

NBER WORKING PAPER SERIES

PRIVATE SAVING AND
PUBLIC POLICY

B. Douglas Bernheim

John Karl Scholz

Working Paper No. 4215

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 1992

This paper was prepared for a conference on "Tax Policy and the Economy," held on November 17, 1992 in Washington D.C. The second author is grateful to the National Science Foundation, which provided research support through Grant Number SES-9211553. We also gratefully acknowledge the work of Robert Avery and Arthur Kennickell, who developed a clean copy of the 1983-1986 Survey of Consumer Finances, and provided extensive documentation. This paper is part of NBER's research program in Public Economics. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #4215
November 1992

PRIVATE SAVING AND
PUBLIC POLICY

ABSTRACT

The evidence presented in this paper supports the view that many Americans, particularly those without a college education, save too little. Our analysis also indicates that it should be possible to increase total personal saving among lower income households by encouraging the formation and expansion of private pension plans. It is doubtful that favorable tax treatment of capital income would stimulate significant additional saving by this group. Conversely, the expansion of private pensions would probably have little effect on saving by higher income households. However, these households are more likely to increase saving significantly in response to favorable tax treatment of capital income. Currently, eligibility for IRAs is linked to an AGI cap, and pension coverage is more common among higher income households than among low income households. The most effective system for promoting personal saving would have precisely the opposite features. Extending tax incentives for saving to higher income households is problematic. We discuss three competing policy options, IRAs with AGI caps, the universal IRA, and the Premium Saving Account (PSA). Our analysis reveals that the PSA system is a more cost-effective vehicle for providing saving incentives to all households, particularly those in the top quintile of the income distribution.

B. Douglas Bernheim
Department of Economics
Princeton University
Princeton, New Jersey 08540-1021
(609) 258-4021
and NBER

John Karl Scholz
Department of Economics
University of Wisconsin-Madison
1180 Observatory Drive
Madison, Wisconsin 53706
(608) 262-5380
and NBER

1. Introduction

Since the mid-1980s, low rates of national saving in the United States have generated an enormous amount of concern among both economists and policy makers. Proposals to address these concerns fall into two broad categories. One category consists of policies designed to increase public saving; the other consists of policies that are intended to promote private saving. The first category is synonymous with deficit reduction, while the second includes tax incentives, pension policy, and strategies for discouraging the use of private debt. Some economists argue that deficit reduction is the most reliable and efficacious method of increasing national saving (see, for example, Summers [1985]), while others maintain that it is essential to restore adequate rates of saving in the private sector (see, for example, Bernheim [1991]). To evaluate the merits of strategies that target private saving, one must resolve two issues. First, aside from the obvious fact that private saving is one component of national saving, are there reasons to be concerned about the rate of private saving? Second, are there any effective and reliable methods of promoting private saving?

This paper investigates several factual matters bearing on both of these questions. Four central findings emerge from our analysis. First, many households do not save enough to provide themselves with adequate financial security, and as a result will be forced to accept significantly reduced standards of living during retirement. This phenomenon is especially prevalent when the head of the household lacks a college education. Second, patterns of asset accumulation among those without college education bear little or no resemblance to the patterns that emerge from standard economic theories. In contrast, those with a college education not only save more adequately for retirement, but also generally behave in a way that more closely

resembles "text book" life cycle planning. Third, consistent with this second finding, employer provided pensions do not appear to displace other personal saving in cases where the head of the household lacks a college education. However, for college educated households, pensions do appear to crowd out private saving. Fourth, it is likely that high income households respond more vigorously to tax incentives for saving than do moderate and low income households.

These findings have important implications for public policy. To the extent that many households prepare poorly for retirement, there is cause to be concerned about the rate of personal saving, per se. Although lower income households may not respond significantly to tax incentives, it should be possible to stimulate rates of saving among this group by encouraging the creation and expansion of private pension plans. For high income households, the implications are reversed: although pensions displace other forms of saving, tax incentives for saving are probably efficacious.

Because eligibility for deductible contributions to Individual Retirement Accounts is subject to an adjusted gross earnings (AGI) cap, lower income households currently receive the most favorable tax incentives for saving. Conversely, households with higher levels of income and education are much more likely to be covered by private pensions. Thus, the current system appears to be designed in a way that minimizes the impact of public policy on personal saving.

Unfortunately, it is difficult to modify the current system in a way that would extend tax incentives to higher income households without raising a host of new problems. The most common proposals are either to drop the AGI cap on IRAs or to design some new, "universal IRA" system without an AGI cap. The efficacy of these proposals is questionable. For many high income households, saving for retirement may already exceed the proposed contribution

limits; in that case, an IRA does not offer any reward for incremental saving. To the extent that IRAs simply generate windfall gains for many wealthy individuals, the system would be perceived as inequitable. Finally, the expansion of IRA eligibility could significantly reduce public saving (increase federal deficits) and thereby defeat the purpose of the proposal.

An alternative method of extending tax incentives for saving to higher income households is through a system of Premium Savings Accounts, or PSAs (see Bernheim and Scholz [1992]). In brief, a household becomes eligible to contribute to a PSA only when its total saving exceeds a minimum threshold (the floor); beyond that point, incremental saving may be placed into a PSA, up to a cap (the ceiling). These floors and ceilings are tied to AGI: higher income households must save more before becoming eligible.

If one believes that it is desirable to provide high income households with tax incentives for saving, does the PSA proposal offer an attractive alternative to universal IRAs? To answer this question, we undertake a comparison of the two proposals. For each proposal, we calculate an index of effectiveness and a measure of windfalls received by higher income individuals. We also assess the relative budgetary costs of these proposals. Our analysis suggests that, relative to a universal IRA system, the PSA proposal would significantly enhance incentives to save among higher income households even as it would reduce both budgetary costs and windfalls to the wealthy.

This paper is organized as follows. Section 2 presents evidence on the adequacy of personal saving. Section 3 examines evidence on the effectiveness of pension policy. Section 4 discusses the impact of tax incentives for saving. Finally, section 5 concludes.

2. The Adequacy of Personal Saving

According to common wisdom, Americans consume too much and save too little. This impression is largely traceable to widely publicized statistics on aggregate personal saving. International comparisons reveal that U.S. households save significantly less than their foreign counterparts. Between 1980 and 1989, Americans saved 7.4% of disposable personal income, compared with 11.4% for OECD Europe and 16.0% for Japan (OECD [1989]). Moreover, since the mid-1980s, the rate of household saving in the U.S. has been well below its historical average (see figure 1).

Although these statistics raise legitimate concerns, they do not provide definitive evidence of a problem. As measured, personal saving excludes capital gains. Thus, households could in principle accumulate wealth at a rapid rate even when measured saving is low. Rates of personal saving can also vary across both time and countries for reasons unrelated to the adequacy of saving considered from the perspective of individual households. To understand this second point, consider the following hypothetical example. Envision two countries, A and B, that are identical in all respects, except that the elderly make up a larger fraction of the population in A than in B. Since households tend to accumulate wealth prior to retirement and decumulate wealth thereafter, one would expect to observe a higher rate of aggregate personal saving in country B. Indeed, in an economy with no growth in either population or productivity, dissaving by retirees could completely offset saving by workers: one could in principle observe no personal saving in the aggregate, regardless of how well individual households prepared for retirement. Thus, ultimately, one can only judge the adequacy of personal saving by examining microeconomic data on the behavior of individual households.

Generally, the available evidence suggests that American workers have prepared poorly for retirement. According to Diamond [1977], during the 1960s roughly 40% of married couples and over 50% of unmarried individuals reported that, after retirement, they received no money income from assets. At age sixty, nearly 30% of middle-class individuals lacked sufficient wealth to replace two years' worth of income. Similarly, Hammermesh [1984] concluded that, during the 1970s, most elderly individuals failed to accumulate resources sufficient to sustain their accustomed standards of living. Indeed, consumption shortly after retirement exceeded the highest sustainable level of consumption by an average of 14%. Hammermesh found that, within a few years of retirement, most households were forced to reduce their expenditures substantially.¹ Hausman and Paquette [1987] also documented substantial declines in consumption for men who involuntarily retired earlier than expected.

Unfortunately, these microeconomic studies of the adequacy of personal saving are now somewhat dated. In addition, they tend to use fairly arbitrary standards of income or consumption replacement to judge the adequacy of wealth accumulation. In this study, we adopt a different strategy. Using an elaborate model of household decision making, we simulate asset accumulation profiles.² We then compare these simulated profiles with actual profiles estimated from more recent household survey data.

The simulation model uses a life cycle planning framework to establish the criteria for household decision making. In this framework, a household's standard of living at any point

¹Other economists have reached somewhat more optimistic conclusions. See Kotlikoff, Spivak and Summers [1982].

²Development of this model was sponsored by Merrill Lynch & Co., Inc. For a detailed description of the model, see Bernheim [1992b].

is taken to be a function of its material consumption per capita, and its overall well-being depends upon both its present and future standards of living. Loosely speaking, the life cycle framework implies that a household should accumulate wealth sufficient to finance a standard of living during retirement that is consistent with its standard of living prior to retirement.

The model takes as inputs certain descriptive data concerning a household, including age, birth cohort, current earnings, pension coverage, education, marital status, and gender (if single). Based on these characteristics, the model imputes an earnings history, a family composition history, and mortality probabilities. The earnings history is extrapolated from cross-sectional age-earnings profiles and is adjusted to reflect the economy's baseline wage growth. Similarly, the family composition history is constructed from estimates of the relationship between household size and various household characteristics. Mortality probabilities are obtained from gender-specific actuarial tables. The model also incorporates important macroeconomic factors, such as interest rates, inflation, economic growth rates, aspects of the federal, state, and local income tax systems, and social security statutes.

The model generates consumption and asset trajectories through an iterative procedure. The first step in this procedure determines the household's asset accumulation plan for the first year of its economic life (taken to be age 26). The choice of a plan is based in part upon forecasts of its future economic prospects.³ Decisions taken in the initial year determine the level of retirement assets that the household carries into the following year. The second step of the procedure determines the household's asset accumulation plan in its second year of economic life. Since the household may learn more about its economic prospects between the

³Forecasts of macroeconomic variables are calculated using simple econometric models.

first and the second years, its forecasts may change. Consequently, the household's second year plan may deviate from the plan that it envisioned in the first year. For example, if during the second year rates of interest rise unexpectedly, the household may decide to take advantage of this development by saving more than it had planned previously. Financial decisions taken in the second year establish the level of retirement assets carried over into the third year. This procedure is repeated, until the current year is reached. The household's asset accumulation plan for the current year supplies the forecast for future asset and consumption trajectories.⁴

It is important to understand that the model describes the accumulation of assets only for retirement. There are, of course, many reasons to save. Households should prepare for the possibility of illness, layoff, disability, death, and other risks for which they are imperfectly insured. In addition, most households accumulate resources to pay for large expenses such as college tuition or the purchase of an automobile. In the current study, no attempt is made to estimate the extent to which households should save for these other reasons.

Figures 2 and 3 depict the output of an illustrative simulation run. This particular simulation was constructed for a household with the following characteristics: age 27 (as of 1991), some college education, married, current earnings of \$60,540, and the primary earner is covered by a private pension. Figure 2 displays, in constant 1991 dollars, this household's

⁴It should be noted that, in each year, the model treats all forecasts of future prospects as if these prospects were known with certainty. Yet, in each decision year, additional information is acquired and forecasts are revised. It would be preferable to employ a simulation model that would explicitly recognize uncertainty concerning future conditions and incorporate this uncertainty into consumer decision making. However, this alternative approach poses considerable technical problems. The simplified approach adopted here probably has the effect of understating the ideal level of asset accumulation since, in the presence of uncertainty, households would also have a precautionary motive for saving.

trajectory of after-tax earned income (including pensions and social security) and its consumption. Note that after-tax earnings rise steeply early in life. Earnings growth continues at a reduced level until the individual reaches age 55, at which point it begins to fall. After retirement, earned income consists of social security and private pension benefits. Since pensions are less than perfectly indexed for inflation, real benefits decline gradually over time.

As a direct consequence of its rapid earnings growth early in life, the household saves nothing for retirement prior to age 30. In fact, during its 20s, the household would prefer to borrow against future income in order to consume more than its current after-tax income. However, the model does not permit these households to obtain loans since they lack collateral. Between ages 30 and 80, the consumption trajectory is relatively flat. This property reflects the household's preference for a stable standard of living. However, during the 30s and 40s, consumption is somewhat elevated relative to the 60s and 70s. This pattern results from changes in household composition: between the ages of 30 and 50, the typical household incurs significant child-rearing costs. Consumption declines rapidly after age 80 until, at age 101, it matches after-tax retirement benefits. Falling survival probabilities drive this end-of-life decline. Since there is a relatively low probability of reaching age 90, the household would prefer to sacrifice the standard of living that it would receive at age 90 (if it survived that long) in favor of a higher standard of living earlier in life.

Figure 3 depicts the associated trajectory of retirement assets. Assets are accumulated at an increasing rate early in life. They reach a peak at retirement, and then steadily decline until they are exhausted at age 100.

Actual asset trajectories are estimated using the 1983 and 1986 waves of the Survey of

Consumer Finances (SCF). The Federal Reserve in conjunction with other Federal agencies sponsored the SCF, and it is recognized as one of the best available sources for data on household balance sheets.⁵

We restrict attention to married couples for which the husband was fully employed and between the ages of 25 and 64 in 1986. A total of 1314 households in the SCF sample satisfy these criteria. Our measure of accumulated net worth includes stocks and mutual funds, bonds, checking and savings accounts, IRA and Keogh accounts, money market accounts, certificates of deposit, profit sharing and thrift accounts, the dollar cash value of whole life insurance, and other financial assets, as well as net equity in property (other than primary residences) and business assets, less credit card debt, consumer debt, and other debt.⁶ This measure of net worth excludes all assets and liabilities associated with homes and vehicles. This is appropriate since households appear to have a strong aversion to paying living expenses during retirement by reducing home equity (see Venti and Wise [1989]); moreover, it seems likely that few individuals save for retirement by accumulating wealth in the form of vehicles.

We divide the sample into two subgroups, based upon whether or not the husband completed college. Our sample includes 474 husbands who completed college and 840 husbands who did not complete college. Education is of interest for two reasons. First, it may be related to differences in behavior, either because education enhances an individual's ability to formulate coherent long range plans, or because those who pursue more education do so precisely because

⁵See Avery and Elliehausen [1988] and Avery and Kennickell [1988] for a more complete discussion of the SCF.

⁶Accumulated wealth for 1983 is expressed in 1986 dollars using the Consumer Price Index.

they are more likely to be concerned about the future. Second, education is highly correlated with income. Adjusted for age, the median earnings of households in which the husband is college educated are roughly 57% higher than the median earnings of households in which the husband is not college educated. Since IRA eligibility is subject to an AGI cap, differences in saving behavior across income categories is of particular interest. Although it might seem more natural to divide the sample by income in order to examine these differences, that approach poses the practical difficulty that income varies with age. For example, a household with earnings of \$50,000 at age 27 is probably much wealthier over the course of a lifetime than a household with earnings of \$55,000 at age 55. For this purpose, one can think of education as a proxy for permanent income.

As a first step in our analysis, we examine changes in wealth between 1983 and 1986. In order to control for differences in resources across households, we focus our discussion on the ratio of the change in wealth to total wage income. We divide our sample into subgroups based upon age (25 to 29, 30 to 34, 34 to 39, 40 to 44, 45 to 49, 50 to 54, 55 to 59, 60 to 64). For each of these subgroups, we calculate the median change-in-wealth-to-wages ratio (adjusting for sampling weights). The use of medians rather than means is important, because the distribution of assets is highly skewed. Although mean wealth is quite high for many population subgroups, this fact tells us very little about the adequacy of saving for the typical household within these groups. Rather, it primarily reflects the extreme behavior of a few unrepresentative households. In contrast, median wealth is not influenced by extreme outliers.

We then simulate asset accumulation trajectories for households that are representative of each population subgroup. The household characteristics (wage income, years of education)

chosen for these simulations are based on within-subgroup population medians. Two simulations are conducted for each population subgroup: one assumes that the primary earner is covered by a private pension, while the other assumes that the primary earner is not covered by a private pension.

When comparing estimated and simulated trajectories, it is important to bear in mind that the simulations focus on preparation for retirement as the sole motive for saving. Unfortunately, when examining the data, one cannot determine whether particular assets were accumulated for retirement or for some other purpose. Aside from excluding residences and vehicles, we make no attempt to discern saving motives. Consequently, the comparison between estimated trajectories and simulated trajectories may provide an overly optimistic picture of the adequacy of retirement saving.

Figure 4 depicts results for the "no college" sample. "Actual" refers to median change-in-wealth-to-wage ratios based upon the SCF; "Sim/no pen" indicates simulated change-in-wealth-to-wage ratios for a representative household without pension coverage for the primary earner; and "Sim/pen" denotes simulated change-in-wealth-to-wage ratios for a representative household with pension coverage for the primary earner. Note that simulated change-in-wealth-to-wage ratios rise steeply with age. This occurs for two reasons. First, wages increase more rapidly than consumption during most of an individual's working life (refer back to Figure 2). Second, reinvested capital income rises as the household accumulates assets. In contrast, the actual change-in-wealth-to-wage ratios do not vary significantly with age. By the time the household reaches middle age, simulated asset accumulation exceeds actual accumulation by a wide margin. Although actual asset accumulation is higher than the simulated profiles at ages

27 and 32, this result is of little consequence, since the data reflect saving for a variety of purposes aside from retirement. Overall, Figure 4 suggests that, between 1983 and 1986, households without a college education saved far less than our simulation model predicts.

Figure 5 depicts results for the "college" sample. The contrast between Figures 4 and 5 is remarkable. In cases where the head of household has completed college, both simulated and actual change-in-wealth-to-wage ratios rise steeply with age. Moreover, simulated asset accumulation tracks actual asset accumulation remarkably well. Taken at face value, Figure 5 suggests that college educated households saved adequately for retirement between 1983 and 1986.

Low rates of asset accumulation do not necessarily imply that households are inadequately prepared for retirement. In principle, a household with high initial assets in 1983 (relative to the simulated trajectory) could save little between 1983 and 1986 and still remain above the simulated asset trajectory in 1986. Conversely, high rates of asset accumulation do not necessarily imply that households are adequately prepared for retirement, since these households may have started out well below the simulated trajectory. To evaluate the adequacy of retirement preparation, one must therefore examine levels of wealth in addition to changes in wealth.

Consequently, as a second step in our analysis, we examine levels of wealth for 1986. We proceed exactly as in the first step, except that we focus on wealth-to-wage ratios, rather than on change-in-wealth-to-wage ratios. Figures 6 and 7 depict the results. These figures reinforce the lessons drawn from Figures 4 and 5. In particular, actual asset trajectories are far below simulated trajectories for the "no-college" sample, while actual and simulated trajectories

track each other rather well for the "college" sample. Indeed, since the majority of college educated workers are covered by private pensions, it appears that actual asset trajectories are actually above simulated trajectories for those with a college education.

Although it is tempting to conclude that inadequate saving is largely confined to those without a college education, this conclusion must be tempered by two considerations. First, as indicated by Figure 1, personal saving declined sharply after the 1983 to 1986 period. Using a sample of relatively young individuals (ages 25 through 44) surveyed in early 1992, Bernheim [1992a] found much more pervasive evidence of inadequate saving. Second, the model probably understates the amount of wealth that each household ought to accumulate. The most obvious reason for this bias is that the simulations envision retirement planning as the sole motive for saving. In addition, it is quite likely that the model overstates mortality probabilities (since it does not make any allowance for the fact that average life expectancy is projected to increase), understates the importance of health and long term care costs for the elderly, and fails to consider the effects of mounting economic pressures that may force Congress and employers to scale back existing retirement benefits.

Before proceeding to the next section, it is important to discuss one potential criticism of this analysis. Some readers may be inclined to argue that our investigation sheds more light on the nature of tastes than on the adequacy of saving for retirement (see, for example, Lazear [1992]). After all, any measure of adequacy is subjective. If a household has chosen to save relatively little, who are we to argue? Presumably, the household has its own best interests at heart.

We are not persuaded by this argument. Even in the context of the traditional life cycle

hypothesis, individuals may face incentives that lead to inefficiently low levels of saving. For example, as individuals save more, they may lose the benefits from social insurance programs (Hubbard, Skinner, and Zeldes [1992]), risk the loss of eligibility for college scholarships (Feldstein [1992]), or reduce assistance from other members of the extended family (Bernheim and Stark [1988]).⁷ These considerations may be particularly important for lower income individuals.

It is also possible that, for some households, the life cycle hypothesis may not adequately characterize saving behavior. For some economic decisions one can argue that, through trial and error, an individual eventually learns to behave in a way that is consistent with utility maximization. This argument, however, is less persuasive in the context of life cycle saving. Each individual accumulates resources for retirement only once; there is no opportunity to learn from one's mistakes. Moreover, the life cycle saving decision is extraordinarily complex, in that it requires an individual to contemplate labor earnings, investment strategies, macroeconomic trends, and a vast assortment of risks, all over a very long time frame. It would be surprising if the average individual, in isolation, with no practice and little or no training, would act as a perfectly rational, farsighted utility maximizer. Manski [1990] discusses the circumstances in which learning from others can take the place of personal experience. Even with good role models and reference groups, however, it is difficult to imagine that households do not deviate from their optimal life cycle consumption profiles.

In recent years, a number of economists have argued against the view that individuals act

⁷Bernheim [1992c] provides a more comprehensive discussion of the factors that could produce an inefficiently low level of saving within the standard life cycle framework.

as if they maximize an intertemporal utility function, and have instead emphasized the importance of behavioral concepts such as habit, mental accounting, and self-control (see, for example, Shefrin and Thaler [1988]). Behavioral theories allow for the possibility that individuals may regret their bad habits and lack of foresight after the fact. Consequently, the notion of inadequate saving has a clear normative meaning within the context of these theories.

The evidence offered in this section suggests that most households without a college education do not behave in a manner consistent with optimal life cycle planning. These households save little relative to their simulated asset trajectories. Moreover, their estimated and simulated trajectories do not even exhibit the same qualitative patterns (refer once again to Figure 4).⁸ Given the behavioral differences between households with and without college degrees, an important question arises: is it possible to design policies that effectively stimulate saving at all levels of education and income?

3. Pension Policy

In recent years, asset accumulation in private pension plans has accounted for a substantial fraction of personal saving (see Bernheim and Shoven [1988] and Bosworth, Burtless, and Sabelhaus [1991]). This observation raises the possibility that policies affecting private pensions may have powerful effects on aggregate personal saving. Whether or not these effects would actually materialize in practice depends upon the manner in which workers would respond to an expansion of private pension coverage. Economic theory suggests that such an expansion

⁸It could be argued that low income individuals save less relative to simulated saving because they discount the future more heavily. Although a higher discount factor would reduce saving, it would not alter the qualitative features of the asset trajectories (unless discounting was high enough to prevent the accumulation of any significant assets for retirement).

would simply crowd out other forms of personal saving. The simulation results presented in section 2 illustrate this principle. However, previous studies of personal saving generally fail to find evidence for the hypothesis that private pensions reduce other forms of personal saving (see, for example, the review in Shefrin and Thaler [1988], especially pages 622-624). Depending upon whether one credits the theoretical analysis or the empirical studies, one can reach dramatically different conclusions about the effect of pension policy on aggregate personal saving.

The analysis of section 2 raises an intriguing possibility: if the behavior of those with college education conforms to the predictions of the life cycle hypothesis, while the behavior of those without a college education does not, then perhaps private pensions do displace personal saving among the college educated, but do not displace personal saving among the rest of the population. In this case, pension policy could be an effective tool for stimulating total personal saving, as long as it is primarily used to provide incentives for expansion of coverage among lower income (less educated) workers.

To investigate the effect of pensions on household saving, we estimate equations that explain the median value of the wealth-to-wage ratio (henceforth WWRAT) as a function of the husband's age (AGE), total household earnings (EARN), and a dummy variable summarizing the husband's private pension coverage (PENS).⁹ We employ a cubic function of AGE to allow for flexible age-wealth trajectories. Since earnings may be related to the shape of the asset trajectory as well as to its absolute level, we also interact EARN with AGE. For similar

⁹Since our object is to explain median wealth rather than mean wealth, we employ quantile regression techniques (least absolute deviations), rather than more traditional regression techniques (least squared deviations).

reasons, PENS is interacted with AGE in most specifications.

For purposes of comparison with the previous literature, it is useful to begin with results for the entire data sample (all households, irrespective of educational attainment). The following estimated equation is consistent with the view that pensions fail to displace other forms of personal saving:

$$\begin{aligned} \text{WWRAT} = & -4.72 + 0.411 \text{ AGE} - 0.115 (\text{AGE})^2/10 + 0.106 (\text{AGE})^3/10^3 \\ & (1.34) \quad (0.098) \quad (0.023) \quad (0.017) \\ & - 0.140 \text{ EARN}/10^4 + 0.581 \text{ EARN} \times \text{AGE}/10^6 - 0.0004 \text{ PENS} \\ & (0.024) \quad (0.045) \quad (0.0379) \end{aligned}$$

Note that the coefficient of the pension dummy is economically trivial and statistically insignificant.

The absence of a relationship between pension coverage and personal saving is sometimes interpreted as providing evidence that standard economic models do not faithfully represent the typical household's decision making process. Proponents of the standard model, however, argue that the absence of a pension effect is a statistical artifact. Pension coverage is not random. A worker who is concerned about retirement may turn down job offers from employers who fail to provide attractive pension benefits. Conversely, a worker who gives little thought to retirement may be unwilling to accept a job that provides pension coverage if this offer entails a reduction in current disposable income. If these hypothetical facts are indeed descriptive of behavior, then those who are inclined to select jobs with pension coverage will tend to save more than those who are inclined to select jobs without pension coverage. For any particular individual, a private pension may displace other forms of saving; however, in the data, this pattern may be obscured by the fact that pension coverage is correlated with the inclination to

save more. In other words, a sample selection effect may offset the saving-displacement effect.

Unless we can determine whether the absence of a pension effect is a behavioral phenomenon or a statistical artifact, we cannot predict the impact of a change in pension policy on personal saving. We suggest a method of distinguishing between these two hypotheses, based upon the following argument. It is certainly possible that, by coincidence, a sample selection effect exactly offsets the saving-displacement effect. However, the magnitude of the saving-displacement effect should vary systematically with age; specifically, the difference between the accumulated assets of those with pension coverage and those without pension coverage should increase with age. Although the sample selection effect may also vary with age, it seems highly unlikely that these two effects would exactly offset each other at every age.

We therefore estimate a new equation, in which an interaction term involving PENS and AGE is added to the list of explanatory variables. The estimated relationship is as follows:

$$\begin{aligned} \text{WWRAT} = & -4.64 + 0.406 \text{ AGE} - 0.114 (\text{AGE})^2/10 + 0.105 (\text{AGE})^3/10^3 \\ & (1.47) \quad (0.108) \quad (0.025) \quad (0.019) \\ - & 0.137 \text{ EARN}/10^4 + 0.576 \text{ EARN} \times \text{AGE}/10^6 - 0.098 \text{ PENS} + 0.310 \text{ PENS} \times \text{AGE}/10^2 \\ & (0.026) \quad (0.049) \quad (0.169) \quad (0.373) \end{aligned}$$

Taken individually, the coefficients of the pension variables lack statistical significance. In addition, the hypothesis that both coefficients equal zero is entirely consistent with the data.¹⁰ In other words, private pension eligibility is not systematically related to either the level or the shape of the asset accumulation trajectory. Although, in principle, this could still reflect the offsetting effects of asset displacement and sample selection, it seems implausible that these

¹⁰The F-statistic for the joint hypothesis that both coefficients equal zero is 0.78. One would obtain an F-statistic of this magnitude or greater roughly 46% of the time, even if the true coefficients were zero.

effects would offset each other at every age. It appears more likely that there is little or no behavioral link between pension eligibility and personal saving.

Of course, the preceding results do not distinguish between households on the basis of education. In light of our previous findings, it is clearly important to make this distinction. We therefore estimate wealth trajectories separately for the two subgroups (college educated, not college educated) described in section 2.

We obtain the following equation for households without a college education:

$$\begin{aligned} \text{WWRAT} = & -3.60 + 0.313 \text{ AGE} - 0.086 (\text{AGE})^2/10 + 0.077 (\text{AGE})^3/10^3 \\ & (1.51) \quad (0.111) \quad (0.026) \quad (0.020) \\ - & 0.110 \text{ EARN}/10^4 + 0.406 \text{ EARN} \times \text{AGE}/10^6 + 0.008 \text{ PENS} - 0.125 \text{ PENS} \times \text{AGE}/10^2 \\ & (0.044) \quad (0.088) \quad (0.176) \quad (0.395) \end{aligned}$$

Note that the coefficients of the pension variables are even smaller, both in terms of economic and statistical significance, than they were in the previous equation (based on the full sample). In order to illustrate the implications of this equation, we extrapolate asset trajectories for hypothetical households with pension coverage and households without pension coverage.¹¹ Figure 8 exhibits these trajectories. Note that pension eligibility bears little or no relationship to the path of the wealth-to-wage ratio.

In contrast, we estimate the following equation for households with a college education:

¹¹For the purpose of this calculation, the household's earnings are taken to be constant at \$30,000. This figure is close to median age-adjusted earnings for households without a college education.

$$\begin{aligned}
\text{WWRAT} = & 5.78 - 0.502 \text{ AGE} + 0.126 (\text{AGE})^2/10 - 0.080 (\text{AGE})^3/10^3 \\
& (4.58) \quad (0.329) \quad (0.077) \quad (0.058) \\
+ & 0.062 \text{ EARN}/10^4 - 0.021 \text{ EARN} \times \text{AGE}/10^6 + 0.899 \text{ PENS} - 2.83 \text{ PENS} \times \text{AGE}/10^2 \\
& (0.048) \quad (0.091) \quad (0.563) \quad (1.25)
\end{aligned}$$

Note that the coefficients of the pension variables in this equation are much more significant, both economically and statistically. The data decisively reject the hypothesis that both of these coefficients equal zero.¹² Note also that the signs of all the other coefficients in the equation for college educated households are exactly the reverse of the signs of these coefficients in the equation for households without college education. Clearly, the behavior of these two groups differs markedly.

In order to illustrate the implications of the estimated equation for college educated households, we extrapolate asset accumulation trajectories for hypothetical households with private pensions and without private pensions.¹³ Figure 9 depicts these trajectories. Note that those individuals who are eligible for pensions accumulate resources at a significantly slower rate than those individuals without pensions. Remarkably, at age 62, the gap between the assets of these two groups is almost identical in magnitude to the gap that emerges from our simulations (figure 7). These patterns are strongly consistent with the view that private pensions displace other personal saving for college educated households. It is unlikely that the observed relationship between pension coverage and saving results from spurious factors, since such

¹²The F-statistic for this hypothesis is 5.60, which is significant at the 99% level of confidence.

¹³For the purpose of this calculation, the household's earnings are taken to be constant at \$50,000. This figure is close to age-adjusted median earnings for households with a college education.

factors would presumably also have produced the same patterns for less educated households.¹⁴ These results suggest that previous studies may have failed to find a significant saving displacement effect simply because they did not distinguish between households on the basis of education.

The contrast between figures 8 and 9 points to a clear and important conclusion: private pensions displace personal wealth accumulation only when the head of the household is college educated. This is consistent with the findings of section 2 on the adequacy of personal saving. Indeed, our evidence tends to support a more general conclusion: college educated households behave in the manner predicted by standard economic theories of saving, while households with less education do not.

It should be emphasized that past and current policies have been more successful at stimulating the expansion of pension coverage among college educated workers than among those with less education. Analysis of the SCF data reveals that 75.2% of husbands with college degrees are covered by private pensions, in comparison to 55.7% of husbands without college degrees. In other words, the current system is quite effective at providing pensions to those who reduce other saving in response, but is substantially less effective at providing coverage to those individuals for whom pensions would represent incremental saving.

Proposals to encourage or require portability of private pensions could have the effect of expanding pension coverage among lower income workers. The absence of portability substantially reduces the benefits of pension coverage for workers with relatively little job

¹⁴It is worth mentioning that there is some evidence of a small sample selection effect: the trajectory for households with pensions starts out slightly above the trajectory for households without pensions. However, this effect is not the dominant pattern in the data.

stability. On average, these workers also have relatively low earnings. Portability would enhance the attractiveness of pensions to workers with little job stability, and thereby increase the likelihood that these workers would obtain pension coverage.

A more drastic proposal would require employers to provide all workers with private pension coverage. Although this requirement would not be related to income, it would have a far greater impact on pension coverage among low income workers than on pension coverage among high income workers (roughly three-fourths of whom are already covered). In response to this requirement, those employers who had not previously provided pension coverage would probably reduce other forms of employee compensation. Our analysis suggests that newly covered workers with high earnings would simply adjust other personal saving to offset this change. In contrast, newly covered workers with lower earnings would reduce consumption. In effect, mandatory pension coverage would force lower income households to increase saving.

Our analysis also raises concern about the recent growth of 401K plans. Participation in these plans is entirely voluntary. Many employers have substituted 401K plans for more conventional plans in an effort to reduce operating expenses. We would not expect this trend to have a significant impact on the total amount of saving done either by or on behalf of high income workers. However, the elimination of compulsory contributions may significantly depress total saving by or on behalf of lower income workers.

4. Tax Incentives for Saving

The most commonly discussed strategies for stimulating personal saving entail reductions in capital income taxation. Economic theory suggests that households will respond to a higher after-tax rate of return by increasing future consumption relative to current consumption.

However, the increase in anticipated future net worth resulting from higher rates of return may actually induce households to save less. Indeed, empirical estimates of the interest elasticity of saving (which measures the sensitivity of saving to the after-tax rate of return) vary widely (see, for example, Boskin [1978], Summers [1981], and Hall [1988]).

Current policies provide saving incentives primarily to lower income households. Eligibility for deductible Individual Retirement Accounts (IRAs), for example, is subject to an AGI cap. The existence of an AGI cap raises an important question: does the interest elasticity of saving vary systematically across income classes?

Simulations based upon the model described in section 2 suggest that higher income households should be much more responsive than lower income households to a change in the after-tax rate of return. According to these simulations, a 35 year old, high school educated couple with pension coverage and annual earnings of \$30,000 in 1991 should have saved roughly 1.5% of its earnings.¹⁵ A permanent, one percentage point rise in the before-tax rate of return modestly increases this figure to 1.6%. In contrast, a 35 year old, college educated couple with pension coverage and annual earnings of \$50,000 in 1991 should have saved roughly 4.8% of its earnings. The same permanent, one percentage point rise in the before-tax rate of return increases this figure by a much larger amount (both absolutely and proportionately), to 5.5%. Similar results hold for couples without pensions. For the representative high-school educated couple, saving would fall from 5.6% of earnings to 4.5% of earnings in response to the higher rate of return; for the college educated couple, it would fall less (both absolutely and

¹⁵In this calculation, the rate of saving is defined as saving above and beyond reinvestment of capital income, divided by after-tax wage income.

proportionately), from 9.2% to 8.8%.¹⁶

The contrast between the simulated responses of college educated individuals and high school educated individuals becomes even more striking when one factors in the statistics on pension coverage mentioned at the end of section 3. Averaging across those individuals with pensions and those individuals without pensions, the simulations imply that saving by college educated households would increase by 10.2% in response to a permanent one percentage point increase in the before-tax rate of return, while the saving of households without college degrees would fall by 4.5%. Consequently, a policy that provides tax incentives for saving exclusively to lower income households excludes those households that are most likely to increase saving in response to the policy; indeed, it is conceivable that such policies could actually reduce aggregate personal saving.

It is important to emphasize that this positive relationship between income and the interest elasticity of saving results from a natural economic consideration, rather than from some peculiar feature of the simulation model. It is natural to assume that, when planning for the future, most households are concerned first and foremost with assuring themselves of some minimum standard of living. As lifetime resources increase, households have more discretion to allocate resources over time in a way that increases consumption above and beyond this minimum

¹⁶These simulations imply that the interest elasticity of saving tends to be higher when the household has private pension coverage. The explanation for this phenomenon is straightforward. An increase in the rate of return reduces the present discounted value of future income; in that sense, it makes the household poorer, and reduces current consumption (this effect was originally noted by Summers [1981]). Since pension income is received after retirement, its present discounted value is more sensitive to the rate of return than is the present discounted value of future (pre-retirement) earnings. Thus, those individuals with pensions are more likely than those individuals without pensions to reduce current consumption in response to an increase in the rate of return.

standard. Saving to provide for minimum consumption is, in effect, saving for a fixed target. It is well known that an individual who saves to achieve some target will reduce saving in response to an increase in the rate of return (see, for example, Bernheim and Shoven [1988]). In contrast, discretionary saving to finance consumption over and above the target responds positively to an increase in the rate of return. For low income households, saving to achieve the minimum consumption target is probably far more important than saving to fund incremental consumption. Thus, target saving dominates the simulated behavior of these households, and produces a low or negative interest elasticity of saving. On the other hand, for high income households, saving to fund incremental consumption is probably far more important than saving to achieve the minimum consumption target. Consequently, incremental saving dominates the simulated behavior of these households and produces a high interest elasticity of saving. In the Appendix, we develop this argument mathematically and demonstrate for a simple model that the interest elasticity of saving rises with income.

Throughout this section, we have assumed that household behavior accords with standard economic theories. The preceding sections call this premise into question. However, this observation does not undermine our conclusion. We have found that the behavior of college educated (high income) households does correspond to the predictions of standard theories; consequently, for this group, it is likely that one would observe a substantial interest elasticity of saving. On the other hand, we have also found that the behavior of households without a college education (those with lower income) does not conform to standard economic theories. Although this finding reduces our faith in the applicability of our simulation results, it does not reverse our conclusions concerning the interest elasticity of saving. The notion that households

will respond to a change in the after-tax rate of return is predicated upon the assumption that households rationally anticipate and plan for future economic contingencies. To the extent that this assumption proves incorrect, there is no particular reason to believe that low income households will respond to a change in the after-tax rate of return.

Most current proposals to provide tax incentives for saving are patterned after IRAs. IRAs were established as part of the 1974 Employee Retirement Income Security Act, to give workers who were not covered by employer-provided pension plans added incentives to accumulate resources for retirement. In 1981, IRA eligibility was extended to all taxpayers. Subsequently, the Tax Reform Act of 1986 curtailed the tax-deductibility of IRA contributions for high income households.

Two prominent current policy initiatives would reverse the direction of the 1986 reforms and extend tax incentives for saving to households in higher income brackets. The Bush Administration's Family Saving Accounts (FSAs) would allow single individuals with incomes below \$60,000 and married couples with incomes below \$120,000 to make contributions of up to \$2,500 per person (not including children) to qualified accounts. The FSA proposal is an example of a "back-loaded" system: contributions are nondeductible, but accumulated funds are not taxed upon withdrawal. An alternative proposal, the Bentsen-Roth "super-IRA," would allow contributions of up to \$2,000 per person (not including children) to either a traditional or back-loaded IRA.

Unfortunately, there are sound conceptual reasons to doubt the effectiveness of extending eligibility for IRA-style accounts to higher income households. First, contributions are capped. Under the current system, a single taxpayer, for example, can make no more than \$2,000 in tax-

deductible contributions. For any taxpayer who would have saved more than \$2,000 in the absence of IRAs, the availability of an IRA does not affect the costs or the benefits that result from an additional dollar of saving, and therefore provides no incentive on the margin for the taxpayer to increase saving. In such cases the IRA constitutes a "give-away" of public funds, and its principal effect is to reduce federal tax receipts. In addition, the IRA may actually induce the taxpayer to spend more on current consumption, since it increases his/her total after-tax resources. For both of these reasons, the IRA would contribute to a lower rate of national saving. These concerns are of little significance for low income households, because few of them would save more than \$2,000 in the absence of the program. It is far more likely that high income households would save more than the contribution limit. Thus, a standard IRA-style scheme may be a particularly ineffective method of providing high income households with tax incentives for saving.

Second, even if a taxpayer would not (in the absence of IRAs) have saved more than the IRA contribution limit in a given year, he/she could take full advantage of the IRA deduction either by drawing down previously accumulated assets or by borrowing. Indeed, the 1991 Tax Guide for College Teachers devotes a full page to the issue, "What if You're Short of Cash to Fund Your IRA?" (pp. 229-230). The Guide describes an IRS private letter ruling that allows households to finance their IRAs by borrowing. Contributions funded either by shifting existing assets or by borrowing do not increase household saving. Instead, they depress national saving by reducing Federal tax receipts, and add to the Federal budget deficit. Once again, high income households, who possess greater wealth, financial sophistication, and access to credit markets, are more likely than lower income households to engage in borrowing or asset shifting,

and thereby defeat the purpose of the program.

Empirical evidence on the efficacy of IRAs is mixed. Gale and Scholz [1992] find little evidence that IRAs stimulated household saving between 1983 and 1986. Venti and Wise [1986, 1987, 1990, 1991] and Feenberg and Skinner [1989] argue that most IRA contributions during this period represent net increases in household saving. Joines and Manegold [1991] conclude that the effects of IRAs on household saving are unlikely to be as large as the estimates of Venti-Wise and may be as small as the estimates of Gale-Scholz.

An alternative proposal to promote household saving, based upon "Premium Saving Accounts" (PSAs), is described in Bernheim and Scholz [1992]. A PSA system would require each taxpayer to save in total some fixed amount (the "floor") before becoming eligible to make contributions to a tax-favored account. For each dollar saved in excess of the floor, the taxpayer would be eligible to contribute one additional dollar to the tax-favored account, up to some limit ("the ceiling"). These floors and ceilings would vary with adjusted gross income (AGI) and certain types of capital income. As with IRAs, capital income accrued on balances held in PSA accounts would be exempt from taxation.¹⁷

The use of both floors and ceilings would create "windows" of program eligibility. Consider, for example, a husband and wife with a combined AGI of \$80,000. They might face a floor of \$8,000, and a ceiling of \$12,000. Should they save less than \$8,000 in the corresponding tax year, they would not be eligible to make any contributions to a tax favored

¹⁷With this essential structure, a PSA system could be either front-loaded or back-loaded. Penalties could be established to lock funds into tax-favored accounts for relatively short periods of time (e.g. seven years), or until some age close to retirement (perhaps age 59 1/2). Accounts could be established for specific purposes (e.g. retirement, purchase of a home, college education), or the accounts could be unrestricted.

account. If, on the other hand, they saved \$9,500, they would be eligible for favorable tax treatment on \$1,500. If they saved more than \$12,000, then they would be eligible to make the maximum contribution of \$4,000 (the difference between \$12,000 and \$8,000).

The most important and distinctive feature of a PSA system is that floors and ceilings would vary with AGI. Eligibility windows could be positioned to maximize, within each income class, the number of households receiving tax breaks on the marginal dollar of saving. Thus, higher income taxpayers would not be deprived of tax incentives for saving; rather, they would simply be required to save larger fractions of their incomes than lower income taxpayers before becoming eligible for the program. It would also be much more difficult for households to take advantage of tax-favored PSA accounts by shifting assets or by borrowing, because eligibility would be based upon total saving. An individual cannot increase total saving by shifting assets from one account to another or by borrowing in order to invest.¹⁸

To implement a PSA system, one needs to measure saving. Bernheim and Scholz [1992] propose the following measure:¹⁹

1. Net purchases of assets (i.e. for assets on which investors receive capital gains and losses, total purchases minus total sales),

plus

2. The January 1-to-January 1 change in cash account balances (for example, bank accounts),

¹⁸The administrative feasibility of monitoring total saving for each taxpayer is discussed in Bernheim and Scholz [1992].

¹⁹Many economists would define saving as the change in the stock of wealth between two points in time. If one adopts this definition, then saving is very hard to measure -- one would need to assess the market value of all assets every year. The definition used in the text represents a compromise between economic logic and administrative feasibility.

minus

3. The January 1-to-January 1 change in total debt (for example, mortgages and consumer credit).

In effect, saving is defined as the incremental personal resources that an individual diverts into investments in any given year, over and above reinvested capital gains.²⁰

If this definition of saving is employed, then it is also important to adjust each taxpayer's eligibility floors and ceilings upward by the amount of capital income other than capital gains. In the absence of such an adjustment, the system would distort investors' choices among assets, causing them to tilt their portfolios towards assets that produce current income, rather than capital gains. See Bernheim and Scholz [1992] for a detailed discussion of this issue.

In the remainder of this section, we evaluate the effects of three distinct strategies for stimulating household saving: an IRA-like program with an AGI cap (henceforth referred to as the "standard IRA" system), an IRA-like program without an AGI cap (henceforth referred to as the "universal IRA" system), and a PSA system. We compare the cost effectiveness of extending tax incentives for saving to higher income taxpayers through universal IRAs and PSAs.

Table 1 contains illustrative eligibility schedules for a PSA system. We selected these particular schedules after examining the empirical distribution of saving. We restricted attention to a class of simple schedules, and chose the schedules that provide maximum saving incentives.²¹

²⁰Note that it is possible to compute this measure of saving without assessing the value of unrealized capital assets since, by definition, unrealized gains are fully reinvested.

²¹An eligibility schedule belongs to this class if the floor is set equal to zero up to some level of AGI, beyond which the floor rises linearly with AGI. We also studied more complex

The schedules define eligibility windows for each level of AGI. In order to facilitate comparisons with IRAs, we have adopted window widths of \$2,000 per year for single households, \$2,250 per year for married couples with one earner, and \$4,000 per year for married couples with two earners. The lower end of the window (the floor) is determined by a two-part calculation. First, compute the value of an algebraic expression involving AGI. We refer to this value as the "unadjusted floor" -- it is identical for all taxpayers with the same level of AGI. Second, add capital income other than capital gains to obtain the "adjusted floor" (or simply, "the floor"). Table 1 indicates, for example, that a dual-earner married couple with an AGI of \$30,000 and no capital income would have a floor of \$0 and a ceiling of \$4,000. In contrast, a couple with an AGI of \$120,000 and dividend and interest income of \$2,000 would have a floor of \$16,362 ($.167 \times \$86,000 + \$2,000$) and a ceiling of \$20,362. Figure 10 graphs the proposed eligibility schedule for married couples. Note that because the typical American household saves very little, the floor is zero for lower income households.

The standard and universal IRA systems differ from the PSA proposal in that the IRA systems anchor the eligibility window at \$0 for all income classes, and no adjustment is made for capital income. The standard system phases out deductible contributions for couples with incomes between \$40,000 and \$50,000 and for single taxpayers with incomes between \$25,000 and \$35,000. The universal system allows all households to make deductible contributions.²²

schedules, but discovered that it was difficult to improve significantly upon the simple schedule described in the text.

²²The IRA-like proposals we simulate are superior to actual IRA schemes, because, in practice, IRA schemes are susceptible to tax arbitrage strategies involving borrowing and asset shifting, which our simulations do not capture.

We will compare these plans on the basis of three criteria. The first criterion is a measure of effectiveness. Specifically, for each plan, we estimate the number of households that would receive a higher after-tax rate of return on the incremental dollar of investment. We refer to these households as the IMPACT GROUP. Our second criteria is a measure of wasteful subsidization. Specifically, for each plan, we estimate the number of households that would make the maximum eligible contribution to a tax-favored account while continuing to receive the unsubsidized after-tax rate of return on the incremental dollar of investment. We refer to these households as the NO-IMPACT GROUP. Our third criterion is also a measure of wasteful subsidization: we calculate the budgetary cost of subsidizing the NO-IMPACT GROUP. We refer to this cost as the GIVEAWAY.

In evaluating these plans, it may be useful to consider other criteria, such as the ratio of the number of households in the IMPACT GROUP to the DOLLARS OF GIVEAWAY. A "bang-for-the-buck" ratio of this type would provide some indication of the cost-effectiveness of each proposal.

To calculate the size of the IMPACT GROUP, NO-IMPACT GROUP, and DOLLARS OF GIVEAWAY, we must predict the extent to which each household would participate in each plan. Our predictions are predicated on three behavioral assumptions. First, we assume that no household would make a contribution to any tax-favored account unless it would contribute to a universal IRA. Second, if a household would make a contribution to a universal IRA, then that same household would also, if eligible, contribute to either a standard IRA or a PSA. The justification for these first two assumptions is compelling: both proposals are identical to the universal IRA system, except that eligibility is more restricted.

Our third assumption concerns the magnitude of contributions. For each alternative proposal, we assume that an eligible household would make the maximum allowable contribution when two conditions are satisfied: the household would contribute to a universal IRA, and its total saving would exceed its eligibility ceiling. This third assumption is more problematic than the others. Saving in a tax favored account may be an imperfect substitute for other forms of saving due, perhaps, to restrictions on early withdrawal. Thus, it is conceivable that some households would contribute less than the maximum amount even when these conditions are satisfied. In practice, it would be very difficult to identify these households with available data.

Obviously, our three assumptions are helpful only if we know whether or not households are inclined to make contributions to universal IRA accounts. During the period of universal IRA eligibility, this inclination can be inferred from actual behavior. Consequently, we base our calculations upon a sample of households surveyed in 1983 and 1986 for which data on IRA participation are available.

More specifically, we use the 1983 and 1986 waves of the SCF. Key variables are constructed as follows. Income is defined as the average of total household income for 1983, 1984, and 1985. Our measure of saving corresponds to the definition proposed above. To calculate the net change in assets exclusive of capital gains or losses, we calculate (by asset category) the average constant contribution needed to generate the balance in 1986, given the observed balance in 1983 and the average rate of return that prevailed during this period. Our asset categories include stocks and mutual funds, bonds, IRA and Keogh accounts, money market accounts, certificates of deposit, profit sharing and thrift accounts, the dollar cash value of whole life insurance, and other financial assets (for a more detailed discussion of the

calculations, see Gale and Scholz, 1992). Changes in cash account balances (saving and checking accounts) and total debt are measured directly.

Table 2 provides background information on the saving propensities of households that contributed to IRA accounts. The top panel shows that, on average across income groups, total saving was positive for only 59.5 percent of participating couples. Among those households that saved, average saving ranged from \$5,693 in the lowest income quintile to \$90,296 in the highest income quintile. These averages, however, are affected by the relatively small number of households that saved very large amounts. The third column of the table provides the saving of the typical (median) participating household with positive total saving in each income quintile. These figures range from \$3,840 in the bottom quintile to \$19,695 in the top quintile. The typical participating couple with positive saving accumulates \$7,128 per year. The corresponding figure for single households is \$3,579.²³

Table 3 compares the effect of the policies on married couples. The top panel shows the size of the IMPACT GROUP. Overall, the PSA system provides incentives to 2.4 million couples, roughly 90% more than the IRA with AGI restrictions and 30% more than the universal IRA. The difference is particularly pronounced in the top income quintile. By definition, the IRA with AGI caps ignores these households. Relative to the universal IRA, the PSA increases the number of couples receiving marginal incentives in the top income quintile by 122.4%.²⁴

²³Burless, Bosworth, and Sabelhaus [1991] and Venti and Wise [1992] present tabulations on household saving from several microdatasets, including the Survey of Consumer Finances. The numbers on saving presented in Tale 2 differ from these other tabulations in that they apply only to those households saving positive amounts.

²⁴Since high income households are more likely to increase saving in response to tax incentives, there is some justification for selecting an eligibility schedule that would maximize

Since, in this sample, over 60% of positive household saving is attributable to households in the top quintile of the income distribution, this improvement is particularly important.

The bottom two panels of Table 3 measure the NO-IMPACT group and the cost of these ineffective subsidies. The calculations show, for example, that the PSA system would reduce the number of households in the NO-IMPACT group by 1.75 million (28.2%) and would reduce Federal expenditures on ineffective subsidies by \$2.0 billion (34.0%), relative to the universal IRA. In terms of cost-effectiveness, the PSA system increases the ratio of the IMPACT GROUP to the GIVEAWAY by 96.5% overall, and by 287.2% (that is, by a factor of almost four) in the top income quintile. The IRA with AGI caps also effectively reduces ineffective subsidies and budgetary cost, but it achieves this reduction by excluding the very households that are most likely to respond to tax incentives.

Table 2 reveals that participating single households saved considerably less than married households. Nevertheless, the gains from adopting the PSA system would still be substantial for single households. Table 4 indicates that the size of the IMPACT group would increase by 15.1% overall, and by 235.4% in the top income quintile, relative to the universal IRA proposal. Moreover, both the size of the NO-IMPACT group and the GIVEAWAY would fall relative to the universal IRA proposal. The result is a 49.7% increase in overall cost-effectiveness (the

the IMPACT GROUP in the top income quintile, rather than the total IMPACT GROUP. This would be accomplished with a schedule that sets the floor equal to zero as long as AGI is below \$45,000, and increases the floor by 17.7 cents for each dollar of AGI over \$45,000. Relative to the universal IRA, this PSA schedule increases the number of couples receiving marginal incentives in the top income quintile by 266.1%. Surprisingly, the use of a schedule that maximizes the IMPACT GROUP for the top quintile, rather than the total IMPACT GROUP, reduces the total IMPACT GROUP only slightly (by roughly 9,000 households). However, it increases the NO-IMPACT GROUP by 9.7%.

ratio of the IMPACT group to GIVEAWAY), and a 551.3% increase in cost effectiveness for the top income quintile, relative to the universal IRA proposal.

These comparisons of IRA and PSA proposals incompletely incorporate behavioral responses. For example, households saving strictly less than the PSA eligibility floor might increase their saving in order to become eligible for PSAs. It is also possible that these proposals will differentially affect saving for psychological reasons. Indeed, those who believe that IRAs significantly stimulated private saving often suggest psychological explanations, such as the following: (1) IRAs were aggressively marketed by financial institutions; (2) IRAs provided taxpayers with an effective way of earmarking funds for retirement, thereby facilitating the division of funds into distinct "mental accounts," some of which are psychologically more difficult to invade; (3) the existence of a sizable early withdrawal penalty effectively locked saving into IRAs, thereby helping households to impose self-discipline; and (4) the IRA eligibility limit provided households with a saving "target." Empirical evidence suggests that the fourth effect was particularly important (many households contributed exactly \$2,000, the widely publicized contribution limit, even in cases where they were actually eligible to contribute more). The PSA system, like a universal IRA program, would preserve all these features. Indeed, the fourth effect would probably be strengthened with PSAs, since the proposal would provide many taxpayers with more ambitious targets.

5. Conclusions

The evidence presented in this paper supports the view that many Americans, particularly those without a college education, save too little. Our analysis also indicates that it should be possible to increase total personal saving among lower income households by encouraging the

formation and expansion of private pension plans. It is doubtful that favorable tax treatment of capital income would stimulate significant additional saving by this group. Conversely, the expansion of private pensions would probably have little effect on saving by higher income households. However, these households are more likely to increase saving significantly in response to favorable tax treatment of capital income. Currently, eligibility for IRAs is linked to an AGI cap, and pension coverage is more common among higher income households than among low income households. The most effective system for promoting personal saving would have precisely the opposite features.

Extending tax incentives for saving to higher income households is problematic. We have discussed three competing policy options, IRAs with AGI caps, the universal IRA, and the Premium Saving Account (PSA). Our analysis reveals that the PSA system is a more cost-effective vehicle for providing saving incentives to all households, particularly those in the top quintile of the income distribution.

Pension policies and tax policies do not exhaust the full range of strategies for stimulating personal saving. One particular class of policies not discussed here merits further attention. An accumulating body of evidence, including that contained in sections 2 and 3 of this paper, suggests that the behavior of many households (particularly those with lower income) are not well-described by traditional economic theories. Consequently, it may be possible to design more effective policies by educating the population or by exploiting the psychology of saving. The Japanese appear to have had considerable success with such a strategy during the post-war period (see Horioka [1988] and Bernheim [1991]). The development of a framework for analyzing policies of this type is an important research priority.

References

- Avery, Robert B. and Gregory E. Elliehausen (1988), "1983 Survey of Consumer Finances: Technical Manual and Codebook," mimeo, Board of Governors of the Federal Reserve System, August
- Avery, Robert B. and Arthur B. Kennickell (1988), "1986 Survey of Consumer Finances: Technical Manual and Codebook," mimeo, Board of Governors of the Federal Reserve System, November
- Bernheim, B. Douglas (1991), The Vanishing Nest Egg: Reflections on Saving in America, Twentieth Century Fund
- Bernheim, B. Douglas (1992a), "Is the Baby Boom Generation Preparing Adequately for Retirement? Summary Report," mimeo, Princeton University, August
- Bernheim, B. Douglas (1992b), "Is the Baby Boom Generation Preparing Adequately for Retirement? Technical Report," mimeo, Princeton University, August
- Bernheim, B. Douglas (1992c), "A Discussion of Two Papers on Saving," prepared for the Economics of Aging Conference in Caneel Bay, mimeo, Princeton University, May
- Bernheim, B. Douglas and John Karl Scholz (1992), "Premium Saving Accounts: A Proposal to Improve Tax Incentives for Saving," mimeo, Princeton University and the University of Wisconsin--Madison, September
- Bernheim, B. Douglas and John B. Shoven (1988), "Pension Funding and Saving," in Zvi Bodie, John B. Shoven, and David A. Wise (eds), Pensions in the U.S. Economy, Chicago: University of Chicago Press and NBER, 85-111
- Bernheim, B. Douglas and Oded Stark (1988), "Altruism Within the Family Reconsidered: Do Nice Guys Finish Last?" American Economic Review 78, December, 1034-1045.
- Boskin, Michael (1978), "Taxation, Saving, and the Rate of Interest," Journal of Political Economy, April, 3-27
- Bosworth, Barry, Gary Burtless, and John Sabelhaus (1991), "The Decline in Saving: Evidence from Household Surveys," Brookings Papers on Economic Activity, 1, 183-241.
- Diamond, Peter A. (1977), "A Framework for Social Security Analysis," Journal of Public Economics, December, 275-298
- Feenberg, Daniel R. and Jonathan Skinner (1989), "Sources of IRA Saving," in Lawrence Summers, ed., Tax Policy and the Economy, MIT Press, 25-46

- Feldstein, Martin (1992), "College Scholarship Rules and Private Saving," NBER Working Paper No. 4032, March
- Gale, William G. and John Karl Scholz (1992), "IRAs and Household Saving," mimeo, UCLA and University of Wisconsin--Madison
- Hall, Robert (1988), "Intertemporal Substitution in Consumption," Journal of Political Economy, 96, 337-357
- Hamermesh, Daniel S. (1984), "Consumption During Retirement: The Missing Link in the Life-Cycle," Review of Economics and Statistics, February, 1-7
- Hausman, Jerry A. and Lynn Paquette (1987), "Involuntary Early Retirement and Consumption," in G. Burtless, ed., Work, Health and Income Among the Elderly, Brookings Institution, 151-175.
- Horioka, Charles Y. (1988), "Why is Japan's Private Saving Rate So High?" in R. Sato and T. Negishi (eds), Recent Developments in Japanese Economics, Tokyo: Harcourt Brace Jovanovich Japan/Academic Press
- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes (1992), "Precautionary Saving and Social Insurance," mimeo, Columbia University, April
- Joines, Douglas H. and James G. Manegold (1991), "IRA and Saving: Evidence From a Panel of Taxpayers," mimeo., U.S.C.
- Kotlikoff, Laurence J., Avia Spivak, and Lawrence H. Summers (1982), "The Adequacy of Saving," American Economic Review, December, 1056-1069
- Lazear, Edward P. (1992), "Some Thoughts on Savings," Hoover Institution Working Paper #E-92-17, June.
- Manski, Charles F. (1990), "Dynamic Choice in a Social Setting: Learning From the Experience of Others," Journal of Econometrics, forthcoming.
- Organization for Economic Cooperation and Development, (1989), Economic Outlook, Historical Statistics
- Shefrin, Hersh M. and Richard H. Thaler (1988), "The Behavioral Life Cycle Hypothesis," Economic Inquiry, 26, October, 609-643
- Summers, Lawrence (1981), "Capital Taxation in a Life Cycle Growth Model," American Economic Review, September

- Summers, Lawrence (1985), "Issues in National Savings Policy," National Bureau of Economic Research Working Paper #1710, Cambridge, Massachusetts, September
- Venti, Steven F. and David A. Wise (1986), "Tax-Deferred Accounts, Constrained Choice and Estimation of Individual Saving," Review of Economic Studies, LIII, 579-601
- Venti, Steven F. and David A. Wise (1987), "IRAs and Saving," in Martin Feldstein, ed., The Effects of Taxation on Capital Accumulation, Chicago: University of Chicago Press and NBER, 7-48
- Venti, Steven F. and David A. Wise (1989), "Aging, Moving, and Housing Wealth," in David A. Wise (ed), The Economics of Aging, Chicago: University of Chicago Press and NBER
- Venti, Steven F. and David A. Wise (1990), "Have IRAs Increased U.S. Saving?: Evidence from Consumer Expenditure Surveys," Quarterly Journal of Economics, 105 (August), 661-698
- Venti, Steven F. and David A. Wise (1991), "The Saving Effect of Tax-Deferred Retirement Accounts: Evidence from SIPP," in B. Douglas Bernheim and John B. Shoven eds., National Saving and Economic Performance, Chicago: University of Chicago Press and NBER, 103-128
- Venti, Steven F. and David A. Wise (1992), "Government Policy and Personal Retirement Saving," Tax Policy and the Economy, 6, 1-41

Appendix: Household Income and the Interest Elasticity of Saving

Consider a two-period model, in which an individual maximizes

$$u(c_1, c_2) = \alpha^{-1} [(c_1 - m)^\alpha + \beta (c_2 - m)^\alpha]$$

subject to

$$c_1 + \frac{c_2}{1+r} = w + \frac{w\mu}{1+r},$$

where c_t denotes consumption in the t -th period of life ($t=1,2$), r is the after-tax rate of return, w is earnings in the first period, $w\mu$ is earnings in the second period (the reader should construe this as pension income, where $\mu < 1$ is the replacement rate), m is minimum consumption, and α is the (constant) elasticity of marginal utility with respect to consumption. Optimization requires

$$c_2 - m = \gamma (c_1 - m),$$

where

$$\gamma = [\beta(1+r)]^{\frac{1}{1-\alpha}}.$$

Substituting this expression into the budget constraint and solving for first period saving ($s = w - c_1$), we obtain

$$s = (w - m) - \frac{(1+r)(w - m) + (w\mu - m)}{1 + r + \gamma}.$$

From this equation, we derive an expression for the interest elasticity of saving:

$$e_s = \frac{r}{s} \frac{ds}{dr} = - \frac{r \left[1 - \left(\frac{\gamma \xi}{(1+r)(1-\alpha)} \right) \right]}{1+r+\gamma},$$

where

$$\xi = \frac{(1+r)(w-m) + (w\mu - m)}{\gamma(w-m) - (\mu w - m)}.$$

It is then apparent that

$$\frac{de_s}{dw} = \left[\frac{r\gamma}{(1+r)(1-\alpha)(1+r+\gamma)} \right] \frac{d\xi}{dw}.$$

It follows that the sign of de_s/dw is the same as the sign of $d\xi/dw$. But

$$\frac{d\xi}{dw} = \frac{m(1+r+\gamma)(1-\mu)}{[(\gamma-\mu)w + m(1-\gamma)]^2},$$

which is positive if and only if $\mu < 1$. Thus, as long as income replacement during retirement is less than complete, the interest elasticity of saving rises with income.

Table 1

Deductible Contribution Formula: Married Couples

If your income is	Deductible Qualified Contribution Floor (Added to Capital Income)	Deductible Qualified Contribution Ceiling (Added to Floor) ²⁵
Less than \$34,000	0	\$2,250 or \$4,000
Greater than \$34,000	$.167 \times (\text{Income} - 34,000)$	\$2,250 or \$4,000

Deductible Contribution Formula: Single Households

If your income is	Deductible Qualified Contribution Floor (Added to Capital Income)	Deductible Qualified Contribution Ceiling (Added to Floor)
Less than \$42,000	0	\$2,000
Greater than \$42,000	$.34 \times (\text{Income} - 42,000)$	\$2,000

²⁵For the purpose of comparison with IRAs, married couples with one earner are allowed to contribute \$2,250 and married couples with two earners can contribute \$4,000. In the actual implementation of this proposal we see no compelling reason to make this distinction.

Table 2²⁶

Panel A: Average Annual Saving of Married Couples, IRA Contributors 1983-1986²⁷

Income Quintiles	Percentage that Save	Dollars	
		Average Amount for HHlds that Save	Median Amount for HHlds that Save
Second	61.5	12,741	6,839
Middle	58.5	11,140	5,432
Fourth	61.7	19,693	9,358
Highest	56.4	90,296	19,695
Full Population	59.5	26,728	7,128

Panel B: Average Annual Saving of Single Households, IRA-Contributors 1983-1986²⁸

Income Quintiles	Percentage that Save	Dollars	
		Average Amount for HHlds that Save	Median Amount for HHlds that Save
Lowest	56.6	2,885	2,322
Second	58.7	5,015	5,841
Middle	70.7	4,402	2,475
Fourth	63.0	8,017	3,518
Highest	60.1	36,212	9,923
Full Population	61.7	11,130	3,579

²⁶Data are from the 1983-1986 Survey of Consumer Finances. Saving is defined in the text, income deciles are given in Table 1.

²⁷The weighted number of couples that contributed to IRAs between 1983 and 1986 is 13,536,814.

²⁸The weighted number of singles that contributed to IRAs between 1983 and 1986 is 3,252,938.

Table 3:²⁹ A Comparison of 3 Saving-Incentive Proposals, Married Couples

<u>IMPACT GROUP (in 1000s)</u>			
<u>Income Quintile</u>	<u>IRA w/ AGI Cap</u>	<u>Universal IRA</u>	<u>PSA</u>
Lowest	559	559	560
Second	344	344	377
Middle	353	550	602
Fourth	0	284	622
Highest	0	102	228
Full Population	1,256	1,840	2,388

<u>NO-IMPACT GROUP (in 1000s)</u>			
<u>Income Quintile</u>	<u>IRA w/ AGI Cap</u>	<u>Universal IRA</u>	<u>PSA</u>
Lowest	1,039	1,039	921
Second	1,262	1,262	1,143
Middle	1,277	1,080	813
Fourth	0	1,420	773
Highest	0	1,416	817
Full Population	3,578	6,218	4,467

<u>ANNUAL GIVEAWAY (in \$1,000,000s)</u>			
<u>Income Quintile</u>	<u>IRA w/ AGI Cap</u>	<u>Universal IRA</u>	<u>PSA</u>
Lowest	465	465	401
Second	813	813	725
Middle	728	1,002	767
Fourth	0	1,631	858
Highest	0	1,950	1,119
Full Population	2,006	5,861	3,870

²⁹Simulations use data from the 1983-1986 Survey of Consumer Finances. Saving and column headings are defined in the text. The PSA schedule is given in Table 1.

Table 4:³⁰ A Comparison of 3 Saving-Incentive Proposals, Single Taxpayers

IMPACT GROUP (in 1000s)			
<u>Income Quintile</u>	<u>IRA w/ AGI Cap</u>	<u>Universal IRA</u>	<u>PSA</u>
Lowest	188	188	199
Second	105	105	150
Middle	136	173	173
Fourth	25	97	38
Highest	0	40	134
Full Population	454	603	694

NO-IMPACT GROUP (in 1000s)			
<u>Income Quintile</u>	<u>IRA w/ AGI Cap</u>	<u>Universal IRA</u>	<u>PSA</u>
Lowest	196	196	185
Second	280	280	236
Middle	312	275	275
Fourth	290	304	263
Highest	0	350	197
Full Population	1,078	1,405	1,155

ANNUAL GIVEAWAY (in \$1,000,000s)			
<u>Income Quintile</u>	<u>IRA w/ AGI Cap</u>	<u>Universal IRA</u>	<u>PSA</u>
Lowest	66	66	62
Second	141	141	117
Middle	168	158	158
Fourth	85	188	163
Highest	0	292	151
Full Population	460	845	650

³⁰Simulations use data from the 1983-1986 Survey of Consumer Finances. Saving and column headings are defined in the text. The PSA schedule is given in Table 1.

Figure 1: The Rate of Personal Saving,

National Income Accounts




Figure 2: Simulated After-Tax Income
and Consumption Trajectories




Figure 3: Simulated Wealth Trajectory




Figure 4: Rates of Asset Accumulation

No College




Figure 5: Rates of Asset Accumulation

College




Figure 6: Total Assets

No College




Figure 7: Total Assets

College




Figure 8: Estimated Wealth Trajectories

No College




Figure 9: Estimated Wealth Trajectories

College




Figure 10: PSA Eligibility Schedule
Married Couples (no capital income)

