THE WELFARE EFFECTS OF PAY-AS-YOU-GO RETIREMENT PROGRAMS: THE ROLE OF TAX AND BENEFIT TIMING

Alan D. Viard

FEDERAL RESERVE BANK OF DALLAS

THE WELFARE EFFECTS OF PAY-AS-YOU-GO RETIREMENT

PROGRAMS: THE ROLE OF TAX AND BENEFIT TIMING

Alan D. Viard

Revised June 2006

THE WELFARE EFFECTS OF PAY-AS-YOU-GO RETIREMENT PROGRAMS:
THE ROLE OF TAX AND BENEFIT TIMING

Abstract: It is well known that pay-as-you-go retirement programs reduce steady-state welfare and the capital stock in dynamically efficient OLG economies. The common two-period OLG model obscures, however, the dependence of these effects on the ages at which taxes are paid and benefits are received. Program changes that shift taxes to older workers or benefits to younger retirees have effects similar to reductions in program size, yielding steady-state welfare gains and increases in capital accumulation while imposing transition costs on current generations. This analysis has policy implications for both tax and benefit timing.

I. INTRODUCTION

The study of two-period OLG models has yielded important insights into the welfare effects of pay-as-you-go retirement programs in dynamically efficient economies. A pay-as-you-go program offers windfall gains during its start-up phase, but lowers steady-state utility because its steady-state rate of return equals the economy's growth rate, which, under dynamic efficiency, is lower than the marginal product of capital. Shutting down or scaling back the program allows future generations to earn higher returns, but imposes transition costs on current generations who have paid into the program but have not yet received full benefits. The future generations' gains and the transition costs are equal in present value. It is well known that these results extend to continuous-time and multi-period OLG models.

I show, however, that the two-period model fails to capture the role of one important factor. The welfare effects of a pay-as-you-go program depend upon its life-cycle timing – the average ages at which each cohort pays taxes and receives benefits. In the two-period model, taxes must be paid in "period one" and benefits must be received in "period two." In contrast, actual programs have flexibility in the allocation of taxes

1

within the working lifetime and benefits within the retirement years. I show that the program's steady-state welfare loss is smaller when taxes are paid at later ages or benefits are received at earlier ages.

These effects arise because the pay-as-you-go program's rate-of-return shortfall is less harmful to each cohort when compounded over a shorter or later time period. Shifting taxes to older ages or benefits to younger ages therefore aids future generations in a manner similar to reducing the size of the pay-as-you-go program. Like a reduction in program size, though, such a timing shift imposes a transition cost on current generations, equal in present value to future generations' gains. Specifically, shifting taxes to older workers boosts lifetime taxes for some of the cohorts working at the time of implementation, while shifting benefits to younger retirees reduces lifetime benefits for some of the cohorts retired at that time.

In a simple calibration of the U.S. Social Security program, a payroll tax exemption during the first 10 years of working life (with a revenue-neutral tax increase on older workers) reduces the program's steady-state welfare loss by about one-sixth. This policy raises lifetime taxes for most of the cohorts working at the time of implementation. Policy changes that raise benefits for younger retirees (with a budget-neutral benefit reduction for older retirees) also generate steady-state welfare gains and transition costs, but of smaller magnitudes.

In section II, I review the familiar analysis of pay-as-you-go retirement programs in two-period OLG models. In section III, I explain the role of tax and benefit timing in a continuous-time OLG model under the assumption that all taxes are paid at a single age and all benefits received at a single age. In section IV, I show that these results generalize

to the more realistic case in which taxes and benefits are paid at multiple ages. I examine

the implications for tax timing in section V and those for benefit timing in section VI.

Section VII concludes.

II. REVIEW OF TWO-PERIOD MODELS

I begin by reviewing the basic properties of pay-as-you-go retirement programs in

the familiar two-period OLG model. Technology is linear, implying fixed factor prices,

and labor supply is inelastic. In each period t, the number of workers is N^t and the per-

worker wage is G^t. The gross-of-principal one-period marginal product of capital is R. I

assume $R > NG$, so that the economy is dynamically efficient.

Consider a simple pay-as-you-go retirement program. In period t, each of the

N^t workers pays lump-sum tax τG^t and each of the N^{t-1} retirees receives lump-sum

benefit τNG^t.[1] The lifetime present-value net burden on each period-t worker is

$$(1) \qquad \tau G^t \left\{ 1 - (NG/R) \right\}.$$

The money's worth ratio, the present value of benefits divided by the present value of

taxes, equals (NG/R).

Since the present value (1) would equal zero if R equaled NG, the program's

internal rate of return is NG, the economy's growth rate.[2] With $R > NG$, the continued

operation of the pay-as-you-go program places a burden on each future generation.[3] The

period-t closed-group liability is the period-t present value of the aggregate burden that

---

[1] As befits a pure pay-as-you-go program, budget balance is assumed to hold in each period. This assumption rules out the temporary surpluses and deficits often posted by (essentially) pay-as-you-go programs, such as the post-1983 U.S. Social Security surpluses.
[2] This result was derived by Samuelson (1958) and Aaron (1966). Also, see Lindbeck and Persson (2003, 79), Feldstein and Liebman (2002, 2257-2258), Geanakoplos, Mitchell and Zeldes (1999, 84), and Auerbach and Kotlikoff (1987, 147-148).

continuation imposes on period-t workers and later generations (equivalently, their gain

from ending the program). Multiplying $\tau G^s\{1-(NG/R)\}$, the burden on each period-s

worker, by cohort size N^s and discount factor R^{t-s} and summing across s from t to

infinity yields,

(2) $CGL_t = \tau(NG)^t$.

Although abruptly ending the program at the beginning of period t benefits

period-t workers and later generations, it imposes a "transition cost" on the period-t

retirees, who have already paid taxes, but lose benefits of $\tau(NG)^t$. The transition cost is

equal to CGL_t, the present value of future generations' gains. This present-value equality

also applies to the initial creation of the program – the present value of the program's

burden on future generations equal the start-up gains of the initial retirees, who receive

benefits without paying taxes. It can be shown that the equality also applies to gradual

and phased-in changes.[4]

As discussed by Lindbeck and Persson (2003, 82), Kotlikoff (2002, 1878-1886)

and Auerbach and Kotlikoff (1987, 148), the pay-as-you-go program also depresses

capital accumulation. With endogenous factor prices, the reduction in the capital stock

results in lower wages and a higher marginal product of capital.

It is well known that the two-period model's basic insights extend to continuous-

time models: the program's steady-state return equals the economy's growth rate, the

program lowers steady-state welfare (under dynamic efficiency), and ending the program

---

[3] In contrast, if $R < NG$ (the economy is dynamically inefficient), the pay-as-you-go retirement program increases all generations' wellbeing. Auerbach, Mankiw, Summers and Zeldes (1989) provide evidence, though, that the U.S. economy is dynamically efficient.
[4] For further discussion, see Lindbeck and Persson (2003, 80-81), Kotlikoff (2002, 1881-1882), Feldstein and Leibman (2002, 2258-2259), and Geanakopolis, Mitchell and Zeldes (1999, 86-87).

imposes a transition cost on current generations equal to the closed-group liability. I now
show, however, that continuous-time models also provide an important role for life-cycle
timing effects that are suppressed in the two-period model.

III. TWO-AGE PROGRAMS IN CONTINUOUS TIME

I examine a continuous-time overlapping-generations economy with linear
technology. At date t, e^{nt} people begin economic life, the age-a population is $e^{n(t-a)}$, and
total population is $e^{nt}P$, where $P \equiv \dfrac{1 - e^{-nL}}{n}$ and L is the length of economic life. At date
t, the per-capita wage equals e^{gt} and national labor income equals $e^{(n+g)t}P$. The marginal
product of capital is $r > n+g$, so the economy is dynamically efficient.

Consider a simple "two-age" pay-as-you-go retirement program that collects taxes
solely at age A_T and pays benefits solely at age $A_B > A_T$. Transfers are equal to a fixed
fraction τ of national labor income. At date t, each of the $e^{n(t-A_T)}$ individuals aged A_T pays
tax of $\tau P e^{gt+nA_T}$ and each of the $e^{n(t-A_B)}$ individuals aged A_B receives benefit $\tau P e^{gt+nA_B}$.
Each individual entering the economy at date s then faces a net lifetime burden with a
date-s present value of $e^{gs}(PVT - PVB)$, where

(3) $PVT \equiv \tau P e^{(n+g-r)A_T}, \quad PVB \equiv \tau P e^{(n+g-r)A_B}.$

A. Steady State Effects of Tax and Benefit Timing

As in the two-period model, the steady-state burden is proportional to the
program's size τ and is increasing in the rate-of-return shortfall r-n-g. But, the burden is
also greater if the tax age A_T is lower and the benefit age A_B is higher. These tax and
benefit timing effects, which are the focus of this paper, are not captured by the two-

period model, in which tax payment and benefit receipt must occur in "period one" and "period two," respectively.

Specifically, the money's worth ratio, the present-value ratio of benefits to taxes, is $e^{(n+g-r)(A_B-A_T)}$. This ratio has a straightforward interpretation. If an individual were required to invest in an asset paying $n+g$ rather than the market return r for an interval of length $A_B - A_T$, this would be ratio of the present value of the payout to the present value of the initial investment outlay. The same effects arise when an individual participates in a pay-as-you-go program with return $n+g$. In either context, a rate-of-return shortfall is more harmful when compounded over a longer interval.

To be more concrete, set $n+g$ equal to .03 (reflecting 1 percent population growth and 2 percent productivity growth) and r equal to .05 (a conservative estimate of the marginal product of capital). Investing at 3 percent rather than 5 percent over a one-year interval is only slightly harmful; the money's worth ratio equals $e^{-.02}$ or .98, so only 2 percent of the investment is lost due to the below-market return. But, investing at such returns over a 10-year interval is considerably more harmful; the money's worth ratio equals $e^{-.2}$ or .82 and the loss is 18 percent. Over a 30-year period, the harm is much greater, with a money's worth ratio of $e^{-.6}$ or .55 and a 45 percent loss.

A pay-as-you-go program imposes similarly small steady-state welfare losses if there is only a one-year gap between taxes and benefits; if, say, Social Security taxes were paid at age 50 and benefits received at age 51, there would be little loss from the below-market returns. The losses are much greater if taxes are paid at age 40 and benefits are received at age 70.

Expression (3) also reveals that, holding fixed the interval $A_B - A_T$ and the program size τ, the absolute burden is smaller if the tax and benefit ages are later. Delaying both the tax and the benefit ages by one year reduces each present value by 2 percent, which leaves the money's worth ratio unchanged, but reduces the size of the net burden by 2 percent. The beginning-of-life present value of the burden is therefore 2 percent smaller if taxes are paid at age 41 and benefits received at age 51 than if taxes are paid at age 40 and benefits are received at age 50. This result can also be understood by considering the investment analogy. The beginning-of-life present value burden of a required below-market investment depends upon the beginning-of-life present value of the amount invested. In this case, delaying each individual's tax by one year while holding τ fixed reduces the beginning-of-life present value of the tax by 2 percent. The one-year delay raises the size of the tax payment by 3 percent (since revenue remains a fixed fraction τ of national labor income, which grows 3 percent each year); the present value of the tax then falls by 2 percent because it is discounted an additional 5 percent.

B. Present-Value Equality Continues to Hold

As is well known, the present-value equality also holds in continuous-time models. In this two-age case, abruptly shutting down the program imposes a transition cost on individuals aged between A_T and A_B, who are denied their benefit despite having paid their tax. Of course, individuals younger than A_T gain (in the same manner as future generations), so their gains must be subtracted to obtain the net transition cost on current generations. As shown in section A of the appendix, the closed-group liability (the welfare gain to future generations) and the net transition cost are both equal to

$e^{(n+g)t} \dfrac{PVT - PVB}{r - n - g}$ for a date-t shutdown. It can be shown that, as before, the equality

also holds for gradual and phased-in changes.

The application of the present-value equality to the two-age program makes intuitive sense. As discussed above, if there is only a one-year gap between the tax and benefit ages, the gains to future generations from a shutdown are small because they are spared only one year of below-market returns. It can readily be seen that the transition cost on current generations is then also small, since only one annual cohort suffers such a cost. Conversely, if the tax and benefit ages are 30 years apart, abolition offers large gains to future generations who are spared 30 years of below-market returns. But, the transition cost is also large because 30 annual cohorts are harmed.

As in the two-period model, the present-value equality also applies to program start-up; the present value of the program's burden on future generations equal the start-up gains of the initial cohorts who receive benefits without paying taxes. The burden imposed on future generations by a program that collects taxes at age 50 and pays benefits at age 51 is small; the start-up bonus offered by its abrupt introduction is also small, because only those aged between 50 and 51 receive benefits without paying taxes. For a program that collects taxes at age 40 and pays benefits at age 70, the future burden is large; the start-up bonus is also large, with everyone aged between 40 and 70 receiving benefits without having paid taxes.

As discussed above, holding $A_B - A_T$ fixed, delaying both ages reduces the program's burden on future generations and hence their gains from its abolition. It can readily be seen that such a delay also reduces the net transition cost on current generations. As previously noted, cohorts younger than A_T gain from abolition; if A_T is

higher, then a larger number of current cohorts have such gains. Because their gains are subtracted in computing the net transition cost on current generations, that cost is smaller.

The present-value equality has another important implication. Because the equality holds for any pair of tax and benefit ages, it also holds for any change from one pair to another. By raising the tax age or lowering the benefit age, policymakers can reduce the program's steady state burden without reducing its size. But, they cannot avoid the transition cost. Abruptly raising the tax age from 40 to 50 imposes a transition cost on workers then aged between 40 and 50; they are taxed again at age 50 under the new rules, after having been taxed at age 40 under the old rules. Abruptly lowering the benefit age from, say, 70 to 60 imposes a transition cost on retirees then aged between 60 and 70; they were too young to receive benefits under the old rules, but are too old to receive benefits under the new rules. As before, the transition cost also arises from gradual and phased-in changes. As in the two-period model, there is no free lunch.

IV. GENERAL CASE

The above analysis assumes that taxes are paid at a single age and benefits are received at another single age. I now show that the conclusions also apply to programs that collect taxes and pay benefits at a variety of ages.[5]

Assume that, at each date t, each of the $e^{n(t-a)}$ individuals aged a pays $T(a)e^{g(t-a)}$ or receives $B(a)e^{g(t-a)}$, subject to the budget constraint

---

[5] Section B of the appendix confirms that the present-value equality holds in this general case.

$$P^{-1}\int_0^L e^{-(n+g)a}T(a)da = P^{-1}\int_0^L e^{-(n+g)a}B(a)da = \tau.^6$$ Then, the present-value burden for an

individual entering the economy at date t is $e^{gt}(PVT - PVB)$, where

(4) $$PVT \equiv \int_0^L e^{-ra}T(a)da, \quad PVB \equiv \int_0^L e^{-ra}B(a)da \quad.$$

Taking a first-order Taylor approximation to the logs of *PVT* and *PVB* with

respect to *r*, evaluated at *r* equal to *n+g*, and using the budget constraint yields

(5) $$\ln PVT \approx \ln(\tau P) - (r - n - g)A_T, \quad \ln PVB \approx \ln(\tau P) - (r - n - g)A_B,$$

where $$A_T \equiv \frac{\int_0^L aT(a)e^{-(n+g)a}da}{\int_0^L T(a)e^{-(n+g)a}da}, \quad A_B \equiv \frac{\int_0^L aB(a)e^{-(n+g)a}da}{\int_0^L B(a)e^{-(n+g)a}da}.$$ (From the budget constraint,

the denominators of the A_T and A_B expressions both equal τP).

Taking the exponential of (5) and substituting into (4) yields an expression

identical to (3). Up to a Taylor approximation error, the analysis is unchanged from the

two-age case, except that A_T and A_B are now weighted average ages of tax payment and

benefit receipt rather than the single ages previously considered.

In the weighted averages that define A_T and A_B, each age is weighted by the

present value of taxes or benefits at that age, using the discount rate *n+g* (the value

around which the Taylor approximation is taken). These weighted averages are

algebraically identical to the bond duration measure of Macaulay (1938, 48-50), which is

prominent in the bond pricing literature. Macaulay duration is a weighted average of the

time remaining until a bond's future payments, with weights given by the present value

of each payment. The economic interpretation is the same in both contexts; just as a

---

[6] In the two-age case, $T(a)$ is $\tau Pe^{(n+g)A_T}$ at A_T and zero elsewhere while $B(a)$ is $\tau Pe^{(n+g)A_B}$ at A_B and
zero elsewhere.

bond's duration governs the sensitivity of its present value (price) to the interest rate, so these weighted average ages govern the sensitivity of the present values *PVT* and *PVB* to the discount rate *r*.

To obtain more specific results and to avoid reliance on a Taylor approximation, I calibrate a stylized representation of the U.S. Social Security old-age and survivor (but not disability) program, as further detailed in section C of the appendix. I continue to set *n* to .01, *g* to .02, and *r* to .05. I assume that individuals work from economic ages 0 to *X* and are retired from economic ages *X* to *L*. I set *X* equal to 42 and *L* equal to 60, corresponding to work from biological ages 20 to 62 and retirement from biological ages 62 to 80. The population parameter *P* is then 45.1.

Under the benchmark policy, the program is financed by an age-uniform payroll tax of rate τ, so the timing of tax payments matches the timing of wages. I fit a quadratic cross-sectional age-earnings profile to recent data. Also, under the benchmark policy, benefits are paid from ages *X* to *L* and remain unchanged in real terms for each cohort throughout retirement.

In this benchmark case, *PVT* equals 29.89τ and *PVB* equals 16.62τ. These values are the same as those in a two-age program with A_T equal to 20.6 (roughly the midpoint of working life) and A_B equal to 49.9 (close to the midpoint of the retirement period). The money's worth ratio is .556. The closed-group liability is 14.7 times annual benefit payments, or about $6.4 trillion in 2005.

Starting from the benchmark policy, I first examine reforms that alter the timing of taxes and then turn to reforms that change the timing of benefits.

V. TAX TIMING CHANGES

Consider a revenue-neutral replacement of the age-uniform payroll tax with a young-worker exemption policy. As detailed in section D of the appendix, such a policy imposes no tax on the earnings of workers below economic age Y and maintains revenue neutrality by taxing earnings from ages Y through X at a rate higher than τ. **Figure 1** depicts the steady-state tax burden PVT for exemption ages from zero (the benchmark policy) to 42 (the extreme policy in which each worker pays her lifetime taxes in a single payment right before retirement). As the exemption age rises, the lifetime present value of the tax burden falls, in accordance with the above analysis.

FIGURE 1
Present Value of Taxes with Young-Worker Exemption
(multiple of annual per-capita tax at date of labor force entry)
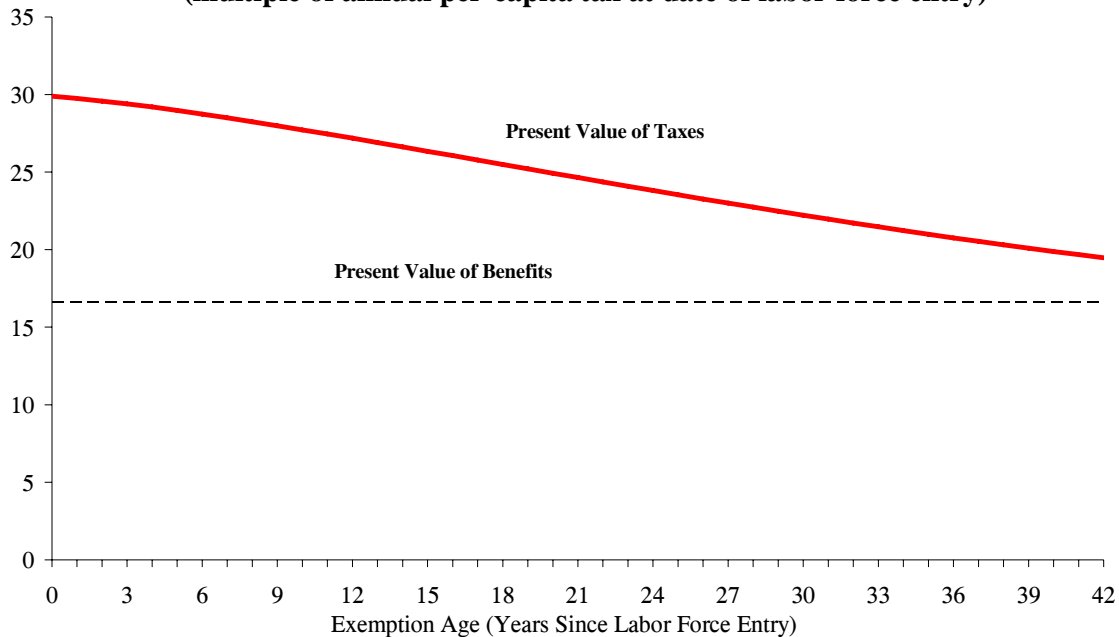
Consider Y equal to 10, so that workers are exempt from taxes during the first 10 years of working life (biological ages 20 through 30), with a revenue-neutral increase in the tax rate on older workers to 1.21τ. Then, PVT is reduced from 29.89τ to 27.72τ,

which is equivalent to raising A_T from 20.6 to 24.4 in a two-age program. Since *PVB* still

equals 16.62τ, the net lifetime loss from the pay-as-you-go system falls from 13.27τ to

11.10τ, a reduction of more than 16 percent. In other words, the steady-state gain from

this young-worker exemption is equal to that attained by scaling back the system by one-

sixth across the board. The reduction in the 2005 closed-group liability is about $1.0

trillion.

In accord with the present-value equality, however, the gain to future generations

is accompanied by a transition cost of the same size on current generations. If exempting

workers from taxes during their first 10 working years yields the same future gains as

shrinking the program by one-sixth, then it must impose the same aggregate transition

costs on current generations. The costs are allocated differently; under the young-worker

exemption, the cost is borne solely by current workers, rather than current retirees.

FIGURE 2
Individuals' Net Losses from Young-Worker Exemptions
(multiple of annual per-capita tax at implementation date)

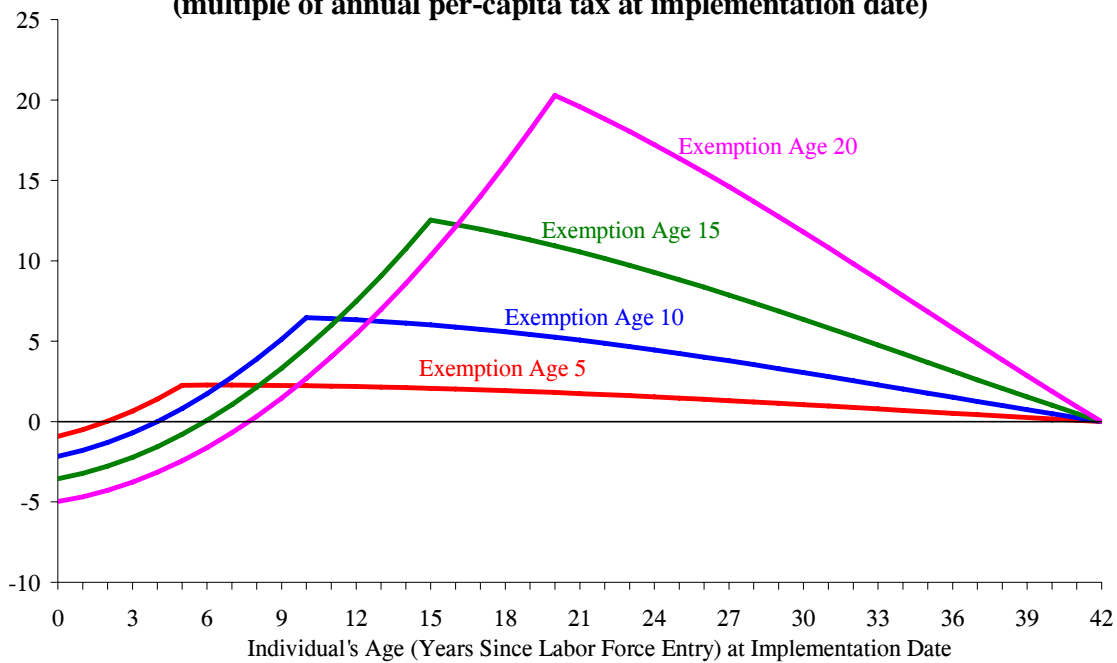Individual's Age (Years Since Labor Force Entry) at Implementation Date

Figure 2 plots the present-value loss (negative if gain) borne by each member of the various working cohorts from the abrupt introduction of young-worker exemptions with exemption ages of 5, 10, 15, and 20. For Y equal to 10, workers aged 0 through 4.0 are net winners and older workers are net losers. In general, the loss is greatest for workers around the exemption age, who obtain no gain from the exemption and who have the longest exposure to the higher rate.

By assuming particular utility and production functions (as detailed in section E of the appendix), it is possible to compute the effects on capital accumulation. For this purpose, I set τ equal to .056, the 2005 ratio of old-age and survivor benefits to national labor income. **Figure 3** plots the increases in the steady-state capital stock resulting from young-worker exemptions for ages 0 through 42. For comparison, the chart shows the 8.0 percent increase that would arise from abolition of the pay-as-you-go program.

FIGURE 3
Increase in Steady-State Capital Stock from Young-Worker Exemption

Percent Increase

Increase in Steady-State Capital Stock from Abolition of Pay-As-You-Go Program

Exemption Age (Years Since Labor Force Entry)
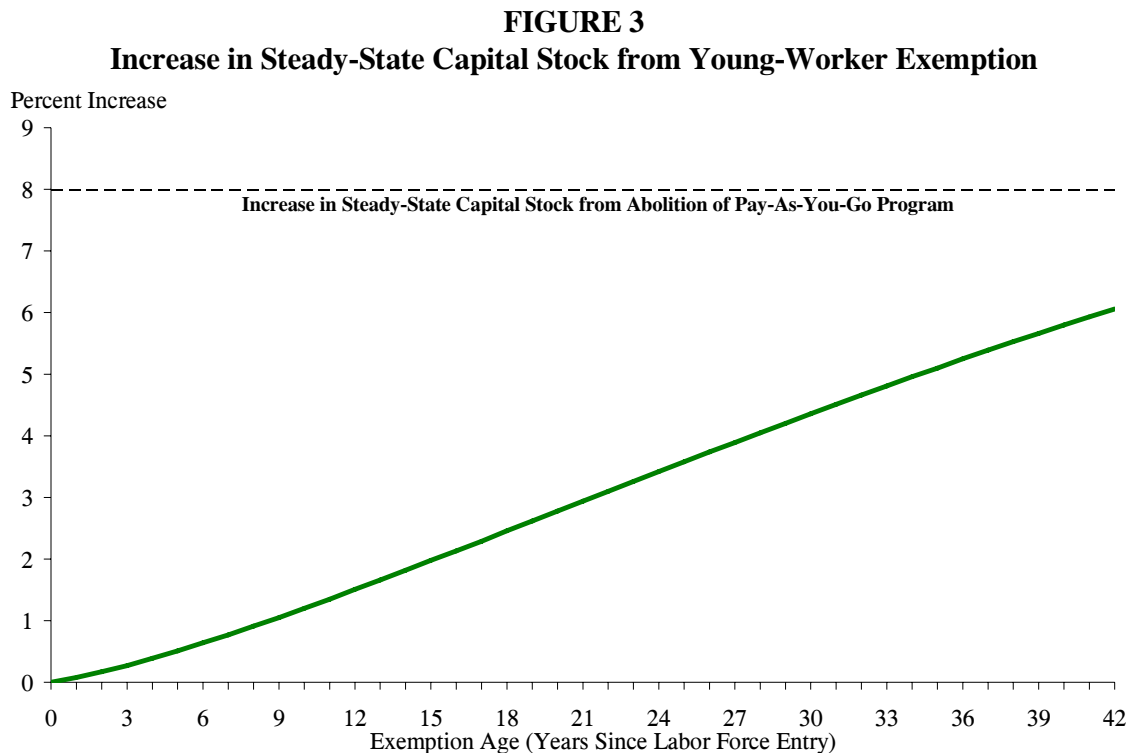
The relative effects of different policies are virtually unchanged from the partial-equilibrium framework. For example, an exemption age of 10 increases the steady-state capital stock by 1.2 percent, which is 15 percent of the increase attained from abolition of the program; recall that this policy yielded 16 percent of the partial-equilibrium steady-state welfare gains offered by abolition.

The capital-accumulation effects follow straightforwardly from the analysis of Seidman and Lewis (2003), who show that any revenue-neutral shift of the tax burden from young to old increases the steady-state capital stock. As they note, a similar point has been made by authors studying the choice between consumption and wage taxation, including Auerbach and Kotlikoff (1987, 58-60) and Summers (1981). The present analysis applies this general insight to payroll taxation, thereby linking the analysis to the literature on pay-as-you-go retirement programs and permitting an extension to benefit timing. This analysis, like much of the pay-as-you-go literature, emphasizes the partial-equilibrium welfare effects, with less attention to the capital-accumulation effects emphasized by Seidman and Lewis. This difference in emphasis is largely a matter of taste, since the two effects are inextricably linked.

Hubbard and Judd (1987) present a separate argument for a young-worker payroll tax exemption based on borrowing restrictions. (Also, see Hurst and Willen (2004)). Hubbard and Judd assume a pre-funded Social Security system that pays the market return r. They note that, if young workers face binding borrowing restrictions, their shadow interest rates exceed r and it is then desirable to delay tax payments. In contrast, the present analysis assumes no borrowing restrictions, so that workers' shadow interest rate equals r, but considers a pay-as-you-go system that pays rate of return $n+g < r$. If

these complementary analyses are combined, the steady-state welfare gain from a young worker exemption is even larger, as borrowing restrictions push workers' shadow rate above r while the pay-as-you-go system offers a return, $n+g$, lower than r.

VI. CHANGES IN BENEFIT TIMING

The analysis in the preceding section considered the steady-state gains and transition costs associated with delaying tax payments. As discussed earlier, qualitatively similar effects can be achieved by accelerating benefit receipt. Although benefit timing changes generally have smaller effects than tax timing changes, the effects can still be significant. I consider three policies to alter benefit timing, with full details in section F of the appendix.

The easiest way to compare the effects of tax and benefit timing changes is to consider the (unrealistic) policy that is analytically parallel to the young-worker exemption; an old-retiree cutoff that eliminates benefits for retirees above age J with a budget-neutral benefit increase for younger retirees. The old-retiree cutoff has much smaller effects than the young-worker exemption. For example, the extreme policy of paying all lifetime benefits at the onset of retirement (J equal to 42) raises PVB only from 16.62τ to 19.48τ, a gain of 2.86τ; the corresponding extreme policy of collecting all lifetime taxes at that same age (Y equal to 42) lowers PVT from 29.89τ to 19.48τ, a gain of 10.41τ, over three times larger.[7] Similarly, denying benefits during the second half of retirement (J equal to 51) raises PVB by 1.28τ, while eliminating taxes during the first half of working life (Y equal to 21) lowers PVT by 5.25τ.

---

[7] The combination of the two extreme policies equates PVB to PVT at 19.48τ, eliminating the steady-state welfare loss. Indeed, this combination shuts down the program, since there are no real effects from paying taxes that are immediately and fully refunded in the form of benefits.

The smaller impact of benefit timing changes is easily explained. Raising A_T by one year and lowering A_B by one year have the same proportional effects, reducing *PVT* by 2 percent and increasing *PVB* by 2 percent. But, since *PVT* is almost twice as large as *PVB* under the benchmark policy, the tax change has the larger absolute effect. Also, because working life is longer than retirement (here, 42 versus 18 years), the tax changes in the above comparisons alter A_T by more than the benefit changes alter A_B. For example, the extreme tax timing policy raises A_T by 21.6 years, from 20.4 to 42, while the extreme benefit timing policy lowers A_B by only 7.9 years, from 49.9 to 42.

I next consider changes in the benefit growth rate during retirement. Under the benchmark policy (as in the actual Social Security system), each individual's real benefits remain constant throughout retirement. A budget-neutral change that hikes initial benefits but then lets benefits fall 4 percent per year throughout retirement raises *PVB* from 16.62τ to 16.96τ. Conversely, a budget-neutral change that reduces initial benefits and then lets them rise 4 percent per year lowers *PVB* to 16.27τ.

Finally, I consider a policy parameter that has been changed in past Social Security reforms, the benefit eligibility age. Relative to leaving the program unchanged, an eligibility-age increase obviously shrinks the program and reduces the steady-state welfare loss. Because the eligibility-age increase delays benefit receipt, however, it reduces the steady-state burden by less than a budget-equivalent across-the-board benefit reduction.

To examine this issue, consider a budget-neutral eligibility-age increase in which benefits are delayed until $V > X$ and are increased for older retirees. As shown in **Figure 4**, denying benefits during the first three years of retirement (*V* equal to 45) lowers *PVB*

from 16.62τ to 16.03τ. Although changes in benefit timing have smaller impacts than changes in tax timing, the effects can still be significant. In 2005, a three-year eligibility-age increase would result in a present-value aggregate loss for future generations of almost $300 billion, compared to across-the-board cuts that lowered aggregate benefits by the same amount.

FIGURE 4
Present Value of Benefits, Under Various Eligibility Ages
(multiple of annual per-capita tax at date of labor force entry)
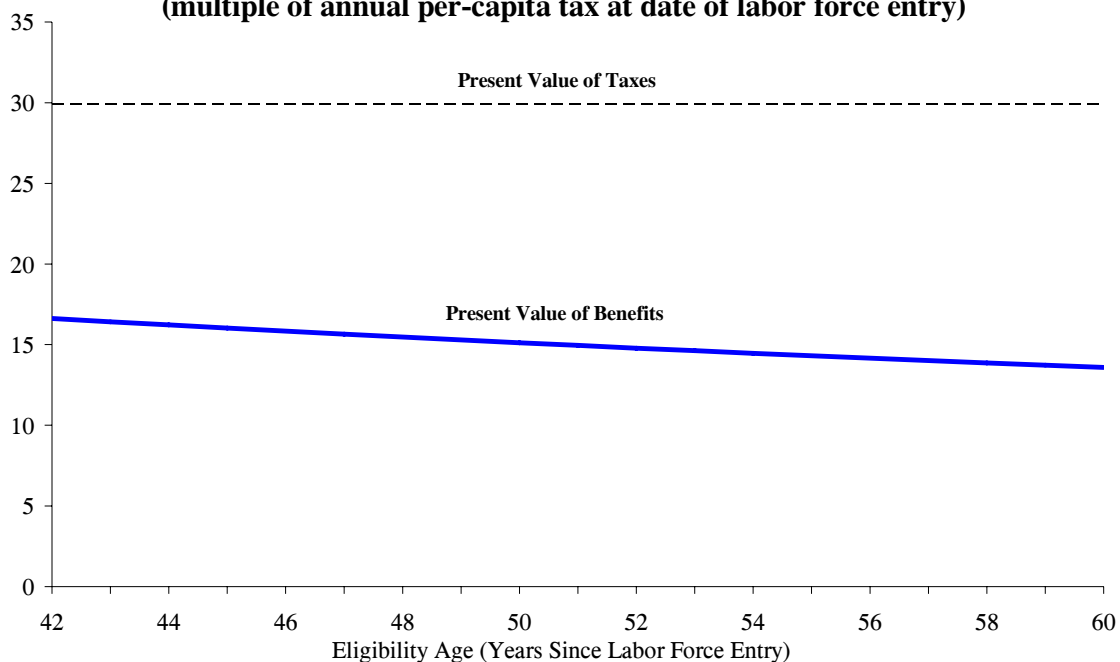
This analysis has assumed known lifetimes. In a world with uncertain lifetimes and imperfect annuitization, the policies considered in this section, particularly the old-retiree cutoff, would harm individuals who enjoy unexpectedly long lifetimes. Such effects would have to be considered in a more complete analysis.[8] Like the smaller impact of benefit timing changes, these effects may suggest that tax timing changes are of greater policy relevance.

---

[8] Feldstein (1990) considers the choice of benefit growth rates in a four-period OLG model, analyzing both the effects considered here and the effects of uncertain lifetimes and imperfect annuitization.

VII. CONCLUSION

In a continuous-time OLG model, a pay-as-you-go retirement program's steady-state welfare loss is higher when taxes are paid earlier or benefits are received later. The larger loss arises because the pay-as-you-go program's rate-of-return shortfall is more harmful to each cohort when compounded over a longer or earlier time period. Policy changes that exempt younger workers from payroll taxes (with a revenue-neutral tax increase on older workers) increase steady-state welfare, but impose transition costs on older workers when implemented. Policy changes that increase benefits for younger retirees (with a budget-neutral benefit cut for older retirees) have similar, but smaller, effects. Policy analyses of proposed changes in pay-as-you-go retirement programs should consider how the changes affect tax and benefit timing.

REFERENCES

Aaron, H.J. "The Social Insurance Paradox." *Canadian Journal of Economics and Political Science*, 32, 1966, 371-374.

Abel, A.B., N.G. Mankiw, L.H. Summers, and R.J. Zeckhauser. "Assessing Dynamic Efficiency: Theory and Evidence." *Review of Economic Studies*, 56(1), 1989, 1-19.

Auerbach, A.J. and L.J. Kotlikoff. *Dynamic Fiscal Policy*. Cambridge, UK: Cambridge University Press, 1987.

Feldstein, M. "Imperfect Annuity Markets, Unintended Bequests, and the Optimal Age Structure of Social Security Benefits." *Journal of Public Economics*, 41(1), 1990, 31-43.

Feldstein, M. and J.B. Leibman. "Social Security," in *Handbook of Public Economics*, v. 4, edited by Alan J. Auerbach and Martin Feldstein. Amsterdam: North-Holland, 2002, 2245-2324.

Geanakoplos, J., O.S. Mitchell and S.P. Zeldes. "Social Security Money's Worth," in *Prospects for Social Security Reform*, edited by Olivia S. Mitchell, Robert J. Myers and Howard Young. Philadelphia: University of Pennsylvania Press, 1999, 79-151.

Hubbard, R.G. and K.L. Judd. "Social Security and Individual Welfare: Precautionary Saving, Borrowing Constraints, and the Payroll Tax." *American Economic Review*, 77(4), 1987, 630-646.

Hurst, E. and P. Willen. "Social Security and Unsecured Debt." National Bureau of Economic Research Working Paper 10282, January 2004.

Kotlikoff, L.J. "Generational Policy," in *Handbook of Public Economics*, v. 4, edited by Alan J. Auerbach and Martin Feldstein. Amsterdam: North-Holland, 2002, 1873-1932.

Lindbeck, A. and M. Persson. "The Gains from Pension Reform." *Journal of Economic Literature*, 41(1), 2003, 74-112.

Macaulay, F.R. *Some Theoretical Problems Suggested by the Movement of Interest Rates, Bond Yields, and Stock Prices in the US Since 1856*. New York: NBER, 1938.

Samuelson, P.A. "An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money." *Journal of Political Economy*, 66(6), 1958, 467-482.

Seidman, L.S. and K.A. Lewis. "The Later You Pay, the Higher the *k*." *Southern Economic Journal*, 69(3), 2003, 560-577.

Summers, L.H. "Capital Taxation and Accumulation in a Life Cycle Growth Model." *American Economic Review*, 71(4), 1981, 533-544.

APPENDIX

Section A: Closed-Group Liability and Transition Cost in Two-Age Model

To obtain the date-*t* closed-group liability, multiply the burden on each date-*s*

entrant from (3) by the number of such entrants e^{ns} and the discount factor e^{-rs} and

integrate across *s* from *t* to infinity, yielding $\tau P e^{(n+g)t} \dfrac{e^{(n+g-r)A_T} - e^{(n+g-r)A_B}}{r-n-g}$.

The net transition cost is computed as follows. For each age *a* between A_T and A_B,

$e^{n(t-a)}$ individuals lose $\tau P e^{g(t-a)+(n+g)A_B}$ at date $t + A_B - a$; the cohort's payments have date-

t present value $\tau P e^{(n+g)t+(n+g-r)(A_B-a)}$. Integrating across *a* from A_T to A_B yields gross

transition cost $\tau P e^{(n+g)t} \dfrac{1 - e^{(n+g-r)(A_B-A_T)}}{r-g-n}$. For each age *a* less than A_T , $e^{n(t-a)}$ individuals

avoid a burden with date-*t* present value $e^{g(t-a)+ra} \tau P \left\{ e^{(n+g-r)A_T} - e^{(n+g-r)A_B} \right\}$. Integrating

across *a* from zero to A_T yields $\tau P e^{(n+g)t} \left\{ e^{(n+g-r)A_T} - e^{(n+g-r)A_B} \right\} \dfrac{e^{(r-n-g)A_T} - 1}{r-g-n}$ or

$\tau P e^{(n+g)t} \dfrac{e^{(r-n-g)A_B} + 1 - e^{(n+g-r)(A_B-A_T)} - e^{(n+g-r)A_T}}{r-g-n}$. Subtracting from the gross transition

cost yields $\tau P e^{(n+g)t} \dfrac{e^{(n+g-r)A_T} - e^{(n+g-r)A_B}}{r-n-g}$.

Section B: Closed-Group Liability and Transition Cost in General Case

To obtain the date-*t* closed group liability, apply discount factor $e^{r(t-s)}$ to the

burden $e^{gs}(PVT - PVB)$ on each of the e^{ns} date-*s* entrants and integrate across *s* from *t*

to infinity to obtain $e^{(n+g)t} \dfrac{PVT - PVB}{r-n-g}$.

The transition cost is computed as follows. The combined date-t present-value loss of the cohorts aged 0 through L as of date t equals

$$e^{(n+g)t}\int_0^L e^{-(n+g)h}\left(\int_x^L e^{r(h-a)}[T(a)-B(a)]da\right)dh\,,\text{ which equals}$$

$$e^{(n+g)t}\int_0^L e^{-ra}\left(\int_0^a e^{(r-n-g)h}dh\right)[T(a)-B(a)]da\ \text{ or }\ e^{(n+g)t}\int_0^L \frac{e^{-(n+g)a}-e^{-ra}}{r-n-g}[T(a)-B(a)]da\,.$$

Since $\int_0^L e^{-(n+g)a}[T(a)-B(a)]da = 0$ from the budget constraint and

$$\int_0^L e^{-ra}[T(a)-B(a)]da = PVT - PVB\,,\text{ this expression can be rewritten as}$$

$$e^{(n+g)t}\frac{PVB-PVT}{r-n-g}\,.$$

Section C: Calibration of Social Security System

An individual of age a (from 0 to X) at date t receives earnings

$$e^{gt}P\frac{1+w_1 a+w_2 a^2}{Z(n+g,0,X)}\,,\text{ where } Z(n+g,0,X)\text{ is a scaling factor. For any discount rate } \rho,\text{ any}$$

starting age α, and any ending age ω, the scalar $Z(\rho,\alpha,\omega)$ is the age-zero present value (at discount rate ρ) of $e^{ga}\left(1+w_1 a+w_2 a^2\right)$ between the specified ages,

$$
\begin{aligned}
Z(\rho,\alpha,\omega) &\equiv \int_\alpha^\omega e^{(g-\rho)a}(1+w_1 a+w_2 a^2)da\\
&=\left[e^{(g-\rho)\alpha}-e^{(g-\rho)\omega}\right]\left\{(\rho-g)^{-1}+(\rho-g)^{-2}w_1+2(\rho-g)^{-3}w_2\right\}\\
&+\left[\alpha e^{(g-\rho)\alpha}-\omega e^{(g-\rho)\omega}\right]\left\{(\rho-g)^{-1}w_1+2(\rho-g)^{-2}w_2\right\}+\left[\alpha^2 e^{(g-\rho)\alpha}-\omega^2 e^{(g-\rho)\omega}\right](\rho-g)^{-1}w_2
\end{aligned}
$$

The $Z(n+g,0,X)$ term in the denominator of the earnings function scales wages to keep the

per-capita wage equal to e^{gt}. So, $T(a)=e^{ga}\tau P\dfrac{1+w_1 a+w_2 a^2}{Z(n+g,0,X)}$ and PVT then equals

$$\tau P\frac{Z(r,0,X)}{Z(n+g,0,X)}\,.$$

Using Social Security Administration worker and earnings data and Census

population data for 2003, I constructed a proxy for per-capita earnings (number of

workers multiplied by median earnings divided by population) for each five-year age

cohort between ages 20 and 60 and for the two-year cohort aged 61 to 62. I regressed this

proxy on a constant, the midpoint economic age of each group (treating biological age 20

as economic age zero), and age squared and rescaled the coefficients to set the intercept

to one, obtaining w_1 equal to .2608 and w_2 equal to -.00511.

The budget constraint then requires that $B(a) = \tau P \dfrac{n+g}{e^{-(n+g)X} - e^{-(n+g)L}}$ for a from X

to L. Then, PVB equals $\tau P \dfrac{n+g}{r} \dfrac{e^{-rX} - e^{-rL}}{e^{-(n+g)X} - e^{-(n+g)L}}$.

Section D: Young-Worker Exemption

The young-worker exemption taxes ages Y through X at rate $\tau \dfrac{Z(n+g,0,X)}{Z(n+g,Y,X)}$.

Under this policy, $T(a)$ equals zero for a less than Y and equals $e^{ga}\tau P \dfrac{1 + w_1 a + w_2 a^2}{Z(n+g,Y,X)}$ for

a from Y through X. Then, PVT equals $\tau P \dfrac{Z(r,Y,X)}{Z(n+g,Y,X)}$.

At date t, for each a between Y and X, each of $e^{n(t-a)}$ workers suffers a loss with

present value $e^{ra+g(t-a)}\tau P \left\{ \dfrac{Z(n+g,0,X)}{Z(n+g,Y,X)} - 1 \right\} \dfrac{Z(r,a,X)}{Z(n+g,0,X)}$ from the higher tax rate. For

each a between 0 and Y, each of $e^{n(t-a)}$ workers of age a suffers a loss with present value

$e^{ra+g(t-a)}\tau P \left\{ \dfrac{Z(n+g,0,X)}{Z(n+g,Y,X)} - 1 \right\} \dfrac{Z(r,Y,X)}{Z(n+g,0,X)}$ from the higher tax rate they will face after

attaining age Y, but has a gain with present value $e^{ra+g(t-a)}\tau P \dfrac{Z(r,a,Y)}{Z(n+g,0,X)}$ from the

exemption enjoyed until age Y.

Section E: Equilibrium Capital Stock

Letting K denote the capital stock divided by national labor income, the steady

state values of K and r satisfy two conditions. First, $K = \dfrac{\kappa}{(1-\kappa)(r+\delta)}$, where κ is the

capital share in a Cobb-Douglas production function for gross-of-deprecation output and

δ is the deprecation rate. Second, K equals the aggregate stock of national saving, which

is the present value of future consumption minus the present value of future disposable

non-capital income,

$$K = \int_0^L e^{-(n+g)h} \left\{ \int_h^L e^{r(h-a)} K(a)\,da \right\} dh = \int_0^L e^{-(n+g)h} \left\{ \int_h^L e^{r(h-a)} [C(a)-W(a)-B(a)+T(a)]\,da \right\} dh,$$

where $K(a)$, $C(a)$ and $W(a)$ are capital holdings, consumption, and pretax labor income of

each worker at age a. (Like $B(a)$ and $T(a)$, they are expressed as a fraction of the per-

capita wage when the individual is of economic age zero.) If the consumer maximizes

$\int_0^L e^{-\lambda a} \log(C(a))\,da$, then $C(a) = \dfrac{\lambda e^{(r-\lambda)a}}{1-e^{-\lambda L}} \int_0^L e^{-rh}[W(x)+B(x)-T(x)]\,dh$. Setting κ to .35, δ

to .04, and τ equal to .056, I solved backwards to find that a time preference rate λ of .018

yields a capital stock consistent r equal to .05 under the age-uniform benchmark policy.

(In 2005, old-age and survivor benefits were 5.6 percent of national labor income,

defined as employee compensation plus two-thirds of proprietors' income). I then

computed the equilibrium values of K and r under program abolition and the young-

worker exemptions.

Section F: Changes in Benefit Timing

A budget-neutral old-retiree cutoff with cutoff age J sets $B(a)$ equal to

$$\tau P \frac{n+g}{e^{-(n+g)X} - e^{-(n+g)J}}$$ for a from X through J and zero for a from J through L. Then, PVB

equals $\tau P \dfrac{n+g}{r} \dfrac{e^{-rX} - e^{-rJ}}{e^{-(n+g)X} - e^{-(n+g)J}}$.

A budget-neutral change that lets each cohort's real benefit grow at rate q sets

$B(a)$ equal to $e^{qa} \tau \dfrac{1 - e^{-nL}}{n} \dfrac{n+g-q}{e^{(q-n-g)X} - e^{(q-n-g)L}}$ for a from X to L. Then, PVB equals

$$\tau P \frac{n+g-q}{r-q} \frac{e^{(q-r)X} - e^{(q-r)L}}{e^{(q-n-g)X} - e^{(q-n-g)L}} .$$

A budget-neutral increase in the eligibility age to V sets $B(a)$ equal to zero for a

from X through V and equal to $\tau P \dfrac{n+g}{e^{-(n+g)V} - e^{-(n+g)L}}$ for a from V through L. Then, PVB

equals $\tau P \dfrac{n+g}{r} \dfrac{e^{-rV} - e^{-rL}}{e^{-(n+g)V} - e^{-(n+g)L}}$.