

The Design of Retirement Saving Programs
in the Presence of Competing Consumption Needs

Andrew Samwick
Dartmouth College and NBER
andrew.samwick@dartmouth.edu
September 21, 2009

Abstract:

This project considers the optimal design of Social Security taxes when households' saving decisions include motives like housing, education, and uncertainty in addition to retirement. At issue is the timing, not the expected present value, of taxes over the working lifetime. A 10-year, revenue-neutral delay in the onset of payroll taxes generates a welfare gain equal to approximately 18 percent of one year of annualized income. This is equivalent to giving a typical worker \$8,000 of assets upon entering the work force. Welfare gains from revenue-neutral payroll tax delays are larger when individuals must also save to overcome down payment constraints on housing purchases near the beginning of their work lives and slightly lower when they must save later in their work lives to finance college educations for their children.

Keywords: Social Security, payroll tax, saving, housing, college

I thank Jun Liang for exceptional research assistance. This research was supported by the U.S. Social Security Administration through grant #10-M-98363-1-01 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium. The findings and conclusions expressed are solely those of the author(s) and do not represent the views of SSA, any agency of the Federal Government, or the NBER. Any errors are my own.

I. Introduction

The standard life cycle model is based on the retirement motive for saving – the need to transfer resources from the working life when income is predictably high to a planned period of retirement, when income is predictably low. Social Security is a government program designed to ensure that households engage in at least some life cycle saving. It changes their budget set in two important ways. First, the system provides benefits for retirees, the disabled, and their survivors in the form of an indexed annuity that is determined by a formula based primarily on age of retirement, lifetime earnings, and marital status. Second, the financing of these benefits (net of income taxes) comes from the Social Security payroll tax, levied on employees and their employers for earnings up to a maximum taxable earnings level. Under current law, the payroll tax is scheduled to be the same proportion of covered earnings in each year. If the Social Security system is actuarially fair, and if households would optimally be saving more than the payroll tax for retirement during their working years, then the system will likely raise welfare through the longevity insurance it provides.

However, it is clear from empirical data on household saving that for many households, these conditions are not met. In particular, the typical household: (1) has a low saving rate or faces liquidity constraints over much of its working life, (2) has a saving rate that increases substantially in the years just prior to retirement, and (3) saves for several non-retirement reasons during those working years, including for housing, to make transfers to children, and as a precaution against uncertainty. Recognizing the first two of these three, Hubbard and Judd (1987) and Hurst and Willen (2007) consider the

further increases in welfare that can obtain from making the Social Security payroll tax depend on age. Specifically, they show that a revenue-neutral delay in the payment of Social Security taxes can raise welfare by allowing households to smooth their consumption more effectively over their working lives.

This paper builds on these prior studies by recognizing the third feature of pre-retirement saving – namely, the presence of competing saving and consumption needs over the working part of the household’s life cycle. The key insights of this project derive from the need to accumulate resources in advance of the purchase of housing and the financing of college educations. For housing, traditional mortgages have down payment requirements that can only be met, or other financing constraints that can be relaxed, when assets are accumulated in advance of the purchase. For college, the student loan market is imperfect, in that students and their families are typically unable to borrow the full costs of attending and must therefore have some resources available. The need to prefund the purchases leads to an interaction between the design of Social Security and saving for these motives through the life cycle budget constraint. For example, making a revenue-neutral change to the payroll tax rates—lowering tax rates early in the life cycle and raising them later in the life cycle—would free up resources while households are young and saving for housing and college educations.

The key contribution of this paper is to build a comprehensive model of intertemporal consumption decisions that incorporates multiple reasons for saving. To the standard life cycle model in which retirement is the only motive, it adds income uncertainty, bequests, and stylized purchases of housing and college expenses. It incorporates both the current payroll tax schedule and a stylized version of the income tax

schedule for a single individual. The model is parameterized to be consistent with empirical studies of income dynamics, asset accumulation, and expenditures on housing and college educations. For a given set of parameters, the model is first solved to obtain optimal consumption rules at each age as a function of assets and income. The outcomes of the model can then be simulated by confronting those optimal consumption rules with random draws from the income process. The outcome of most importance is the individual's expected lifetime utility, and differences in expected utility are normalized as an equivalent variation: what fraction of annualized income provided in a lump sum at the beginning of life would enable the individual to achieve the same level of welfare as the revenue-neutral delay in the onset of payroll taxes?

The results of those comparisons depend broadly on the amount of saving that the individual is doing under the baseline. As a central estimate, a 10-year, revenue-neutral delay in the onset of payroll taxes generates a welfare gain equal to approximately 18 percent of one year of annualized income. This is equivalent to giving a typical worker \$8,000 of assets upon entering the work force. Welfare gains from revenue-neutral payroll tax delays are larger when individuals must also save to overcome down payment constraints on housing purchases near the beginning of their work lives and slightly lower when they must save later in their work lives to finance college educations for their children.

The remainder of the paper is organized as follows. Section II provides background on the importance of non-retirement motives for saving, using data from the Surveys of Consumer Finances from 1989 – 2007. This background motivates the development in Section III of a model that includes a number of these motives, including

uncertainty, bequests, and specific consumption needs such as housing and college educations. The parameterization of the model is discussed in Section IV, and Section V illustrates the age-consumption and age-asset profiles that emerge from the model in the baseline cases. Section VI introduces both specific consumption needs and revenue-neutral payroll tax shifts and presents the main analytical results. Section VII discusses caveats and extensions to the model and concludes.

II. Background

This section presents some basic data on saving levels, horizons, and motives from the Surveys of Consumer Finances to motivate why the standard life cycle model must be augmented to allow for both uncertainty and specific consumption needs.

Figure 1 provides evidence on the low and late saving behavior of typical households. It shows the median financial assets by age group in the Surveys of Consumer Finances from 1989 – 2007, measured as months of the median household income for that age group (e.g., a value of 12 indicates the financial assets of the median household are equivalent to the income of the median household measured over one year). The bottom three curves show that for households under age 55, financial assets increased relative to income between 1989 and 2001 but have remained flat thereafter. For the youngest two age groups, there has been essentially no increase in financial assets relative to income over the 18 year period. A typical household age 50 (the midpoint of the 45-54 group) in 2007 had about 10 months of income in financial assets. This is not a particularly large nest egg heading into years when retirement becomes more

common. The increase in financial assets relative to income for those aged 55-64 started in 1995 and continued through 2007. The typical household age 60 now has about 16 months of income in financial assets (still not a particularly large nest egg). The key feature of the graph is that the gaps between curves get wider at higher ages – the pace of saving picks up later in the life cycle.

Why do households seem to do their retirement saving at low rates *and* in their later working years? Figure 2 provides one possible explanation. It shows the distribution by age of households' self-reported saving horizons in the 2007 Survey of Consumer Finances.¹ The categories have been stacked for each age group so that they total 100 percent. For households under age 25, over 80 percent report a saving horizon of a few years or less. This percentage declines steadily at older ages, reaching a low of 48 percent for households age 46 to 55. It is simply implausible to argue that retirement is an important determinant of saving decisions for young households whose saving horizon ends decades or more before retirement could be expected to occur.

Figure 3 suggests additional reasons why households' attentions are elsewhere for a good portion of the life cycle. It shows the distribution by age of households' self-reported primary reasons for saving from the 2007 Survey of Consumer Finances. The individual responses have been aggregated into the categories shown. The bottom category in each age group shows the percent of households reporting retirement as the primary motivation for saving. It increases steadily through the working years, from 10 percent for those under 25 to 59 percent in the group age 56-65. The next most prominent reason for saving is against uncertainty, starting at 24 percent for the youngest

¹ The exact question is, "In planning (your/your family's) saving and spending, which of the following is most important to [you/you and your (husband/wife/partner)]: the next few months, the next year, the next few years, the next 5 to 10 years, or longer than 10 years?"

households and falling to about 18 percent over the working years. Other identifiable reasons include housing, special purchases (e.g., consumer durables like vehicles and home furnishings), and for the family, which are relatively more important at younger ages than in the immediate pre-retirement period. All three are prominent for households under age 25. “Family” becomes particularly important for those up to age 45, when saving for sending children to college often must be done.

The explicit study of saving motives has been largely excluded from the economics literature on saving and the design of retirement income programs. When economists study household saving decisions, the standard procedure is to focus on one motive, such as retirement, and make strong simplifying assumptions about the other motives so that they can be relegated to the background. But a simple list of the various reasons why households save gives an indication of how many different decision problems the typical household will have to face. In Samwick (2006), I suggest that for a typical person, life-cycle circumstances that influence saving may include:

1. A period of independent living while young, in which the focus is less on saving and more on keeping credit card and educational debt to a minimum;
2. The formation of a household, with the addition of children and its concomitant demands on disposable income;
3. The purchase of a first home, in which down payment constraints of 5 to 20 percent of the purchase price may force households to begin saving for their first explicit target (Engelhardt, 1996);
4. The college education of children, with often complicated incentives involving means-testing through financial aid formulas (Feldstein, 1995);
5. Several sources of income risk, including career progress, unemployment, and disability (Hubbard, Skinner, and Zeldes, 1995; Engen and Gruber, 2001; and Chandra and Samwick, 2008);

6. Retirement, with its own risks of health and longevity; and
7. The possibility of an intentional bequest at the end of life (Ameriks et al., 2007).

The analysis below is designed as a first attempt to add saving to pre-fund specific needs – housing (motive #3) and college educations (motive #4) – into a model that has been used primarily to consider saving for retirement (motive #6) and uncertainty (motive #5). The model also permits a bequest (motive #7), but this exists more for completeness than as a focus of the analysis.

III. Model Specification and Solution Method

This section presents a life cycle model of consumption that encompasses precautionary, targeted, and bequest motives for saving in addition to the retirement motive on which the standard model is based. The basic structure of the model is that in each period of life, the individual chooses a value of consumption, C_s , as a function of the two state variables in the model, current assets, A_s , and current income, Y_s . The individual's value function in period t , $V_t(A_t, Y_t)$, is defined as:

$$V_t(A_t, Y_t) \equiv \max_{\{C_s\}} E_t \sum_{s=t}^T (1+\delta)^{t-s} \left(\prod_{q=t}^s p_q \right) (p_s u(C_s) + (1-p_s)v(A_s))$$

such that :

$$(1) \quad \begin{aligned} a) \quad & u(C) = \frac{C^{1-\gamma}}{1-\gamma} \\ b) \quad & v(A) = \phi_1 \left(\frac{(\phi_2 + A)^{1-\gamma}}{1-\gamma} \right) \\ c) \quad & X_s = A_s + Y_s - h_s(Y_s) - z_s(A_s, Y_s) \\ d) \quad & A_{s+1} = (1+r)(X_s - C_s) - g(Y_s, X_s - C_s) \\ e) \quad & A_s \geq 0 \quad \forall s \end{aligned}$$

The value function is equal to the expected discounted utility of consumption in each period from the current period t to the final period T , discounted by a factor of $1/(1+\delta)$ each period. The rate of time preference is δ and is similar to an interest rate in governing the utility tradeoff across periods. The individual is subject to mortality risk in each period, where p_s is the probability of living through the period conditional on having survived to that period. Thus, $1 - p_s$ is the probability of dying that period, and the $\prod p_q$ term represents the probability of surviving until that period from the initial period, t . The survival probabilities used in the solving the model are based on the life table for the total population of the United States in 2004, as published in Arias (2007).

The within-period utility function is assumed to be additively separable in consumption and all other factors that affect utility, so that these factors can be omitted from the optimization problem. Utility is derived from either of two sources: the optimal consumption choice made for an individual that survives the period and the amount of the bequest, equal to current assets, if the individual does not survive. The utility of consumption each period is assumed to take the Constant Relative Risk Aversion

(CRRA) form: $u(C) = C^{1-\gamma}/(1-\gamma)$, where γ is the coefficient of relative risk aversion.

With a utility function such as CRRA that has a convex marginal utility function (i.e. $u'''(C) > 0$), there is a precautionary motive for saving, and greater uncertainty in the income process will induce greater saving.² Following the literature, the utility of a bequest in a given period takes the same CRRA form, with two modifications.³ The first is to include a parameter, ϕ_1 , that determines how much utility the individual gets from a dollar of bequest relative to a dollar of consumption. The second is to include a parameter, ϕ_2 , which affects the marginal utility of the bequest – a higher value of the parameter reduces the marginal value of a bequest.

Equation (1c) defines the concept of “cash on hand” that is available to finance consumption and income taxes each period. To obtain cash on hand, X_s , assets are augmented by pre-tax labor income, Y_s , but reduced by payroll taxes and other consumption needs. The payroll tax, $h_s(Y_s)$, is based on the formula in place in 2009. A payroll tax of 12.4 percent for Social Security is levied on earnings up to a maximum taxable earnings limit of \$106,800.⁴ A further 2.9 percent payroll tax for the Health Insurance portion of Medicare is levied on all earnings. As the other elements of the budget constraint do not include the employer’s share of this tax, only half of the full payroll tax paid on the individual’s behalf is deducted from Y_s . This function is subscripted by the year, which is not necessary in the baseline case. However, in

² The use of the CRRA utility function is standard in both the empirical and theoretical literature on precautionary saving. CRRA utility means that a consumer remains equally willing to engage in gambles over a constant proportion of current wealth as wealth increases. An alternative, and perhaps more realistic assumption, might be that the consumer will accept larger proportional risks as wealth increases. See Kimball (1990) for a discussion and derivation of the key results for precautionary saving.

³ See Chen and Smetters (2009) and Ameriks et al. (2007), for example.

⁴ The 12.4 percent tax includes coverage for disability insurance, but the impact of disability is not modeled in this paper. See Chandra and Samwick (2008) for a similar model that incorporates the risk of disability.

simulations that alter the timing of payroll taxes over the working life, this function will be modified so that the amount of payroll taxes paid on a given level of income varies with age.

As discussed in the previous section, individuals may be subject to specific consumption needs as they progress through their working lives. The two that are included specifically in the analysis below are housing and children's educations. These are modeled as deductions from the amount of resources available for other consumption at specific ages, $z_s(A_s, Y_s)$. At this stage, three simplifications have been imposed on the way these needs are modeled. First, the individual is assumed to have no choice as to whether to pay for these consumption needs. Second, the individual is assumed to know precisely the years when these consumption needs will arise. Third, the consumption needs are modeled as functions of assets and labor income rather than fixed dollar amounts. A richer model, a subject of work in progress, would remove these simplifications and allow for more general results.

Equation (1d) shows how assets accumulate from one period to the next. Cash on hand is used to finance consumption and the income taxes, $g(\cdot, \cdot)$. The income tax liability is a function of two arguments. The first is labor income and the second is the amount of saving, or $X_s - C_s$. For simplicity, the portfolio decision is restricted to a single riskless asset paying a return, r , each period. Thus, capital income is just $r*(X_s - C_s)$. The individual's taxes are calculated based on the 2009 tax schedule for a single taxpayer who does not itemize deductions and receives all capital income as interest or dividends rather than capital gains. Payroll taxes are assumed to be paid as the labor income is earned, prior to the consumption decision each period. Since income taxes depend on capital

income and thus the outcome of the consumption decision during the period, they are assumed to be paid at the end of the period.

The last part of Equation (1) is the liquidity constraint. Equation (1e) requires assets to be positive in each period – the individual cannot borrow against future income to finance current consumption. This is a simplification that nonetheless acknowledges the credit constraints that prevent individuals from borrowing too heavily against future income outside of a secured or collateralized relationship.⁵

The processes that describe income uncertainty and the evolution of current income are as follows:

Before retirement:

$$a) \ln(Y_s) = \ln(P_s) + u_s$$

$$b) u_{s+1} = \rho u_s + \varepsilon_{s+1}$$

$$c) \varepsilon_s \sim i.i.d. N(0, \sigma^2)$$

(2) At retirement:

$$d) Y_{s+1} = PIA(Y_s, P_s, \dots, P_1)$$

After retirement or disability:

$$e) Y_{s+1} = Y_s$$

In this model, the individual retires at a planned date that is known from the beginning of the working life. A richer model would include a disutility of work (if not the risk of involuntary retirement due to health or other reasons) and a choice over the

⁵ The outcomes of the model are not greatly affected by allowing a fixed amount of unsecured borrowing. It also imposes the liquidity constraint directly, rather than including a much higher rate for borrowing that would discourage but not prohibit large amounts of unsecured borrowing. See Hurst and Willen (2007) for an analysis of consumption and Social Security with a richer modeling of credit constraints.

retirement age based on economic factors. Prior to retirement, the natural log of current income is equal to the natural log of permanent income (P_s) plus a shock to income (u_s) that follows an AR(1) process. The innovations to that AR(1) process are assumed to be independently and identically drawn from a normal distribution with mean zero and variance σ^2 .⁶

At retirement, income is given by the Social Security benefit formula, denoted here by a function PIA, for the primary insurance amount that a worker earns based on earnings covered by the system. The benefit formula used is based on the 2009 schedule, with one modification. The actual Social Security benefit formula is based on an average of the highest 35 years of earnings over an individual's career. Without some simplification, this formula would require two additional state variables in the model, corresponding to the average as of the current period and the lowest earning year still included in the average.

To avoid additional state variables, the benefit can be based only on the state variables and any fixed parameters of the optimization problem. The latter includes permanent income in each period, $\{P_s\}$. In the year prior to retirement, Equation (2a) shows that the cumulative income shock, u_s , can be recovered as $\ln(Y_s/P_s)$. The benefit formula makes the simplifying assumption that u_s evolved through a series of annual shocks of equal size.⁷ This pseudo-history is then used in the benefit formula, which in 2009 provided a replacement rate of 90 percent on average earnings up to \$8,928, 32 percent on average earnings above that level but below \$53,796, and 15 percent on

⁶ In the simulations, the mean of the shock to income, not the log of income, is normalized to be one in all periods.

⁷ Algebraic manipulation of Equation (2b) shows that this shock is $\varepsilon = u_n(1-\rho)/(1-\rho^n)$, where n is the final year of work.

covered earnings above the second level. After retirement, income is unchanged at this new level and is no longer uncertain. In the analysis below, the model is set up so that the individual enters the workforce at age 21, retires at age 65, and lives to a maximum age of 100.

The solution method for stochastic optimization problems with multiple state and control variables is discussed in detail in Carroll (2001). The solution begins in the last period of life, T , when the problem is trivial because the household simply consumes all of its assets and after-tax income (or dies, leaving a bequest of A_T), yielding an optimal value for C_T as a function of the state variables A_T and Y_T . These solutions generate the value function, $V_T(A_T, Y_T)$, and the partial derivative, $V_T^A(A_T, Y_T)$, which represents the marginal value of an additional dollar in assets at the beginning of period T . Moving back to the period $T-1$ problem, we can rewrite the objective function as:

$$(3) \quad V_{T-1}(A_{T-1}, Y_{T-1}) \equiv \max_{\{C_{T-1}\}} p_{T-1}u(C_{T-1}) + (1 - p_{T-1})v(A_{T-1}) + \beta E_{T-1}[V_T(A_T, Y_T)]$$

The problem in period $T-1$ is a special case, since there is only mortality risk and no income uncertainty. More generally, given the function $V_{t+1}(A_{t+1}, Y_{t+1})$ and the associated partial derivative, the problem at period t is:

$$(4) \quad V_t(A_t, Y_t) \equiv \max_{\{C_t\}} p_t u(C_t) + (1 - p_t)v(A_t) + \beta E_t[V_{t+1}(A_{t+1}, Y_{t+1})]$$

These one-period problems have first-order conditions given by:

$$(5) \quad u'(C_t) - \left(\frac{1 + r(1 - g_2(Y_t, X_t - C_t))}{1 + \delta} \right) E_t [V_{t+1}^A(A_{t+1}, Y_{t+1})] = 0$$

The first term in the first-order condition is the marginal utility of an additional dollar of consumption in period t .⁸ The second term is the expected discounted value of saving that dollar to be used in period $t+1$. The dollar grows by the after-tax interest rate and has a marginal value of V^A at that time. In this expression, $r \cdot g_2(Y_t, X_t - C_t)$ is the marginal tax on another dollar of saving, i.e. the derivative of the tax liability function with respect to its second argument. The expected marginal utility of a dollar of assets at time $t+1$ is discounted back to period t utility by a factor of $1/(1+\delta)$. The difference between the marginal utility of consumption and the expected marginal utility of assets in the next period is zero at the optimal level of consumption.⁹

The solution of the optimization problem is a series of consumption rules that determine consumption in each period as a function of assets and income. Once the optimal consumption rules have been obtained, the model can be simulated forward by specifying initial values of the state variables, drawing random shocks to income, and applying the consumption rules to generate distributions of asset balances in each period. In the simulations below, the model is evaluated using the average outcomes based on 5,000 independent random draws of the income profile.

⁸ Note that the survival probability does not pre-multiply this marginal utility because the consumption decision happens only if the individual survives.

⁹ The solution method is complicated by the liquidity constraint. The constraint that A_{t+1} cannot be negative implies that the maximum amount of consumption in the prior period is such that $X_t - C_t - g(Y_t, X_t - C_t) = 0$.

The key outcome of the model is a value for the expected value of $V_1(A_1, Y_1)$, computed as the average value of the program across the 5,000 income profiles and a starting asset value of zero at the beginning of the work life. For a given set of preference parameters, $\{\delta, \gamma, \phi_1, \phi_2\}$, we can measure the gain in expected utility from changing the individual's budget constraint by comparing the value of $V_1(A_1, Y_1)$ before and after the change. A natural metric is equivalent variation, measured here by κ in the equation:

$$(6) \quad V_1^1(A_t, Y_t) = V_1^0(A_t + \kappa, Y_t)$$

The equivalent variation indicates how much the individual would have to be compensated to remain under the original budget constraint to achieve the level of expected utility available under the new budget constraint.¹⁰ For simplicity, κ is assumed to be available at the beginning of the working life, where a liquidity-constrained individual gets the most bang-for-the-buck for an additional dollar. While κ could be measured as a simple dollar value, the most natural units for it are as a percentage of labor income. Since income changes from year to year, the equivalent variation is measured as a percentage of the equivalent annual income (EAI), defined as:

¹⁰ Given the functional form for the bequests, specifically the fact that ϕ_2 is not zero, the value function is not homothetic in consumption and bequests. This precludes the use of a measure of equivalent variation as in Chandra and Samwick (2008) or Hurst and Willen (2007) that multiplies consumption and would be interpreted as "a percentage of consumption in each period."

$$(7) \quad EAI = \frac{E_1 \sum_{s=1}^T (1+r)^{1-s} \left(\prod_{q=1}^s p_q \right) (p_s Y_s)}{\sum_{s=1}^T (1+r)^{1-s} \left(\prod_{q=1}^s p_q \right) (p_s)}$$

By construction, the actuarial present value of receiving the EAI in each period of life is equal to the expected actuarial present value of the labor income and Social Security income actually received by the individual.

IV. Model Parameters

The baseline model consists of assumptions about the income process, preference parameters, the size of specific consumption needs, and the potential changes in the timing of payroll taxes by age. The variety of parameter configurations are chosen to illustrate the differences in the timing of saving over an individual's pre-retirement years. Specifically, the greater the individual's desire to consume at or above current income early in the life cycle, the more valuable will be opportunities to delay the payment of taxes, even if the delay generates higher taxes later in the working life.

One factor that determines the desire to borrow early in the life cycle is the slope of the age-earnings profile. When earnings are expected to rise rapidly, there is less reason to save for life cycle reasons. There are a number of prior studies that have estimated the parameters of labor income processes, using a variety of datasets and methods. For example, Cocco, Gomes, and Maenhout (2005), hereafter CGM, used the Panel Study of Income Dynamics to estimate an age profile for total family non-capital

income (i.e. including both spouses, if present, and income from social insurance and transfer programs) separately for different education groups. The profile (here for the lowest education group) adjusted to constant 2009 dollars is based on a cubic polynomial in age and is graphed in Figure 4.

Other authors, most notably Murphy and Welch (1990), hereafter MW, have cautioned that the failure to include higher order terms in the polynomial for age may understate early-career income growth and overstate late-career income declines. They recommended a fourth-order polynomial in age. They also focused on individual earnings, rather than total family income. The age-earnings profile from their study (averaged across all education groups), based on log earnings for a sample of white males in the Current Population Surveys, is also graphed in Figure 1.¹¹ The difference in the slope of the age-income profile is noticeable. Individuals facing the MW profile are more likely to face binding liquidity constraints, while those facing the CGM profile can be expected to begin saving earlier in their working lives.

Figure 4 depicts the average age-income profiles, or the profile of permanent income $\{P_s\}$ in the model. Around this average is uncertainty, captured by an AR(1) process with parameters of $\rho = 0.95$ and $\sigma = 0.15$. These parameters have become standard in the precautionary saving literature since the work of Hubbard, Skinner, and Zeldes (1995). The leading alternative is a random walk model in which $\rho = 1$, as in CGM. Either assumption would be fine for the present purpose, which is to compare the simulations for individuals who face realistic amounts of income uncertainty to

¹¹ Murphy and Welch present results for log weekly earnings. The MW and CGM profiles are both adjusted to constant 2009 dollars and initiated at a level of income at age 21 such that they generate replacement rates for a typical worker with medium earnings in the Social Security system, as provided at <http://www.ssa.gov/OACT/TR/2009/lr6f10.html>.

simulations for individuals facing no uncertainty and thus without a precautionary motive for saving.

The second key factor that determines the individual's to desire early in the life cycle is the size of the discount rate, δ , relative to the interest rate, r . If the two rates are similar, then the term that pre-multiplies the expected marginal utility of assets in Equation (5) will be close to unity, and the individual will seek to balance consumption across the two periods. As δ rises relative to r , the marginal utility of consumption must fall to keep the first-order condition from failing. Given the concavity of the utility function, this requires current consumption to rise. Thus, higher discount rates lead to higher current consumption and a greater desire to borrow in the present.

To capture this difference, the discount rate takes one of two values. For simulations of a "patient" consumer, δ is assumed to be 0.03. For simulations of an "impatient" consumer, δ is assumed to be 0.08 percent. The two values are consistent with estimates of discount rates in Samwick (1998). A patient consumer begins saving for retirement early in the lifecycle, while an impatient consumer typically delays saving for retirement for some portion of the lifecycle.

There are three additional preference parameters. The coefficient of relative risk aversion, γ , is assumed to be 3. In a CRRA model, this results in an intertemporal elasticity of substitution of $1/3$. This assumption is commonly found in prior studies of individual level saving decisions. Values chosen to match macroeconomic data are often much higher but implausible as a description of individual-level behavior. Higher values generate both more precautionary saving and a greater aversion to allowing consumption to vary over the lifecycle, even if the variation is anticipated. The last two preference

parameters pertain to the utility of bequests. In the simulations, ϕ_1 is set at a value of 1 and ϕ_2 is set at a value of \$24,000. These values are chosen to de-emphasize the bequest motive for saving relative to prior literature, such as Ameriks et al. (2007), in which bequest motives were the central focus. Specifically, ϕ_1 is lower and ϕ_2 is higher, such that in the simulations below we observe individuals spending down their assets steadily at the end of life rather than maintaining or building a large bequest.

For the purposes of the simulations, the specific consumption needs of the individual are calibrated to a typical household's experience. The individual is assumed to have one child to put through college. Mathews and Hamilton (2002) report that in the 2000 Census, the average age of a mother at the birth of a first child was 25 and the average age of a mother at the birth of any child was 27. Choosing 18 as the child's age upon enrollment, the individual is assumed to incur college costs for four years beginning in at age 45. According to Sutton (2009), the College Board reports that annual tuition and fees for two-year public colleges were \$2,402 for 2008-2009, \$6,585 for four-year public colleges, and \$25,143 for four-year private colleges. However, actual costs to the student's family depend on the family's financial resources through the financial aid process.

For the analysis below, it is assumed that college expenses are as determined by the Expected Family Contribution (EFC) in the standard financial aid formula for a family of 3 with 1 child in college. In brief, the formula specifies an EFC that is increasing in both income and assets. With income for these purposes defined as labor income, the two key determinants of the EFC are the state variables in the model, making

it straightforward to implement.¹² Using the EFC in this manner captures the positive relationship between financial resources and the demand for education. Figure 5 shows the relationship between the EFC and income for asset levels up to \$400,000. The EFC is zero for income levels below \$27,000. Assets do not increase the EFC until income reaches \$50,000. Since the most expensive colleges now have total costs of about \$50,000 per year,¹³ annual costs are capped at this amount regardless of income or assets.

The individual is assumed to make the transition to homeownership in the twelfth year of the working life, or age 32. Data on housing ownership and income from the 2007 Survey of Consumer Finances is presented in Bucks et al. (2009) and reproduced in Table 1. The rows of the table distinguish age categories, and the columns present data on the probability of owning a home, the median value of the home conditional on owning one, the probability of having a mortgage, the median mortgage debt conditional on having a mortgage, and median income. The last two columns show the ratios of median home value and mortgage debt to median income by age group. For households under age 65, the ratio of house values to income varies between 4.0 and 5.3, while the ratio of mortgage debt to income falls from an initial value of 4.1 to 2.0 over the working life. Given these ratios, along with typical down payment requirements of 10 – 20 percent of the purchase price, a reasonable figure for the amount of the down payment is 50 percent of a current year's income. The required consumption amount is therefore 50 percent of income in the 12th year of the working life that begins at age 21.

Finally, the simulations below distinguish between the existing payroll tax, which does not vary by age, and alternatives in which the payroll tax rates are lower at younger

¹² A detailed grid of the EFC for various asset and income combinations can be viewed at <http://www.finaid.org/calculators/quickefcchart.phtml>.

¹³ See, for example, <http://www.campusgrotto.com/most-expensive-colleges-for-2008-2009.html>.

ages and higher at older ages. Following Hubbard and Judd (1987) and Hurst and Willen (2007), these changes are implemented by delaying the onset of the payroll tax until specified ages, and then implementing it at a sufficiently high rate such that the alternatives have the same expected tax revenue collections as the existing system. Two alternatives are considered: one in which payroll taxes are delayed for 10 years and one in which payroll taxes are delayed for 20 years.¹⁴

The key parameters are therefore the increases in the payroll tax rates when they are eventually implemented. The parameters differ for the two age-income profiles, with the shift being greater for the CGM profile because its early years have relatively higher income. For the CGM profile, a delay of 10 years requires the payroll tax rates to be increased by a factor of 1.4211 and a delay of 20 years requires a factor of 2.3803. Focusing on the employee's share of the payroll tax, this requires the uniform rate (under the maximum taxable earnings) to rise from 7.65 percent to 10.87 and 18.21 percent, respectively. For the MW profile, the corresponding factors are 1.2673 and 1.9634 and the corresponding tax rates are 9.70 and 15.02 percent, respectively. Given the partial equilibrium nature of the model and the absence of a labor/leisure decision, changes in the timing of the employer's share of the payroll tax are not analyzed in the present analysis.

¹⁴ Note that the equivalence holds only in expected value. An individual whose income shocks are positive early in life and negative later in life will pay lower taxes under the delayed payroll tax system. The opposite is true for a worker whose income shocks are negative early in life and positive later in life.

V. Consumption and Asset Profiles by Age

This section illustrates the average consumption and asset profiles that result from the solving and simulating the model in the prior sections over a range of parameter groups. There are eight groups, based on all possible combinations of the discount rate being 0.08 (Impatient) or 0.03 (Patient), the income profile being steep (MW) or shallow (CGM), and income being certain ($\sigma = 0$) or uncertain ($\sigma = 0.15$). Figures 6 and 7 show the consumption and asset profiles for the Patient individual, respectively, and Figures 8 and 9 show these profiles for the Impatient individual. In the cases where income is uncertain, each profile shown is the average across the 5,000 random income draws.

In Figure 6, the consumption profiles for the individuals facing the CGM income process are higher than those facing the MW income process, since the latter provides less lifetime income. Part of that incremental lifetime income is consumed in each year. Common to both consumption profiles is the upward sloping portion in the earliest years of the working life – these are years in which the liquidity constraint binds. Because of the comparative steepness of the MW profile, the upward sloping portion extends longer into the working life. For both income profiles, the presence of income uncertainty generates lower consumption early in the working life, as the individuals now need to accumulate wealth to buffer consumption against income fluctuations. These dashed consumption profiles peak later – something that is a consequence of the need to save for precautionary reasons and not desirable even for individuals who are patient. All of the consumption profiles continue to decline through retirement. The flat parts indicate a

period in which assets have been exhausted and the individual optimally sets consumption equal to retirement income.

Figure 7 shows the age-asset profiles corresponding to the consumption profiles in Figure 6. Four aspects of the graph are noteworthy. First, all of the individuals begin their working lives with the liquidity constraint binding. Those facing uncertainty begin saving before those with certain incomes, and, within each pair, those facing the shallow CGM profile begin saving before those facing the steeper MW profile. Second, the precautionary motive for saving is as large a contributor to asset accumulation as the retirement saving motive, as those facing income uncertainty accumulate nearly twice the assets on the eve of retirement as those not facing uncertainty.¹⁵ This fraction is broadly consistent with the estimates in Carroll and Samwick (1998). Third, all of the asset profiles eventually fall to zero before the end of life, in reverse order of when the profiles rose above zero at the beginning of the working life. When assets are zero, the individual simply consumes his after-tax Social Security benefits. Finally, the peaks in the asset profiles are between 2 and 5 times the amount of annual income on the eve of retirement, far too large to be consistent with the data presented in Figure 1, where the typical household had no more than 1.5 years of income in assets.

Figure 8 shows the consumption profiles for the impatient individual. Compared to the patient individual, the period of growing consumption at the beginning of the working life, associated with binding liquidity constraints, lasts for more years. The peak in consumption is also higher, consistent with a higher value placed on current as

¹⁵ Making these comparisons on the eve of retirement understates the importance of precautionary motives relative to retirement motives, as the need for retirement saving is greatest at this age but all of the income uncertainty has been resolved. Consistent with the survey responses in Figure 3, in which the importance of uncertainty as a motive for saving started high and faded over the working life, the heights of the dashed curves are more than twice as high as the height of the solid curves at younger ages.

opposed to future utility. Precautionary saving, as indicated by the gap between the dashed and solid curves, is also less apparent. Since the shocks to income persist over time before fading away, a lower emphasis on the future translates into less precautionary saving. Finally, at the end of the life cycle, consumption falls to equal income more rapidly after retirement.

Each of these differences is also apparent in Figure 9, which shows the age-asset profiles for the impatient individual. Compared to the patient individual, asset accumulation by the eve of retirement is only half as large. As in the case of the patient individual, precautionary motives explain about half of the total asset holdings on the eve of retirement. Most importantly, the asset values at the peak accumulations are only 1.3 – 3 times the value of income on the eve of retirement. These ratios are much more in line with those shown in Figure 1. The need to assume a high degree of impatience to match the saving patterns of the typical household – here reflected by the ratio of median assets to median income – is one of the key aspects of the “Buffer Stock” model of saving proposed by Carroll (1992, 1997).

VI. Consumption Needs and Alternative Tax Schedules

As modeled above, the impact of specific consumption needs on individual welfare is negative. Purchasing a house or sending a child to college requires some pre-funding. This pre-funding must come at the expense of other consumption that the individual would like to enjoy. Payroll taxes depress consumption whenever they are levied. For liquidity constrained individuals, payroll taxes are particularly harmful to

expected utility because they force a household that would like to borrow to do additional saving. As shown by Hubbard and Judd (1987) and Hurst and Willen (2007), revenue-neutral delays in the timing of payroll taxes can increase lifetime welfare. This section considers the interaction between such revenue-neutral shifts and the presence of specific consumption needs.

Figure 10 shows the average consumption profiles for two individuals: one facing a down payment constraint equal to 50 percent of his income at age 32 and one facing college costs as given by the EFC in Figure 5 at ages 45 – 48. The other parameters of the consumption problem match the impatient individual facing the CGM income process with no uncertainty. The outer envelope of the two curves matches the consumption profile for that individual shown in Figure 8. In each case, consumption begins to decline a few years prior to the age at which the specific consumption need begins. After those needs are met, consumption reverts back to the profile that prevailed in the absence of the specific needs. The welfare loss due to the consumption need is captured by these reductions in other consumption.

Table 2 shows the magnitude of these losses for each of the eight parameter combinations whose consumption and asset profiles were graphed in Figures 6 – 9. The first column identifies the parameter combination with a letter A – H. The next three columns show the parameter assumptions for that group: discount rate, income profile, and standard deviation of income shock. The fifth column shows the expected present discounted utility of consumption and bequests for that parameter group – the value of $V_1(0, Y_1)$. When the specific consumption needs are introduced, expected utility falls from this level. The final three columns show the equivalent variations – the amount of

assets that would have to be given to the individual at the beginning of life as a share of EAI – to enable him to achieve the level of utility in the absence of the consumption need. Separate results are shown for just the housing need, just the college need, and both needs.

The key pattern that emerges in Table 2 is that welfare losses that result from specific consumption needs are larger for parameter groups in which the individual was doing more saving in the baseline, whether for life cycle or precautionary reasons. Equivalent variations are higher for the patient compared to the impatient individual and for the CGM profile compared to the MW profile. Greater income uncertainty increases the welfare loss for the patient individual and for the impatient individual in the case of the housing need. (For the college need, higher uncertainty results in slightly lower welfare losses for the impatient individual.) Welfare losses range from near zero to over 40 percent of EAI, depending on the parameter group and which consumption needs are imposed. The housing need generates a larger welfare loss, largely because it occurs earlier in life. There is some evidence that the welfare losses compound – the welfare loss of the two needs together is larger than the sum of the welfare losses of the needs taken individually when the latter are not too small.

Figure 11 shows the impact of revenue-neutral delays in the timing of the payroll tax on the consumption profile of the impatient individual facing the CGM income profile without uncertainty or specific consumption needs. With a 10-year delay, payroll taxes are zero for the first 10 years and then jump to a level 1.4211 times their current rates. With a 20-year delay, the payroll tax rate increases by a factor of 2.3803 compared to current rates. These delays enable the consumption profiles to start out higher and

remain there until the payroll taxes commence. After that time, consumption is lower for the remainder of the life cycle, the more so the longer the delay. Given the rising income profile, compounded by impatience, this is a shift that makes the individual better off.

Table 3 shows the equivalent variations from a 10-year revenue-neutral delay in the payroll tax for the same eight parameter groups. Focusing on Parameter Group A in the first row of the table, an impatient individual facing the MW income profile and no income uncertainty would value the 10-year delay as equivalent to starting his working life at age 21 with an initial wealth of 17.26 percent of his equivalent annual income. For individuals facing the MW profile, the equivalent variation is about 17 percent of EAI regardless of whether they are patient or impatient or whether they face uncertainty in their income process or not.

For individuals facing the CGM income profile, the equivalent variation is much higher for the impatient individual at 28.56 ($\sigma = 0$) or 26.07 ($\sigma = 0.15$) percent of EAI and somewhat lower for the patient individual at 9.96 ($\sigma = 0$) or 12.79 ($\sigma = 0.15$) percent of EAI. Recall that compared to the MW profile, the CGM profile starts out at higher income levels. The payroll tax delay is potentially more valuable to both the impatient and patient individuals. However, the impatient individual will take full advantage of the delay, using it to boost consumption to the point where the liquidity constraint is binding for nearly the whole delay. In contrast, the patient individual will boost consumption only for a shorter period and begin to save some of it after a few years. Overall, a simple average of the 8 parameter groups gives an equivalent variation of 18.19 percent of EAI. Applying that average to the average EAI (\$40,536 for MW, \$51,227 for CGM) yields \$8,346. Implementing a 10-year revenue neutral delay in the payroll tax would generate

a welfare gain comparable to giving each individual about \$8,000 upon entering his working years.

The next column of Table 3 shows the equivalent variations when the individuals have a need to pre-fund a down payment on a house at age 32. For each parameter group, the equivalent variation is higher than in the first column when the housing need was absent. The payroll tax delay is more valuable when the housing need is present because it defers most early tax payments that would come from low- (or no-) saving individuals until after the housing purchase is made. Averaging across the eight different parameter groups, the equivalent variation is now 19.66 percent of EAI, an increase of 1.47 percentage points or 8 percent of the baseline value of 18.19 percent.

The final column of Table 3 shows the equivalent variations when the individuals have a need to pay for four years of college expenses beginning at age 45. The equivalent variations are positive but somewhat less than when the college need was absent. Taking the simple average across the eight groups, the equivalent variation is now 17.72 percent of EAI, a reduction of 0.47 percentage points or 3 percent of the baseline value when the consumption need was absent. What accounts for this reduction? Recall that the payroll tax deferral is not presented as an option – it changes the individual’s budget constraint in the same way that a loan at an interest rate of 3 percent would change it. If the individual wanted to save most of the payroll tax reduction, he would have to pay taxes on the incremental saving and not achieve the full 3 percent rate of return. Since the college needs occur 15 years after the deferral ends and the “loan” begins to be repaid, there is considerably more saving of the payroll tax deferral involved with the college needs than with the housing needs.

Table 4 presents the analogous results for a revenue-neutral payroll tax delay of 20 years. When neither consumption need is present, the equivalent variations are higher than for the 10-year delays in Table 3 for seven of the eight parameter groups. The simple average of those groups yields an equivalent variation of 21.18 percent, an increase 2.99 percentage points or 16 percent compared to the 18.19 percent for the 10-year delay. As with the 10-year delay, the introduction of the housing need increases the value of the 20-year payroll tax delay. Because of the longer delay, the impact is even larger – the average equivalent variation rises from 21.18 to 24.70 percent of EAI, an increase of 3.52 percentage points or 17 percent. Similarly, the welfare gains are in all but one case positive when the individual faces the college need though reduced compared to when the need was absent. The average equivalent variation is now 19.53 percent of EAI, a reduction of 1.65 percentage points or 8 percent compared to the baseline of 21.18 percent.

VII. Conclusion

The Social Security system relies primarily on the payroll tax for financing, a tax that is applied uniformly to earnings regardless of the age of the worker. This paper considers the scope for welfare gains from revenue-neutral payroll tax delays in the presence of potential liquidity constraints and non-retirement reasons for saving, including income uncertainty, housing down payments, and college costs. As in prior studies such as Hurst and Willen (2007) and Hubbard and Judd (1987), the welfare gains are found to be positive, with a simple average estimate of 18.19 percent of equivalent

annual income in initial assets for a 10-year delay and 21.18 percent for a 20-year delay. Given the values of equivalent annual income for the income profiles used here, the equivalent variation for the 10-year delay translates into the same welfare gain as giving each individual about \$8,000 in assets at the beginning of his working life.

Extending the prior literature, this paper shows that the welfare gains are of comparable magnitudes when income uncertainty and thus a precautionary motive for saving are added to the analysis. It also shows that in the presence of an asset constraint early in the life cycle, modeled here as the need to accumulate a down payment for a housing purchase, the welfare gains to revenue-neutral payroll tax delays are somewhat higher. The average welfare gains rise to 19.66 and 24.70 percent of annualized income for the 10- and 20-year delays, respectively.

For consumption needs that arise later in the life cycle, such as paying college expenses for children, the welfare gains from revenue-neutral payroll tax delays are generally still positive but smaller than without those costs. The need to save most of the payroll tax deferrals for an extended period of time in a taxable account erodes some of their value. On average, the welfare gains are 17.72 and 19.53 percent of annualized income for the 10- and 20-year delays, respectively.

The simplifications made in the analysis above provide a number of opportunities to enhance the model in future work. One important simplification was to model the two consumption needs as mandatory and welfare reducing. An obvious extension is to model the gains that individuals receive from experiencing them and make them voluntary. For example, the down payment would translate into part ownership of an illiquid asset (the house) and be available for liquidation at some price or as part of the

bequest. The college expenditures would generate children with higher earning potential, some of which might flow back to the parent or lessen the urgency of the bequest motive. This extension could magnify or shrink the welfare gains described in the paper, depending on the way it was modeled.

The analysis has also focused on the statutory payroll tax rate paid by the individual. This is done for consistency with the income processes as they were estimated in other studies and with the income tax schedule present in the model. Including the employer contribution for payroll taxes would require considering more formally the incidence of the payroll tax. Including some portion of the employer's share of the payroll tax in the revenue-neutral delays would magnify their welfare impact. Other modifications to the analysis could include more subtle tax shifts than periods of full exemption and could consider individuals facing income profiles that are on average much higher or much lower than the typical individuals presented here.

Another important extension is to consider the implications of the payroll tax delays on the labor/leisure tradeoff. In addition to Hurst and Willen (2007) and Hubbard and Judd (1987) who focus on the saving margin, there is a prior literature that considers whether tax rates should depend on age to encourage younger workers to spend more time in the labor market. The basic insights are presented in Kremer (2003) and more fully developed in Lozachmeur (2006), Viard (2007), and Weinzierl (2008). Such an extension is particularly relevant in the case of large changes to the payroll tax rate at higher ages, where labor supply might be very sensitive to tax rates given the option of retiring early. In all cases, including the labor supply decision would impact the federal government's budget through its collection of not just payroll but income taxes as well.

In the present analysis, with the permanent income profile fixed, income tax collections do not vary much according to the timing of payroll taxes because total lifetime income does not change.

References

- Ameriks, John, Andrew Caplin, Steven Laufer, and Stijn Van Nieuwerburgh (2007). “The Joy of Giving or Assisted Living? Using Strategic Surveys To Separate Bequest and Precautionary Motives,” National Bureau of Economic Research, Working Paper No. 13105, May.
- Arias, Elizabeth (2007). “United States Life Tables, 2004,” National Vital Statistics Reports 56, No. 9, December. [http://www.cdc.gov/nchs/data/nvsr/nvsr56/nvsr56_09.pdf]
- Bucks, Brian K, Arthur B. Kennickell, Traci L. Mach, and Kevin B. Moore (2009). “Changes in U.S. Family Finances from 2004 to 2007: Evidence from the Survey of Consumer Finances,” *Federal Reserve Bulletin*, 95, A1-A55.
- Carroll, Christopher D. (1992). “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence,” *Brookings Papers on Economic Activity*, 2, 61–156.
- Carroll, Christopher D. (1997). “Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics*, 112, 1 – 56.
- Carroll, Christopher D. (2001). “Lecture Notes on Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems,” Manuscript, Johns Hopkins University, April.
- Carroll, Christopher D. and Andrew A. Samwick (1998). “How Important Is Precautionary Saving?” *Review of Economics and Statistics*, 80, 410 – 419.
- Chandra, Amitabh and Andrew A. Samwick (2008). “Disability Risk and the Value of Disability Insurance,” in David M. Cutler and David A. Wise (eds.) *Health at Older Ages: The Causes and Consequences of Declining Disability Among the Elderly*. Chicago: University of Chicago Press, 295 – 336.
- Chen, Ying and Kent Smetters (2009). “Optimal Portfolio Choice over the Life Cycle with Social Security,” Manuscript, University of Pennsylvania, May.
- Cocco, João F., Francisco J. Gomes, and Pascal J. Maenhout (2005). “Consumption and Portfolio Choice over the Life Cycle,” *Review of Financial Studies*, 18, 491-533.
- Engelhardt, Gary V. (1996). “Consumption, Down Payments, and Liquidity Constraints,” *Journal of Money, Credit, and Banking*, 28, 255-271.
- Engen, Eric M. and Jonathan Gruber (2001). “Unemployment Insurance and Precautionary Saving,” *Journal of Monetary Economics*, 47, 545-579.
- Feldstein, Martin S. (1995). “College Scholarship Rules and Private Saving,” *American Economic Review*, 85, 552-566.

- Feldstein, Martin S. and Andrew A. Samwick (1992). "Social Security Rules and Marginal Tax Rates," *National Tax Journal*, 45, 1-22.
- Hubbard, R. Glenn and Kenneth L. Judd (1987). "Social Security and Individual Welfare: Precautionary Saving, Borrowing Constraints, and the Payroll Tax," *American Economic Review*, 77, 630-646.
- Hubbard, R. Glenn, Jonathan S. Skinner, and Stephen P. Zeldes (1995). "Precautionary Saving and Social Insurance," *Journal of Political Economy*, 103, 360-399.
- Hurst, Erik and Paul Willen (2007). "Social Security and Unsecured Debt," *Journal of Public Economics*, 91, 1273-1297.
- Kimball, Miles S. (1990). "Precautionary Saving in the Small and the Large," *Econometrica*, 58, 53-73.
- Kremer, Michael (2003). "Should Taxes Be Independent of Age?" Manuscript, Harvard University.
- Lozachmeur, Jean-Marie (2006). "Optimal Age-Specific Income Taxation," *Journal of Public Economic Theory*, 8, 697 – 711.
- Mathews, T.J. and Brady E. Hamilton (2002). "Mean Age of Mother, 1970 – 2000," *National Vital Statistics Reports* 51, No. 1, December.
[http://www.cdc.gov/nchs/data/nvsr/nvsr51/nvsr51_01.pdf]
- Samwick, Andrew A. (1998). "Discount Rate Heterogeneity and Social Security Reform," *Journal of Development Economics*, Vol. 57, No. 1, 117-146.
- Samwick, Andrew A. (2006). "Saving for Retirement: Understanding the Importance of Heterogeneity," *Business Economics*, 41, 21-27.
- Sutton, Chavon (2009). "Parents pay college costs amid recession: study," *Reuters*, August 20. [<http://www.reuters.com/article/domesticNews/idUSTRE57J5UA20090820>]
- Viard, Alan D. (2007). "The Trouble with Taxing the Young," *American Enterprise Institute, Tax Policy Outlook No. 3*, October.
- Weinzierl, Matthew (2008). "The Surprising Power of Age-Dependent Taxes," Manuscript, Harvard University.

Table 1
Housing and Income by Age

	Homeownership		Mortgage		Income	Ratios to Income	
	Percent	Median	Percent	Median	Median	Home	Mortgage
Less than 35	40.7	175	37.3	135	32.9	5.3	4.1
35–44	66.1	205	59.5	128	51.4	4.0	2.5
45–54	77.3	230	65.5	107	54.5	4.2	2.0
55–64	81.0	210	55.3	85	45.2	4.6	1.9
65–74	85.5	200	42.9	69	27.8	7.2	2.5
75 or more	77.0	150	13.9	40	22.6	6.6	1.8

Source: Survey of Consumer Finances, as tabulated by Bucks et al. (2009)

Note: Medians are conditional on ownership

Table 2
Equivalent Variation for Accepting Consumption Needs

Parameter Group	Discount Rate	Income Profile	Income Shock	Expected Utility	Consumption Needs		
					Housing	College	Both
A	8%	MW	0%	-0.0118	4.29%	0.41%	4.70%
B	8%	MW	15%	-0.0150	5.26%	0.35%	5.60%
C	8%	CGM	0%	-0.0044	15.41%	2.85%	19.14%
D	8%	CGM	15%	-0.0059	17.37%	2.36%	20.21%
E	3%	MW	0%	-0.0181	6.77%	1.73%	8.89%
F	3%	MW	15%	-0.0242	8.47%	1.94%	10.96%
G	3%	CGM	0%	-0.0091	25.19%	11.58%	41.84%
H	3%	CGM	15%	-0.0124	26.03%	12.92%	40.55%

Notes:

- 1) Expected utility is the expected present discounted value of the utility of consumption and bequests for an individual facing the given income process, a uniform payroll tax, and no additional consumption needs.
- 2) Equivalent variations are the amount that an individual facing the specified consumption needs would need in first-period assets in order to be as well off as an individual not facing that need.
- 3) Equivalent variations are specified as percentages of Equivalent Annual Income, defined as the amount that if received each year, would have the same expected present value as the specified income process.
- 4) MW refers to the income profile generated from Murphy and Welch (1990). CGM refers to the income profile from Cocco, Gomes, and Maenhout (2005). Equivalent annual income is equal to \$51,227 (CGM) or \$40,536 (MW).

Table 3
Equivalent Variation for 10-Year Revenue-Neutral Payroll Tax Exemption

Parameter Group	Discount Rate	Income Profile	Income Shock	Consumption Needs		
				Neither	Housing	College
A	8%	MW	0%	17.26%	17.71%	17.20%
B	8%	MW	15%	17.35%	18.01%	17.31%
C	8%	CGM	0%	28.56%	30.78%	27.97%
D	8%	CGM	15%	26.07%	28.42%	25.70%
E	3%	MW	0%	17.00%	17.89%	16.83%
F	3%	MW	15%	16.54%	17.67%	16.37%
G	3%	CGM	0%	9.96%	13.08%	8.71%
H	3%	CGM	15%	12.79%	13.71%	11.68%
Average				18.19%	19.66%	17.72%

Notes:

- 1) Equivalent variations are the amount that an individual facing an age-invariant payroll tax would need in order to be as well off as an individual facing 10-year revenue neutral delay in payroll taxes.
- 2) Equivalent variations are specified as percentages of Equivalent Annual Income, defined as the amount of income that if received each year, would have the same expected present value as the specified income process.
- 3) MW refers to the income profile generated from Murphy and Welch (1990). CGM refers to the income profile from Cocco, Gomes, and Maenhout (2005). Equivalent annual income is equal to \$51,227 (CGM) or \$40,536 (MW).
- 4) Under the 10-year delay, payroll taxes are zero for 10 years and then increase by factors of 1.4211 (CGM) or 1.2673 (MW) relative to the current age-invariant rates for the remainder of the working life.


Table 4
Equivalent Variation for 20-Year Revenue-Neutral Payroll Tax Exemption

Parameter Group	Discount Rate	Income Profile	Income Shock	Consumption Needs		
				Neither	Housing	College
A	8%	MW	0%	20.32%	22.32%	20.07%
B	8%	MW	15%	21.40%	23.71%	21.25%
C	8%	CGM	0%	35.71%	42.15%	34.29%
D	8%	CGM	15%	36.20%	41.40%	35.00%
E	3%	MW	0%	17.63%	20.46%	17.17%
F	3%	MW	15%	19.92%	22.99%	19.26%
G	3%	CGM	0%	2.84%	7.11%	-2.85%
H	3%	CGM	15%	15.43%	17.42%	12.07%
Average				21.18%	24.70%	19.53%

Notes:


- 1) Equivalent variations are the amount that an individual facing an age-invariant payroll tax would need in order to be as well off as an individual facing 20-year revenue neutral delay in payroll taxes.
- 2) Equivalent variations are specified as percentages of Equivalent Annual Income, defined as the amount of income that if received each year, would have the same expected present value as the specified income process.
- 3) MW refers to the income profile generated from Murphy and Welch (1990). CGM refers to the income profile from Cocco, Gomes, and Maenhout (2005). Equivalent annual income is equal to \$51,227 (CGM) or \$40,536 (MW).
- 4) Under the 20-year delay, payroll taxes are zero for 20 years and then increase by factors of 2.3803 (CGM) or 1.9634 (MW) relative to the current age-invariant rates for the remainder of the working life.

Figure 1: Median Financial Assets to Median Family Income by Age




Source: Author's tabulations of SCF data presented in Bucks et al. (2009), Tables 1 and 6

Figure 2: Primary Horizon for Financial Decisions by Age, 2007



Source: Author's tabulations of the Survey of Consumer Finances, 2007

Figure 3: Primary Reasons for Saving by Age, 2007



Source: Author's tabulations of the Survey of Consumer Finances, 2007

Figure 4: Average Labor Income Profiles






Figure 5: Expected Family Contributions for College Costs of One Child




Source: Author's tabulations of data from www.finaid.org

Figure 6: Average Consumption Profiles, Patient Individual




Notes: The Patient individual has a time preference rate of 3 percent. The four profiles differ according to whether income is uncertain with a standard deviation of 0 or 15 percent and whether the income profile is steep as in Murphy and Welch (1990) or shallow as in Cocco, Gomes, and Maenhout (2005).

Figure 7: Average Asset Profiles, Patient Individual




Notes: The Patient individual has a time preference rate of 3 percent. The four profiles differ according to whether income is uncertain with a standard deviation of 0 or 15 percent and whether the income profile is steep as in Murphy and Welch (1990) or shallow as in Cocco, Gomes, and Maenhout (2005).

Figure 8: Average Consumption Profiles, Impatient Individual




Notes: The Impatient individual has a time preference rate of 8 percent. The four profiles differ according to whether income is uncertain with a standard deviation of 0 or 15 percent and whether the income profile is steep as in Murphy and Welch (1990) or shallow as in Cocco, Gomes, and Maenhout (2005).

Figure 9: Average Asset Profiles, Impatient Individual



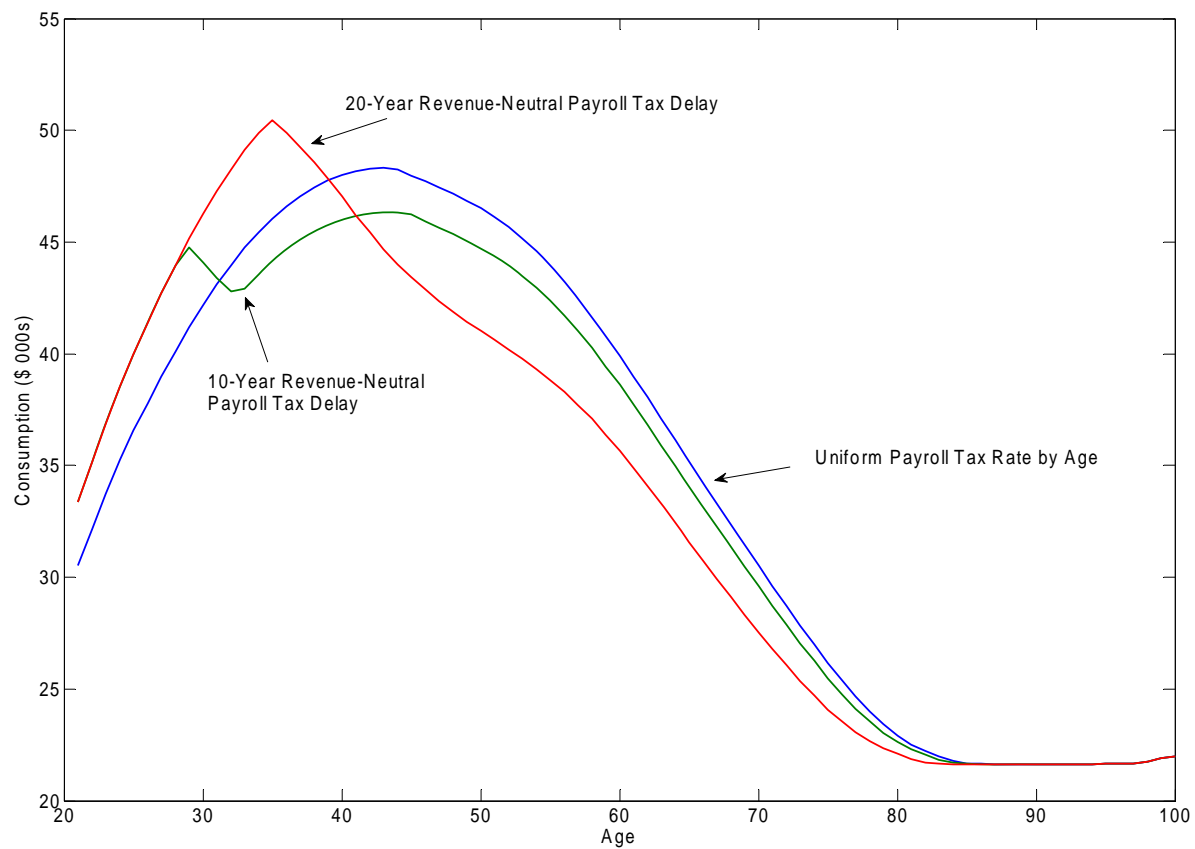
Notes: The Impatient individual has a time preference rate of 8 percent. The four profiles differ according to whether income is uncertain with a standard deviation of 0 or 15 percent and whether the income profile is steep as in Murphy and Welch (1990) or shallow as in Cocco, Gomes, and Maenhout (2005).

Figure 10: Average Consumption Profiles, Impact of Consumption Needs



Notes: Figure shows the consumption profiles for an impatient individual facing the CGM income profile and no income uncertainty.

Figure 11: Average Consumption Profiles, Impact of Payroll Tax Shifts



Notes: Figure shows the consumption profiles for an impatient individual facing the CGM income profile and no income uncertainty.