

**WOULD A PRIVATIZED SOCIAL SECURITY SYSTEM
REALLY PAY A HIGHER RATE OF RETURN?**

BY

**JOHN GEANAKOPLIS, OLIVIA S. MITCHELL
AND
STEPHEN P. ZELDES**

COWLES FOUNDATION PAPER NO. 1002



**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
AT YALE UNIVERSITY**

Box 208281

New Haven, Connecticut 06520-8281

2000

Would a Privatized Social Security System
Really Pay a Higher Rate of Return?
*John Geanakoplos, Olivia S. Mitchell, and
Stephen P. Zeldes*

AS THE U.S. Social Security system has matured, the rate of return received by participants has fallen. In the coming years, around the time the baby boom generation retires, the system will experience a budget shortfall. First, tax revenues will fall short of promised benefits, requiring spending the interest earnings, and then the principal, of the trust fund. Eventually, the trust fund will be depleted. This projected insolvency will necessitate benefit cuts or tax increases, leading to further declines in the rates of return individuals can expect.

Many advocates of reform suggest that an answer to this problem is to privatize Social Security. They argue that the creation of individual accounts invested in private capital markets, and especially in the stock market, will produce better rates of return for individuals than will Social Security. For example, Stephen Moore of the Cato Institute recently claimed that “privatization offers a much higher financial rate of return to young workers than the current system . . . if Congress were to allow a 25-year-old working woman today to invest her payroll tax contributions in private capital markets, her retirement benefit would be two to five times higher than what Social Security is offering.”¹⁴ Presidential candidate Steve Forbes criticized Social Security because “the average worker retiring today receives a lifetime return of only about 2.2 percent on the taxes he has paid into the system. Contrast this with

This paper draws on a longer paper by the authors (Geanakoplos, Mitchell, and Zeldes, forthcoming). Useful comments were provided by Jeff Brown, Martin Feldstein, and Kent Smetters.

14. Moore (1997)

the historic 9–10 percent annual returns from stock market investments. . . . The advantages of an IRA-type approach are overpowering.”¹⁵

Our goal in this paper is to challenge the following popular argument: projected returns to Social Security are low relative to expected returns on stocks and bonds, and therefore everyone would receive higher returns and be better off if the United States moved to a privatized system where individuals could directly invest their contributions in stocks and bonds. We argue that for households with access to diversified capital markets, privatization without prefunding would not increase Social Security returns, when properly measured. Privatization together with prefunding would eventually raise the rate of return to future generations of participants, but at the cost of a lower rate of return to early generations.

The real economic benefit to privatization is the diversification that would be made available to constrained households that cannot participate on their own in diversified capital market investment portfolios. The improved rate-of-return argument in favor of privatization thus has no force unless there are constrained households without access to diversified capital markets. However, this group of people is not generally recognized to be key to the popular argument just quoted. Indeed, in the presence of such constrained households, young market-savvy workers, who according to the popular argument should find privatization most appealing, will probably get lower returns as a result of privatization.

We begin by defining what privatization means and by distinguishing that concept from diversification and from prefunding. Next, we ask whether projected returns on Social Security are in fact below those anticipated from U.S. capital markets.¹⁶ Finally, we ask whether low Social Security returns are a valid reason to support a privatization that does not involve prefunding. We conclude that several valid rationales can be offered to support privatization, but taken by itself, the low rate of return on Social Security is not among them.

What Do We Mean by Social Security Privatization?

To begin, it is useful to draw a clear distinction between three terms that are often confused: *privatization* versus *diversification* versus *prefunding* of Social Security. By privatization we mean replacing the current mostly unfunded defined-benefit Social Security old-age program with a defined-contribution system of individual accounts held in individual workers' names. By diversification

15. Steve Forbes, "How to Replace Social Security," *Wall Street Journal*, December 18, 1996

16. Although we focus on the rate-of-return concept of money's worth in this paper, additional measures are evaluated in more detail in our longer study (Geanakoplos, Mitchell, and Zeldes, forthcoming).

we mean investing funds (either from the personal accounts or from the Social Security trust fund) into a broad range of assets. These assets might include U.S. private sector stocks and bonds, and foreign securities, in addition to the government bonds now used exclusively by the Social Security trust fund. At the present time, the focus is on diversifying into stocks. By prefunding we mean reducing the sum of the system's implicit and explicit debt (table 4-4).¹⁷

In the public debate these terms are often linked, but they are conceptually different. It is easy to see why all three categories nevertheless appear together in the public mind. Suppose the Social Security system had begun as a forced saving plan in which all workers were obliged to set aside money for their retirement, which would be put into private accounts invested in a balanced portfolio of stocks and bonds. Then from the beginning the system would have been privatized, diversified, and completely prefunded.

The situation is quite different now from what it was in 1935 when the U.S. Social Security system began. Social security systems in the United States and in most developed economies have amassed substantial unfunded liabilities; assets on hand are insufficient to pay for the present value of benefits that have been accrued and promised to workers and retirees based on contributions already made.¹⁸ In the United States, for instance, the unfunded present value of Social Security promises totals about \$9 trillion.¹⁹ Although this Social Security debt is implicit rather than explicit, it is economically significant. The United States will surely not decide to eliminate this implicit debt by ignoring all of the promised benefits.

Our tripartite decomposition is intended to emphasize that Social Security reform could change any one of these three categories without changing the other two. For example, the trust fund could invest in stocks as well as bonds, thus diversifying without privatizing. Alternatively, workers could be given private accounts in which the money was always invested in government bonds, thus privatizing without diversifying. This is the case in Mexico's new individual account system, where the government has required that all pension assets be invested only in inflation-indexed bonds.²⁰ It is also possible to raise system

17. There is debate over whether prefunding should refer to a change that reduces the total government debt rather than just the Social Security debt. Here we assume that any changes in Social Security do not change the non-Social Security debt or deficit.

18. The concept of unfunded liability is different from that of actuarial imbalance. Actuarial imbalance is defined as the present value of expected benefits over some period (often seventy-five years) minus the present value of expected tax receipts over the same period, minus the current value of the trust fund. The United States currently has a seventy-five-year actuarial imbalance of about 2.2 percent of payroll a year, or about 2.9 trillion dollars in present value. See Goss (forthcoming).

19. Goss (forthcoming).

20. Mitchell and Barreto (1997).

Table 4-4. *Differentiating Privatization, Prefunding, and Diversification of Social Security*

- *Privatization*: Replace existing Social Security system with a system of individual accounts held and managed by individuals.
- *Prefunding*: Raise contributions or cut benefits so as to lower the sum of explicit and implicit debt associated with the system.
- *Diversification*: Invest Social Security funds into a broad range of assets, including equities.

		<i>Privatization</i>	
		<i>No</i>	<i>Yes</i>
<i>Prefunding</i>	<i>No</i>	<ul style="list-style-type: none"> • Current system <p style="text-align: center;"><i>Diversification</i></p> <div style="border: 1px solid black; padding: 5px; width: fit-content; margin: auto;"> <p><i>No</i>: Current system <i>Yes</i>: Borrow, invest proceeds in equities through trust fund</p> </div>	<ul style="list-style-type: none"> • Create individual accounts • Issue recognition bonds • Perpetually roll over principal and enough interest to keep path of debt same as that of unfunded liability under current system <p style="text-align: center;"><i>Diversification</i></p> <div style="border: 1px solid black; padding: 5px; width: fit-content; margin: auto;"> <p><i>No</i>: Require individual accounts to hold bonds <i>Yes</i>: Permit individual accounts to hold equities and bonds</p> </div>
	<i>Yes</i>	<ul style="list-style-type: none"> • Raise taxes/cut benefits to decrease unfunded liability <p style="text-align: center;"><i>Diversification</i></p> <div style="border: 1px solid black; padding: 5px; width: fit-content; margin: auto;"> <p><i>No</i>: Invest trust fund in bonds <i>Yes</i>: Invest trust fund in equities</p> </div>	<ul style="list-style-type: none"> • Create individual accounts • Issue recognition bonds • Raise taxes/cut benefits to make path of debt lower than path of unfunded liability under current system <p style="text-align: center;"><i>Diversification</i></p> <div style="border: 1px solid black; padding: 5px; width: fit-content; margin: auto;"> <p><i>No</i>: Require individual accounts to hold bonds <i>Yes</i>: Permit individual accounts to hold equities and bonds</p> </div>

Privatization, prefunding, and diversification are distinct concepts. It is possible to have any one, without either of the other two.

funding without involving individual accounts; taxes could be raised or benefits cut and the proceeds could be put into a central trust fund. Singapore's national Provident Fund, for example, is a nonprivatized, prefunded system where the central government collects taxes sufficient to generate substantial assets, which it then invests on the system's behalf. Conversely, people could be given individual accounts without prefunding of benefit promises. A privatized but unfunded pension system has recently been established in Latvia, where payroll taxes are collected by the government, which then credits workers' so-called notional accounts with paper returns on contributions. Chile is the best-known example of a country whose program is both prefunded and privatized: here workers hold assets in individually managed accounts, and debt under the old system is being reduced over time.

Likewise in the United States, privatization without prefunding is quite possible; a Social Security reform that created a national 401(k)-type system of private accounts could be implemented with no change in Social Security debt. For example, the Social Security system could issue new explicit debt (recognition bonds), guaranteed by the federal government, with payouts set exactly equal to the benefit promises that have accrued to date under the current system. These bonds would be given to current participants in lieu of their accrued future benefit payments. All new contributions to Social Security would then go directly into private individual accounts.²¹ If the recognition bonds were paid off in full with new government tax receipts, the debt would eventually disappear when the last of today's workers finally died. However, the Social Security system could instead borrow again in the future by issuing new bonds to meet the recognition bond coupon payments, and then again and again to meet the payments on these new bonds. Subject to some limits, the government could choose how much interest and principal to roll over and how much to pay off. If the government were to choose to keep the path of explicit debt equal to the path of unfunded liabilities under the current system (the implicit debt), the result would be a Social Security system that was privatized without being prefunded.

In what follows we shall examine the claim that a privatized, diversified Social Security system could deliver higher returns without any additional prefunding. We do so in three steps. First, we analyze returns in a privatized system that confines investments to government bonds. Next, we examine returns in a privatized and diversified system in which investments in stocks are permitted. Finally, we analyze returns in a privatized and diversified system and allow for the possibility that there are some constrained households who do not currently have access on their own to diversified capital markets.

21. More gradual transitions to a system of individual accounts that leave (implicit plus explicit) debt unchanged are also possible.

Table 4-5. *Annual Inflation-Adjusted Returns on Stocks and Government Bonds, 1926-96*

Percent		
<i>Asset</i>	<i>Arithmetic average real return</i>	<i>Standard deviation</i>
S&P 500	9.4	20.4
Long-term government bond	2.4	10.5
Intermediate-term government bond	2.3	7.1
Short-term T-bill	0.7	4.2

Source: Authors' calculations based on data from Ibbotson & Associates (1998).

Are Projected Rates of Return on Social Security Lower Than Those on U.S. Capital Markets?

A starting point for projecting future returns on U.S. capital markets is to examine historical returns. Table 4-5 reports the historical average of inflation-adjusted (that is, real) returns on stocks, bonds, and Treasury bills, as well as their variability. The average annual real return on stocks (as proxied by the S&P 500) was 9.4 percent; the corresponding return on intermediate-term government bonds was 2.3 percent.²² Whether these provide reasonable forecasts for future years depends in part on one's judgment about whether the past will predict the future.


Cohort-specific rates of return under Social Security, from a study by Dean Leimer, are presented in figure 4-1.²³ Historical data on current and past workers and retirees, as well as projections of future contributions and benefits, are used to compute rates of return, which we sometimes call by their more precise name, internal rates of return (IRRs). We note first that prospective (internal) rates of return depend on a host of predictions, including mortality, population growth, and real wage growth.²⁴ Second, the prospective IRR data are not those

22. These are arithmetic averages of annual returns from 1926 through 1996 taken from Ibbotson and Associates (1998). The 1994-96 Social Security Advisory Council projected that in the future stocks would earn a 7 percent real return, compared with 2 percent real return on bonds.

23. Leimer (1994).

24. The internal rate of return (IRR) is defined as the interest rate that equates the present value of taxes paid to the system and the present value of benefits received, by cohort. Two other "money's worth" measures sometimes reported are the present value ratio (PVB/PVT), or the present value of benefits divided by the present value of taxes, and the net present value (NPV), or the present value of benefits received minus the present value of taxes paid. The IRR and its companion money's worth measures are appealing because they seek to summarize in a single measure how a household might evaluate the complex multiyear stream of Social Security payments and

Figure 4-1. *Estimated Real Internal Rates of Return on Social Security Contributions*



Source: Leimer (1994) tax increase balanced budget scenario.

that would follow from forecasting taxes and benefits under *current* Social Security rules, since that system faces insolvency or actuarial imbalance. Instead the IRR figures for the future assume that taxes will be increased sufficiently so that the system does not run out of money.²⁵

Figure 4-1 shows that early cohorts under the program received very high IRRs in real terms. Workers born in 1876 received a real return of over 35 percent a year. Workers born in 1900 received a 12 percent real return. The figure reveals that IRRs fell over time for subsequent cohorts. Leimer estimated the

receipts. However the measure's simplicity belies the extensive assumptions and calculations needed to arrive at a single summary number. For example, in order to conclude that workers born in 1930 anticipate a Social Security IRR of 4 percent, it is necessary to compute what that cohort paid in payroll taxes over all years of work (Leimer 1994). Not all those born in 1930 have retired as yet, so future earnings profiles, taxes, and retirement patterns must be forecast. Each group's tax payments must then be compared against the stream of Social Security benefits actually paid out to people of that birth year. Of course, many 1930-cohort members are still living, so future benefit payments and mortality patterns must again be estimated. For details see Geanakoplos, Mitchell, and Zeldes (forthcoming).

25. Leimer also provides estimates under the assumption that benefits are reduced to maintain solvency.

return to be 5.7 percent for those born in 1920, and about 2 percent for those born 1950–70. For workers born in 1975, he forecasted that IRRs would be around 1.8 percent, dropping down to 1.5 percent for those born in 1998.

Young or middle-aged workers listening to the Social Security reform debate today could reasonably ask how their anticipated IRR from Social Security would compare with what could be earned by investing in U.S. capital markets. Leimer's data indicate IRRs of about 1.5 percent for future cohorts; theoretical models of a pure pay-as-you-go Social Security system in steady state suggest that the IRR would be expected to equal the growth of the wage base, which is currently forecast to be about 1.2 percent.²⁶ Either way, it is clear that projected IRRs are well below expected returns from investment in equities based on historical averages, and fall short of average government bond returns in table 4-5 as well. Thus, the popular rate of return argument in favor of privatization seems to apply whether or not the investments are diversified.

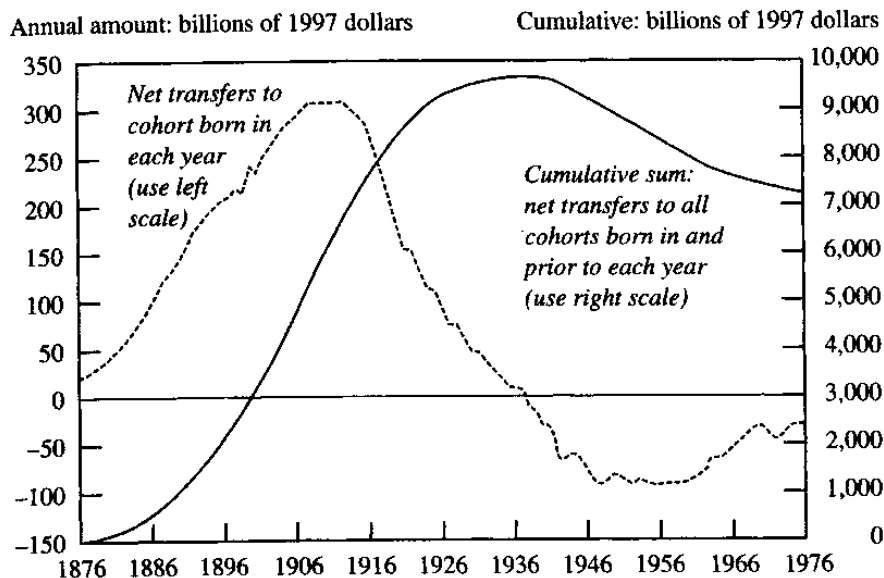
Why Are Projected Social Security Returns Low?

Projected Social Security returns are low, not because of waste or inefficiency, but because the system developed as a primarily unfunded, pay-as-you-go system. In a pure pay-as-you-go system, all contributions received by the system are paid out in the same year as benefits to someone else, and no trust fund is accumulated. This means that there are some early beneficiaries who receive benefits even though they have not made any contributions. The U.S. Social Security system was not started as a pure pay-as-you-go system, because only those individuals who contributed some money to the system were eligible for benefits. Nevertheless, the accumulated trust fund was minimal, so that it was still the case that the present value of benefits received by those retiring soon after 1940 far exceeded the present value of contributions they had made.

The key to understanding why IRRs must fall in our nearly pay-as-you-go Social Security system is to exploit the connection between net present values (NPV) and internal rates of return (IRR). Whenever the NPV for a cohort is positive, the IRR for the cohort will be greater than the market rate, and vice versa. Therefore, stating that early generations received a positive net transfer from Social Security is equivalent to saying that they received above-market rates of return on their contributions.²⁷

26. For a discussion of the assumptions used to estimate future wage base growth rates as well as future Social Security benefit and tax paths, see Advisory Council on Social Security (1997).

27. Comparing the present value of a cohort's benefits and taxes is an application of the generational accounting approach to measuring fiscal policy. See, for example, Auerbach, Gokhale, and Kotlikoff (1994).

Figure 4-2. *Social Security Net Intercohort Transfers*

Sources: Leimer (1994) tax increase balanced budget scenario and authors' calculations. All figures are present values as of 1997.

In figure 4-1, we saw that the real rates of return for early cohorts (birth-year 1876–1900) ranged from 12 to 37 percent.²⁸ Figure 4-2 presents estimates of cohort net present values, derived from Leimer's study, corresponding to these cohort rates of return.²⁹ The dotted line (left scale) is the net present value, in 1997 dollars, of all contributions and benefits for individuals born in the indicated year. As expected for the beginning of a pay-as-you-go system, the present value of benefits exceeded taxes for the first wave of retirees. Each birth-year cohort between 1880 and 1900 (those retiring approximately 1945–1965) received a lifetime transfer (in 1997 present value dollars) of between 40 billion and 240 billion dollars. The solid line in figure 4-2 (right scale) is the cumulative sum of all net transfers received by cohorts born prior to and including the indicated birth year. Cohorts born through 1900 received a cumulative net transfer of about 3 trillion dollars.

28. Recall that these are real or inflation-adjusted returns; nominal returns would obviously be even higher.

29. Leimer (1994) used the interest rates on trust fund assets (essentially intermediate-term U.S. Treasury bonds) to convert current dollars into 1989 dollars. We convert the series into 1997 dollars using the corresponding interest rates between 1989 and 1997.

In part because the Social Security tax rates started out low in 1937 (the initial tax rate was 2 percent) and gradually rose over time, the positive transfers to retirees continued well past the first wave of beneficiaries. Those born between 1901 and 1917 received net transfers with present value of about \$4.9 trillion. In addition, the twenty age cohorts born between 1918 and 1937 are scheduled to receive a further net transfer of \$1.8 trillion, in present value. Adding up the \$3 trillion, \$4.9 trillion, and \$1.8 trillion, we arrive at an estimate of the total net transfer to the first sixty age cohorts of \$9.7 trillion, or roughly \$10 trillion.

What do the positive net subsidies and high returns for early cohorts have to do with the low returns forecasted for current and future cohorts? It can be shown that the present value across all cohorts (from the beginning of the system forward for the infinite future) of the net transfers received must sum to zero.³⁰ Inevitably, cohorts born after 1937 must give up in aggregate the whole \$10 trillion, discounted back to 1997 dollars.³¹

In other words, since past cohorts received positive net transfers, some present and future cohorts must receive negative net transfers. The connection described above between net transfers and rates of return means that we can

30. To see this, first consider a system that is pay as you go, so that in each year aggregate taxes are equal to aggregate benefits, that is, benefits minus taxes equals zero. Hence, the present value of each year's benefits minus taxes must equal zero, and therefore the sum across all years of these present values must equal zero as well. So long as the present value of all benefits across all years and the present value of all taxes across all years are each finite, which must necessarily hold if the interest rate is on average higher than the rate of growth of the economy, we can rearrange and regroup benefits and taxes however we like. Grouping together benefits and taxes for each birth cohort, it must be, as claimed, that the present values of net transfers (benefits minus taxes), cohort by cohort, sum to zero.

This is also true with prefunding, and even with trust fund borrowing, provided that the trust fund assets are invested at the same interest rate used to compute present values. In this case, the present value of all benefits paid up to and including any year T is equal to the present value of all taxes up to and including year T plus the present value of the year T trust fund. Allowing T to go to infinity and assuming that the trust fund is not allowed to grow increasingly negative at rate r or faster and that the government would not want it to grow increasingly positive at rate r or faster, the present value of net transfers (benefits minus taxes) made to all birth cohorts must equal zero. See Geanakoplos, Mitchell, and Zeldes (forthcoming) for more details.

The difference between a prefunded system and a pay-as-you-go system is that in the latter, the early cohorts must get positive transfers, leaving negative transfers for later cohorts, whereas in a prefunded system, the early cohorts may not get any positive transfers. Our current Social Security system has very little prefunding. The correspondingly small trust fund indeed earns a bond rate of return, so our analysis is relevant.

31. It is important to note that this discounting is at the real rate of interest. If, for example, the real rate of interest is 2.3 percent a year and inflation is 3 percent, and if the \$10 trillion were paid back in one lump sum in 30 years, then it would cost \$48 trillion in 2027 dollars. The relentless power of compounding interest is what makes the burden of the initial Social Security transfers many years ago loom so large today.

translate this statement about net transfers into a statement about rates of return. Because past cohorts received rates of return greater than market rates, current and future cohorts must receive rates of return lower than market rates.³²

How big the negative transfer must be for any one cohort born after 1937 depends on how the total transfer is spread across all the cohorts. Observe that according to (our modifications of the numbers in) Leimer, the cohorts born between 1938 and 1977 are scheduled to receive negative transfers of about \$2.5 trillion, reducing the total projected net transfer to cohorts born before 1978 to approximately \$7.2 trillion. Of course much of this \$2.5 trillion transfer has not yet occurred. The young born in 1975, for example, have just begun to work, and so have not had time to make any significant transfers through the Social Security system. Roughly half of the working careers of cohorts born between 1938 and 1977 have passed by 1997. We might guess therefore that subtracting the present value of Social Security taxes already paid by cohorts born between 1938 and 1977 from the present value of Social Security benefits already accrued as a result of these taxes gives about half of \$2.5 trillion. That would mean that the sum of the net tax burden on all contributions made after 1997 would equal \$9.7 trillion minus \$1.25 trillion, or about \$8.5 trillion. In fact, the number must be equal to the unfunded liability, which was calculated independently by Stephen Goss at just under \$9 trillion, thus providing a check on our numbers.³³

There is a simple way of estimating the transfers that might be made each year after 1997. Suppose the total negative transfer of \$9 trillion is spread equitably among all the cohorts, in the sense that every cohort is (or will be) asked to give up the same percentage of its earnings each year of its working career. In a steadily growing economy, the required annual transfer in year t would then be approximately equal to $(r - g) \times$ (unfunded liabilities at end of year $t - 1$), where r is the riskless real rate of return and g is the growth rate of the economy (approximately equal to the sum of population growth and technological improvement).³⁴ Assuming that r is about 2.3 percent (as indicated in

32. If the system had instead started as and remained a fully funded system, then early participants would have received market returns, that is, zero net transfers, as would all current and future workers.

33. Goss (forthcoming).

34. The transfer in period $t + 1$ ($TRANS_{t+1}$) is the amount by which contributions in period $t + 1$ exceed the present value of the benefits accrued as a result of those contributions. Define UL_t as the unfunded liability at time t . It can be shown (Geanakoplos, Mitchell, and Zeldes, forthcoming) that the change in the unfunded liability between t and $t + 1$ ($\Delta UL_{t,t+1}$) is equal to $(r \times UL_t) - TRANS_{t+1}$. In a steady-state economy, the unfunded liability must grow at rate g , i.e., $\Delta UL_{t,t+1} = g \times UL_t$. This implies that the transfer must be exactly $(r - g) \times UL_t$, as claimed in the text.

table 4-5), and that g is about 1.2 percent, and using \$9 trillion as the current unfunded liability, transfers in 1997 must be on the order of \$100 billion. Measured as a percentage of Social Security taxes paid, which were about \$400 billion in 1997, this transfer represents about 25 percent of annual payroll taxes.³⁵

Spreading the remaining start-up cost of our pay-as-you-go Social Security system evenly over all cohorts requires each cohort to give up about 25 percent of every annual contribution, or 3 percentage points of the current 12.4 percent payroll tax. In other words, implicit interest payments explain why young workers may expect only 75 percent of their taxes back in (the present value of) benefits over their lifetimes.³⁶ As long as the Social Security debt is spread out over all subsequent cohorts, returns on Social Security will be lower for every cohort than returns paid by bonds.

Are Low Social Security Returns a Valid Reason to Support Privatization?

The second step in the causal argument supporting privatization seems the easiest and most straightforward—so much so that it is often taken for granted. If projected returns on Social Security are significantly lower than those offered in U.S. capital markets, doesn't it immediately follow that we would all be better off if we were allowed to invest Social Security contributions directly in private securities? Frequent arguments in the popular press and some studies by advocates of privatization suggest that this is the case.³⁷ Yet, this conclusion is misleading for two reasons: it ignores transition costs (that is, how do we eliminate the implicit Social Security debt); and it does not account for changes in risk borne by participants.

35. This number is, of course, only an approximation. To get it we assumed that the real interest rate would stay constant at 2.3 percent and that the wage base would grow steadily at 1.2 percent. Among other things, both assumptions ignore demographic changes. In the current baby boom era, both real interest rates and growth rates are higher than we assumed. If anything, however, it looks like actual $r - g$ is higher than we presumed, which would imply that the necessary transfers might be even more than we suggest.

36. As a check, consider a hypothetical worker who pays level real Social Security taxes for forty years and receives level real benefits for the next twenty years at an internal rate of return of 1.2 percent (an estimate of g). If instead, his contributions were reduced by 25 percent but his benefits remained the same, then his internal rate of return would be 2.12 percent, which is close to our estimate of r of 2.3 percent.

37. See, for example, Steve Forbes, "How to Replace Social Security," *Wall Street Journal*, December 18, 1996, and Beach and Davis (1998) There are sites on the World Wide Web that calculate for users the benefit stream they will likely receive from Social Security and compare it to the income stream attainable from investing their Social Security contributions in private capital markets. See, for example, the Cato Institute web site at www.socialsecurity.org/calc/calculator.html.

Transition Costs

Let us begin by ignoring issues related to risk, such as diversification. Since the rate of return on bonds is greater than the rate of return of Social Security, it would be easy to increase the return for future cohorts by simply ignoring past contributions and not paying any benefits accrued under the current system. The old Social Security system could be shut down, and all new Social Security tax receipts could be put into private accounts invested in bonds. Future cohorts would be able to earn market returns. But then the entire \$10 trillion cost of subsidizing the first sixty cohorts would in effect be borne by the current middle-aged and old, who would then have paid into the system for years and received nothing in return.

Alternatively, one could shut down the old system and privatize but continue to pay all the Social Security benefits accrued to date, based on past contributions. As noted earlier, recognition bonds could be issued to workers and retirees for the full amount of the unfunded liability (that is, \$9 trillion). If the government did not default on these bonds, new taxes would have to be raised to pay interest on the recognition bonds. Further, if the new taxes were set to keep the path over time of recognition bond debt the same as the path of implicit debt under the current system, then the new taxes would correspond exactly to the transfers we mentioned in the last section.³⁸ In other words, it can be shown that the new taxes raised would eliminate all of the higher returns on individual accounts.³⁹ Let us see why.

Consider a steady-state economy growing at the constant rate g with market interest rate r , and with a pay-as-you-go Social Security system begun sometime in the distant past. As pointed out earlier, the implicit tax paid through Social Security in each year is $(r - g) \times$ (the unfunded liability at the end of the previous year). If suddenly at date t Social Security were privatized and transition costs were ignored (for example because accrued benefits were never paid), the IRR on the privatized system in which individual accounts were all held in marketed bonds would be equal to r . If instead recognition bonds were issued, their market value at date t would have to be equal to the unfunded liability in that year. In order to keep the debt growing at the same rate g as the economy, taxes in the next year would have to be raised in the amount $(r - g) \times$ (unfunded liability). The extra taxes needed to finance the payments on the recognition bonds would thus be identical to the transfers made each year in the old Social Security system. By choosing the tax rates appropriately, the tax

38. In a steady state, the required path of debt would keep constant the ratio of outstanding recognition bond debt to GDP.

39. For a further discussion of this result, see Geanakoplos, Mitchell, and Zeldes (forthcoming).

burden could be made to fall exactly on the same people who were contributing more to Social Security than they were receiving in benefits.⁴⁰ Aside from the transfers, participants in the current pay-as-you-go Social Security system are in effect earning the bond rate of return on their money. In a privatized system in which households invested their forced saving in bonds, they would have to pay in new taxes exactly what they paid before in transfers.

In other words, the rate of return on a privatized system in which all private investment is in marketed bonds, net of new taxes needed to pay the appropriate interest on the recognition bond debt, would be identical to the low rate of return on the current system. This is true regardless of how large the difference is between r and g —this difference could be 1.1 percent, 3 percent, 10 percent, or even higher.⁴¹ That is why it is fallacious to assume that all future participants in a privatized Social Security system could earn returns equal to those forecast for U.S. capital markets. We thus agree that returns to the current Social Security system are low, but we argue that there is no costless way of improving them for all current and future workers.

Suppose, alternatively, that taxes were raised disproportionately on current cohorts, and receipts were used to buy out some of the recognition bond debt, so that the system reform increased the degree of funding.⁴² Then later cohorts would face lower taxes to pay the interest on the remaining recognition bond debt. For these later cohorts, returns on Social Security contributions, net of recognition bond taxes, would be higher than under the current system. But for current workers, returns on a privatized system would be lower than under the current system. Thus, transition costs mean either that returns under a privatized system would be the same as the current system (privatization with no prefunding), or that returns in the privatized system would be lower than under the current system for some or all of the currently alive cohorts and then higher for later cohorts (privatization with prefunding). The debate should be focused on whether this is a trade-off worth making, not on whether there is a free lunch.⁴³

40. There are different ways of measuring accrued benefits, and each method would require a different tax scheme to make taxes in a privatized system just equal to transfers in the current Social Security system. We give one example. Suppose accrued benefits are defined on an equal present value basis, that is, suppose each dollar of contributions brings the same present value of accrued benefits (discounted back to the point when the dollar is contributed). Suppose this ratio is .75. Then a proportional tax of $(1 - .75) \times 12.4$ percent = 3.1 percent would leave everyone exactly as well off in a privatized system as he or she was in the current pay-as-you-go Social Security system.

41. Higher r makes privatized returns higher but, as we have just seen, it also increases the interest burden of the unfunded liability.

42. A potential advantage of prefunding is that it would, according to most economists, increase national saving. There is no reason to believe that privatization without prefunding would necessarily increase national saving. See, for example, Mitchell and Zeldes (1996).

43. As argued, for example, by Feldstein (1997).

Risk Adjustment

Advocates of Social Security reform might be tempted to say that people would invest their individual accounts in stocks—earning a higher rate of return—rather than in bonds. By now readers should be suspicious of arguments that promise something for nothing. After all, we just saw that although bonds earn a higher return than the Social Security system, privatizing and investing in bonds would give no higher return once the new taxes needed to redeem the outstanding liabilities of the current system were properly accounted for.

IRR computations usually overlook the fact that workers and retirees bear different risks in a privatized versus a publicly run Social Security system. It is important to keep in mind that when one asset is riskier than another, it should have a higher equilibrium expected market return; this explains why stocks on average should earn a higher return than bonds. As a result, safer returns from a conventional Social Security system cannot be directly compared with riskier returns anticipated from an equities-based system, unless risk adjustment is done to make the programs comparable. By risk adjustment, we mean adjusting downward the expected value of risky returns (for example, stock returns) to recognize the increased risk associated with these securities.

Households that already hold both stocks and bonds in their portfolios and that have the expertise to buy and sell stocks and bonds should value an additional dollar of stocks the same as an additional dollar of bonds, even though stocks have a much higher expected rate of return. If they valued an additional dollar's worth of stocks more than an additional dollar's worth of bonds, they would have already sold some of their bonds and bought stocks with the money.⁴⁴ For this type of household, the risk-adjusted rate of return on an additional dollar of stocks is identical to that on bonds. Put differently, when considering a Social Security change that alters a small fraction of such a household's portfolio, it would be appropriate to compare the current Social Security system to a privatized system in which the individual accounts were required to invest only in government bonds.⁴⁵ And we have already seen in that case that privatization does not bring higher returns, correctly measured.

For a larger change in Social Security investment policy, such households could simply use their non-Social Security portfolio to completely offset the change in Social Security. For example, imagine a middle-aged household

44. This is not to say that such a household would put its first dollar of saving into bonds instead of stocks, but rather after its (perhaps considerable) investment in stocks, the next dollar of stock would increase well-being no more than another dollar of bonds.

45. In practice, government bond returns are not equal to Social Security returns, and neither is riskless. We ignore these issues here.

paying \$5,000 a year in Social Security taxes (including employer contributions), which initially were being placed directly into a personal account (run by the government) in the U.S. Treasury bond market. Suppose the household was also saving outside of the Social Security account an additional \$10,000 a year, \$9,000 in stocks and \$1,000 in treasury bonds. Overall, the household is putting 60 percent in stocks (\$9,000 out of \$15,000), and 40 percent in bonds (\$6,000 out of \$15,000). If the government suddenly switched the entire Social Security account into stocks, this household could restore the 60-40 split of its total saving by putting \$4,000 into stocks and \$6,000 into bonds out of its private saving. In fact, no matter how the Social Security account is invested, this household will be no better or worse off. The appropriate risk-adjusted return to such a household from a Social Security account, no matter how it is invested, is the bond return. For a household that is already diversified, the correctly measured return from a privatized and diversified Social Security system is no higher than the current system.⁴⁶

The real economic benefit to privatization comes from the attendant diversification that would be made available to households that cannot participate on their own in diversified capital markets. According to economic theory, every household whose income is uncorrelated with stock returns should include some stock in its portfolio in order to take advantage of the higher returns stocks provide compared to bonds and other safe assets. The first dollars invested in stocks would definitely raise risk-corrected returns, because they would bring extra returns without adding much risk. Subsequent dollars invested in stock continue to raise returns, but at the cost of more and more risk, eventually lowering risk-adjusted returns. How much money should optimally be invested in stocks depends on the wealth and risk tolerance of the household.

Households that are constrained in their private portfolio from holding sufficient amounts of stock would be helped if a portion of their Social Security returns coincided with equity returns. What type of households would fit into this category and how do we recognize them? First, there are households that will not accumulate any wealth outside of Social Security, but who would like to borrow money to invest in the stock market if they could find a lender to loan

46. This argument holds even if there is an inexplicably large equity premium, as some economists claim there is. Suppose the excess return on stocks above bonds is much more than can be justified by the risk differential, because many households are irrationally underinvesting in equities. In that case, we might be tempted on paternalistic grounds to force these irrational households to hold more stock. However, this is no simple matter. If Social Security were simply privatized, these households would likely choose not to hold any equities in their Social Security accounts for the same reasons they held too few stocks in the rest of their portfolio. If households were instead forced into equities in Social Security, they would likely undo this by reducing their holdings of equities in the rest of their portfolio.

them money at the bond rate.⁴⁷ Second, there are households that will invest 100 percent of their private wealth in the stock market and that would like to invest even more in the stock market if they could borrow at the bond rate. Finally, there are households that will accumulate wealth outside of Social Security, but will choose not to invest in the stock market either because they are uninformed about the relative returns and risks or because it is not worth it to them to incur the fixed costs of becoming a stock market investor. Each of these groups would benefit from owning individual accounts invested in the stock market.⁴⁸ This provides us with an economic rationale for privatization and diversification.⁴⁹

There remains the issue of how many "constrained" households there are, that do not have sufficient access to the stock market. The second group is probably small. The first and third groups are a subset of the people who own no stocks. Fully 59 percent of the U.S. population (in 1995) did not hold stocks in any form.⁵⁰ Some of the nonstockholders, however, are young and have not yet accumulated much Social Security or private wealth. What we really want to know is what fraction of the population is unlikely to hold stock in the future, when they are accumulating Social Security wealth. Although we do not have direct estimates of this, about 50 percent of the population 44 to 54 years old and 60 percent of those 55 to 64 years old held no stock in 1995. But stock-ownership has risen in the past and is likely to continue to rise in the future, so this is probably an upper bound on the fraction of the population that will not hold stock in the absence of Social Security privatization and diversification.

47. Given the chance, these "liquidity constrained" households would prefer to use borrowing to increase current consumption. If borrowing could be used only to purchase other assets, however, these households would choose to borrow and purchase stocks.

48. This assumes that the fixed costs would be lower under an individual account system.

49. This raises the issue of whether diversification into stocks is better achieved through privatized individual accounts or through central trust fund investments. If constrained investors tend to be irrational or myopic in their investment decisions, then it would be more advantageous for them if the Social Security trust fund itself undertook stock investments on their behalf. Unconstrained investors could still keep control of all their asset holdings (inside and outside of their Social Security funds) by compensating in their private accounts for whatever the trust fund did that was not to their taste, as we saw above. However, households such as those who have chosen optimally to hold small amounts of stocks, perhaps because they are risk averse, may be forced to hold too much extra in stocks in their Social Security accounts. These households would be made worse off. The more heterogeneous are constrained households in their tolerance of risk, and the wiser we think constrained households would be in their investment choices, the more attractive is privatization as a means of achieving diversification. The more homogeneous are constrained households in their tolerance of risk, and the more myopic we think constrained households would be in their investment choices, the more attractive is trust fund investment as a means of achieving diversification.

50. This is based on calculations using the 1995 Survey of Consumer Finances. It includes stocks held directly, through mutual funds, and through defined-contribution pensions. See Kennickell, Starr-McCluer, and Sunden (1997) and Ameriks and Zeldes (in progress).

Also, not all those without stock would benefit substantially by holding stocks. For some, their wage income or private business income might be sufficiently positively correlated with the stock market that they are better off not holding any equities. Others, who hold zero stock because they are very risk averse and face a small fixed cost, would only see a small benefit from increased holdings of stock in their Social Security account. Overall, we would guess that significantly fewer than half of households would benefit substantially, and therefore experience substantially higher risk-adjusted returns, from increased Social Security investment in stocks.

If there is a large number of constrained households, then Social Security diversification would have macroeconomic consequences as well. The most important general equilibrium effect of diversification would be an increase in the demand and thus in the price of stocks and a corresponding decline in their expected return. As the value of the stock market increased, current holders of stocks, which disproportionately include the wealthy and older workers who have had time to accumulate stock, would see their wealth increase. Young workers would find that they earned smaller returns on their future stock purchases than they would have in the absence of diversification.⁵¹

While privatization and diversification may indeed bring benefits to some constrained households, and perhaps some indirect benefits to the economy as a whole, the presence of these households is not part of the currently popular rate-of-return argument for privatization. According to that argument, all households, and especially the youngest (25-year-old workers), are supposed to see higher returns from privatizing Social Security. Yet when properly measured, rates of return for unconstrained households will stay the same or actually decline for young unconstrained households.

Another risk consideration is that the government-run Social Security system provides insurance functions that an individual-account system probably would not provide, including insurance for shocks to earnings, length of life, disability, and inflation. For instance, the benefit formula is structured to provide a higher rate of return to low lifetime earner households than to high lifetime earner households. This is justified on the grounds that a social insurance plan can pool over the entire population certain risks that are difficult to insure privately, particularly disability, unemployment, and poverty. To the extent that

51. Social security diversification would bring some indirect benefits to the economy if there were many constrained households. Unconstrained households would end up holding less stock, because some of it would be in the hands of Social Security accounts held by constrained households who could not buy stock previously. Thus unconstrained households would bear less risk. They would be inclined to shift the mix of investment projects undertaken toward more risky ones. This might in turn raise future GDP.

private equivalents for these forms of social insurance would be more costly or nonexistent, privatizing the program would increase participants' risk exposure. Consequently, to the extent that social insurance affords benefits that private insurance could not, this raises the risk-adjusted return on a government-run Social Security system.

There is another risk factor that must be accounted for as well. A publicly run defined-benefit program incurs political risk. This occurs because current workers cannot effectively contract with and bind as-yet-unborn cohorts of taxpayers to pay for them when they grow old. As a result, baby boomers feel unsure that they will wield the political clout to extract rising payroll taxes from their children and grandchildren, to support them when they are old. This is a risk that cannot be traded, so there is no way for those more willing to bear this risk to trade with those less willing to bear it. Although the simple models described above do not have a good way to price this type of risk, correcting for political risk would probably work in the direction of increasing the risk-adjusted returns and relative attractiveness of a privatized system.

Conclusion

In this paper we argue that privatization, diversification, and prefunding are distinct and should be considered as such. We also argue that it is worthwhile separating a move toward system privatization into three steps. Suppose that individual accounts are created, but that all benefits accrued to date from past Social Security contributions are honored, and that the funding level of the system is held constant (because new explicit debt is issued to cover the unfunded liability of the present Social Security system). In the first step, consider a world with no uncertainty in which individuals are only allowed to invest the individual accounts in government bonds. Although the market return on bonds is greater than that projected on Social Security, once transition costs (to pay interest on the new debt) are accounted for, the rate of return on the two systems would be identical. We estimate that paying for the transition costs would require about a 25 percent tax on all payments into new accounts. This surtax would wipe out all of the extra returns attainable by holding bonds in people's individual accounts. Our 25 percent number depends on forecasts of future growth rates g of the economy, and future real interest rates r . If those forecasts of r and g were to change, the tax would have to change. But the result that all of the gains to privatization would disappear in extra taxes to pay off the unfunded liability is true regardless of the specific values of r and g .

Next, consider adding idiosyncratic risk (for example, shocks to earnings and length of life). If the privatized system contains no special insurance provisions

and the private market is unable to provide these, then the risk-adjusted rate of return under privatization would be lower than under the current system.

Third, consider adding aggregate risk and allowing households to invest their individual accounts in equities. For households that already held both stocks and bonds elsewhere, and were thus unconstrained in their portfolio choice, the risk-adjusted rate of return would be no higher than the risk-adjusted rate of return if the individual accounts were held entirely in bonds. For the others, namely, those who do not currently have as much invested in equities as they might like, the risk-adjusted rate of return on a portfolio with equities would be higher as a result of diversification into stocks, thus providing an important rationale for diversification. Nevertheless, this return would not be as much higher as a naive comparison of expected returns would suggest. If this group is large, there will also be general equilibrium effects that will likely lower the expected return on stocks, thus making worse off young households who would have invested in stocks anyway.

Finally, what if the system's funding were increased at the same time as individual accounts were created? Although the increase in funding would likely raise the rate of return on Social Security for future generations (and raise national saving in the process), it must be kept in mind that it would do so only at the cost of lower returns to current cohorts.

We argue that the Social Security reform debate should focus on the trade-off between returns for current cohorts versus future cohorts; the risk-bearing benefits to the economy of enlarging the population that holds a diversified portfolio to include constrained households; and whether diversification is better implemented via privatized personal accounts or through trust fund investments.

We recognize that privatization has several other important benefits, including increased portfolio choice, reduced political risk, possibly reduced labor supply distortions, and an intangible increased sense of ownership and responsibility. It is also possible that privatization may improve the political feasibility of implementing prefunding or diversification, thus increasing the chance of achieving the benefits to future generations and constrained households described earlier.

Indeed some find these overwhelming reasons to favor privatization. These positives, however, must be balanced against the loss of social risk pooling mentioned above, somewhat higher administrative costs, and perhaps the social and personal costs of permitting workers to make "unwise" portfolio choices.⁵²

52. See Mitchell (forthcoming) for an analysis of administrative costs.

In any event, our main message is that the popular argument that Social Security privatization would provide higher returns for all current and future workers is misleading, because it ignores transition costs and differences across programs in the allocation of aggregate and household risk.