This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Risk Aspects of Investment-Based Social Security Reform

Volume Author/Editor: John Y. Campbell and Martin Feldstein, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-09255-0

Volume URL: http://www.nber.org/books/camp01-1

Publication Date: January 2001

Chapter Title: Asset Allocation and Risk Allocation: Can Social Security Improve Its Future Solvency Problem by Investing in Private Securities?

Chapter Author: Thomas E. MaCurdy, John B. Shoven

Chapter URL: http://www.nber.org/chapters/c10590

Chapter pages in book: (p. 11 - 40)

1

Asset Allocation and Risk Allocation
Can Social Security Improve Its Future Solvency Problem by Investing in Private Securities?

Thomas E. MaCurdy and John B. Shoven

Policy makers widely accept that social security faces a long-run solvency problem. The social security trustees publish a seventy-five-year forecast of the OASDI system's finances every year. Since the 1983 amendments, the forecast has gotten consistently more ominous. Figure 1.1 shows the intermediate-assumptions forecast for the trust fund balances in 1983, 1985, and 1998.

The trustees' 1998 report has the OASDI trust fund completely depleted in 2032 (when the youngest baby boomer, born in 1964, will be sixty-eight) rather than 2063 (when the same individual will be ninety-nine years old), the forecast that immediately followed the 1983 social security amendments. The report has as the intermediate or "best-guess" forecast that the system's income (not counting net interest income on the special issue bonds held in the trust fund) will fall short of its costs beginning in 2013. In 1999, the surplus was projected to be 1.52 percent of covered payroll, but this turns negative in 2013, and, by the end of the trustees' forecasting period (2075), the annual shortfall under the intermediate set of assumptions is forecast to be 6.43 percent of covered payroll. That means that, if we operated the system strictly on a pay-as-you-go basis, the payroll tax rate would have to be increased by 6.43 percentage points to achieve balance in 2075. The 2075 OASDI deficit (with the intermediate-assumptions

Thomas E. MaCurdy is professor of economics at Stanford University, senior fellow of the Hoover Institution, and a research associate of the National Bureau of Economic Research. John B. Shoven is the Charles R. Schwab Professor of Economics at Stanford University and a research associate of the National Bureau of Economic Research.

The authors have benefited from the comments and help of John Campbell, Victor Fuchs, Jim Poterba, Sylvester Schieber, and Steve Zeldes. The authors acknowledge the fine research assistance of Stanford Ph.D. students Davide Lombardo and Clemens Sialm.

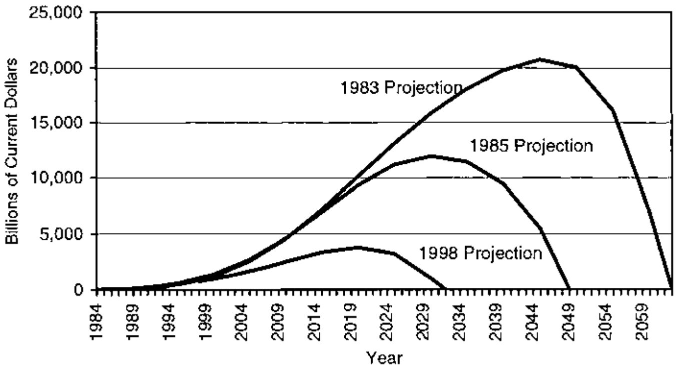

Fig. 1.1 Three forecasts of the OASDI trust fund balances

forecast) amounts to 4.26 percent of GDP. One way to look at the financial solvency problem of social security is that the intermediate forecast is for fourteen years of modest and declining surpluses, followed by ever-growing deficits as far as the eye can see. Against that we have a relatively small trust fund (currently $760 billion) that generates roughly $40 billion of interest income but that is now forecast to be completely depleted by 2032, as shown in figure 1.1. Naturally, the financial situation for OASDI is much worse under the trustees' high-cost set of assumptions. With them, the date of OASDI trust fund exhaustion is 2022. The date when payroll-tax receipts first fall short of benefit payments is 2006, and the eventual shortfall of income to costs in 2075 is 16.04 percent of taxable payroll.

The social security trustees summarize the seventy-five year outlook for OASDI by computing the long-run actuarial balance. In 1998, the seventy-five year actuarial balance was in deficit by 2.19 percent of covered payroll with the trustees' intermediate assumptions. What that means is that, had the payroll tax been immediately increased in 1998 by 2.19 percentage points and the increase maintained for the next seventy-five years, the life of the trust fund would have been extended to the seventy-five-year horizon under the intermediate set of economic and demographic assumptions. Even under this hypothetical scenario, the payroll-tax proceeds would be less than benefit payments under the currently legislated benefit structure beginning in roughly 2020. Further, the seventy-five-year actuarial balance would last for exactly one year; with every passing year, one fewer surplus year would be in the seventy-five-year window, and one more deficit year would be included. The immediate payroll-tax increase needed to fix the solvency problem of the current OASDI system permanently would be significantly greater than 2.19 percent. Steven Goss (1999), of the Office of the Actuary of the Social Security Administration, has estimated that the permanent or open-ended actuarial deficit is 4.7 percent of covered payroll under the intermediate set of assumptions.

All this illustrates the dimensions of the problem of returning the existing OASDI system to long-run solvency. Operating within the existing structure of social security, there are only two obvious paths: raise taxes or cut benefits, or both. If we take Goss's 4.7 percent of covered payroll figure for the perpetual actuarial deficit, permanently fixing the finances of the system (under the intermediate set of economic and demographic assumptions) will require immediate tax increases or benefit cuts totaling 38 percent. If the implementation of these steps is delayed, as it almost certainly will be, then the benefit cuts and tax increases will need to be larger. Neither tax increases nor benefit cuts of this magnitude are economically or politically attractive. Naturally, people are looking for a more palatable way out of the social security solvency problem.

One natural place to look is the investment returns earned on the trust fund. Currently, the trust fund is exclusively invested in special nonmarketable U.S. government bonds. When these bonds are issued, their interest rate is set at the average interest rate of marketable Treasury bonds with a maturity of four years or more. The bonds have one special feature, which offsets their nonmarketability: they are redeemable at par at any time. This feature is not generally available on publicly held bonds and notes, with the exception of U.S. government savings bonds. The special issue bonds are certainly safe, with no price risk and with the principal and interest fully backed by the U.S. government, but Treasury interest rates are somewhat less than what is offered on AAA corporate bonds and trail the average total return earned on common stocks by a wide margin.

The question that we address in this paper is whether a significant fraction of the whole solvency problem could reliably be solved by having the Social Security Administration invest the OASDI trust fund in higher-yielding private securities. The analysis in MaCurdy and Shoven (1992) suggests that such a strategy might yield an improvement in social security's finances.[1] Of course, there is a more fundamental question regarding whether society as a whole would benefit from this new asset allocation. Another question that we address is how risky such a change in the trust fund asset allocation would be and who would bear that risk. Related to this question, we discuss the feasibility of social security delivering a true defined-benefit pension program to its participants.

1. Our paper investigated the effects of defined-contribution accumulators consistently acquiring stocks or bonds for careers between 1926 and 1989. We showed that stock accumulators consistently ended up with more money than bond accumulators for all careers longer than twenty years in the period 1926–89. The result is weakened by the fact that we used only the actual historical pattern of financial returns and that there are only three completely independent twenty-year runs of data in the period 1926–89. We extended the data to 1876–1990 in an unpublished manuscript. There we found that the career length had to be forty years before stock accumulators always ended up ahead of bond accumulators. Once again, the data limitations do not allow one to predict with confidence the likelihood of stocks outperforming bonds for long careers. In the present paper, we try to get around the data limitations by using the bootstrap statistical procedures.

1.1 The Proposals

We focus on two of the reform proposals that rely on investing the central trust fund in higher-yielding private securities in order to maintain the general benefit structure of social security. However, our analysis is applicable to any plan that attempts to make progress on the solvency of the system by simply reallocating portfolios toward higher-yielding securities, including the plan outlined by President Clinton in his 1999 State of the Union Address. The two plans of this type that we describe in detail are the "maintain-benefits" (MB) proposal of the 1994–96 Advisory Council (often referred to as *the Bob Ball plan*) and the plan offered recently by Henry Aaron and Robert Reischauer (1998). Among the measures that both plans advocate, the asset reallocation is credited with the largest effect on reducing or eliminating the seventy-five-year actuarial deficit.

The details of the two plans are shown in table 1.1. There are many similarities between them. The numbers in the tables come directly from Aaron and Reischauer (1998, table 6-1) and from volume 1 of the *Report of the 1994–96 Advisory Council on Social Security* (SSA 1997, table 1, app. II). The figures are based on estimates of the Office of the Actuary of the Social Security Administration, assuming the intermediate alternative II economic and demographic conditions. The numbers for the same action may differ slightly for the two plans since the estimates for the Aaron-Reischauer plan were made roughly two years later than those for the MB plan. Further, the plans differ slightly in how they implement each measure. Even with that said, both proposals have the same goal—to deal with the seventy-five-year actuarial deficit within the context of the current structure of benefits. Both assume that the measurement of the CPI will be changed in ways that lower official inflation and cause the program's costs to grow more slowly.[2] Both advocate covering all newly hired state and local workers, increasing the number of years included in the AIME (average indexed monthly earnings) formula from thirty-five to thirty-eight, and taxing all social security benefits received over and above the employees' own actual contributions. Both proposals characterize their changes in the personal income taxation of social security benefits as putting them on the same basis as private pension income. However, their plans involve significantly higher taxes than the Schieber-Shoven PSA-2000 plan, which also claims to put social security and private pension benefits on the same tax footing (Schieber and Shoven 1999). The PSA-2000 plan taxes half of social security benefits on the argument that half of social security contributions were made with before-tax money (the

2. Many wage contracts are officially or unofficially tied to the CPI. It is not clear that the Office of the Actuary took the feedback of changing the official CPI on wages into account. If not, the gain from the restructuring of the CPI may be exaggerated.

Table 1.1 Contribution of Each Feature of the Aaron-Reischauer Plan and the Maintain-Benefits Plan to the 75-Year OASDI Actuarial Deficit (all numbers are % of covered payroll)

	Aaron-Reischauer Plan	Maintain-Benefits Plan
Projected long-term deficit	2.19	2.17
Effect of assumed changes by BLS in CPI measurement	−.45	−.31
Cover all newly hired state and local employees	−.21	−.22
Increase the number of years in the AIME formula from 35 to 38	−.25	−.28
Tax 85% of social security benefits	−.36	−.31
Increase the OASDI payroll tax in 2045 by a combined 1.6 percentage points	N.A.	−.22
Redirect revenue for taxation of benefits from hospital insurance to OASDI starting in 2010	N.A.	−.31
Raise the normal retirement age to 67 by 2011, and then index further increases to improvements in mortality	−.49	N.A.
Raise the age of early retirement eligibility to 64 by 2011, and then index further increases to improvements in mortality	−.23	N.A.
Gradually reduce spouse's benefit from 50 to 33.33%; raise benefits for surviving spouses to 75% of couple's combined benefit	+.15	N.A.
Gradually invest part of the trust fund assets in private stocks and bonds	−1.20	−.82
75-year deficit with plan	−.85	−.24

Note: N.A. = not applicable.

employer contribution) and half with after-tax dollars (the employee contribution). Essentially, Schieber and Shoven's plan would have half the contributions treated like a Roth IRA (after-tax contributions with tax-free withdrawals) and half like a regular IRA (before-tax contributions with taxable withdrawals). In contrast, the Aaron-Reischauer and Ball plans tax 85 percent or more of social security benefits.

The MB plan contains two features that Aaron and Reischauer did not choose to copy. One is a 1.6 percentage point payroll-tax increase in 2045, and the other is a redirection of money from Medicare to OASDI. The latter feature of the MB plan seems ill advised as the long-run finances of Medicare are in worse shape than are those of OASDI. Similarly, the Aaron-Reischauer plan contains several features not in the MB plan. These include accelerating and extending the increase in the age of normal retirement, advancing the age of early retirement, and readjusting the spousal and surviving spouse benefits. The Office of the Actuary of Social Security estimates that, both plans if enacted immediately and completely, would more than eliminate the seventy-five-year actuarial deficit of the system.

The common feature of the two plans shown in table 1.1 is that the biggest contributor to eliminating the long-run solvency problem is investing part of the OASDI trust fund in private securities. The two plans differ in the details of how they would do that. In fact, the MB plan in the 1994–96 Advisory Council report suggests that further study be given to the idea before implementation. Nonetheless, when the Advisory Council scored the plan to see whether it achieved the goal of eliminating the seventy-five-year actuarial deficit, it included the provision that the trust fund would begin investing in stocks in the year 2000 and that the proportion of trust fund assets in stocks would gradually rise until it reached 40 percent in 2015. The assumed real rate of return on the stock portion of the portfolio in the MB plan is 7.0 percent. The Aaron-Reischauer plan is to have the trust fund balances that exceed 150 percent of one year's benefits gradually invested in common stocks and corporate bonds. Since the OASDI trust fund balances currently exceed the 150 percent of annual payout criterion, the switch to private securities would begin immediately under the Aaron-Reischauer program. Remarkably, Aaron and Reischauer estimate that 55 percent of the whole actuarial deficit of 2.19 percent of payroll would be eliminated by this asset reallocation alone. They also estimate that the combination of the asset reallocation and the effect of the changes that the Bureau of Labor Statistics is making in the construction of the CPI eliminates over 75 percent of social security's projected long-term deficit. Any 75 percent cure that is this painless deserves careful scrutiny.

1.2 The Net Transaction

The net transaction involved in having the central OASDI trust fund invest a portion of its assets in common stocks is an asset swap. Social security or the federal government sells additional government bonds to the public and uses the proceeds to acquire common stock from the public. In the social security context, the system is selling one set of assets (U.S. government bonds) in order to acquire another set of assets (a diversified portfolio of common stocks) of equal value. Of course, this transaction can be examined without reference to social security's financial problems at all. The real issue is whether the government can improve the welfare of taxpayers (or social security participants) by issuing and selling bonds and using the proceeds to buy common stocks.

There are at least two reasons to be skeptical about the advantages of the net transaction being discussed. First, the total capital stock and wealth in the economy (at least to a first approximation) are unaffected by the asset swap. Therefore, the level of GDP and national income is unchanged. If social security or the government can systematically improve its financial position by making this exchange, the private sector, which presumably is on the other side of the transaction, is systematically losing. It is hard to imagine that the politicians and bureaucrats running social security are systematically getting the better of the pension fund managers and institutional investors who are buying the bonds from the government in exchange for common stocks. It is possible, of course, that the exchange could be played out on international capital markets so that Americans as a whole could conceivably end up as net winners (or losers) in the transaction. Second, there is the matter of risk. While stocks have a much higher expected return than government bonds, they also involve much higher risks (particularly given the fact that the current special issue bonds have the feature that they can be redeemed at par at any time). The two plans discussed above are very vague as to how much additional risk the system would be assuming and how that risk would be shared among taxpayers and social security participants. As we discuss in section 1.5 below, there is a real question about whether it is feasible to maintain a universal coverage defined-benefit pension plan funded with risky securities.

Having the government exchange bonds for stocks in the hope of relieving the solvency problems of social security is a form of financial engineering. The recent history of such schemes is not promising. Savings and loans sell short-term liabilities (certificates of deposit and demand deposits) and acquire higher-yielding long-term mortgages. The savings-and-loan crisis of the 1980s was brought about when short-term interest rates peaked in 1980 in the mid-teens when the savings-and-loan mortgage portfolio yielded 6–8 percentage points less. The usual yield relation did not hold, and the resulting bankruptcies and bailouts cost taxpayers massive sums of money. The Orange County debacle resulted from a similar failed attempt to exploit the shape of the yield curve. The recent massive losses of hedge funds resulted from asset swaps that failed to perform as expected. Before the U.S. government engages in exchanging bonds for stocks, careful analysis is clearly warranted.

1.3 Analyzing the Stocks-for-Bonds Swap

To gauge the riskiness of the net transaction of selling government bonds and buying corporate stocks, we use market returns from 1954 to 1997 and simulate what would have happened had the government completed the transaction. We separately look at what would have happened with the government selling ten-year and twenty-year bonds. We assume that market returns would have been unaffected by the government transaction. This almost certainly favors the exchange strategy that in reality might very well drive up interest rates and depress equity returns. The counterfactual simulations have the government selling bonds in the past, buying the S&P500 stock portfolio with the proceeds, and then paying all the bond payments (interest and principal) from the resulting stock portfolio. The strategy is deemed successful if all the payments on the bonds can be made with the stock funds with money left over. The strategy is deemed to have been a failure if the stock portfolio is unable to generate sufficient cash to make the required bond payments. In the discussion presented below, we initially explore how the strategy would have worked with the actual time series of returns generated by the stock market (and the actual historical interest rates on government bonds). The interest rates are those published in the *Economic Report of the President* (1998) and computed by the Board of Governors of the Federal Reserve System. The stock market returns are the "large company stocks" (i.e., S&P500) total return series in Ibbotson Associates (1998). The problem with this historical approach is that there are not many independent ten- or twenty-year periods in our data set (1954–97). In fact, there are only four completely independent ten-year runs of data and two of twenty years. We could add more data by examining pre-1954 information, and we do so in some of the analyses in this section. There is a serious question about whether adding data before 1954 helps or hurts when assessing the viability of a proposed strategy for the twenty-first century. The problem is that the additional data are likely to be drawn from a different regime and, therefore, may do more harm than good to the analysis. After examining the actual performance of the stock market in the period 1954–97, we report on extensive bootstrap simulations of what could have happened using the same annual return data, but now with the order of the returns randomly scrambled according to a bootstrap statistical approach.

First, we examine the case of the government borrowing money for ten years at the actual historical interest rates on U.S. government bonds and investing in the S&P500. Figure 1.2 shows the money left over at the maturity of the bonds (per dollar borrowed) after all bond payments had been made out of the stock account. The net cost of the original asset swap is zero, so any residual money is pure profit, and any residual shortfall is loss. Given our definition of success and failure, the asset-swap strategy would have failed seven times out of the thirty-four counterfactual experiments. The limitations of the actual historical record are very apparent in the results, however. Basically, the strategy of the government issuing ten-year bonds and buying the S&P500 would have worked from 1954 through 1964 and again from 1975 through 1987. However, it would have failed in seven of the ten years 1965–74. This reflects the limitations of observing only four completely independent ten-year sets of data. On an annual basis, the failure rate is seven of thirty-four, or 20.6 percent, but a more informative way of reporting it is that the failure rate is roughly one in four (from the fact that we had four non-overlapping stretches of data and the strategy worked in three of the four). It should also be noted that our criterion for success is very modest. We are not requiring that stocks

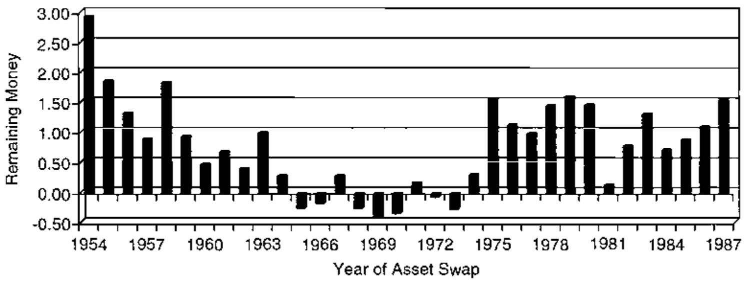

Fig. 1.2 Net money generated at the completion of a swap of ten-year U.S. government bonds for common stocks

match the 7 percent real yield assumed in scoring the MB plan. Rather, we are simply requiring that the stock returns were sufficient to pay off the bonds, which had low or even negative ex post real returns (especially in the 1970s).

While we have been focusing on the nontrivial chance of failure of the asset-swap strategy, it should also be noted that figure 1.2 indicates that it would indeed have worked twenty-seven of thirty-four times. Further, when it works, the average gain is quite substantial. The numbers in the figure represent the amount of money left in the stock account (per dollar borrowed) at the maturity of the bonds. The overall average amount left over for the thirty-four experiments is 78.5 cents per dollar borrowed; the average over the twenty-seven years with positive outcomes is almost \$1.05. Of course, it is well known that, on average, stocks outperform bonds by a wide margin. What is less well known is that borrowing money for ten years to buy stocks is very risky—even with the favorable interest rates available to the government.

To acquire a better sense of how such a government asset-swap policy might work in the future, we have used the same set of 1954–97 stock returns as the basis for generating simulated sequences of returns using a moving-block bootstrap method to allow for autocorrelation. For each year 1954–97, we calculate what would have happened had the government borrowed additional money with ten-year bonds (at the actual prevailing interest rate). The stock market return in the year of the borrowing is taken as the actual return in that year. However, the succeeding real returns for the next nine years are randomly chosen in blocks of three-year sequences or "blocks" (with replacement) from the set of all realized real returns. For example, in order to calculate how this strategy would have worked in 1971, we compute the terms on the borrowed money according to the prevailing nominal interest rates in 1971. Further, we assume that the return on the stock market for the first year was the actually observed 1971 return. For the returns for the next nine years, we randomly choose three dates (with replacement, meaning that the same date could be chosen more than once) between 1954 and 1995. If we choose 1958, 1988, and 1956, for instance, then the ten-year string of real returns under this simulation would be (R_{71}, R_{58}, R_{59}, R_{60}, R_{88}, R_{89}, R_{90}, R_{56}, R_{57}, R_{58}), where R_t is the total real return of the S&P500 for year t. The fact that we are choosing three-year blocks of data should be apparent in this sample vector of returns. The real returns chosen via the bootstrap technique are converted into nominal returns using actual historical inflation rates. For instance, in the counterfactual experiment of borrowing money in 1971, we use the actual inflation rates for 1971–80 to generate the nominal returns resulting from the real returns chosen by the bootstrap process. For each asset-swap year, we examine one thousand such random sequences of returns.[3] The failure rate for the strategy differs by the year of the swap because of different interest rates and because of different first-year stock market returns. The results of this bootstrap simulation are shown in figure 1.3.

Implicitly, we have assumed that there is stability in the underlying probability distribution determining real stock returns during the period 1954–97 and that each observation was equally likely. Of course, it is certain that these assumptions are not entirely correct; it is also certain that many low-probability events never occurred in the observation period. Finally, we are assuming that the same underlying probability distributions that generated the 1954–97 observations will generate stock market returns in the future. In total, this is a strong set of assumptions. Nonetheless, the implied results provide important insights. According to figure 1.3, the overall failure rate is predicted to be 22.1 percent, meaning that the asset-swap policy failed 9,723 times in the 44,000 simulations we did for the combined forty-four years of analysis. The expected failure rate for the most recent ten-year period was 44.7 percent, owing to the higher real interest rates of this period. Note that these failure rates are most likely conservative since they still are based on the robust returns realized in U.S. equity markets in the period 1954–97. These failure rates may therefore form something of a lower bound on the riskiness of the asset-swap strategy going forward.[4]

The patterns of failure probabilities shown in figure 1.3 stand in sharp contrast to the success or failure patterns shown in figure 1.2 above, which uses the actual realized time series of returns. The reasons for the divergence are easily explained. The accelerating inflation of the 1970s was not fully incorporated into government interest rates until 1981 and 1982. The

3. Our preliminary analysis indicates that selecting block sizes 1, 2, 3, and 5 would not change the main findings reported here.

4. Once again, we are concentrating on the failure probability as a measure of the riskiness of the asset swap. The asset swap does work more than 50 percent of the time and often generates large profits.

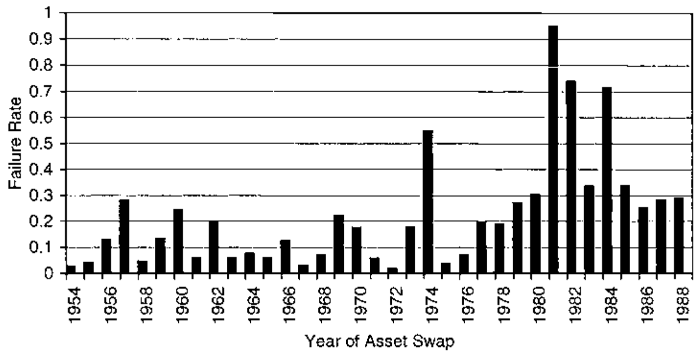

Fig. 1.3 Failure rate by year of asset swap for ten-year bonds, bootstrap simulations with three-year blocks, 1954–88

ex ante analysis of figure 1.3 shows that to be the period with the highest chance of failure for the ten-year bonds-for-stocks swap (particularly with the poor performance of the stock market in 1981). However, ex post the stock market has performed extraordinarily well from 1982 to the present. As a result, the asset-swap strategy would have worked in 1981 and 1982, even though our bootstrap simulations show that it had a relatively low probability of doing so. On the other hand, the ex ante failure probabilities in 1965–72 were not particularly high; nonetheless, in fact, the strategy would have failed in six of those eight years, as shown in figure 1.2. The point is that the returns that were actually experienced were just a single set of draws from an economic and financial system generating risky returns.

We perform the same analysis for twenty-year bonds. In this case, the limitations of the data are even more severe since there are only two completely independent sets of twenty-year observations in the 1954–97 data set. Therefore, it is difficult to know how to interpret the fact that the asset-swap strategy would have worked for every year 1954–77 using actual interest rates and the actual time series of stock market returns. This says very little about the underlying chance of failure with a future policy of swapping twenty-year government bonds for a diversified portfolio of common stocks. The same bootstrap approach is used in this case as with the ten-year bonds. Now, instead of a single historical set of stock market returns, we can generate thousands of simulated sequences of returns generated from a random selection of the actual annual observations. The results of doing that are shown in figure 1.4.

The pattern of failure probabilities is primarily determined by the actual pattern of nominal interest rates on government bonds and by the actual first-year return on the stock market. Therefore, it is not surprising that the general pattern of failure rates is similar in the twenty-year bond case of figure 1.4 and the ten-year bond case of figure 1.3. The high failure

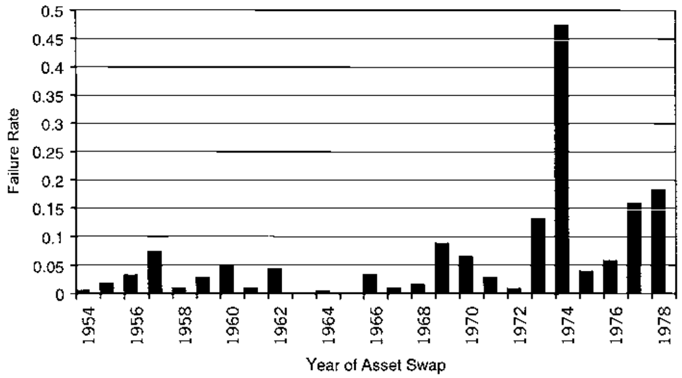

Fig. 1.4 Failure rate by year of asset swap for twenty-year bonds, bootstrap simulations with three-year blocks, 1954–78

probability (47.3 percent) for the asset-swap strategy in 1974 is due to the fact that the real return on the S&P500 in 1974 was −34.5 percent. If stocks lose more than one-third of their value immediately after bonds are sold to buy stocks, the chances of making the swap work are greatly diminished.

For twenty-year bonds, the overall predicted failure rate during 1954–78 is 6.24 percent. In the most recent ten-year period (1969–78), the failure rate is predicted to be 12.27 percent, which we view as a considerable chance of failure. This suggests that, were the government to borrow money by issuing twenty-year government bonds (assuming that it could do so without raising interest rates), and were the real returns on the stock market generated by the same process that produced the returns of 1954–97 (although with the ordering of the returns randomly scrambled), the chance that the government would have to borrow additional money when the bonds mature is about one in eight. The asset swap has a considerable chance of being counterproductive in terms of the finances of the federal government in general and social security in particular.

The historical record on interest rates also represents only a single set of realizations of what might have happened. Future real (or nominal, for that matter) interest rates may be higher or lower than what we have observed in the past. To add some robustness to our results, we examine what would have happened to the finances of the government had it issued inflation-indexed bonds and bought stocks with the proceeds. In our first experiment, we set the real interest rate on ten-year inflation-protected bonds at 3.5 percent. This is actually somewhat less than the recently observed interest rate on this type of security. Each year, the principal on the bonds is marked up to reflect realized inflation. The bonds are assumed to pay 3.5 percent interest each year on the revised (i.e., inflation-adjusted) principal amount and at maturity return the original investment marked up by cumulative inflation. The question we want to ask is how likely it is that an S&P500 stock portfolio can generate sufficient returns to pay off the claims of such inflation-indexed bonds.

Once again, we can look at how this asset swap would have fared historically could it have been accomplished without affecting the actually observed stock market returns. In this case, we have used the observed stock market returns and inflation rates for the period 1927–97 and examined counterfactual asset swaps for the sixty-one years 1927–87. Basically, we have six full and completely independent (i.e., nonoverlapping) runs of ten-year data. Figure 1.5 shows the results for this case with ten-year inflation-indexed bonds. The strategy would have failed eighteen times out of the sixty-one years examined, for a failure rate of 29.5 percent. The failures are strongly autocorrelated, however, because of the overlapping returns. If the swap strategy fails in year *t*, it is quite likely that it will also fail in year $t + 1$ since the two experiments involve nine common years of stock market returns. Not surprisingly, figure 1.5 shows that the asset-swap strategy would have had mixed success in the period 1927–40, followed by twenty-three years of complete success, followed by eleven years in a row of failure and, finally, by thirteen years of success. Some of the failures are quite substantial. In seven of the 1964–74 years, the government would have had to borrow a second time an amount roughly equal to its initial issuance of bonds to be able to pay the principal of the maturing bonds.

The autocorrelation in the success or failure of this strategy calls into question one aspect of the MB and Aaron-Reischauer plans that we have

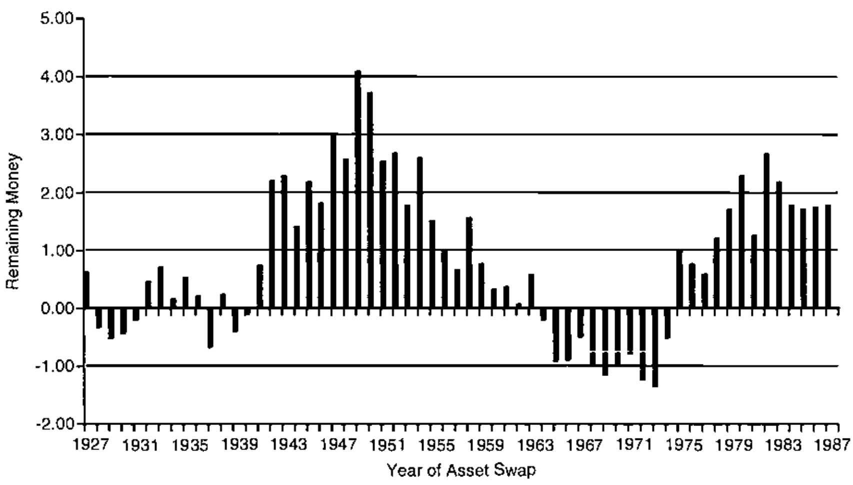

Fig. 1.5 Net money generated at the completion of a swap of ten-year inflation-indexed U.S. government bonds for common stocks

not yet modeled—namely, the fact that both plans advocate that the asset swaps take place continuously. Both plans would phase in the private securities investments of the OASDI trust fund only gradually. One could hope that some kind of "dollar-cost-averaging" phenomenon would make this less risky than making the asset reallocation all at once. However, the eleven failures in a row from 1964 through 1974 indicate that considerable risk remains even with such a gradual transition. Had social security engaged in this particular asset swap (ten-year inflation-indexed bonds for a portfolio of the S&P500) from 1964 through 1974, the financial crises that the system faced in 1977 and again in 1983 would have been considerably worsened.

We have also simulated how the asset swap of ten-year inflation-indexed bonds for common stocks would perform in the future if future returns and inflation rates were drawn from the 1927–97 set of realizations drawn randomly with a bootstrap technique. The predicted overall failure rate was 29 percent. That is, the stock portfolio failed to generate the necessary 3.5 percent real return required to pay off the obligations of the inflation-indexed bonds 29 percent of the time. Had we defined failure as the stock investments falling short of the 7 percent real rate of return assumed in scoring the MB plan in table 1.1, the failure rate would, of course, be much higher.

The results for twenty-year inflation-indexed bonds are similar. Figure 1.6 shows the counterfactual history for asset swaps from 1927 through 1977. The strategy would have failed eighteen times out of the fifty-one years, for a failure rate of 35.3 percent. Even more telling is the fact that it would have failed fifteen years in a row from 1959 to 1973. The strategy is very successful from 1932 through 1954 and again in 1975–77. This is not surprising since we know that the average real return on the stock market in 1927–97 has been well above 3.5 percent. The failure or limited success of the strategy for twenty years in a row (1955–74), however, is a big problem for its advocates. The overall failure rate when the 1927–97 inflation rates and stock market returns are scrambled with a bootstrap technique of choosing the observations is 25.4 percent.

Table 1.2 sums up what we have learned thus far. We have examined a total of four different cases for asset swaps based on ten-year bonds and the same four cases for twenty-year bonds. The bootstrap simulations are probably more informative regarding the future chances of success for the exchange than are the numbers based on the actual historical time series of returns. That is because the future will not be a precise replay of the past. The bootstrap simulation technique gives us a better handle on the probability of future success. If real stock market returns are generated in the future from the same distribution that generated the actual 1954–97 returns, and if real interest rates on bonds approximate 3.5 percent, then the failure rate of a strategy of swapping bonds for stocks is about 29

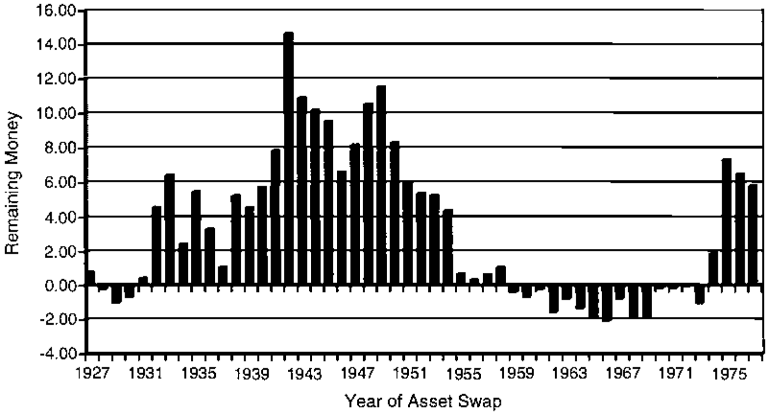

Fig. 1.6 Net money generated at the completion of a swap of twenty-year inflation-indexed U.S. government bonds for common stocks

Table 1.2 Summary of Failure Rates (%) of a Strategy of Swapping Bonds for S&P500 Stocks

Bond Maturity	Actual Interest Rates; Actual 1954–97 Time Series of Stock Market Returns	Actual Interest Rates; Bootstrap Simulations of Real Returns Chosen from 1954–97 Realizations	3.5% Inflation-Indexed Bonds; Actual 1927–97 Time Series of Inflation and Stock Market Returns	3.5% Inflation-Indexed Bonds; Bootstrap Simulations of Real Returns Chosen from 1927–97 Realizations
10 years	20.6	22.1	29.5	29.0
20 years	0	6.2	35.3	25.4

percent for ten-year bonds and 25 percent for twenty-year bonds. These figures refer to the final column of results in table 1.2. It is important to note, however, that all variations of the asset-swap simulations that we have conducted imply a considerable risk of failure.

1.4 Tax Considerations

We have ignored personal income tax issues thus far, but they deserve some consideration. The federal government collects taxes on the returns of stocks held in private hands and would recapture some of the interest payments on the bonds through tax proceeds. If the average marginal tax rate on the returns is the same for the stocks and bonds, then it is a fairly minor correction to the results presented above. For instance, assume that on the "other side" of the transaction are defined-contribution pension funds. Further, assume that, as people see that the social security trust fund is moving out of bonds and into stocks, they are willing to move in the opposite direction in their own IRA, 401(k), or other defined-contribution accounts (without a change in interest rates). If the average marginal tax rate of the pension plan holders will be 25 percent when the money is drawn out, then the government effectively owns one-fourth of the assets in the plan. If the government makes money on the swap on its own account, it loses one-fourth of its profits owing to its participation in the defined-contribution accounts (which took the other side of the asset swap) via the tax system. The failure-rate analysis of the previous section need not be modified, although the absolute size of the gains and losses is reduced by 25 percent. If the other side of the transaction consisted of corporate accounts backing defined-benefit pension promises, then the gain or loss in the funding of the pension plan presumably would be reflected in the stock of the corporation offering the defined-benefit plan. Stockholders will eventually pay capital gains taxes on their stock (sooner rather than later if the stocks are held by an actively managed mutual fund), and, once again, the government will find itself with roughly a 20 or 25 percent position on the opposite side of the swap that it thought it was engaged in. Also, once again, the failure-rate analysis of the previous section is unaffected.

We think that pensions are the most likely "other side" of the asset-swap transaction. We have also examined the case in which privately held assets are exchanged with the government. If the average marginal tax rate on the interest payments of the government exceeds the blended average marginal tax rate faced by dividends and short- and long-term capital gains from the stock returns, then the failure probabilities are slightly improved by tax considerations. The government recaptures a higher percentage of its interest payments through the tax system than it did on the return on common stocks. The effect is fairly minor, however. We were probably exploring the upper limit of the effect in assuming that the average marginal income-tax rate on interest on the bonds was 33 percent, whereas the average marginal tax rate on long-term capital gains was 20 percent. Both numbers are almost certainly on the high side, but the difference may be reasonable. With such parameters, the projected failure rate of issuing fixed-interest-rate ten-year bonds turns out to be 25.1 percent. The twenty-year-bond failure rate falls from 6.2 to 5.4 percent. The failure rate for ten-year bonds remains 21.1 percent even if somehow the stocks had paid zero taxes and the bonds faced an average marginal tax rate of 33 percent. The worst case for the government is the opposite of that, of course. If the bonds are somehow held in completely tax-free accounts and the stocks are fully taxable (33 percent on dividends and short-term capital gains and 20 percent on long-term capital gains), then the predicted failure rate with ten-year fixed-interest bonds jumps to 37.9 percent.

Our conclusion regarding taxes is that they are an important second-order consideration but that the results of the previous section—that the government issuing bonds and using the proceeds to buy stocks is very risky—is not materially altered by including the effects of taxes.

1.5 Interpretation of Failure Results

We feel that the most appropriate of our asset-swap simulations in terms of predicting the likelihood of future failure of such a policy are those that use the 3.5 percent real-interest-rate inflation-protected bonds and generate real stock returns with the bootstrap procedure. These results are relevant even if the government actually swaps nominal bonds for stocks because the real interest cost of such bonds is far more likely to approximate 3.5 percent than the negative real interest rates that prevailed in the 1970s. The result is that the asset-swap strategy has about a 25 percent chance of failure with twenty-year bonds. This means that there is a 25 percent chance that the switch into stocks will make the solvency problems of social security worse rather than better after twenty years. Of course, there is about a 50 percent chance that the strategy will not work as well as assumed by its proponents (and listed in table 1.1). That is because the 7 percent real rate of return assumed in that table is roughly the median long-run real rate of return for stocks.

Someone might take the view that, if stocks are behind after twenty years, the government can simply borrow the shortfall (or more) and buy yet more stocks. With a long enough time horizon, stocks are bound to beat out bonds, right? First, this is not quite right. It is true that the probability of a shortfall is lower the longer the time horizon. Of course, the magnitude of the shortfall in the unlikely event that stocks have a bad forty- or fifty-year run can be enormous if each intermediate shortfall is covered with more borrowed money. Second, and probably more important, the balance-sheet position of social security is economically and politically important. The social security crises of 1977 and 1983 were caused by the impending exhaustion of the OASDI trust fund. Of course, the government could have bailed out the fund by borrowing money and handing it over to social security. But that action was not seriously considered. Our point is that the trust fund would have been in even worse shape in the late 1970s and early 1980s had it been investing in stocks in the 1960s instead of in the special issue Treasury bonds. In that case, the benefit cuts and tax increases would have been even more severe than they were. One would not have been able to get through that period with the argument that we can just borrow more money and buy more stocks and sooner or later the strategy will pay off. The extra risks of stocks translate directly into riskier future benefits and taxes.

The failure probabilities that we have estimated are only as good as the modeling of the underlying process generating real stock market returns. In a number of respects, we have erred on the side of favoring the strategy. First, we have assumed that the government could engage in the swap without changing interest rates and stock market rates of return. Second, we have assumed that future stock market returns will be generated from the observed returns from 1954–97 (1927–97 in some cases). The average real rate of return on stocks from 1954–97 was 9.7 percent. The average from 1927–97 was 9.6 percent. Still, relative to either international comparisons or longer historical returns (Siegel 1994), these average returns are quite high.

For most of the results that we have presented, we have worked with three-year blocks of stock market returns. We have assumed that each of the three-year blocks observed from 1954 to 1997 is equally likely to be repeated, independent of the previous returns drawn. There is a literature on long-run mean reversion that indicates that high returns are more likely to follow low returns (and vice versa) (see, e.g., Poterba and Summers 1988; Fama and French 1988; and Campbell, Lo, and MacKinlay 1997, chap. 2). Were this true, then long-horizon investing would be safer than if returns were drawn independently from previous realizations. The statistical evidence on long-run mean reversion is not very powerful for the same reason that we have had to downplay our actual counterfactual results—there simply are not many long runs to observe in order to investigate such long-run phenomena. The empirical studies that have been conducted cannot rule out that returns are actually serially independent. With that said, all that we can say is that, if returns were generated by a process exhibiting long-run mean reversion, then the asset-swap strategy would fail less often than we have estimated. Our simulations of the asset-swap strategy using one-, two-, three-, and five-year blocks of data showed roughly the same failure probabilities. This indicates that, to the extent that it operates in under five years, mean reversion is at most a relatively weak phenomenon.

Our overall conclusion is that the simplifications in this study have not biased the results against the asset-swap policy. The relatively high mean return that we assume for equities must be at least as important an assumption as the lack of mean reversion in returns over periods greater than three years. If asked the likelihood that investing some of the central trust fund of social security in equities would worsen its finances in ten or twenty years, our best-guess answer would be that there is a 25–30 percent chance.

1.6 Is a Universal Defined-Benefit Retirement Plan Feasible?

Our conclusion thus far has the usual "no-free-lunch" ring to it. Unfortunately, one cannot eliminate the actuarial deficit of social security by swapping bonds for stocks. If it were possible, then we would also do well to eliminate the cost of the defense budget with the same maneuver. This raises the question of who bears the risks in the current social security system and who would bear the additional risks were the trust fund to invest in private securities as advocated by the MB and the Aaron-Reischauer plans. Even more fundamentally, it raises the issue of whether it is feasible for the government to offer a defined-benefit plan with universal coverage.

At one extreme, consider a defined-benefit plan offered by a financially strong company such as General Motors. The company can promise its workers a particular or "defined" benefit in retirement and fund that obligation with private stocks and bonds. The worker can accumulate a vested, safe retirement benefit (although not safe from the effects of inflation), while the firm, or, more accurately, its stockholders, bears the investment risks of the underlying portfolio. The point is that it is quite transparent how a safe benefit was created from risky investments—the risky residual claim is borne by the General Motors stockholders.

At the other extreme, consider what would happen if a self-employed person decided that he would like to provide himself a defined-benefit pension funded with stocks and bonds. Is it feasible? Can a person invest, say, $30,000 in stocks and bonds and promise himself or herself a safe retirement benefit of $3,000 per year? The answer is no since the employer bears the risk of the assets not being adequate to fund the benefits. But a self-employed person is both employee and employer. The investment risks have not been transferred at all, and, therefore, the do-it-yourself defined-benefit strategy cannot work. The only way for a self-employed person to end up with a defined-benefit-like retirement claim is to contract with an insurance company or some other third party that will provide the benefit and assume the investment risks.

Now consider a universal coverage defined-benefit social security retirement-pension system funded with risky private securities. Is it possible to provide safe retirement benefits as in the General Motors example, or is it more like the example of the self-employed person who cannot make meaningful retirement-benefit promises without the help of a third party? If everyone is in the social security system, then collectively everyone bears the investment risk, and it makes little sense to call the program a *defined-benefit* scheme. We can think of only two ways to maintain the defined-benefit nature of the system. One is to divide the population by age cohort. Conceivably, the elderly can be promised safe benefits with the young playing the role of the insurance company and bearing the investment risks of the underlying portfolio. The future payroll taxes or the future benefits of the young will be more risky if the trust fund engages in the kind of asset swap that we have been analyzing. One of the drawbacks of a defined-benefit structure is that the risks imposed on the young are not very transparent. If we maintain the defined benefits for both the old and the young, then the taxes faced by the young will be even more volatile. The point is simply that the considerable risks involved in investing in private securities, which we have documented above, have to be borne by someone. The defined-benefit nature of the system does not work very well if the investment risks are borne by those paying the personal income tax. The reason is that there is a nearly complete overlap between social security participants and those paying the personal income tax. In order to create a safe asset from a risky portfolio, the risks must be shifted to someone other than the holder of the supposedly safe claim. Shifting the risks from social security participants to taxpayers is not enough of a shift to accomplish this.

At least conceptually, the investment risks could be shifted internationally. Shiller (1998) makes this point in his theoretical paper on risks and social security. The Europeans, for instance, could play the role of the insurance company and offer Americans a safe defined-benefit pension system by assuming the investment risks themselves. There is a question of what terms they would demand to hold all the risks. Further, while it is clear that there are advantages of the young and old and the Europeans and Americans all sharing in the risks of global financial markets, it is not at all clear why one group (the young or the Europeans) should hold all the risk so that another group (the elderly Americans) can enjoy a safe, defined-benefit retirement program. The optimal risk-sharing arrangement would have everyone bearing some risk, and, thus, any optimal universal pension plan would not be strictly of a defined-benefit type.

One problem with investing the central OASDI trust fund in private securities that we have not emphasized is the danger that the investment decisions would be subject to political pressure and, therefore, would not manage risk efficiently. There would undoubtedly be pressure to invest in only American companies or only unionized companies or to allocate the investments evenly across the fifty states or not to invest in tobacco companies or those doing business in Burma. Aaron and Reischauer are confident that certain organizational structures that they propose (including the creation of a "Social Security Reserve Board") could insulate fund management from political control by elected officials. We are not so sure. It is quite possible that capital would be less efficiently allocated with government ownership of equities, meaning that the asset swap could reduce aggregate future GDP.

1.7 Conclusions

In this paper, we have examined the risk of having the central OASDI trust fund invest in private securities. We have noted that that net transaction is an asset swap with the government selling bonds to the public and using the proceeds to buy a portfolio of common stocks of equal value. The asset reallocation does not increase saving, wealth, or GDP. We examine whether it actually would improve the finances of the social security system. We adopt a convention that the asset swap is deemed successful if the stock portfolio generates sufficient cash to pay off the interest and principal of the bonds and still have money left over and a failure if the stock portfolio is insufficient to do that and the bond repayments require another round of borrowing. We have looked at several cases, including ten- and twenty-year fixed-interest-rate bonds and ten- and twenty-year inflation-indexed debt. The predicted failure rate from the bootstrap simulations of future returns ranges from 6 to 29 percent. Further, failure would be autocorrelated with long strings of possible annual failures. For instance, when we examined the counterfactual issuance of twenty-year inflation-indexed bonds, we found that the asset-swap strategy failed fifteen years in a row between 1959 and 1973. Clearly, this policy cannot reliably reduce the actuarial deficit of social security.

Individual accounts and defined-contribution pension plans certainly involve significant risks. An individual who shifts from bonds to stocks runs the same or greater risk that the move will be counterproductive after ten or twenty years. However, the defined-contribution plans have the advantage that their risks are straightforward. Further, people who want to minimize risks can invest in such safe assets as inflation-indexed bonds. The analysis of this paper has shown that an OASDI trust fund invested in private securities would generate risk for Americans, but the precise incidence of that risk would likely remain ambiguous.

References

Aaron, Henry J., and Robert D. Reischauer. 1998. *Countdown to reform: The great social security debate.* New York: Century Foundation Press.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay. 1997. *The econometrics of financial markets.* Princeton, N.J.: Princeton University Press.

Economic report of the president, 1998. 1998. Washington, D.C.: U.S. Government Printing Office.

Fama, Eugene F., and Kenneth R. French. 1988. Permanent and temporary components of stock prices. *Journal of Political Economy* 96, no. 2:246–73.

Goss, Stephen. 1999. Measuring solvency in the social security system. In *Prospects for social security reform,* ed. Olivia Mitchell, Robert J. Myers, and Howard Young. Philadelphia: University of Pennsylvania Press.

Ibbotson Associates. 1998. *Stocks, bonds, bills, and inflation: 1998 yearbook, market results for 1926–97.* Chicago.

MaCurdy, Thomas E., and John B. Shoven. 1992. Stocks, bonds, and pension wealth. In *Topics in the economics of aging,* ed. David A. Wise. Chicago: University of Chicago Press.

Poterba, James M., and Lawrence H. Summers. 1998. Mean reversion in stock prices. *Journal of Financial Economics* 22:27–59.

Schieber, Sylvester J., and John B. Shoven. 1999. *The real deal: The history and future of social security.* New Haven, Conn.: Yale University Press.
Shiller, Robert J. 1998. Social security and institutions for intergenerational, intragenerational, and international risk sharing. NBER Working Paper no. 6641. Cambridge, Mass.: National Bureau of Economic Research, July.
Siegel, Jeremy J. 1994. *Stocks for the long run.* Chicago: Irwin.
Social Security Administration (SSA). 1997. *Report of the 1994–96 Advisory Council on Social Security.* Vol. 1, *Findings and recommendations.* Washington, D.C.
———. Various years. *Annual report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Disability Insurance trust funds.* Washington, D.C.: U.S. Government Printing Office.

Comment Stephen P. Zeldes

As the debate over social security reform progresses, it is becoming increasingly obvious that many of the key issues relate to uncertainty. One question that arises is in what ways the benefit payments to the elderly should be linked to the outcomes of key economic and demographic variables in the economy—such as mortality, real wages, birthrates, unemployment, interest rates, and stock returns. A second question is what techniques should be used to evaluate, on an ex ante basis, reform proposals that involve risky outcomes.

Consider first the current evaluation techniques used by the Office of the Actuary at the Social Security Administration (SSA). The actuaries estimate the most likely outcome for a set of economic and demographic variables—these are referred to as *intermediate-cost* assumptions. The *actuarial balance,* projected path of the trust fund, and *money's-worth* measures (e.g., net present value and internal rates of return to participants) are based on these assumptions. These provide a reasonable assessment of the expected outcome for the system and for individuals. They do not attempt to incorporate the riskiness of any variables. The SSA incorporates uncertainty only by describing two alternate scenarios: *high* and *low* cost. While this does provide some information about the dispersion in possible outcomes, it is an incomplete and inadequate way of dealing with uncertainty, for the following reasons. First, this technique does not assign probabilities to the different outcomes. Second, it ignores correlations between variables, both economic and demographic. Third, when examining reforms, it ignores general equilibrium effects—the effects of these reforms on the underlying demographic and economic variables.

The SSA is careful to point out that its goal is to examine expected values and that it is not attempting to account for risk. However, many

Stephen P. Zeldes is the Benjamin Rosen Professor of Economics and Finance at the Graduate School of Business, Columbia University, and a research associate of the National Bureau of Economic Research.

who use money's-worth and other SSA statistics inappropriately interpret them as corresponding to welfare or utility rankings when in fact they do not. Since most of the reform packages that have been proposed involve investment in the stock market, through either individual accounts or a central trust fund, it is particularly problematic to use non-risk-adjusted money's-worth measures to rank the attractiveness of alternative reforms, both versus the current system and versus each other (for more on this issue, see Geanakoplos, Mitchell, and Zeldes [1999]).

MaCurdy and Shoven examine the potential risks and benefits to social security of undertaking an open market operation selling Treasury bonds and using the proceeds to purchase stocks, to be held either in the trust fund or in individual accounts. This, as well as analyzing any reform plan that involves stock market investment, requires an understanding of the equity premium: the difference between the return on equities and the return on Treasury bills or bonds. How we think about the equity premium and to what we attribute it are key to understanding many of the issues raised at this conference.

Since the stock market has a higher expected return than the bond market, this change in investment policy improves both the actuarial balance of the system and the money's worth as currently measured. It also typically increases the risk. Evaluating by how much the risk increases and assessing the trade-off of whether the higher return is worth the higher risk are difficult to do and in general require using a stochastic general equilibrium model. Such an approach has been followed in many of the papers included in this volume. These models have the advantages that they remedy the three problems described above; that is, they assign probability distributions, allow for correlations between variables, and incorporate the general equilibrium effects that can be crucial to the analysis. They also, by necessity, oversimplify the world that they are modeling—sacrificing realism at the expense of tractability.

MaCurdy and Shoven have an alternative approach with a more modest goal. Rather than analyzing the entire general equilibrium model, they analyze important properties of the joint distribution of stock and bond returns.

First, consider the basic facts. Table 1C.1 shows the means and standard deviations of annual real returns on stocks and government bonds from 1926 to 1997. The average real return on the S&P500 was 9.7 percent, as compared to 2.6 and 2.3 percent on long- and intermediate-term Treasury bonds, respectively, and 0.7 percent on short-term Treasury bills. The higher historical return on stocks is seen even more dramatically in figure 1C.1. One dollar invested in December 1925 would have grown (in constant dollars) to $203.19 if invested in the S&P500, to $4.34 if invested in long-term government bonds, and to $1.58 if invested in three-month Treasury bills.

The higher average return on stocks is clear. What about the risk? Look-

Table 1C.1 Annual Inflation-Adjusted Returns on Stocks and Government Bonds, 1926–97

Asset	Arithmetic Average Return (%)	SD (%)
S&P500	9.7	20.5
Long-term government bond	2.6	10.5
Intermediate-term government bond	2.3	7.0
Short-term Treasury bill	.7	4.2

Source: Calculations are based on data from Ibbotson Associates.

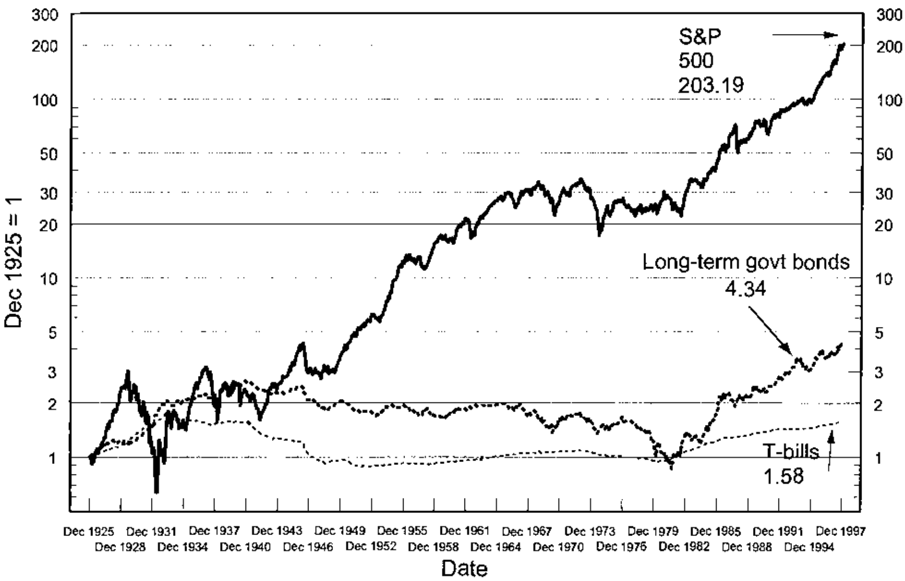

Fig. 1C.1 Stocks, long-term bonds, and Treasury bills, cumulative real returns, 1926–97

ing again at table 1C.1, we see that the standard deviations of annual real returns were 20.5 percent on stocks, 10.5 percent on long-term bonds, 7.0 percent on intermediate-term bonds, and 4.2 percent on Treasury bills. The higher standard deviation on stocks suggests that stocks are riskier than bonds, but it is not sufficient evidence, for a few reasons. First, it is possible that stocks dominate bonds in every period: for example, $r^S = r^B + \varepsilon$, where r^S is the annual return on stocks, r^B is the annual return on bonds, and ε is a random variable that has high variance but is always nonnegative (i.e., ε is ≥ 0 with probability one). In general, we need a stochastic model together with a utility function to evaluate riskiness. But, if stocks dominate bonds in this sense, then no such model is necessary. A simple glance at the data on annual returns tells us that this is not the case. There are many years when the return on stocks is negative and the return on bonds is positive. A second possibility is that risk is best measured by looking at long-horizon returns rather than annual returns. If there is sufficient mean reversion in stock returns, it is possible that, when examining long returns (e.g., twenty-year returns), stocks dominate bonds in the sense described above. An examination of twenty-year rolling annual returns indicates that there was only one twenty-year period in which stocks did not outperform both long-term bonds and Treasury bills, and even in this one period the difference was small. Some observers (e.g., Siegel 1998; and Glassman and Hassett 1998) argue that this type of evidence implies that stocks are not more risky than bonds in the long run.

There is much debate in the literature about what can be learned from the historical data about the ex ante expected return (in the past and in the future), the riskiness of stocks relative to bonds (in the past and in the future), and whether the higher expected return is/was a fair compensation for higher risk (in the past and in the future).

Table 1C.2 is my attempt to summarize the many views on this question. Where in table 1C.2 one stands is likely to be related to how attractive one finds reform proposals that involve social security investments in the stock market.

The first row of table 1C.2, the "faith-in-markets" view, argues that any equity premium that exists is fair compensation for the riskiness of stocks. Mehra and Prescott (1985) accept that expected excess returns were high in the past and that stocks were riskier but argue that they were not sufficiently risky, on the basis of a standard model of the economy, to justify the higher risk. One could interpret this result as implying either that the markets have it wrong and are providing too much compensation for stock market risk or that the standard model is wrong and is incorrectly estimating the amount of risk or expected return of the stock market. A variety of papers have attempted to explain this "equity-premium puzzle," either by arguing that ex post returns were much higher than what was expected ex ante or by arguing that with alternative utility functions risk was actually much higher than Mehra and Prescott estimated. MaCurdy and Shoven do not address the question of whether the equity premium is correctly compensating for risk, but they do challenge the view that stocks are less risky than bonds in the long run (i.e., the view represented in the rows at the bottom of the table).

MaCurdy and Shoven point out that one possible problem with simply examining long-period returns in the past is that there are not many data points, only three and a half independent twenty-year periods between 1926 and 1997—far too few to summarize the distribution of twenty-year returns adequately. They try to get around this problem of limited data by bootstrapping: they assume that the annual returns are representative of

Table 1C.2 Views about the Equity Premium

	Very High Expected Excess Return on Stocks?		Stocks Riskier than Bonds in Long Run?		Equity Premium Correct Compensation for Risk?	
View	Past	Future	Past	Future	Past	Future
"Faith-in-markets" view	Yes		Yes		Yes	Yes
Mehra-Prescott	Yes		Yes		No	
"Good-luck" theory	No (ex ante much less than ex post)		Yes		Yes	
Preference-based solutions to equity-premium puzzle	Yes		Yes		Yes	
MaCurdy-Shoven	Yes		Yes			
Siegel	Yes		No		No	
Glassman-Hassett	Yes (not quite as high as ex post)	No (once Dow hits 35,000)	No	No	No	Yes
Naive view supporting stocks for social security	Yes	Yes	No	No	No	No

the underlying probability distribution, and they draw observations randomly from the sample with replacement. They do this in two ways: first, by drawing single years of stock-return data with replacement and, second, by drawing three-year blocks of stock-return data with replacement. They then construct a thousand observations of ten- and twenty-year stock returns, match these with interest rates on ten- and twenty-year government bonds from the first year of each period, and use these constructed data to calculate a "failure rate," that is, the percentage of times for each starting year that stocks fail to outperform bonds.

I have two main criticisms of the technique used by MaCurdy and Shoven. First, they use nominal rather than real returns when performing their bootstraps. Normally, when computing excess returns on stocks relative to bonds, it makes no difference whether real or nominal returns are used as the inflation rate drops out of the differential. But with their bootstrapping method it does matter because the dates for the stock returns do not correspond to the dates covered by the interest rates. For example, the ten-year nominal interest rate in 1981 was high, reflecting high inflation expectations. MaCurdy and Shoven compare this nominal return to a series of nominal stock returns, the first year of which is 1981, but the other years of which are drawn randomly from the sample, thus likely including many years with low expected inflation. Although stock returns are unlikely to compensate for unexpected inflation, they are likely to compensate for expected inflation. Using nominal returns incorporates a mismatching that might explain why their calculated failure rate for stocks outperforming bonds is so high in the early 1980s (when expected inflation is high) and so low in the 1950s (when expected inflation is low). It would be worth reestimating their results using real returns rather than nominal ones to see whether the results are sensitive to this change.

Second, and more important, their results are unlikely to persuade those who argue that stocks are not risky over the long run. Although I do not believe this view, let me temporarily play devil's advocate. The argument that stocks are not more risky than bonds over long holding periods, even though they are over short holding periods, depends crucially on long-term mean reversion in stock returns. In other words, strings of bad returns must tend to be followed by strings of good returns. Because the bootstrap method randomly draws three-year blocks of data, it will by construction not pick up these long-term correlations. In their simulated data, a three-year period of low returns is, by construction, as likely to be followed by periods of low returns as it is by periods of high returns. Therefore, it cannot possibly capture the long-term mean reversion that is crucial to reducing the long-run riskiness of stocks.

To see this more clearly, suppose that we were looking at the riskiness, not of stocks, but of thirty-year zero-coupon bonds. For simplicity, assume that inflation was always zero. Because of real-interest-rate risk, the one-year-holding-period returns on these bonds are risky. However, high one-year returns in one part of the sample must be followed by low one-year returns later in the sample, and the thirty-year real return on these bonds is completely riskless. Were the MaCurdy-Shoven bootstrap technique applied to these bonds, it would miss the long-term mean reversion, and it would lead to the erroneous conclusion that thirty-year real returns on these bonds are very risky.

With that said, let me emphasize that the MaCurdy-Shoven conclusion that there is risk in stocks relative to bonds, even in the long run, is probably correct. Certainly, a market test would lead to this conclusion. Consider a $1 million portfolio in the S&P500. Were we to ask a large insurance company how much it would charge to guarantee that, thirty years from today, the portfolio would be worth at least as much as the amount in an initial $1 million portfolio invested in thirty-year bonds, it would probably charge a substantial amount and would certainly not provide this option for free. This is, of course, subject to the critique that insurance companies do not understand the risks of the market either.

Suppose that we accept the MaCurdy-Shoven conclusion that stocks are riskier than bonds. The next questions are whether they are sufficiently riskier to justify the high equity premium and whether social security could benefit by doing a swap of stocks for bonds. These are questions that MaCurdy and Shoven do not address. As described earlier, they require a stochastic general equilibrium model. It may very well be that there are tails of the distribution of stock returns that involve low returns and that the marginal utility of consumption is very high in those states of the world. Ignoring or downplaying those states may make the equity premium look like a giveaway.

The question of whether social security could benefit from an open market swap is not fully answered by examining the riskiness of stocks. Even if stocks completely dominated bonds over long horizons and participants in the asset markets did not realize this, we would need to balance the losses of the sellers of equity against the gains to social security participants. On the other hand, even if the equity premium is a fair compensation for the risk borne by current stock market participants, there may still be a role for social security to improve risk sharing. This could be done, for example, by providing exposure to current stock market returns to those without it: adults who are nonstockholders, children, and those not yet born (for more on this, see, e.g., Geanakoplos, Mitchell, and Zeldes [1998, 1999] and most of the papers in the present volume).

Overall, this is a helpful paper that contributes in an important way to the debate on the relative riskiness of stocks and bonds in the long run and the attractiveness of investing social security funds in the stock market. This paper and others in this volume appropriately focus the debate on the risk-return trade-offs inherent in such a change.

References

Geanakoplos, John, Olivia S. Mitchell, and Stephen P. Zeldes. 1998. Would a privatized social security system really pay a higher rate of return? In *Framing the social security debate: Values, politics, and economics,* ed. R. Douglas Arnold, Michael J. Graetz, and Alicia H. Munnell. Washington, D.C.: Brookings.

———. 1999. Social security money's worth. In *Prospects for social security reform,* ed. Olivia S. Mitchell, Robert J. Meyers, and Howard Young. Philadelphia: University of Pennsylvania Press.

Glassman, James K., and Kevin A. Hassett. 1998. Are stocks overvalued? Not a chance. *Wall Street Journal,* 30 March, A18.

Mehra, Rajnish, and Edward C. Prescott. 1985. The equity premium: A puzzle. *Journal of Monetary Economics* 15, no. 2:145–62.

Siegel, Jeremy J. 1998. *Stocks for the long run: The definitive guide to financial market returns and long-term investment strategies.* 2d ed. New York: McGraw-Hill.

Discussion Summary

Robert Shiller remarked that long-term bonds are quite risky because of the institutional feature that a single person, currently Alan Greenspan, can affect their value substantially, for instance, in response to political pressure. *Stephen Zeldes* concurred and noted that it is therefore important to consider three-asset models, with short-term bills, long-term bonds, and stocks.

James Poterba noted that dynamic hedging could be used to at least partially hedge the risk of holding equities. He also wondered whether the criticism of the discussant, Zeldes, concerning the use of nominal rather than real returns in the bootstrap simulation method would have significant implications. He noted that the correlation between equity returns and inflation is low. In response to this discussion, *John Shoven* noted that Zeldes's suggestion to repeat the simulation exercise with real returns would be incorporated in the final version of the paper.

John Campbell remarked that the Siegel view has an internal inconsistency. If stocks are less risky in the long run owing to mean reversion, then one must also accept that expected stock returns are time varying. But, in that case, one cannot logically recommend a buy-and-hold strategy, as Siegel does. Instead, mean reversion in stock returns and the associated time variation in expected equity returns imply that the optimal portfolio strategy involves some market timing. One cannot simultaneously assert that stocks are less risky in the long run owing to mean reversion and that a buy-and-hold strategy is optimal since the latter requires independently and identically distributed (i.i.d.) returns or constant expected returns. Campbell recommended that MaCurdy and Shoven be more explicit about their view on this data-generating process in the paper since it matters so much for the results.

David Cutler asked whether these experiments would not correspond more to borrowing against tax revenues than to using long-term debt since the rate of return seems to be based on wages, not on real interest rates in the economy. *Martin Feldstein* concurred but also noted that the authors are substituting stocks for bonds in the portfolio explicitly. There are two distinct policies: one in which assets are simply being swapped, without changing taxes (as suggested by Aaron and Reischauer and by Ball), and one that involves raising taxes.

When asked by *Feldstein* whether he had discussed the paper with Aaron, *Shoven* replied that Aaron was not pleased with the results. *Feldstein* responded that this was not surprising. He added that their argument would be that, although there is a small probability of a shortfall, they choose to take this risk on behalf of the taxpayer. Feldstein concluded that Shoven is well hedged, having written papers about the benefits of investing the trust fund in equities (e.g., MaCurdy and Shoven 1992) and now showing that this also involves substantial risk.

Shoven responded that one has to make the distinction between a defined-benefit and a defined-contribution plan in this respect. While a defined-contribution plan is typically perceived as risky, the defined-benefit-based plan is purported to be safe, yet is not, and becomes, moreover, significantly more risky when investing in equities.

Shoven acknowledged that any form of long-run mean reversion is not captured by the bootstrap experiment that takes an i.i.d. view. He noted that Poterba's question about the possibility of using dynamic hedging to intervene in the face of adverse conditions was worth pursuing.

Finally, Shoven emphasized that the political rhetoric appears to describe the current system as involving both defined benefits and defined taxes. This is misleading and even more incorrect when investing the trust fund in risky securities.

Reference

MaCurdy, Thomas E., and John B. Shoven. 1992. Stocks, bonds, and pension wealth. In *Topics in the economics of aging,* ed. David A. Wise. Chicago: University of Chicago Press.