achieve

³⁷See Figure A.1(b) for the marginal tax rates applicable at various income levels.

³⁸See A.1(b).

the result by under-saving now and distorting their investment portfolios toward stocks! I emphasize again that this result is dependent upon my particular modeling assumptions, but it demonstrates the striking ways in which nonlinear taxation of retirement investments can distort behavior and lead to inefficiencies.

A.2 Stochastic Tax Rates

The previous two subsections assumed that tax rates for the general population were constant over time. This is, of course, not realistic. Tax rates have fluctuated significantly over time. In fact, the potential variability in future tax rates is often cited as one of the main factors at play for many taxpayers in deciding between traditional and Roth types of accounts. In this section, I analyze the decision taxpayers must make when prevailing tax rates are stochastic.

My approach is simple. I use the Monte Carlo model developed in Section A.1, but for a randomly chosen 50% of the paths, I assume that tax rates increase by 50% during retirement years.³⁹ This is a substantial risk of increase, as well as a substantial increase, but it is also 35 years in the future from the perspective of the taxpayer optimizing his lifetime utility. Such a change over the course of 35 years may not be so unlikely. In any event, the purpose of choosing values here is just to illustrate the effect that tax risk has on retirement planning choices.

The results of my analysis for the equivalent constant consumption expression for the maximum utility solution are shown in Figure A.2(a). Even with such large risk of a future tax change, the optimal utility obtained by choosing the traditional account is still at least as great as that for the Roth account.⁴⁰ The optima for the traditional and Roth accounts do converge for wages over about \$100,000, however.

Even more interesting is the result for the portfolio investment decision in retirement. Figure A.2(b) shows that traditional account investors shift their portfolios even further into stocks, taking on additional stock risk and return to compensate them from exposure to the risk of future tax changes.

B Appendix: Model and Simulation Parameters

In this appendix I provide details about several of the underlying parameters for computing income taxes and Social Security benefits for the model of Appendix A.

³⁹Thus, what is currently the 10% tax bracket would become a 15% tax bracket. Similarly, 25% would become 37.5%, and so forth.

⁴⁰Note that the optimal utility for the Roth account is not affected by the future tax change, since the change by assumption occurs only during retirement years.

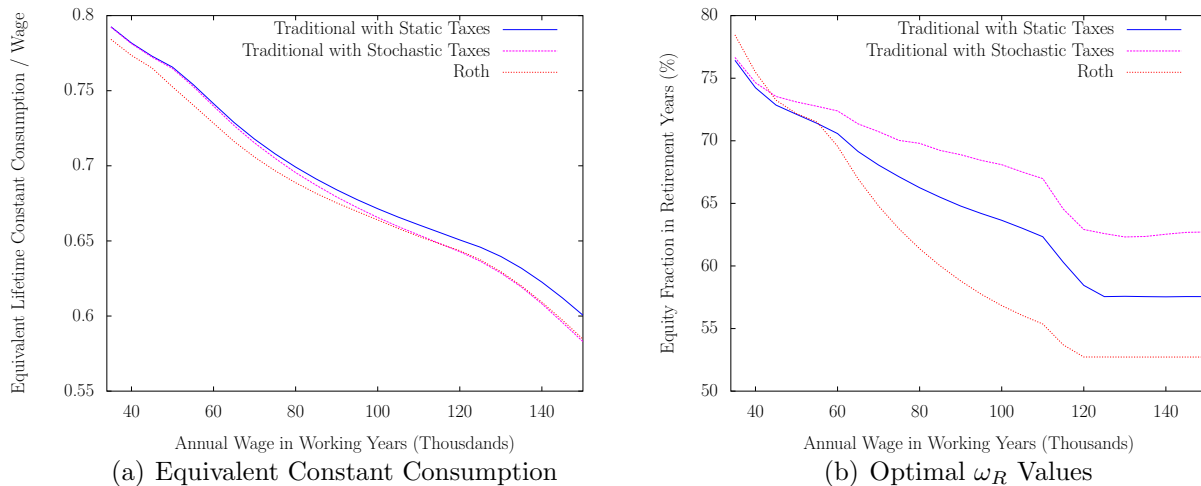


Figure 3: The diagram on the left shows the optimal choice of ω_R , the fraction of investment in equity during retirement. The diagram on the right shows the optimal choice of contribution amount. For both diagrams, the red and blue lines represent the optimal values in the case of Roth and traditional accounts, respectively.

The tax rates used are those applicable to single taxpayers in 2013. The details of the bracket cutoffs and particular rates are listed in Table B. In computing taxable income, I assume a standard deduction and a personal exemption of \$6,100 and \$3,900, respectively. These are the amounts for a single taxpayer in 2013.⁴¹

Rate	Bracket Endpoints	
	Low	High
10%	0	8,925
15%	8,925	36,250
25%	36,250	87,850
28%	87,850	183,250
33%	183,250	398,350
35%	398,350	400,000
39.6%	400,000	

Table 2: Income tax brackets for a single individual in 2013.

I account for Social Security during working years by withholding 6.2% of the lesser of the total wage paid and \$110,000, which was the maximum wage subject to social security withholding in 2012. Because the model assumes that the taxpayer works for the same wage for 35 years, I use that wage to compute the “average index monthly earnings,” *AIME*, for

⁴¹I do not apply the personal exemption phase-out because I do not consider taxpayers with high enough income to be in the phase-out range.

purposes of determining Social Security benefit payments. Specifically, I calculate

$$AIME = \frac{1}{12} \min(\$110,000, Wage),$$

where $Wage$ is the annual wage amount. I then calculate the annual Social Security benefit in accord with the formula for 2012⁴²

$$\begin{aligned} \text{SS Benefit} = & 12 \times [0.90 \times \min(AIME, 791) \\ & + 0.32 \times \min(\max(0, AIME - 791), 4,768 - 791) \\ & + 0.15 \times \max(0, AIME - 4,768)]. \end{aligned}$$

Finally, I calculate the portion of the Social Security benefit that is taxable as⁴³

$$\begin{aligned} \text{SS Taxable} = & \min[0.85 \times SSBen, \\ & \min(0.5 \times SSBen, \\ & \quad 0.5 \times \min(9,000, \max(OI + 0.5 \times SSBen - 25,000, 0))) \\ & + 0.85 \times \max(OI + 0.5 \times SSBen - 34,000)], \end{aligned}$$

where OI is other income during the year, all from traditional retirement account benefits in the model, and where $SSBen$ is the Social Security benefit, computed above.

I do not account for medicare withholding because it does not have an impact on retirement income and so it effects traditional and Roth account holders equally.

⁴²See the worksheet available at <http://www.ssa.gov/pubs/EN-05-10070.pdf>.

⁴³See the instructions for the 2012 IRS Form 1040, p. 29.