Pension Systems and Aggregate Shocks

KARSTEN JESKE

The author is an economist in the macropolicy section of the Atlanta Fed's research department. *He thanks Thomas Cunningham, Juan Rubio-Ramírez, and Ellis Tallman for helpful comments.*

To some analysts the prospective imbalances between pay-as-you-go pension system benefits and the tax base present a great economic challenge in the decades to come.[1] It is currently estimated that the U.S. Social Security Trust Funds will run out of money in 2041, after which either benefit cuts or major tax hikes would have to occur.[2] Other countries face similar or even tougher and shorter-term problems. For example, in Germany wages reported to the defined benefit system are already taxed at a rate of about 20 percent of gross income, compared to 12.4 percent in the United States, and are expected to rise to a staggering 28 percent by 2035. Sinn (1999) calculates that the current value of all implicitly promised future benefits amounts to a number roughly 250 percent of current German annual gross domestic product (GDP), an overwhelming figure compared to the current government-debt-to-GDP ratio of 60 percent.

In the United States and in most Western democracies, proposals to cut benefits to the extent necessary to save social security are politically infeasible. Raising contributions is—from an economic point of view—undesirable because proportional payroll taxes have distortionary effects on both labor supply and savings decisions. In light of these impending funding problems, politicians and academics alike are calling for a reform of social security. In the United States there is a near consensus that such reform is necessary, but there is also controversy about what the reform should look like. One suggestion, put forward by the President's Commission to Strengthen Social Security, is a partial privatization of social security (President 2001). In this proposal current social security surpluses could be used to fund private, individual retirement account (IRA)-style accounts, and the private savings could make up for future benefit cuts.

In the privatization debate, two opposing perceptions seem to hinder productive discussion. First, social security is considered a low-risk vehicle for the provision of old-age income. Second, the exact size of the social security funding problem is extremely difficult to forecast. Minute changes in the growth assumptions of the forecasting models lead to large changes in the long-term forecasts for social security feasibility. In fact, a few economists argue that with only slightly larger annual growth rates than the conservatively chosen rates used by the Social Security Administration, social security will not face any funding problem whatsoever.[3] For example, Robert Reich, a former trustee of the Social Security Trust Fund, argues that "The actuary's projections are based on the pessimistic assumption that the economy will grow only 1.8 percent annually over the next three decades. Crank the economy up just a bit, to a more realistic 2.2 percent a year, and the fund is nearly flush for the next seventy-five years" (Reich 1998).

Unfortunately, the argument also works in the opposite direction; just slightly smaller growth rates than the ones predicted lead to even more catastrophic scenarios than the ones already discussed. To see how sensitive the trust fund finances are with respect to aggregate variables, one need look only at the calculations of the Social Security Administration. Under the benchmark assumption, the value of the trust fund (in current dollars) will be $6.7 trillion in 2030, and the value decreases until 2041, when the fund is depleted. Increasing the wage growth rate from 1.1 percent to 1.6 percent and the labor force growth rate from 0.2 percent to 0.6 percent would not only almost double the trust fund value in 2030 to more than $12 trillion— almost $7 trillion in today's dollars—but would also ensure that the fund is never depleted over the horizon of eighty years. On the other hand, slightly lower growth rates of 0.6 percent for wages and –0.3 percent for the labor force would cause the trust fund to be depleted by 2029. The implicit confidence interval for the estimated trust fund value in 2030 then covers everything between a slight trust fund liability up to a staggering $7 trillion surplus in today's dollars, more than four times the federal budget for fiscal year 2002.

The results of the model simulations show that, in the long run, privatization makes every generation better off, even if a large aggregate shock occurs.

From an economist's perspective, the two perceptions of low-risk social security and the extreme sensitivity of social security finances with respect to unpredictable economic fundamentals contradict each other. Social security cannot be completely riskless, as the uncertainty about the predicted funding problem demonstrates. The viability of social security crucially depends on what such volatile variables as productivity growth, fertility, and immigration turn out to be over the next decades. From a macroeconomic perspective, a pay-as-you-go (PAYGO) system therefore implies a substantial amount of risk, contrary to the amount that proponents of social security would admit.

Moreover, the factors that tend to drive the performance of a PAYGO system are the same that drive returns on financial market assets, and they push in the same direction. For example, lower productivity growth indeed has a negative effect on the returns to physical capital and therefore reduces financial market returns. At the same time, however, lower productivity growth also jeopardizes a PAYGO system because the system's promised benefits become more difficult to finance if wage growth rates are lower than expected.

PAYGO returns also tend to be lower than financial market returns. If, in addition to this low return, PAYGO also has potentially high risk and a high correlation with financial market returns, then—from the perspective of a Sharpe (1964) and Lintner (1965) capital asset pricing model (CAPM)—a PAYGO system may be viewed as a very undesirable asset. To determine whether and to what degree a PAYGO system is undesirable, one must design a model economy in which aggregate shocks affect not only financial market returns but also a PAYGO system. For policy analysis, this model can help to evaluate proposals. For example, privatization may look unattractive because it exposes the retirement income of a representative generation to a substantial amount of risk. If, at the same time, however, social security faces a symmetric risk profile, then privatization appears more attractive.

This article provides such a model in order to address the sensitivity of different retirement schemes to large aggregate shocks, such as a major drop in productivity growth or demographic shocks like a baby-boomer generation. The workhorse model used in the analysis is a so-called life-cycle economy in which agents work when they are young and their old-age consumption is financed by a combination of a PAYGO pension system and private savings. The model is then used to determine how different retirement schemes perform under different kinds of shocks.

The results of the model simulations show that, in the long run, privatization makes every generation better off, even if a large aggregate shock occurs. The intuition for this result is that, under a privatized pension system, savings are higher, and this higher savings level increases the aggregate capital stock because private savings are more desirable and affordable if both benefits and contributions are lower. This higher capital stock increases welfare by an amount high enough to insulate all future generations even from large aggregate shocks. Aggregate risk is mainly a concern for the period immediately after a social security reform.

A Simple Life-Cycle Model

This section introduces a simple model that allows one to assess how sensitive different retirement systems are with regard to a variety of aggregate

shocks. The model is extremely stylized and abstracts from many real-world matters in order to be analytically and computationally tractable. Many simplifications, however, are performed in such a way as to give a PAYGO system the best possible chance to perform well relative to private retirement accounts. Consequently, the model consistently underestimates gains from privatizing social security; therefore the potential welfare improvements presented here serve as a lower bound.

The model is an overlapping generations (OLG) model; people live for a maximum number of periods, and, to lend the model a greater degree of realism, in each period they have a given probability of dying. That is, in every period there is a distribution of agents of different ages. At the beginning of each period a new generation of young agents enters the economy, and in each of the existing generations a specified fraction of agents dies and leaves the economy.

Just as in the real world, agents have a hump-shaped labor productivity profile; that is, people start with a relatively low labor productivity associated with lower wages and then accumulate human capital until their labor productivity peaks at about the age of fifty, after which productivity slowly decreases until age sixty-five. The model assumes that retirement is mandatory at age sixty-six; that is, productivity (and therefore labor income) falls to zero, and people must live on their private savings and a government-sponsored pension plan thereafter.

In the model simulations in the next section, it is assumed that one period is six years. People enter the labor force at the age of eighteen and may live for up to fifteen periods, that is, to an age of 108; the survival probabilities are matched to the probabilities computed from data from the National Center for Health Statistics. For the computations and the quantitative results presented later in the article, this more sophisticated and realistic multigenerational model is used, but most of the intuition works just fine with a simpler version of the model with only two generations alive at any given time.

In this simpler model, even though the economy has infinitely many periods, individuals live for only exactly two periods; that is, there is no probability that a person will die before reaching the second period. (This assumption will be relaxed in the more sophisticated model used in the numerical examples.) The precise timing is illustrated in Figure 1. In period 1 there is an initial old generation (generation 0) that lives only for exactly one more period and then dies and a young generation (generation 1) that lives for two periods. In period 2, generation 2 is born and serves as the young generation whereas generation 1, the previously young generation, is now the old generation. More generally, in period t there are two generations alive: Generation $t - 1$ is the old generation and generation t is the young. For example, in period 3 generation 2 is old and generation 3 is young.

Introducing social security or even expanding an existing social security system is beneficial only for the initial old and middle-aged generations.

While they are young, people work, receive labor income, and pay a payroll tax to finance the government-sponsored retirement scheme for the old generation. When they are old, people cannot work but must finance their consumption with the government pension and private savings. In the notation used throughout the article, subscripts denote time and superscripts denote the generation's birth period. For example, c_t^t, c_{t+1}^t is consumption of a generation t agent in time periods t and $t + 1$.[4] The budget constraint of this agent in period t takes the form

$$(1) \quad c_t^t + k_{t+1}^t = (1 - \tau_t)w_t.$$

The terms on the left-hand side are the expenses of an agent. He can dedicate his income to either consumption, c_t^t, or private savings, k_{t+1}^t, that pay off principal and interest next period. The right-hand side is net labor income, w_t, after paying the proportional

1. It is important to note that the U.S. social security system does not fit the definition of a pay-as-you-go plan in the narrow sense. The current surpluses due to the baby-boomer generation are saved in the Social Security Trust Fund rather than being used to lower contribution rates.
2. The Social Security Trust Fund is actually two funds—the Old Age and Survivors Insurance Trust Fund and the Disability Trust Fund. This article will refer to these two funds as a single fund.
3. Biggs (2000) has the opposite view. He addresses uncertainties about the estimates of a wide range of economic fundamentals underlying the Social Security Administration's computations, such as fertility, longevity, and productivity, and concludes that the estimates tend to be at most reasonable and sometimes on the optimistic side from a historical point of view.
4. This model assumes that agents in each cohort are identical, so there is no within-generation heterogeneity of wealth.

FIGURE 1

Timing in the Overlapping Generations (OLG) Model

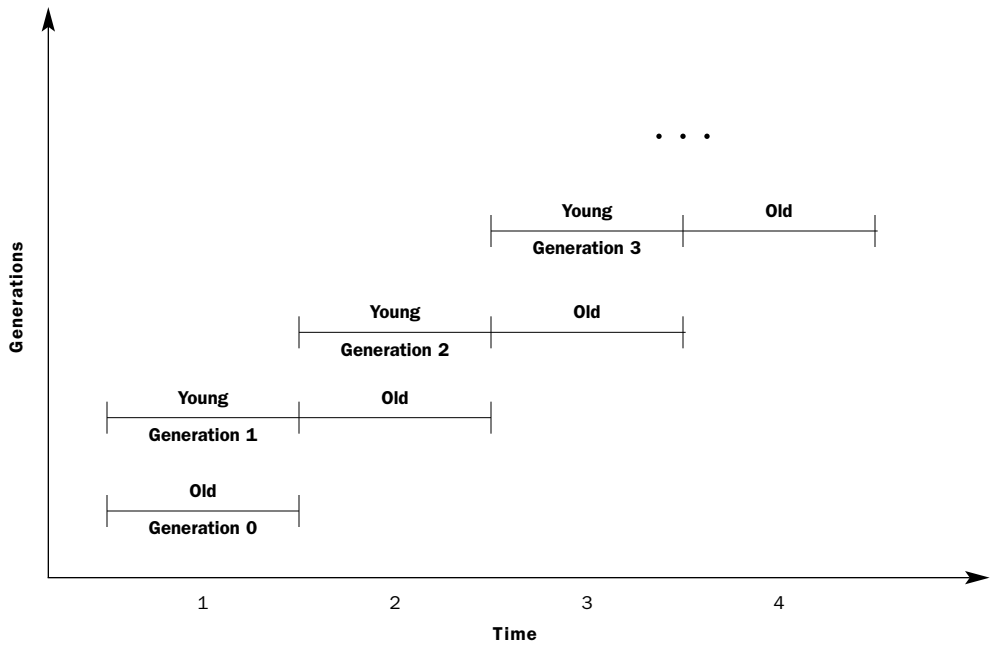

payroll tax, τ_t. After working for one period, the agent retires in period $t + 1$ and has two sources of income:

$$(2) \quad c^t_{t+1} = \delta w_t + (1 + r_{t+1})k^t_{t+1}.$$

The first term on the right-hand side is the government-sponsored retirement scheme, promising to pay a fraction δ of an individual's wage as a social security benefit. This social security system is a defined-benefits system because it promises a fixed replacement ratio of δ of the previously earned wage. The second term is private savings plus the earned interest. Agents take wages and interest rates as given and maximize the objective lifetime utility function $U = u(c^t_t) + \beta u(c^t_{t+1})$ subject to the two budget constraints (1) and (2).

On the aggregate level there is both population and productivity growth. It is easy to see that both growth rates play a vital role for the feasibility of social security. The higher the population growth rate, the lower the retiree-to-worker ratio, and the higher the productivity growth rate, the higher the wage growth rate and thus the easier it is to finance promised benefits out of the payroll tax base. Let g_λ denote the population growth rate, which is assumed to be fixed for now, and λ_t, the size of the generation born in period t. Then, by definition, $\lambda_{t+1} = (1+g_\lambda)\lambda_t$. Output ($Y_t$) is produced by combining labor (L_t) and capital (K_t) according to a production function $Y_t = A_t K_t^\alpha L_t^{1-\alpha}$, where A_t is total factor productivity. Market clearing dictates that the amount of labor used in production is equal to the amount of labor from the young generation, and the amount of capital is equal to the amount of savings the old generation accumulated during its working years. That is, output is equal to $Y_t = A_t(\lambda_{t-1}k_t^{t-1})^\alpha \lambda_t^{1-\alpha}$, where productivity, A_t, grows at rate g_A; that is, $A_{t+1} = (1 + g_A)A_t$.

In the model it is assumed that the government balances its budget period by period, and in order to do so it sets the payroll tax so that the tax revenue is exactly equal to the payments to retirees. (Later in the article, when aggregate shocks are considered, this assumption implies that the tax rate adjusts in order to finance the predetermined benefits to the retirees.[5]) The government budget constraint then implies that the payroll tax revenue, $\lambda_t \tau_t w_t$, is equal to the promised benefits to the old generation, $\lambda_{t-1}\delta w_{t-1}$;[6] that is, the equilibrium payroll tax is equal to $\tau_t = \delta(\lambda_{t-1}/\lambda_t)(w_{t-1}/w_t)$. The intuition that supports this result is that with no population or wage growth, benefits are equal to contributions. If there is wage or population growth, then a fixed benefit level can be achieved with lower contributions because the tax base is increasing.

Alternatively, one could assume that contributions into the pension system are constant and that the government distributes the proceeds to the

current retirees. In this case, if there are fluctuations in wages, payroll taxes as a percentage of wages would remain constant and benefits would adjust—the classic example of a defined-contributions pension system. However, benefit cuts as a reaction to aggregate fluctuations are considered politically difficult to implement. Hence, it is more realistic to use the defined-benefits scheme to model social security in the United States because it is a closer approximation to reality.

There are two main simplifying assumptions in this model. First, labor supply is exogenous, that is, agents' provision of labor does not depend on the wage or the payroll tax. In a more sophisticated model where agents decide how much labor to supply, social security financed by payroll taxes has a distortionary effect because people would work less. Ruling out this effect makes a PAYGO system more attractive than in reality, and the gains from privatizing social security would be underestimated in the model relative to the real world.

The second assumption is that the private savings have to be invested in the domestic economy; that is, no international portfolio diversification is possible. Therefore, the effect on financial asset returns is likely to be overstated in the model, again making social security more attractive than in a more realistic, yet currently computationally intractable, model. This article, therefore, makes the best possible case in favor of social security, and the gains from privatization it depicts are likely a lower bound on the actual gains in a more realistic framework.

Balanced Growth Paths

When numerical simulations are performed, initial conditions in the model matter. In particular, the economy starts with an old generation that owns capital, and the researcher therefore has a choice about what the capital endowment of this initial old generation should be. The preferred choice in economics is to begin in a long-run equilibrium and see how a shock affects the economy, in particular what the path back to the long-run equilibrium looks like. Without population and productivity growth, this long-run equilibrium, called steady state, is an equilibrium in which all model variables stay constant over time. Since there is growth in both population and productivity, however, there is evidently no such steady state in this economy. Instead, there is a long-run equilibrium in which the growth rates rather than the levels stay constant for all variables (but may differ across variables). This long-run equilibrium is also called a balanced growth path. Here, this balanced growth path can be computed easily by noting that the tax, τ, must remain constant and the variables in equation (1)—namely, consumption when young, savings, and wages—all must grow exponentially at the same rate. This result, together with the fact that the wage is simply the marginal product of labor, implies that

Social security reform has beneficial effects in the long run, but in the short run a large portion of the population will be worse off.

$$(3)\quad 1+g_k = 1+g_w = (1+g_A)[(1+g_\lambda)(1+g_k)]^{\alpha}(1+g_\lambda)^{-\alpha},$$

and therefore growth rates of wages and per capita savings are given by

$$(4)\quad 1+g_k = 1+g_w = (1+g_A)^{1/(1-\alpha)}.$$

Equation (4) means that independent of the population growth rate, both wages and savings grow at identical rates equal to approximately $1/(1-\alpha)$ times the productivity growth rate.

5. There are three reasons why this article abstracts from the possibility of a trust fund and instead considers a PAYGO system in which the government adjusts payroll contributions period by period. First, a trust fund is only an option if temporary surpluses are saved to finance future deficits. In two of the three aggregate negative shocks considered here, there are no temporary surpluses and hence no room for a trust fund. Second, in the one case that does display temporary surpluses, namely, the baby-boomer shock, it turns out that the negative welfare effects become even more pronounced with a trust fund; thus, in order to give social security the best possible chance to do well compared to privatized social security, this article assumes there is no trust fund. Third, the bulk of the actual trust fund savings in the United States was accumulated after 1987, well after the first baby boomers entered the labor force, so that the trust fund can all but partially smooth out the future deficits. In other words, because the trust fund build-up occurred so late, the actual U.S. social security system is not far from being a pure PAYGO system.
6. Notice that the promised benefits depend on the lagged wage the current retirees earned when they were young.

FIGURE 2

Increasing Benefits: Average Lifetime Consumption Relative to No Policy Change

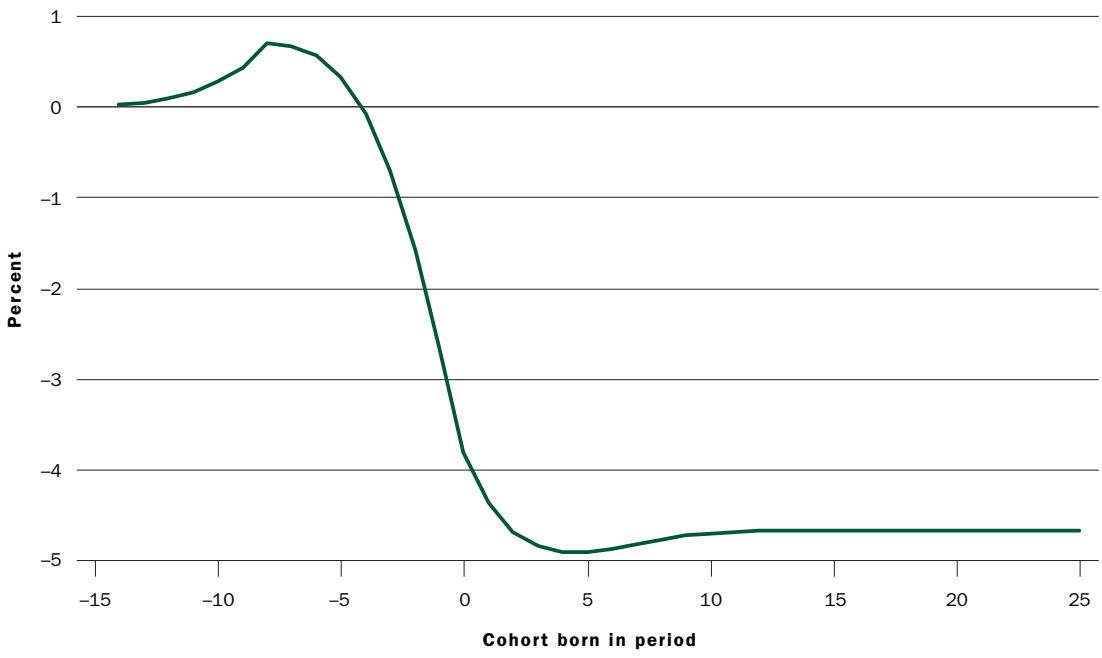

The Facts about Social Security

Equipped with the tools from the OLG model, the discussion now turns to some of the facts about pension systems that are well known in the academic literature but possibly less well understood in the popular media.

The internal rate of return. Without aggregate fluctuations, the internal rate of return of defined benefit systems for a representative generation is equal to the rate of population growth plus the growth rate of real wages. This result, from Samuelson (1958), can be easily derived with the two-period OLG model. According to the calculations from above, an agent's flow of payments into the social security system is $\delta w_t/[(1 + g_\lambda)(1 + g_w)]$ when young and $-\delta w_t$ when old. Consequently, independent of the value for the replacement ratio, δ, the internal rate of return on this flow is $(1 + g_\lambda)(1 + g_w) - 1$, which is approximately equal to the sum of the growth rates of population and wages. The intuition for this result is that the higher the growth rates for wages and population, and thus the higher the growth rate of the payroll tax base, the easier it is to finance social security. For every dollar of old-age benefits, only $1/[(1 + g_\lambda)(1 + g_w)]$ dollars of contributions have to be paid.

This result holds even if there is a positive probability of death before reaching old age. Suppose a fraction π of agents die after their first period of life so they naturally do not receive old-age benefits. In this case the expected benefits of a representative generation are reduced by π, and at the same time the contributions are cut by the same fraction because fewer old people receive old-age benefits, leaving the internal rate of return unchanged. The results from this stylized economy with only a two-period horizon for every generation carry over to more realistic settings in which agents live longer than two periods.

Financial market returns. In the real world, the potential average yield for social security is equal to population plus wage growth, an amount that is smaller than the average return on private savings. This observation spurred the entire privatization debate because an individual worker could invest her social security contributions in the financial market and, in expected terms, achieve a higher level of benefits. For example, for the period 1960 to 2001, real wage plus population growth averaged less than 3 percent per year, and it is expected to average less than 2 percent per year for the next three decades, according to the Social Security Administration's estimates. Individuals born in 2000 can even expect less than 1 percent return on their social security contributions. On the other hand, the long-term real return on financial assets is significantly higher—in the neighborhood of 5.5 percent annually, according to Feldstein and Ranguelova (2001), or 4.6 percent return from a diversified port-

FIGURE 3

Increasing Benefits in Period 0: Changes Relative to No Reform

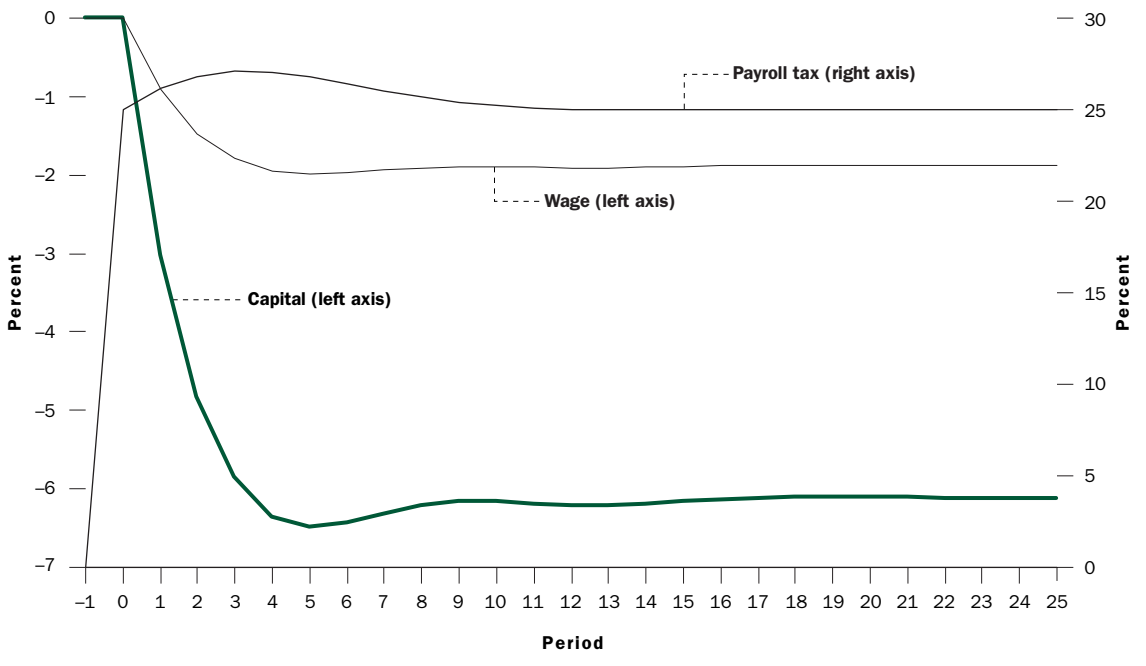

folio considered by the President's Commission to Strengthen Social Security.

Introducing or enhancing a social security system. Introducing social security or even expanding an existing social security system is very beneficial for the initial old and middle-aged generations, but the young generation at the time of the program's introduction and all generations born in the future are worse off.[7] In particular, social security crowds out private savings and therefore leads to a reduction in the level of capital. To demonstrate this result, a more realistic version of the model above is used in which generations live for up to fifteen periods. In the simulation it is assumed that in period 0 pensions increase by 25 percent; that is, the benefit formula is adjusted unexpectedly to increase the ratio of pension benefits to lifetime earnings by one quarter.[8] Notice that this policy affects not only generations born in period 0 and thereafter but also all other current workers, namely, generations –7 to –1, as well as current retirees in generations –14 to –8.

The results of the simulation are demonstrated in Figure 2, which plots the average lifetime consumption of all cohorts affected by the policy change.[9] The ten cohorts that entered the labor force between periods –14 and –5 post a gain in average lifetime consumption. These are the seven cohorts currently retired who get higher benefits without ever paying higher contributions and the three cohorts of workers prior to retirement who profit from higher retirement benefits but must pay higher contributions only for a very limited amount of time. All other cohorts, whether currently alive or born in the future, suffer substantial losses of lifetime consumption on the order of about 4.5 percent.

The reason for this result lies in the response of macroeconomic variables to the policy change. Figure 3 plots aggregate capital, wages, and the

7. The scope of this article is social security reform, exactly the opposite of introducing or enhancing social security, but for completeness this well-known result is included.
8. In this version of the model, since people work for a maximum of eight periods, the calculation of benefits is more complicated than in the two-period OLG model, where the benefit formula consisted of only one replacement ratio, δ. Specifically, the more complex model must define how retirement benefits depend on a total of eight past wages and how they develop during the maximum of seven periods in retirement. It is assumed that retirement benefits stay constant (in real terms) during retirement and benefits are a share of the average lifetime earnings.
9. Average consumption in this context is defined as follows. Suppose an individual born in period t consumes a sequence $(c_t^t,\ldots,c_T^t)$; then average consumption $c_{average}^t$ is defined as the constant consumption profile that makes the individual indifferent between the actual and the constant profile. In the two-period example, $u(c_{average}^t) + \beta u(c_{average}^t) = (1 + \beta)u(c_{average}^t) = u(c_t^t) + \beta u(c_{t+1}^t)$.

payroll tax relative to an economy without policy change over time. A higher level of social security triggers a drop in aggregate capital relative to an economy with unchanged social security benefits. The reason for this drop is that the aggregate capital stock is equal to the sum of all savings, and with higher retirement benefits, private savings become both less desirable as people rely more on social security and less affordable as young and middle-aged workers receive smaller net wages because of higher payroll taxes. With the drop in capital comes a drop in wages (making private savings even less affordable), which in turn causes the payroll tax to initially overshoot to more than 25 percent above the initial level, because now the promised benefits to older cohorts must be financed out of a smaller payroll tax base. It is interesting to note that, in the long run, lifetime consumption drops by about 4.5 percent—much more than the 2 percent wage drop. More than half the drop in lifetime consumption is due to the low internal rate of return on the increased payroll tax contributions.

The longer ago privatization took place, the more likely it is that all cohorts alive will be better off under privatized social security. There may be arguments against privatization, but aggregate risk is not one of them.

Theoretically, there is one positive effect for all generations coming from the introduction or expansion of social security. If there are no markets for annuities, private savings have a disadvantage that, with a random lifespan, there is the risk of outliving one's savings. This risk can be eliminated by social security. Since the expansion of social security in this economy substantially reduces welfare in the long run, the risk-sharing effect from social security must be small, however. This effect is in line with the results from Storesletten, Telmer, and Yaron (1999), who show that the negative effect from crowding out private savings is indeed larger than the positive effect from risk sharing.

The effects of social security reform. Social security reform has beneficial effects in the long run, but in the short run a large portion of the population will be worse off. Naturally, a reform can take many forms, but ultimately all reforms involve lower benefits. This article assumes that a reform takes the form of an immediate one-time reduction of benefits. Alternatively, one could assume a delayed or a stepwise reduction, in which case the initial burden of the transition would be spread over more generations. The long-term effects, however, would be unchanged. The result of initial burden and long-term gains is the precise flip side of the previous result as shown in Figure 4. Cutting benefits hurts the currently retired cohorts (–15 through –9) and even some of the older cohorts currently working but close to retirement. In the long run, though, agents are far better off—by about 4.5 percent of average consumption—because lower retirement benefits in conjunction with a lower tax burden encourage more private savings and therefore more capital accumulation. With more productive capital in the economy, wages are higher, and the increase in lifetime income, coupled with the higher returns on private savings compared to the pension system, makes future generations substantially better off.[10] If retirement benefits are cut, however, all retirees and even workers who are near retirement have to suffer substantial losses. Even a delayed reform in which benefits will be lower in the future penalizes those generations that have to pay contributions into the social security system but receive only reduced benefits when they retire.

The fact that private accounts yield higher returns than social security is therefore often misunderstood as a miraculous way of saving the system by offering higher returns from private accounts. It is often ignored that, if such privatization took place, current retirees' benefits—being a burden on current and future tax payers—would have to be cut or would have to be financed through the general tax base or government debt. In other words, the problem of privatization is the unfunded liability to pay for current retirees if current workers start investing in their private accounts.

Related Literature

The model presented here is neither new nor the only one that tries to address the issue of aggregate fluctuations and social security. Auerbach and Kotlikoff (1987) set up a similar large-scale OLG model to address a wide array of interesting policy questions, including tax reform and social security reform. The model used here can be thought of as a simplified version of the Auerbach and Kotlikoff model that looks at a new set of experiments, namely, the performance of different pension systems if a large aggregate shock occurs.

De Nardi, Imrohoroglu, and Sargent (1999) do a numerical exercise to study the effects of the demo-

FIGURE 4

Partial Privatization through Decreasing Benefits: Average Lifetime Consumption Relative to No Reform

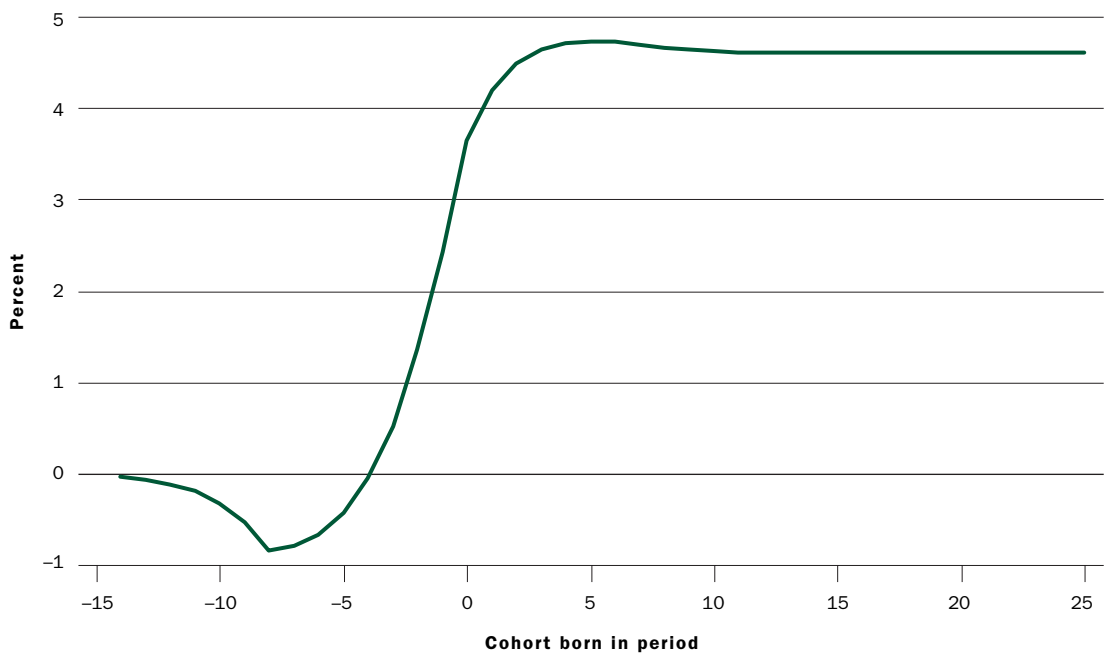

graphic shock on social security that will be caused by retirements among the baby-boomer generation during the coming decades. As one of the main points of the paper, the authors show that, without reform of the public pension system, contributions would have to increase substantially. Higher taxes, however, magnify the distortions in the economy to such a degree that the growth forecasts used by the Social Security Administration seem too optimistic from the perspective of a general equilibrium model because social security taxes discourage agents from both working and saving. In other words, De Nardi, Imrohoroglu, and Sargent quantify a potential feedback effect from higher payroll taxes into the growth forecasts that has been ignored by the Social Security Administration and show that this effect can be substantial.

As mentioned above and demonstrated by Samuelson (1958), introducing a social security system makes a few initial cohorts better off at the cost of making all other future generations worse off. Krueger and Kubler (2002) use an OLG model to test whether this result still holds if aggregate macroeconomic shocks occur—for example, shocks to productivity. Krueger and Kubler point out that it is theoretically possible that all generations may be better off with the introduction of social security, contradicting conventional wisdom. The main ingredients in their model are the assumptions that returns to labor and physical capital are imperfectly correlated and that the social security system is designed as a defined-contribution system. The latter assumption means that the contributions are fixed and benefits vary over time because they are a fixed proportion of the payroll tax base. This setup is quite different from the assumption used in this paper, in which taxes adjust to finance promised benefits, and from the current U.S. social security setup. The welfare-improving

10. It is important to note that, theoretically, it is possible for OLG models to display so-called dynamically inefficient equilibria, as first pointed out by Diamond (1965). The inefficiency involves overaccumulation of savings. In such a case, reducing social security may actually reduce welfare even in the long run because it involves even more overaccumulation of capital. For this reason, Imrohoroglu, Imrohoroglu, and Joines (1995) find in their economy, which displays this inefficiency, social security replacement ratios should optimally be above zero. In the real world this inefficiency, however, appears to be nonexistent because it would involve observing real interest rates on capital that lie below real GDP growth rates, which do not occur in the United States.

FIGURE 5

Baby Boomers in Periods 3–5: Deviation in Average Consumption Relative to No Shock

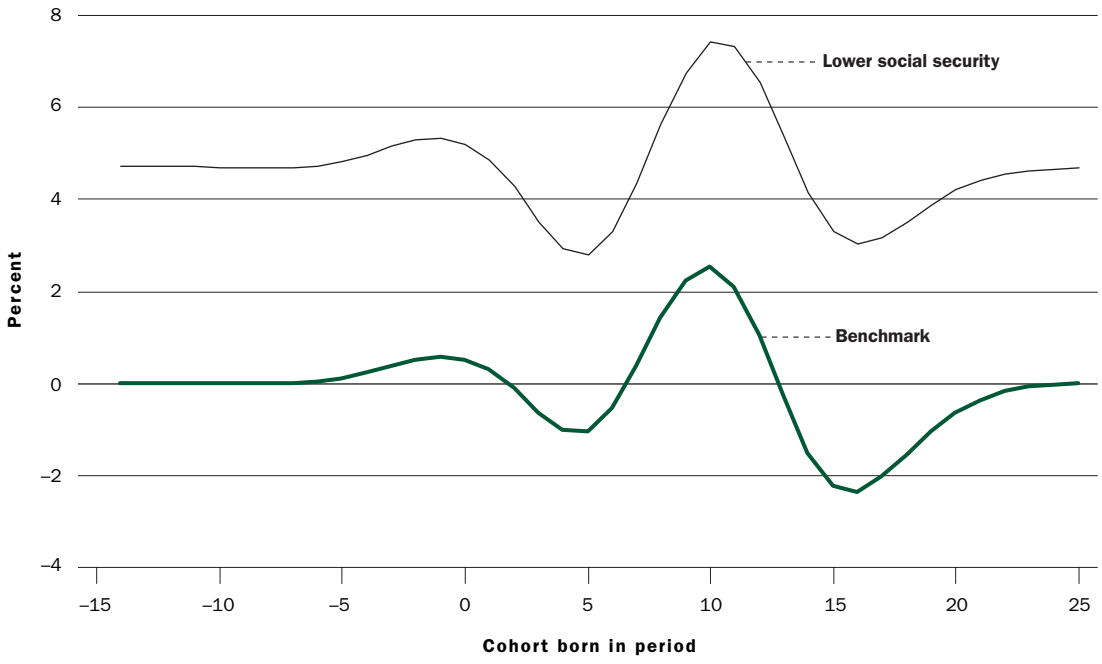

property of social security in Krueger and Kubler's environment stems from the fact that retirees put a value on the asset called labor because the returns of labor and capital are not perfectly correlated and, therefore, diversification gains are possible. Retirees, however, have no labor endowment left, so giving them a claim to labor income through the social security system improves their welfare.[11] Still, there is the well-known negative effect of social security crowding out private savings both through removing incentives on private savings and the tax distortion, but the former effect theoretically could be larger than the latter. Krueger and Kubler do, however, point out that in an economy modeled to match the U.S. economy the negative effect of crowding out savings dominates the risk diversification.

Computations

Using the more complex version of the model, with fifteen generations alive at any given period and having realistic survival probabilities, this section determines the impact of three different aggregate shocks to the welfare of all current and future generations. Since now one period is only six years, the groups entering the labor market each period will be labeled cohorts rather than generations.

Two sets of experiments are conducted, each of which examines the following three aggregate shocks:

1. A baby-boomer generation—three cohorts that are larger relative to both their parents and children. In the computations using the fifteen-period OLG model, it is assumed that in period 0 it becomes apparent that the cohorts entering the labor force in periods 3–5 are 40 percent larger than in the benchmark economy. The three-period gap (equal to eighteen years) accounts for the years between the birth of a cohort and the time the cohort actually enters the labor market.
2. A permanent drop in the productivity growth rate by one-half starting in period 0, causing the long-run wage growth rate to drop from 1.1 percent—the benchmark growth rate used by the Social Security Administration—to only 0.55 percent.
3. A permanent drop in the labor force growth rate from 0.2 percent per year—the benchmark labor force growth rate used by the Social Security Administration—to –0.3 percent per year, which is the pessimistic demographic scenario used in the calculations of the Social Security Administration. Again, it is assumed, for the same reason as before, that the change takes effect in period 3.

The first set of experiments looks at the three alternative shocks' impact on cohorts' welfare under different payroll tax levels when the economies start down their respective balanced growth paths.

FIGURE 6

Effect of Baby Boomers on Factor Prices and Payroll Taxes

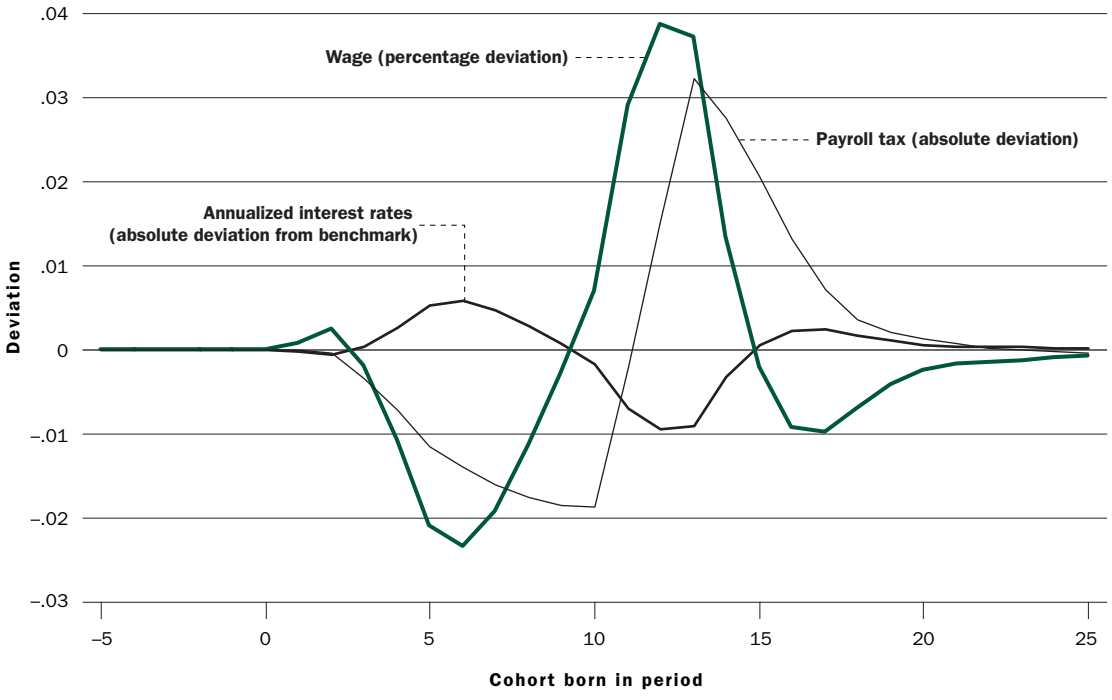

In particular, this analysis looks at one benchmark economy with a benefits structure generating a payroll tax of 12.5 percent, about the same as the payroll tax in effect in the United States, and another economy with both benefits and contributions reduced to three-quarters of the benchmark economy. This second economy is on its long-run balanced growth path and consequently has a higher capital stock than the benchmark economy. One could think of this economy as one that partially privatized social security many periods before and is now on its balanced growth path with lower social security benefits and higher private savings.

The results are plotted in Figures 5–9. In both economies the shocks have an impact on welfare, mostly negative, but people in the economy with the lower payroll tax are uniformly better off; that is, if asked in which economy it would rather live, every cohort would prefer the one with lower payroll taxes.

Specifically, Figure 5 plots the average consumption for both the benchmark economy and the lower-social-security economy relative to the benchmark economy without a shock. In the economy with lower social security coverage, every cohort is better off than its respective cohort in the benchmark economy. Cohort-by-cohort lifetime consumption could lie between 4 percent and 5 percent higher than in the benchmark economy.

Quite interestingly, in the benchmark economy not all cohorts are worse off because of the arrival of the baby boomers. One could say that the parents of the baby boomers gain about 0.5 percent of average lifetime consumption; the baby boomers themselves lose 1 percent, the children of the boomers gain substantially, more than 2 percent, and the biggest losers are the grandchildren of the baby boomers, who lose more than 2 percent of average lifetime consumption. The intuition for this result comes from the general equilibrium structure of the model, in which factor prices (namely, wages and interest rates) and payroll taxes must adjust to macroeconomic shocks.

Figure 6 plots the response of factor prices and payroll taxes to the arrival of the baby-boomer generation. The baby boomers increase the amount of labor input available in the economy starting in period 3. Therefore, the parents of the boomers benefit because their savings yield higher interest as the amount of labor increases. According to the

11. This outcome, of course, raises the question of why a government has to get into the business of providing risk diversification. If there really is a diversification gain, a private market could do the same job without causing the distortions of private savings.

FIGURE 7

Lower Productivity Growth: Deviation in Average Consumption Relative to No Shock

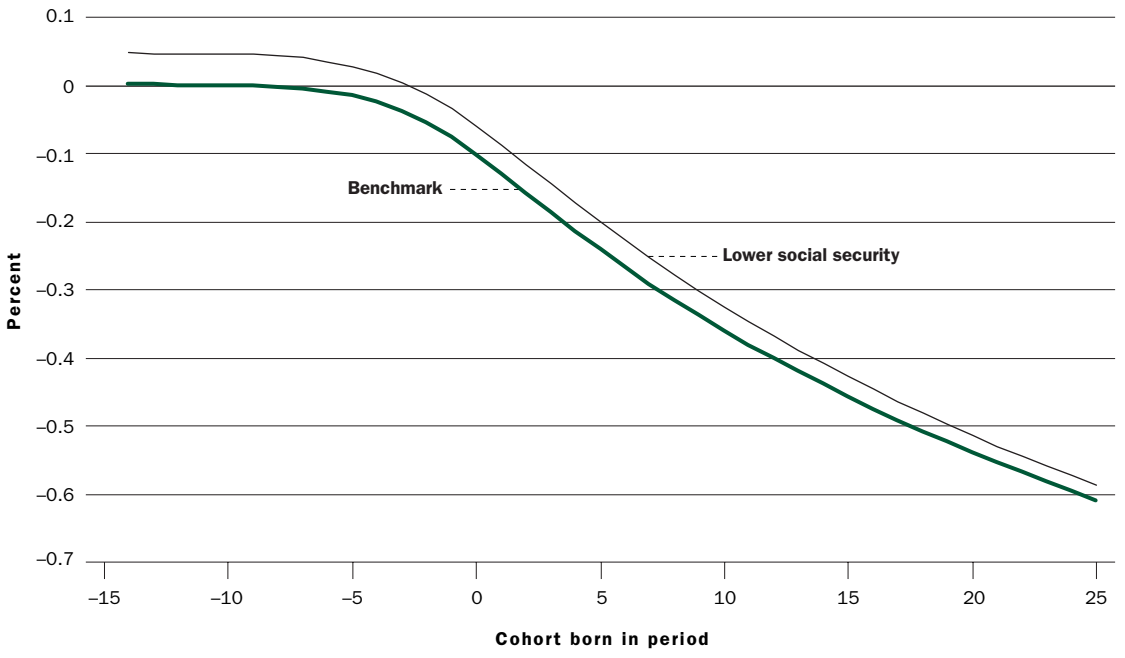

simulations, interest rates are higher in this economy relative to the no-shock economy until period 9, during which the parents of the baby boomers accumulate the bulk of their savings, and the first couple of periods of their retirement. The baby boomers themselves suffer because the larger supply of labor drives down wages until period 9—for almost as long as cohorts 3–5 work. At the same time, interest rates are lower from period 10 on, the start of the baby boomers' withdrawal phase.

The children of the baby boomers gain substantially because they benefit from two developments. First, when they enter the labor force in periods 8–10, the boomers themselves are still working, driving down the payroll tax until period 11 and thus helping to alleviate the future payroll tax hikes. The baby boomers also drive up the aggregate capital stock through their savings, so when they retire they substantially drive up the wages of their children during periods 10–15, covering most of the working years of generations 8–10 and therefore causing large gains for these cohorts. This result is consistent with Bohn's (1999) for the same reason. The relatively small generation of the baby boomers' children posts a net gain because the factor price effect is larger than the fiscal effect coming from higher payroll taxes. By the time the grandchildren enter the labor force, this positive wage effect has reversed, and payroll taxes are still high, driving down their average consumption.

This result is intriguing because conventional wisdom suggests that the post-baby-boomer generations are worse off because they have to pay high payroll taxes to finance the boomers' retirement. The children of the baby boomers, however, post a net gain because their higher wages make up for the payroll tax hike. The large losers are the baby boomers and their grandchildren if general equilibrium effects are taken into account. Notice also that these results are true in a pure PAYGO system, in which contribution rates adjust period by period. If the government were to start a trust fund at the arrival of the baby-boomer generation, the path of the payroll tax would be smoothed out. This fund would even magnify the welfare effects, in particular for the baby boomers and their children. With a PAYGO system the lower payroll tax during the working years of the baby boomers alleviated the factor price effect. A trust fund would eliminate this alleviating effect, causing the baby boomers to be hit even harder. The children of the baby boomers, on the other hand, who already benefit from the movement in factor prices, would get an additional advantage from the trust fund because part of the burden of financing the baby boomers' retirement would now be financed by the baby boomers themselves.

FIGURE 8

Lower Population Growth: Deviation in Average Consumption Relative to No Shock

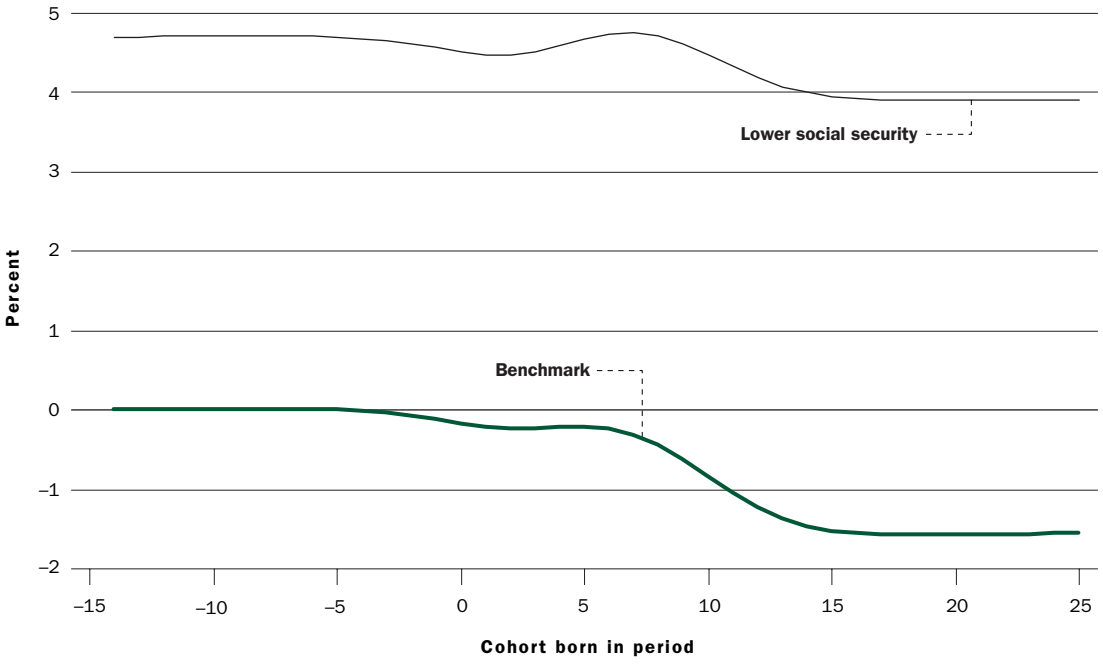

Figure 7 plots the effect of lower productivity growth on cohorts' welfare. All cohorts are uniformly better off in the economy with lower social security. All cohorts are negatively affected by the productivity slowdown, but agents in the economy with lower payroll taxes and higher private savings are shielded better from the adverse effects of the shock.

Finally, in the economy with lower population growth plotted in Figure 8, again every cohort is better off living in an economy with less social security. In addition, the long-run effect in the benchmark economy is more pronounced than in the economy with less social security; cohorts 15–25 lose about 1.6 percent of average consumption in the benchmark economy and only 0.8 percent in the lower-social-security economy. This result comes as no surprise because the long-run rate of return of social security is reduced by exactly the drop in the population growth rate (as demonstrated earlier) while the rate of return on private savings is not affected as much in the long run.[12] Hence, people living in an economy with higher social security levels get penalized more by the adverse demographic shock because their old-age income depends more on social security, an asset whose return is more negatively affected by the demographic shock than private savings are.

Another interesting dimension is the internal rate of return on both private savings and social security in the event of a shock. Figure 9 plots those internal rates of return in the benchmark economy when productivity growth drops in period 0. As expected, the drop in internal rates of returns for the first couple of cohorts is higher for private savings than it is for the social security system. It seems that indeed the social security system provides a better safety net from aggregate shocks than private savings do. However, this result is deceiving because in the long run the drop in the social security yield is higher, about 50 basis points, whereas the yield for private savings recovers for later cohorts and is only about 30 basis points below its original level. In other words, the relative safety of social security for the first couple of cohorts comes at the price of future generations losing a much larger share of their social security yield.

Policy Issues

What is the policy relevance of these results? Evidently, the U.S. economy could not jump

12. The reason is that lower population growth indeed lowers the capital returns because fewer workers are alive, but a large part of the decline is offset by the equilibrium effect of less capital accumulation coming from smaller cohort sizes.

immediately from the regime with high payroll taxes to the one with low payroll taxes and high private savings because a large investment would be necessary to build up the higher capital stock of the latter economy. The analysis shows that, after a partial privatization, in the long run people are better shielded against aggregate shocks. Less social security indeed makes old-age income more sensitive to an aggregate shock, but, since the aggregate capital level is so much higher in the low-payroll-tax economy, this effect is sufficiently cushioned that people are better off with more private savings. Put differently, a concern about a large aggregate shock many years from today would be a reason in favor of privatization, not an objection against it.

Consequently, aggregate risk can be an issue in the privatization debate only over the short horizon right after privatization occurred, that is, before the economy reached its new balanced growth path with higher capital levels. This is the only point at which social security has any chance to beat private accounts on welfare grounds—right after the privatization, when the aggregate capital stock is still low relative to its long-term path and retirees are exposed to the aggregate shock to a higher degree if their benefits are lower.

In the second set of experiments the economy is shocked early on during the privatization, before the new balanced growth path is reached. The three shocks outlined earlier are introduced in each of three new economies, in which the social security reform in the form of a 25 percent reduction of benefits is introduced in periods 0, –5, and –10, respectively. The aggregate shock thus occurs at various stages of the reform, namely, at the same time as the reform, five periods after the reform, and ten periods after the reform, respectively. The welfare effects of the reform combined with the shock are presented in Figures 10–12. The format of these figures is different from that in Figures 5, 7, and 8, which plot the percentage deviation from a no-shock, no-reform economy for the benchmark economy and an economy with lower social security, each of which experienced a shock. Figures 10–12 plot the percentage deviation in the three reform economies with a shock relative to a benchmark economy that also experienced a shock. One could view Figures 10–12 as Figure 4, where an aggregate shock occurred at various stages of the privatization process; that is, a negative number indicates that a cohort is worse off compared to no reform given that a shock occurred, and, vice versa, a positive

FIGURE 9

Internal Rates of Return after a Drop in Productivity Growth in Period 0

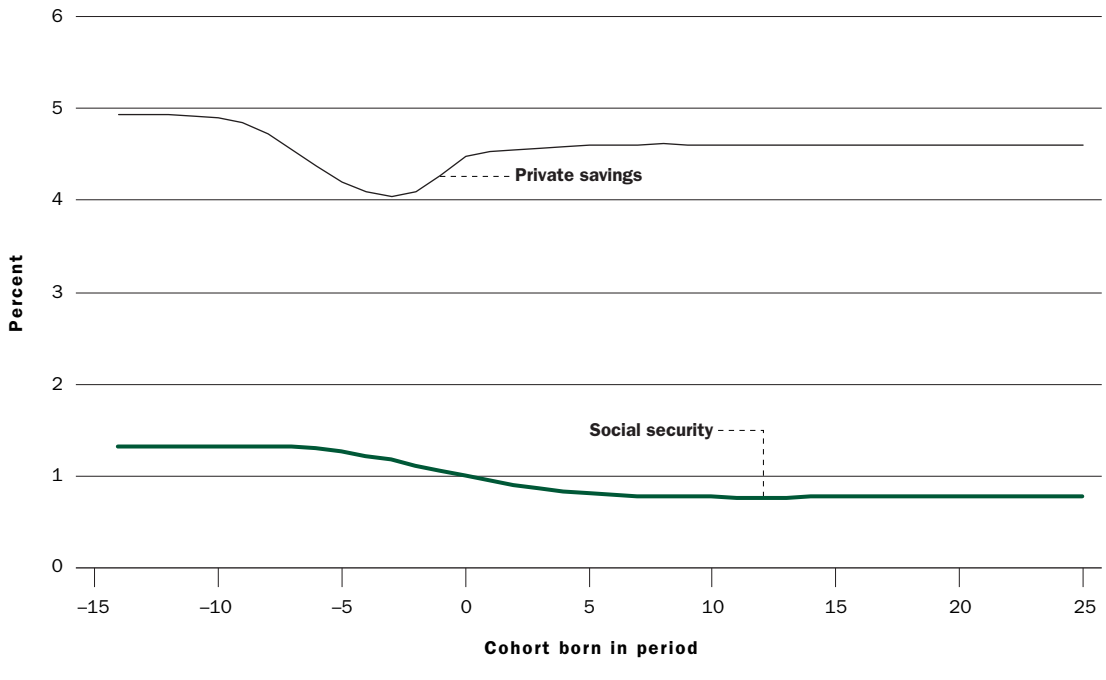

FIGURE 10

Baby Boomers in Periods 3–5: Gain of Partial Privatization

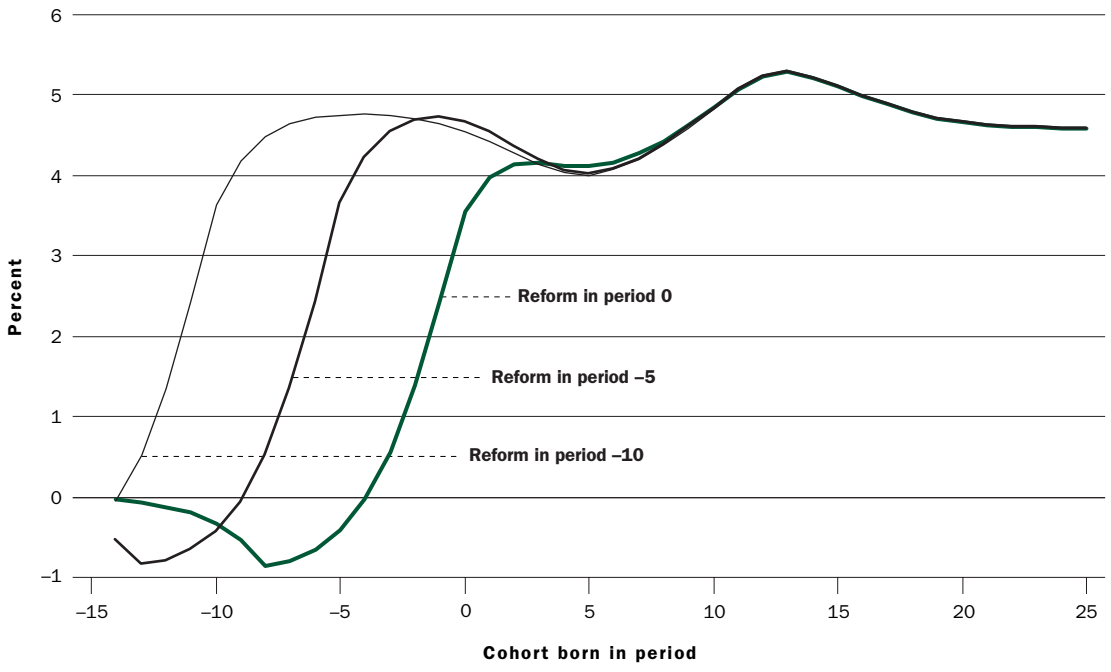

FIGURE 11

Lower Productivity Growth Starting in Period 0: Gain of Partial Privatization

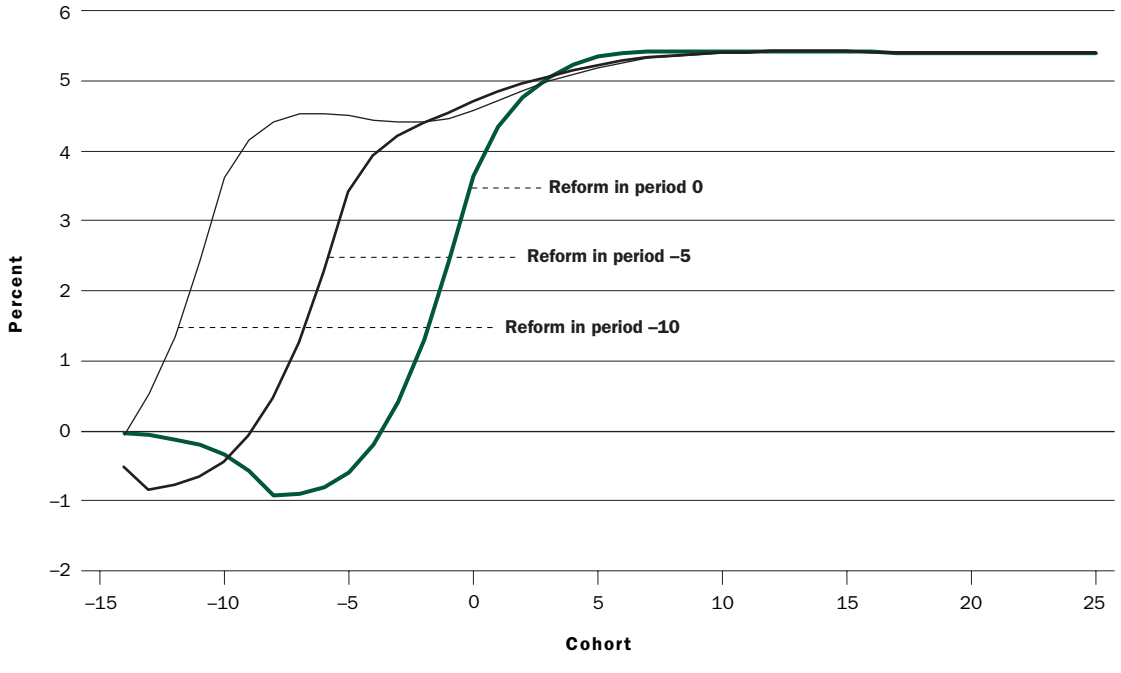

FIGURE 12

Lower Population Growth Starting in Period 3: Gain of Partial Privatization

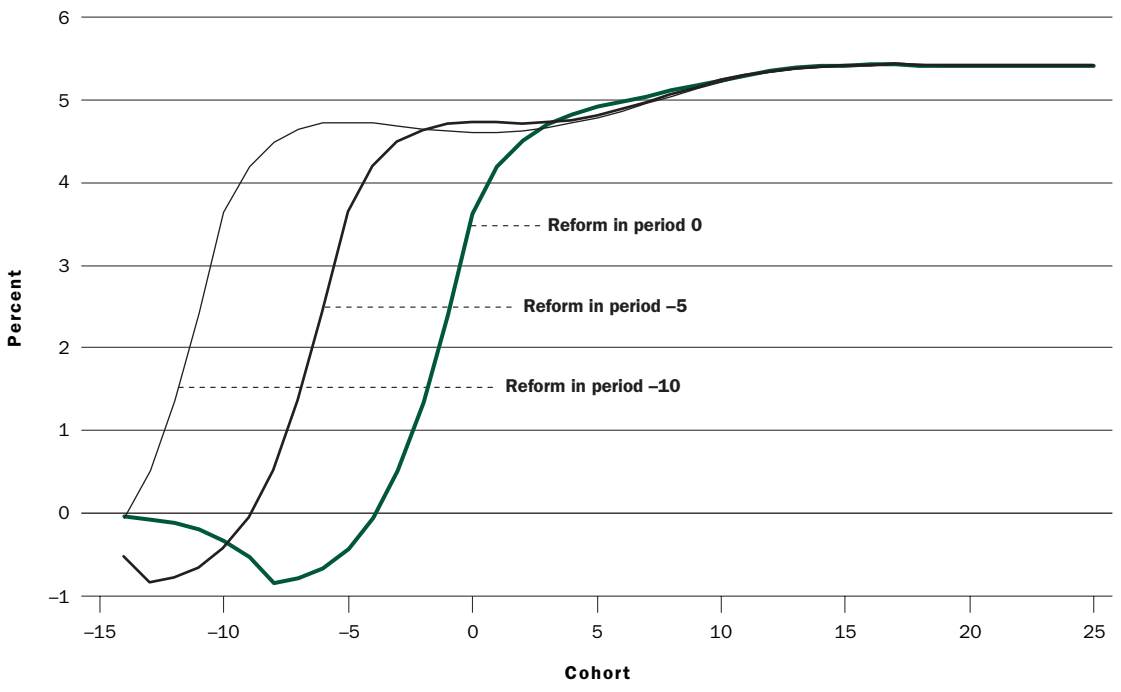

number indicates a cohort is better off with privatization relative to no reform.

The welfare effects in this set of experiments differ from those in the previous sets. Before, every single cohort was better off with lower social security independent of the shock. Now, there are in fact cohorts that are worse off—precisely the older cohorts in the scenarios in which privatization took place in the current period and five periods ago.

The interesting question is whether the older cohorts are worse off because of the reform or the shock. The effects on consumption in Figures 10–12 look strikingly similar to those in Figure 4; if the reform takes place in period 0, in each case the welfare effects on cohorts –14 to 0 are very similar to the ones in Figure 4. If the reform occured five or ten periods ago, then the curves in Figures 10–12 are close to being shifted by five and ten periods, respectively. Consequently, the welfare effects seem to be due to the reform, not the shock. In other words, once privatization is agreed upon, with its consequences on welfare to the initial old cohorts, adding a shock to the economy will reduce their welfare even further, but, quantitatively, this effect is not much larger than it would have been without the privatization. The benefit of the reform is that future generations, who are normally hardest hit by the aggregate shock, benefit a great deal from the privatization.

Conclusion

Implementing a social security system seems like increasing social welfare out of nothing: One generation of initial old agents never has to pay contributions, but they receive benefits. What looks like the modern version of an alchemist's dream of turning lead into gold is in fact a rather costly endeavor for all future generations because the internal rate of return on their contributions and benefits is almost certainly lower than that of financial assets. There is no free lunch after all! Nevertheless, low returns do not necessarily imply that an asset is unattractive. In a standard Sharpe (1964) and Lintner (1965) CAPM model, the attractiveness of an asset is determined by both the expected return and the variance-covariance structure. People perceive a defined-benefits plan such as social security as a relatively riskless asset.

This article establishes a more cautious view on this matter. Social security is risky as well, and, even worse, it is risky for exactly the same reasons that financial assets are. The determinants of implicit social security internal returns are productivity and labor force growth, precisely the same as those for long-term capital market returns. It is certainly true that in the event of a negative shock, say, to productivity growth, the internal rates of return on the private savings of living cohorts drop by more than the

return on social security contributions. In the long term, however, social security returns take a greater hit than returns on physical capital. In other words, the security in pension system returns is deceiving because it pertains only to current retirees, not to future ones. This seems to be a general feature of PAYGO systems: They look attractive to current cohorts, but future cohorts may suffer substantially.

If an economy with a low payroll tax reaches its balanced growth path, then all people in that economy are better off than the people in an economy with higher payroll taxes if an adverse shock hits the economy. Low payroll taxes encourage more private savings in the form of a larger aggregate capital stock, and this cushion of savings ensures that people are better off than they would otherwise be with a higher payroll tax. This result is true for all three scenarios computed here. If an aggregate shock occurs during the transition to lower benefits, certain generations of current retirees and people about to retire suffer a stronger loss than they otherwise would under an unchanged PAYGO system. However, the effect is rather small compared to the welfare loss coming from the reform.

All future generations, on the other hand, in particular those that normally get hit the hardest by aggregate shocks, benefit greatly from the reform. Their average consumption can be far higher than it would be in an economy without privatization if an adverse shock occurs. Moreover, the longer ago privatization took place, the more likely it is that all cohorts alive will be better off under privatized social security. There may be arguments against privatization, but aggregate risk is not one of them.

REFERENCES

Auerbach, Alan, and Lawrence Kotlikoff. 1987. *Dynamic fiscal policy*. Cambridge, New York, and Melbourne: Cambridge University Press.

Biggs, Andrew. 2000. Social security. Is it a crisis that doesn't exist? The Cato Project on Social Security Privatization. SSP No. 21, October.

Bohn, Henning. 1999. Social security and demographic uncertainty: The risk-sharing properties of alternative policies. National Bureau of Economic Research Working Paper 7030, March.

De Nardi, Mariacristina, Selahattin Imrohoroglu, and Thomas J. Sargent. 1999. Projected U.S. demographics and social security. *Review of Economic Dynamics* 2 (July): 575–615.

Diamond, P.A. 1965. National debt in a neoclassical growth model. *American Economic Review* 55 (December): 1126–50.

Feldstein, Martin, and Elena Ranguelova. 2001. Individual risk in an investment-based social security system. National Bureau of Economic Research Working Paper 8074, January.

Imrohoroglu, Ayse, Selahattin Imrohoroglu, and Douglas H. Joines. 1995. A life-cycle analysis of social security. *Economic Theory* 6 (June): 83–114.

Krueger, Dirk, and Felix Kubler. 2002. Intergenerational risk sharing via social security when financial markets are incomplete. Stanford University manuscript, December.

Lintner, John. 1965. The valuation of risky assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics* 47 (February): 13–37.

President's Commission to Strengthen Social Security. 2001. Strengthening social security and creating personal wealth for all Americans. Final report, December. <www.commtostrengthensocsec.gov/reports/final_report.pdf> (December 2002).

Reich, Robert B. 1998. The sham of saving social security. *The American Prospect Online*. <www.prospect.org/columns/reich/rr980600.html> (December 2002).

Samuelson, Paul A. 1958. An exact consumption-loan model of interest with or without the social contrivance of money. *Journal of Political Economy* 66 (December): 467–82.

Sharpe, William. 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance* 19 (September): 425–42.

Sinn, Hans-Werner. 1999. The crisis in Germany's pension insurance system and how it can be resolved. National Bureau of Economic Research Working Paper 7304, August.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron. 1999. The risk-sharing implications of alternative social security arrangements. *Carnegie-Rochester Conference Series on Public Policy* 50 (June): 213–59.