NBER WORKING PAPER SERIES

ASSESSING THE EFFECTIVENESS
OF SAVING INCENTIVES

R. Glenn Hubbard
Jonathan S. Skinner

NBER Working Paper 5686

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 1996

ASSESSING THE EFFECTIVENESS
OF SAVING INCENTIVES

ABSTRACT

In this paper, we argue that there is more to be learned from recent research on the

effectiveness of targeted saving incentives than is suggested by the wide variation in empirical

estimates. First, we conclude that characterizations of "all new saving" or "no new saving" are

extreme; IRAs and 401(k) plans appear to stimulate moderate amounts of new saving. Second, we

suggest a cost-benefit approach to ask: What is the incremental gain in capital accumulation per

dollar of foregone revenue? We find that for quite conservative measures of the saving impacts of

IRAs or 401(k)s, the incremental gains in capital accumulation per dollar of lost revenue are large.

R. Glenn Hubbard   Jonathan S. Skinner
Graduate School of Business   Department of Economics
Columbia University   Dartmouth College
609 Uris Hall   Rockefeller College
New York, NY 10027   Hanover, NH 03755
and NBER   and NBER

The past two decades have witnessed the introduction of a number of public policies designed to provide incentives for household saving. Individual Retirement Accounts (IRAs) were first permitted in 1974 to provide a tax-preferred saving program for employees without pension plans. The tax advantage was the ability to defer paying taxes on the "pension" contribution until the assets were withdrawn at retirement. Despite the fact that half of all workers were eligible to contribute (because they had no pension plan at that time), fewer than one percent of taxpayers contributed.

The Economic Recovery Tax Act of 1981 expanded IRA eligibility to allow virtually all working taxpayers to contribute, and IRA limits were increased. In one year, IRA contributions rose from $5 billion to $28 billion. Over the 1982-1986 period, households contributed more than $170 billion to IRA accounts; by 1986, IRA contributions were about one-fifth of aggregate personal saving. The Tax Reform Act of 1986, however, excluded higher-income taxpayers with employer-provided pensions from making tax-deductible contributions. Total tax-deductible contributions fell by 62 percent in 1987, and have remained low since then.

Another targeted saving program, the 401(k) plan, has become prominent in recent years. While 401(k) plans became available since 1978, it was only after the Department of the Treasury clarified rules for their use in 1981 that they attracted substantial interest. Like traditional IRAs, 401(k) plans involve tax-deductible contributions, no taxes paid on accumulated interest until the funds are withdrawn, limits on annual contributions, and

restrictions on early withdrawals. There are differences, however. A 401(k) plan is available only to employees of organizations that elect to sponsor such plans. Employee contributions to a 401(k) occur through regular payroll deductions, whereas IRA contributions may be made at the employee's discretion. Finally, employers can (and often do) supplement employee contribution rates to a 401(k).

A recent outpouring of research has analyzed the effectiveness of IRAs and 401(k) plans in stimulating saving. Some researchers have found large and significant positive effects of IRAs and 401ks on saving behavior (see Hubbard, 1984; Venti and Wise, 1986, 1987, 1988, 1991; and Poterba, Venti, and Wise, 1994, 1995); the evidence favoring this view is marshalled in Poterba, Venti, and Wise (1996) in this symposium.[1] Others examine the same data and find little or no saving effects of 401(k)s and IRAs, as in the Engen, Gale, and Scholz (1996) contribution to this symposium. In this paper, we first revisit the tangled debate over the effectiveness of saving incentives, and suggest there is good reason to believe that the truth lies somewhere between the extremes of "no new saving" and "all new saving."

More fundamentally, even if one could broker an agreement among the warring factions on the *magnitude* of how IRA and 401(k) contributions affect personal saving, further questions need to be answered. Suppose, for example, one settles on an intermediate estimate that 26 cents of every dollar in IRA contributions represents new saving. Does this mean that IRAs are a rousing success? A complete failure? One cannot answer this question without knowing something about the *cost* of the program, in terms of foregone tax revenue. We

develop a cost-benefit approach which focuses on the incremental gain in long-term capital accumulation per dollar of foregone government revenue from offering the savings incentive program. For even quite conservative measures of the saving effects of IRAs or 401ks, this approach estimates that the incremental gains in capital accumulation per dollar of lost revenue are generally large.

Finally, even cost-benefit analysis of this sort does not allow one to judge whether saving incentives are a success, which we mean an improvement in welfare. To make this judgment, one must first isolate potential market failures that cause people to save too little in the first place. One than can assess whether such failures are sufficiently serious to justify the cost of targeted saving programs. We consider economic arguments that might justify having targeted saving incentives in the first place-- like social benefits from increasing the size of the capital stock, relaxing the intertemporal distortion of consumption and saving decisions made by households, or improving the financial health of households to reduce the government's cost of welfare spending for impoverished elderly. While it is difficult to attach a precise estimate to the value of overcoming these potential market failures is difficult, the most compelling rationale for saving incentives, in our judgment, is that such incentives help to overcome what appears to be inadequate financial planning for retirement.

Individual Retirement Accounts: What Do We Know?

Households who contribute to IRA accounts tend to be wealthier, older, and have higher incomes than those who do not. A detailed picture of the typical IRA contributor, based on 20,000 observations from the Survey of Income and Program Participation, can be taken from Venti and Wise (1991). In 1985, for example, four percent of households with incomes less than $20,000 and with household heads between the ages of 25 and 34 enrolled in IRAs. Holding income constant at less than $20,000, this fraction rises to 18 percent among those between the ages of 55 and 64. Holding age constant, IRA contributions rise dramatically with income. For example, among those age 55-64 and with income of $20,000 to $40,000, contribution rates were 50 percent (compared to 18 percent for those with income less than $20,000, as noted above); for people of the same age with income over $40,000, contribution rates were above 70 percent. Contributors also tend to hold more wealth than noncontributors. In 1983, at the outset of the IRA program, median non-IRA wealth of contributors (age 65 and younger) was $13,500 (Venti and Wise, 1992). By 1986, median non-IRA wealth of contributors had risen to $21,695 (Gale and Scholz, 1994), largely because IRA contributors also tend to save more in non-IRA assets. By contrast, the median noncontributor held only $3,000 in liquid assets. IRA contributors clearly have a greater taste for saving than noncontributors.

Individual Retirement Accounts and Household Saving: Theory

Much of the analysis by economists of households' consumption and saving
decisions is conducted using versions of the life-cycle model. In its most
basic form, the model implies that households save during their working lives
to finance retirement consumption. The pattern of saving over an individual's
lifetime depends on the rate of return to saving, that individual's
preferences over present and future consumption, and the time profile of
earnings. The current generation of life-cycle models adds two features to
the basic approach: imperfect markets for lending, so that households face
limits on their ability to borrow against future resources to finance current
spending, and imperfect markets for insurance, so that uncertainty over, *inter
alia,* future length of life, earnings, or medical expenses, can generate
"precautionary saving" by households.[3]

In the context of the life-cycle model, a savings incentive like an IRA or
a 401(k) plan raises the rate of return for saving done through the mechanism
of that account or plan. However, economic theory teaches that the incentive
raises a household's total saving only if the higher rate of return affects
the household at the margin -- that is, for an incremental dollar of saving.
Roughly three-fourths of all contributors in any given year deposit the full
IRA limit in their account. Several commentators have used this fact as *prima
facie* evidence that IRAs could not generate new saving, because they offer no
marginal incentive to save after the limit is reached (see, for example,
Burman, Cordes, and Ozanne, 1990; and Gravelle, 1991).

5

But this conclusion is too quick. An analysis of consumption (and saving) decisions over a lifetime requires a focus on *lifetime* limits, not annual limits. From a lifetime perspective, the relevant limit on IRA contributions is not the annual limit of, say, $2,000 or $4,000, but the lifetime limit.[4] Gale and Scholz (1994) demonstrated that only 30 percent of IRA contributors contribute at the limit for each of three years, implying that the remaining 70 percent of IRA contributors faced a marginal incentive in at least one of the three years. This evidence suggests that the IRA limits are binding for few households, even in the relatively short term. Remember, even if no contributions are made during the intervening year, the IRA provides a marginal, if unused, incentive.[5]

Some economists eschew the assumptions of life-cycle model, and focus instead on psychological issues of self-control and myopic consumption behavior (Thaler, 1994). This focus suggests that households are not optimizing life-cycle agents, responding to marginal saving incentives as they make lifetime consumption and retirement plans. Instead, they are myopic decisionmakers who have trouble saving for retirement and who respond to programs that encourage self control in setting aside assets future consumption. In this view, IRAs and 401(k)s motivate saving both because of the immediate reward of the tax deduction (including the pleasure of denying the IRS its due) and the fact that money is placed "off limits" for current consumption. There is some evidence consistent with this view. For example, taxpayers are far more likely to contribute to an IRA if they owed money to the IRS in excess of taxes withheld (Feenberg and Skinner, 1989). Apparently,

6

taxpayers would rather write a check for $2,000 to an IRA than a check for $800 to the IRS.[6] Similarly, the life-cycle model predicts that optimizing agents should contribute early in the year to their IRA to maximize tax benefits. However, roughly 40 percent of IRA contributors during the 1984 tax year filed quite late, actually in calendar year 1985 *(Statistics of Income,* 1984).[7]

Both the life-cycle and the behavioral saving models suggest that IRAs and 401(k)s have at least the potential to promote saving even in the short term. However, the magnitude of such an effect can only be determined by looking at empirical evidence.

Individual Retirement Accounts and Household Saving: Evidence

Asessing how much IRAs affect saving in the short term is more difficult than it might first appear. A complete analysis would require a significant amount of information about households: their taxable assets and tax-favored assets, along with earnings, age, and demographic characteristics like marital status or number of children that affect consumption and saving decisions. Households are also likely to have different underlying preferences for saving that are not observable. One means of controlling for different household preferences is to use panel data on the same households over time, thus tracking particular households, but even this approach will not help if preferences about saving vary over time. Some of the differing opinions of savings incentives--and the ambiguities in results--reflect data limitations

that have constrained the ways in which economists have been able to examine effects of saving incentives on household saving.

Before delving into econometric issues, it is useful to consider what people *say* when asked about how they funded their IRA contributions. In a 1983 survey of IRA contributors, Johnson (1985) found that: "About half of the respondents said they would have saved it anyway. About 10 percent said they would have spent it all, while about 40 percent said they would have spent some and saved some." He estimated that of the $32 billion in IRA contributions in 1983, $10 billion comprised new saving, or 31 cents per dollar of IRA contribution. Of course, economists are trained to treat such survey evidence with skepticism unless supported by structural or econometric studies. While the range of econometric estimates of effects of IRAs on saving is broad, we argue below that the empirical evidence is roughly consistent with this survey data.

At the high end of the range of effects are the results of Venti and Wise (1986, 1987, 1988, 1990, 1991). They model the choice among three goods: consumption, tax-favored (IRA) saving, and taxable (liquid) saving. They reason that if IRA saving is a perfect substitute for taxable saving, then the individual will immediately shift taxable saving into IRAs, because IRAs offer the higher net-of-tax rate of return. If, however, IRAs are imperfect substitutes, for other forms of saving, then some IRA contributions will come not at the expense of taxable saving, but at the expense of current consumption. In this case, IRA contributions represent new saving. In evidence from a series of papers, Venti and Wise estimate that 45-66 percent

8

of the increase in IRA contributions comes at the expense of current

consumption, while about 30 percent comes from the tax subsidy, and between 3

and 20 percent comes from a reshuffling of existing saving. (Note that we

subtract the tax break -- the 30 percent above -- in assessing what part of

the IRA is new saving. Hence our reported net saving effect could be -30

percent--for example, if the individual funded 70 percent or the IRA with

shuffled saving, and spent the remaining 30 percent from the reduction in tax

liability.)

Why might IRA and non-IRA saving be imperfect substitutes? Gale and

Scholz (1994) make an important advance by focusing explicitly on the

illiquidity of IRA balances. A household may be concerned that, at some

future point, its savings will be locked up in an IRA when the funds are

needed, perhaps to respond to a a medical emergency or a decline in future

income. The earlier approach of Venti and Wise implicitly treats both saving

and consumption as "goods." By contrast, in the approach of Gale and Scholz,

saving is not an end in itself, but a means to the end of future consumption.

To sort out the explanations for the observed relationships among IRA

contributions and savings, Gale and Scholz (1994) derive the implied saving

function for a particular set of household preferences, where saving is a

function of wealth and age. Their model is also more general than that of

Venti and Wise by allowing for a difference in tastes saving by IRA

contributors compared to noncontributors. Gale and Scholz compare saving

behavior of contributors who are at the IRA limit with contributors who are

not at the limit -- assuming that both groups have a common tasted for savings

9

-- to identify the effect of changes in the IRA contribution limit on national saving. Their estimates show that IRAs have a negative, or at best, zero effect on saving for the sample as a whole. These results suggest that IRA contributions come almost entirely from saving that would have been done in the absence of any incentives.

What's going on here? The intuition behind the Venti and Wise result is that many households, even those with very high income, do not contribute to IRAs. According to the logic of the Venti-Wise model, if IRAs and taxable saving were perfect substitutes, then everyone should contribute. However, even among high-income households, roughly one-fourth do not contribute. Hence IRAs must be imperfect substitutes for non-IRA saving, which implies that IRA contributions are coming from reducing current consumption and increasing overall savings. However, in the Gale and Scholz approach, the fact that some households do not contribute to IRAs is interpreted as evidence that those households have little or no taste for saving. If so, then (as Gale and Scholz suggest) the Venti and Wise results may be biased upward. IRA contributors save more not because of the existence of an IRA program, but because they like to save, in both IRA and non-IRA vehicles.

In contrast, the Gale and Scholz (1994) estimates that IRAs have no impact on saving are probably biased downward, given the extreme fragility of their result. In the Gale and Scholz estimates, they exclude households who reported more (in absolute value) than $100,000 in saving. Using this same exclusion criterion, Poterba, Venti, and Wise (1996) reprogrammed the Gale and Scholz econometric model, and mimicked the Gale and Scholz benchmark result

10

that IRAs have zero (or negative) effects on total saving for this same $100,000 exclusion rule. However, when Poterba, Venti, and Wise reduced the exclusion limit to $90,000, or increased it to $110,000, thereby adding or subtracting just a few observations, the estimated coefficient flipped around -- in both cases -- implying that IRAs were entirely *new* saving.

Because there are so many problems with estimating specific models of IRA contributions, a number of authors have turned to longitudinal studies of saving behavior, using repeated samples over a number of years to assess the extent to which households "reshuffle" existing saving into IRAs. The basic idea is to use saving and IRA information on the same (or similar) households over time. In one such study, Feenberg and Skinner (1989) used the IRS/University of Michigan longitudinal survey of taxpayers over the period from 1980 through 1984. They found that, even after controlling for initial assets in 1980-81, taxpayers who contributed to IRAs also saved substantially more in non-IRA assets. While this evidence weighs against the simplest story of shuffled saving between taxable balances and IRA balances, it cannot be interpreted as proof that IRAs generate new saving. One cannot control for all possible reasons leading to a change in the taste for saving. If a household decided to increase its overall saving because of impending retirement, for example, it might be expected to do so in a variety of investments including IRAs, even if IRAs have no independent effect on their tastes or saving choices.

In a different test of the hypothesis that IRA contributions represent new saving, Joines and Manegold (1995) compare assets and income of new IRA

11

contributors with those who purchased IRAs before the expansion of eligibility in 1982, also using the IRS/University of Michigan taxpayer panel. The thrust of the Joines-Manegold test is the following: If IRA contributions are new saving, then new contributors in 1982 should increase their saving by more than continuing contributors. They find that the marginal effects on saving of increasing the limit on IRA contributions by one dollar are 26 cents or 29 cents of new saving. In addressing a slightly different question, they find that 19 cents to 26 cents out of each dollar of the typical IRA contribution is financed by new saving. Their confidence intervals, however, are wider—between -16 cents to 54 cents in the first case, and -8 cents to 60 cents in the second. [8]

In another recent salvo on this subject, Attanasio and De Leire (1994) compare saving behavior of households just opening an IRA ("new" contributors) with that of households previously making contributions ("old" contributors). Their approach is therefore similar in spirit to Manegold and Joines -- in that they compare only new IRA contributors with old IRA contributors -- but they have more complete data on both changes in assets and changes in consumption for the two groups. Using data from Consumer Expenditure Survey, they test to see whether newly contributing households decrease their non-IRA assets or their consumption to fund contributions. They find that new contributors do not have slower consumption growth, but do experience slower growth of non-IRA assets, which they interpret as supporting the claim that IRA contributions largely represent reshuffled, not new, saving.

However, this interpretation evaluates evaluating only the reshuffling that takes place in the first year for a "new" contributor. Even if one assumes that nearly all IRA contributions by new contributors are shuffled, one still finds that the old contributors -- in the steady state, the vast majority of all contributors -- are barely shuffling at all. When we reinterpreted the Attanasio and De Leire results in this light, we found as much as 49 cents of new saving per $1 contribution to an IRA account (Hubbard and Skinner, 1995).

We have argued that the econometric studies finding very large saving effects are probably biased upward, and the econometric studies finding very small or negative saving effects are probably biased downward. Combined with survey data and other studies suggesting an intermediate impact of IRAs on saving, we believe that a conservative estimate of the effect of IRAs on personal saving would be about 26 cents per dollar of IRA contribution.

401k Plans: What Do We Know?

Enrollees in 401(k) savings plans look much different from IRA contributors. Employees with low levels of income are far more likely to participate in a 401k saving plan (when their employer offers such a plan), then they are to have an IRA. Among workers eligible for a 401(k) plan making between $15,000 and $20,000 annually in 1993, the participation rate was 55 percent, rising to 83 percent for workers earning more than $50,000. Among younger workers at firms with 401(k)s, contribution rates in 1993 were 55

percent for 21-30 year-olds, rising to 67 percent for 31-40 year-olds, and
over 70 percent for older groups to age 65 (Yakoboski, 1993). In other words,
the participation rate for 401(k)s is far higher, especially among low-income
workers and younger workers, than was the participation rate for IRAs.

Estimating how 401(k) plans affect household saving should be easier than
it has proven for IRAs. Individuals who contribute to IRAs are likely to be
more favorably disposed toward saving than those who do not contribute, which
makes the task of distinguishing the marginal effect of IRAs on saving
difficult. By contrast, some firms offer 401(k) plans to employees, and others
do not. It is more appealing to assert that two different groups -- those who
are eligible for 401(k)s and those who are not eligible -- are households that
share common characteristics, the saving behavior of workers eligible and not
eligible for 401(k)s.

Poterba, Venti, and Wise (1994) compare the saving behavior of workers
eligible and not eligible for 401(k)s. They include in their sample the many
workers who are eligible to contribute to a 401(k) plan, but choose not to do
so, to avoid the criticism that individuals who choose to contribute to
401(k)s were eager savers anyway. One finding illustrates the flavor of their
results: In 1984, median financial assets excluding 401(k) and IRA balances,
for those households earning in the $40,000 to $50,000 income range, were
roughly the same for the two groups. Those assets remained generally unchanged
between 1984 and 1991. Between 1987 and 1991, however, median financial
wealth of those eligible for 401(k) plans rose dramatically, largely because
of 401(k) contributions. (Unfortunately, no information is available in the

14

Poterba-Venti-Wise data for 1984 on 401(k) balances.) Assuming that the two groups -- 401(k)-eligible and 401(k)-ineligible households with equal incomes -- hold similar tastes toward saving, and assuming the composition of these workers did not change by much between 1987 and 1991, one might conclude that 401(k)s are entirely new saving.

One possible problem with this conclusion is that, if firms replace defined benefit pension plans with 401(k) plans, workers may show an increase in assets without any effects on overall (pension plus individual) saving. The actual amount of such substitution is likely to be small, however. Papke (1995) estimates that there is almost no substitution of 401(k) for defined contribution plans among large firms, but that among smaller firms, the introduction of a 401(k) plan increases the probability of a termination of a defined benefit pension plan by about 9 percentage points. However, such small firms account for fewer than one-fourth of all employees with 401(k) plans.

An additional difficulty with the Poterba, Venti, and Wise (1994) comparisons is that the 1984 survey data do not include information about existing 401(k) plans. Engen, Gale, and Scholz (1996) for example argue that many of these early 401(k) plans were simply converted from taxable thrift plans organized prior to 1981, but shifted into 401(k)s to take advantage of their tax-preferred status. Hence the apparent rise in wealth affect 401(k)-eligible households between 1984 and 1987 could be illusory. (This point does not affect comparisons between 1987 and 1991, of course.) One problem with assessing this bias is that while we know that a large fraction of 401(k)

15

plans were converted from thrift plans (Engen, Gale, and Scholz, 1996), we know nothing about the size of the balances shifted.

A third difficulty is that firms whose employees are eager savers might also be the ones most likely to implement a 401(k) plan. Then workers eligible for 401(k) plans would be systematically different from those not eligible, a point stressed by Engen and Gale (1995) and by and Engen, Gale, and Scholz (1996). This hypothesis is very hard to test because one cannot compare saving behavior of the two groups after the 401(k) plans have been implemented, since such comparisons would be contaminated by the "treatment" of having offered the 401(k). The likelihood of this "self selection" of 401(k)s by employees who are eager savers is plausible for small firms, though unlikely for workers at very large firms.

A fourth problem is the "dilution" effect of comparing 401(k) contributors and noncontributors over time. For example, suppose that there are "eager" savers and "causal" savers in the population. The earliest participants in 401(k) plans are likely to be the eager savers, so that in 1987, a high proportion of 401(k) participants would be eager savers. By 1991, however, causal savers would account for a larger fraction of 401(k) participants (Bernheim, 1994b). In other words, given the expansion of 401(k) accounts between 1984 and 1991, the typical 401(k) contributor by 1991 may be less inclined toward saving than the typical contributor circa 1984, so the pool of savers is "diluted." For example, Engen, Gale, and Scholz (1994) found a *decline* in total financial wealth among 401(k) contributors between 1987 and 1991; they interpret this to mean that 401(k) contributions were entirely

16

offset by declines in other aspects of financial wealth. A more likely

explanation, however, is that the pool of 401(k) contributors became diluted

during this period. The fraction of 401(k) contributors who also have an IRA

account -- a rough indicator of a prior taste for saving -- declined from 48

percent in 1987 to 37 percent in 1991.

Poterba, Venti, and Wise (1996a) and Engen and Gale (1995) both find that,

when the sample is separated into two groups--those who had an IRA account in

1987 and those who did not--there was an increase in financial assets among

401(k) contributors in each group.[10] This latter comparison does not prove,

of course, that 401(k) contributions increase net wealth, because other

aspects of dilution may be biasing these comparisons (see Bernheim, 1996).

Econometric estimates of how 401(k)s affect saving behavior are bedeviled by

the same problem encountered in the IRA researchs--the difficulty in

controlling for unobservable tastes for saving in the population.

Fifth, recent research by Engen and Gale (1995) suggests a different path

by which 401k assets could be shuffled -- through home equity lines of credit.

Their results suggest that the rise in 401k contributions since 1987 have been

matched nearly dollar for dollar by an increase in home equity debt, implying

that people with 401(k) plans are much more likely to use their home equity

line of credit. In fact, developments in the banking industry have made it

increasingly easy to tap into home equity. However, the Engen-Gale finding

has been criticized, first by Bernheim (1996) on a priori grounds that the

percentage of 401(k)-eligible households with home equity lines of credit (or

who bought a house) is too small to give rise to the very large extent of the

housing equity offset. Poterba, Venti, and Wise (1996) offer an empirical criticism ; they find no difference in the *percentage* changes in housing equity between the two groups.[11]

Overall, the evidence supports the view that 401(k) balances have not been offset by a decumulation of financial assets such as stocks, bonds, and checking accounts. Whether 401(k) contributions are offset by decreases in home equity is, however, still an open question. Our reading of the evidence is that 401(k)s largely represent new saving, if only because there is so little in the form of other financial assets, or home equity among low income and younger 401(k) contributors. We nonetheless recognize that the precise fraction of 401(k)s representing new saving is still under debate.

A Cost-Benefit Approach to Saving Incentives

Even if targeted saving incentives have only moderate effects on saving, a puzzle remain. Suppose that a particular saving incentive generates only four cents of new saving per dollar of contribution to the savings plan. Is this a successful program? The correct answer is: It depends on the cost. If this program loses only one cent of tax revenue per dollar of contribution, then the answer might well be yes -- after all, the policy results in four dollars in new saving per one dollar in revenue cost.

For the IRA program, we can capture this benefit-cost intuition by defining a

ratio:
$$\left[\frac{\Delta\ Private\ Capital\ Accumulation\ per\ \$1\ IRA}{\Delta\ Net\ Tax\ Revenue\ per\ \$1\ IRA} \right] .$$

Both the numerator and denominator are stocks rather than flows, and are defined for a particular time period after the initial IRA contribution. For example, suppose that the taxpayer is in the 36 percent tax bracket, and that 26 cents of the IRA contribution represents new saving, as estimated by Joines and Manegold (1995). Recall that the 26 cents of new saving is *in addition* to the 36 cent tax break which is also deposited in the IRA.[12]   The growth in net capital accumulation in the equation above would therefore be 62 cents (36 cents saved through reduced tax liability plus 26 cents of new saving), divided by the revenue loss of 36 cents. The benefit-cost ratio for the first year after the IRA contribution is therefore 62/36, or 1.72. In other words, there is an increase in private saving of $1.72 per $1 loss in government revenue. If the IRA program were financed through deficit spending, the net impact of the IRA on capital accumulation in the first year would be $0.72 per dollar of revenue loss—or the increase in private saving ($1.72) less the increase in government debt ($1.00).

However, the benefit-cost ratio in just the first year is misleading. IRAs lose additional revenue over time, because taxes are postponed on funds that would have been saved in taxable form, but IRAs then generate revenue when funds are withdrawn. In calculating the benefit-cost ratio, we therefore focus on the change in the stock of private wealth accumulated over the period for which the IRA is held, divided by the accumulated tax revenue loss over the same period.

Such calculations require assumptions about interest rates, tax rates, the length of time for which the IRA is held, and the tax treatment of the saving had it been saved in a taxable form. Because most of the estimates from existing studies are based on data from the 1982-1986 period, we use the tax regime for that period in our benchmark calculations. We assume a holding period of 22 years--which corresponds to buying the IRA at age 50, and cashing it out at age 72[13]--for an initial marginal tax rate of 36 percent (Joines and Manegold, 1995), a final retirement tax rate of 28 percent, an average tax rate on interest and dividend income of 32 percent, and a 60 percent exclusion for capital gains. The representative portfolio, whether invested in an IRA or taxable assets, is assumed to be 29 percent in equity initially, with the remainder in a combination of long-term and short-term bonds an aggregate portfolio consistent with 1985 data (EBRI, 1994).[14] During the period from 1900 to 1990, the geometric mean of the nominal return in the stock market was 9.35 percent, and the geometric mean of a portfolio with one-half short-term bonds and one-half long-term bonds was 4.0 percent (Siegel, 1992).

Assumptions about the discount rate used for government debt are crucial in these evaluations; if we use the low yield on government debt during this period, saving incentives exhibit very large (or even self-financing) effects on capital accumulation, largely because of the arbitrage that occurs when the government can borrow at a low rate of interest, but tax the higher equity returns of the IRA or 401(k) investors.[15] Instead, we use a higher nominal discount rate for government debt of 5.55 percent; this corresponds more

closely to the historical returns on stocks and bonds noted above, with a 29 percent share of equity and 71 percent share of bonds.

Table 1 presents calculations of our measure of the additional private capital accumulation per dollar of foregone tax revenue, for a wide range of estimates. The first row in Table 1 shows how the marginal impact of IRAs on capital accumulation depends on assumptions about the fraction of IRA contributions that are new saving. When there is no new saving from the IRA contribution -- in other words, 64 percent of the IRA is funded by existing saving, and 36 percent funded by the reduction in tax liability --- an IRA program leads to an increase in private saving of only $0.22. Under the assumption that the IRA is debt-financed, the net national capital stock would fall by 78 cents (22 cent increase in private saving, one dollar reduction in government saving). At a compromise estimate of 26 cents in new saving, as suggested by Joines and Manegold (1995), the implied increase in private capital accumulation is $2.21 cents per dollar devoted to the IRA program. Thus, even for a deficit-financed IRA, the net capital stock increases by $1.21. A relatively modest saving effect of IRAs can translate into a substantial "bang for the buck" in terms of capital growth per dollar of foregone tax revenue. The estimated effects are even larger when the marginal saving effect is 40 cents per $1 IRA contribution ($4.31 increase in the private capital stock) or 60 cents ($12.01 increase in the private capital stock).

However, this calculation omits a potentially important effect: The increased supply of loanable funds provided by IRAs will likely be used by

	New Private Saving per Dollar of Revenue Loss					
	0 Cents	10 Cents	19 Cents	26 Cents	40 Cents	60 Cents
Baseline	$0.22	$0.81	$1.51	$2.21	$4.31	$12.01
Include corporate income tax revenue	0.22	0.97	2.33	4.84	self-financing	self-financing
Current tax rates and portfolio share	0.04	0.63	1.35	2.09	4.45	15.51

Table 1: Change in Net Capital Accumulation Per Dollar Increase in Government Revenue Lost on Individual Retirement Accounts
Source: Authors' calculations.

corporations for increased investment, which in turn will generate income and corporate tax payments. Feldstein's (1995) analogous calculations to measure the dynamic revenue loss of the IRA program include this corporate tax effect. We include the effect of corporate taxation in our model by assuming that only equity investments are subject to the 34 percent marginal corporate tax rate used in Feldstein (1995). Because combined (corporate plus individual) tax revenue losses are smaller in this scenario, the predicted impact on private capital accumulation of one dollar in tax revenue is $4.84 at the benchmark saving effect of 26 cents per dollar of IRA contribution. For sufficiently high contributions to new savings, the IRA becomes self-financing; as Feldstein notes, it can actually generate revenue rather than losing revenue.

The tax regime has changed substantially since the mid-1980s. The third row of Table 1 repeats the calculation using more current tax parameters. In these scenarios, we assume a marginal tax rate of 28 percent for contributors, a 24 percent marginal tax rate at retirement, no exclusion for capital gains, and an average 26 percent tax rate on dividends and interest. To reflect the increasing aggregate share of equities in IRAs (EBRI, 1994), we also assume that plan assets are divided equally between stocks and bonds. The estimated incremental impact is quite similar ($2.09) to the pre-1986 tax rules for our assumed midpoint estimate of 26 cents of new saving.

To summarize, IRAs need not stimulate very substantial amounts of new saving per lost dollar of revenue to generate favorable *marginal* increases in the capital stock per dollar of initial revenue loss. The intuition is that even if the aggregate effects of a given IRA program are not large -- in terms

of overall increases in net saving -- the revenue costs can be even smaller, especially once the offsetting effects of higher corporate taxes are taken into account.

Similarly, estimating the effectiveness of 401(k) plans depends on how they affect capital accumulation per dollar of foregone revenue. However, since there are fewer estimates of how 401(k)s affect saving behavior, developing a benchmark estimate is more difficult. On the one hand, if 401(k)s did not crowd out any other types of saving, the incremental private capital accumulation per dollar of revenue cost would be $94, which is likely to pass nearly any threshold of effectiveness.[16] If sorting or reshuffling accounts for fully half of the observed increase in wealth accumulation, the 401(k) program still generates a net increase of $3.60 in private saving per dollar of revenue cost.

A Welfare-Theoretic Approach to Savings Incentives

Suppose for the sake of argument that by raising taxes by $1.00 and using the revenue to expand the IRA program, private saving would rise by $2.21. (By raising taxes and then distributing that money as a tax break, no public dissaving is created, so the entire impact of the plan is on private savings, as in the first row of Table 1.) This increase in the capital stock is not manna from heaven; rather, it is the consequence of households consuming less today in anticipation of consuming additional resources in the future (at retirement). Why fund through distortionary taxes a program which shifts households away from their presently favored level of consumption to one that

23

favors retirement consumption to a greater extent? To offer an economic justification for the existence of saving incentives, one must identify a distortion that the saving incentives are designed to overcome. We consider several possibilities.[17]

A High Social Value of Capital Accumulation. To argue for substantial external effects of increased capital accumulation, one must appeal to models in which capital or investment yield positive external effects on productivity or output, as in the models of Romer (1986) or King and Rebelo (1990). Others have noted the close correlation between saving and investment rates,[18] and between investment rates and Solow residual measures of productivity growth (see Schultze, 1992, page 242). Hence the notion that a larger capital stock yields social external benefits is certainly a valid one, but it is difficult to quantify. One problem with this rationale is that current saving incentives are not well-suited to this purpose. They include restrictions on contributions and the forced withdrawal of assets at older ages, mechanisms not designed to entice the wealthiest households -- those who account for the bulk of the nation's saving -- to save much more.

Reducing the Distortion Between Current and Retirement Consumption. Standard life-cycle models predict that the tax on interest income distorts consumption at retirement years (see, for example, Feldstein, 1978). Shifting one dollar of current consumption to the future at the gross return should provide a first-order welfare gain approximated by the wedge between the gross and net return. However, the IRA and 401(k) program is a leaky bucket in

24

effecting this transfer from current to retirement consumption since to the extent revenue is lost because of partial shuffling.

We use this intuition to consider the example in the section above, using the baseline parameter assumptions and a marginal saving effects of 26 cents per dollar of IRA contribution. First, we assume that consumption in the year of the IRA contribution, say 1996, declines by 26 cents per dollar of IRA contribution. Based on our calculations of the returns to this IRA, consumption at the end of the 22-year holding period, in 2020 rises by 94.5 cents after all taxes are paid.[19] By the first-order conditions, the individual is just indifferent to consuming 26 cents today, or putting that 26 cents in the bank to accumulate at the taxable rate of return, yielding 70.2 cents in the future period. The private individual gain from the the IRA is an extra 24.3 cents in the year 2020 (the 94.5 cent benefit less the (accumulated) foregone cost of consuming today, equal to 70.2).

While the individual is better off, the government loses revenue. The tax loss, accumulated up to the year 2020, is 55.6 cents per dollar of contribution. Raising that revenue is likely itself to entail a deadweight cost, which Ballard et al. (1985) calculate to be about 30 cents per dollar of revenue. Hence the total resource cost of funding the IRA program under these assumptions is 55.6 x 1.3, or 72.3 cents, which exceeds the benefit to the individual of 24.3 cents; hence the saving incentive would fail the cost benefit test judged solely on reducing tax distortions.

A larger saving effect would, of course, make the IRA a more efficient means of encouraging saving. The breakeven point for justifying the IRA

program on the basis of reducing intertemporal distortions is approximately 46 cents of new saving per dollar of contribution. When the corporate tax wedge is included in these calculations, however, the IRA program attains the breakeven point at about our benchmark estimate of 26 cents.

Keeping the Elderly Off Welfare Programs. Welfare programs such as Supplemental Security Insurance (SSI) and Medicaid are designed to assist elderly with limited assets and income. Encouraging households to contribute money into IRAs and 401(k)s could save the government money in the long-term by reducing the chance that individuals would qualify for means-tested welfare programs (see Hubbard, Skinner, and Zeldes, 1995). It is difficult, however, to place a value of the incremental reduction in future government expenditures because households participate in IRAs or 401(k)s today. Another problem with this explanation for saving incentives is that the programs are typically voluntary rather than mandatory. Those most likely to end up on welfare at retirement are probably also those least likely to contribute to any new pension or saving program.

Myopia and Self-Control. We have thus far restricted our attention to individuals facing well-defined, dynamically consistent utility functions. As Bernheim (1996) emphasizes, some available evidence indicates that people stumble through their planning for retirement with little idea of what they require at retirement and, perhaps, little motivation to meet those requirements. For example, Bernheim (1994a) suggests that saving rates on average are only one-third what they should be for households to consume during retirement at levels commensurate with their pre-retirement consumption

26

patterns.   If households made dynamically inconsistent plans (in the sense of Laibson, 1994; or Posner, 1995), there may be an intrinsic value to retirement saving programs that assist in self control. In this case, encouraging people to save helps to offset an "individual failure" or time-inconsistency in planning for the future, which could well yield substantial individual and social benefits.

The difficulty or inability of many individuals to save enough for their retirement may well be the most persuasive justification for encouraging saving incentives. While intuitive, such benefits are difficult to quantify. If one cannot describe preferences in a dynamically consistent way, it is harder still to attach dollar-equivalent values to the shift in the allocation of consumption. As noted above, one problem with viewing IRAs and 401(k) plans as a way to encourage self control is that such programs are voluntary, so that the people who have the most trouble saving for retirement may be the ones least likely to enroll.

Conclusions

We have followed the pattern of much of the recent studies of targeted saving incentives by focusing primarily on short-term effects of IRA and 401(k) programs. We find that, even under conservative assumptions about the extent to which contributions to saving incentives represent new saving, saving incentives generate substantial net capital accumulation over time per dollar of foregone tax revenue. However, the long-term effects are arguably more relevant in assessing the desirability of a permanent targeted saving

27

program. Life-cycle simulation exercises that attempt to quantify the magnitude of IRA and 401(k) programs on the long-term capital stock performed in Engen and Gale (1993) and Engen, Gale, and Scholz (1994) estimate very high benefit-cost ratios of increased private capital stock per dollar of lost revenue (or, in their case, per dollar of increased government debt) -- five dollars of increased capital stock per dollar of tax revenue loss associated with expanding IRA contribution limits, and $17 of expanded capital stock per dollar of tax revenue loss associated with expanded 401(k)s. The resources are careful to qualify these long-term predictions, noting for example that the period of transition is quite lengthy, taking nearly 50 years (after a short-term decline in net capital accumulation) before these benefits are realized.[20] Nevertheless, their exercise emphasizes the importance of focusing on the long-term steady-state impact of these saving programs, which might be quite different from their short-term or transitory effects.

Given more than a decade of data on the impact of targeted saving incentives on saving behavior, it is somewhat surprising that economists still disagree on the fundamental question of whether such incentives work. One reason why disagreements remain is that economists are just beginning to realize how little is understood about consumption and saving behavior, and in particular about the wide variation in saving behavior among people who are of similar age, education, and income. As research on saving incentives provide a better picture of their effectiveness at influencing saving behavior, we hope and expect that it will develop a better picture of why households save.

Acknowledgements

References

Attanasio, Orazio, and Thomas De Leire, "IRAs and Household Saving Revisited: Some New Evidence," Working Paper No 4900, National Bureau of Economic Research, October 1994.

Bernheim, B. Douglas. "Personal Saving, Information, and Economic Literacy: New Directions for Public Policy," in *Tax Policy for Economic Growth in the 1990s*. Washington, D.C.: American Council for Capital Formation (1994a).

Bernheim, B. Douglas. "Comments and Discussion," *Brookings Papers on Economic Activity* (1994b:1): 152-166.

Bernheim, B. Douglas. "Comment," in David A. Wise, ed., *Studies in the Economics of Aging*. Chicago: University of Chicago Press, 1994c.

Burman, Leonard, Joseph Cordes, and Larry Ozanne, "IRAs and National Saving," *National Tax Journal* 43 (September 1990): 259-284.

Employee Benefit Research Institute, "IRAs and Keogh Assets," *EBRI Notes* 15 (October 1994).

Engen, Eric M., and William G. Gale, "IRAs and Saving in a Stochastic Life-Cycle Model," Mimeograph, Brookings Institution, 1993.

Engen, Eric M., and William G. Gale, "Debt, Taxes, and the Effects of 401(k) Plans on Household Wealth Accumulation," Mimeograph, Brookings Institution, October 1995.

30

Engen, Eric M., William G. Gale, and John Karl Scholz, "Do Saving Incentives Work?" *Brookings Papers on Economic Activity* (1994: 1): 85-151.

Feenberg, Daniel R., and Jonathan Skinner. "Sources of IRA Saving," in Lawrence H. Summers, ed., *Tax Policy and the Economy*, vol. 3. Cambridge: MIT Press, 1989.

Feldstein, Martin, "The Welfare Cost of Capital Income Taxation," *Journal of Political Economy* 86 (April 1978): S29-S52.

Feldstein, Martin, "The Effects of Tax-Based Saving Incentives on Government Revenue and National Saving," *Quarterly Journal of Economics* 110 (May 1995): 475-494.

Feldstein, Martin, and Daniel R. Feenberg, "Alternative Tax Rules and Personal Saving Incentives: Microeconomic Data and Behavioral Simulations, in Martin Feldstein, ed., *Behavioral Simulation Methods in Tax Policy Analysis*, Chicago: University of Chicago Press, 1983.

Gale, William G., Testimony before the United State Senate, Committee on Finance. Mimeograph, Brookings Institution, February 2, 1995.

Gale, William G., and John Karl Scholz, "IRAs and Household Saving," *American Economic Review* 84 (December 1994): 1233-1260.

Gravelle, Jane G., "Do Individual Retirement Accounts Increase Savings?" *Journal of Economic Perspectives* 5 (Spring 1991): 133-148.

Hassett, Kevin A., and R. Glenn Hubbard, "Tax Policy and Investment," Mimeograph, Columbia University, March 1996.

Hubbard, R. Glenn, "Do IRAs and Keoghs Increase Saving?" *National Tax Journal* 37 (March 1984): 43-54.

Hubbard, R. Glenn, and Kenneth L. Judd, "Social Security and Individual Welfare: Precautionary Saving, Borrowing Constraints, and the Payroll Tax," *American Economic Review* 77 (September 1987): 97-112.

Hubbard, R. Glenn, and Jonathan S. Skinner, "The Effectiveness of Saving Incentives: A Review of the Evidence," Mimeograph, Columbia University, November 1995.

Hubbard, R. Glenn, Jonathan S. Skinner, and Stephen P. Zeldes, "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving," *Carnegie-Rochester Conference Series on Public Policy* 40 (June 1994): 59-125.

Hubbard, R. Glenn, Jonathan S. Skinner, and Stephen P. Zeldes, "Precautionary Saving and Social Insurance," *Journal of Political Economy* 103 (April 1995): 360-399.

Johnson, Alfred P., "Individual Retirement Accounts Help Boost Saving in the U.S." Testimony to the Committee on Finance, U.S. Senate, in Tax Reform Proposals (99-246, part XIII): 129-149. 1985.

Joines, Douglas H., and James G. Manegold, "IRA and Saving: Evidence from a Panel of Taxpayers," Mimeograph, University of Southern California, February 1995.

Kennickell, Arthur B., Martha Starr-McCluer, and Annika E. Sunden, "Saving and Financial Planning: Some Findings from a Focus Group," Mimeograph, Board of Governors of the Federal Reserve System, January 1996.

King, Robert G. and Sergio Rebelo, "Public Policy and Economic Growth: Developing Neoclassical Implications," *Journal of Political Economy* 98 (October 1990, Part 2): S126-150.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," Mimeograph, Harvard University, 1994.

Papke, Leslie, "Does 401(k) Introduction Affect Defined Benefit Plans?" *National Tax Association Proceedings* (1994): 173-77.

Posner, Richard A. *Aging and Old Age.* Chicago: University of Chicago Press, 1995.

Poterba, James M., Steven F. Venti, and David A. Wise, "401(k) Plans and Tax-Deferred Saving."In David A. Wise, ed., *Studies in the Economics of Aging.* Chicago: University of Chicago Press, 1994.

Poterba, James M., Steven F. Venti, and David A. Wise, "Do 401(k) Contributions Crowd Out Other Personal Saving?" *Journal of Public Economics* 58 (1995): 1-32.

Poterba, James M., Steven F. Venti, and David A. Wise, "Do Retirement Saving Programs Increase Saving?," *Journal of Economic Perspectives*, this issue.

Romer, Paul, "Increasing Returns and Long-Run Growth," *Journal of Political Economy* 94 (October 1986): 1002-37.

Siegel, Jeremy J. "The Real Rate of Interest from 1800-1990: A Study of the U.S. and U.K. *Journal of Monetary Economics*, Vol. 29, No. 2 (April 1992): 227-252.

Schultze, Charles L., *Memos to the President*. Washington, DC: Brookings Instititution, 1992.

Thaler, Richard H., "Psychology and Savings Policies," *American Economic Review* 84 (May 1994): 186-192.

Venti, Steven F., and David A. Wise, "Tax-Deferred Accounts, Constrained Choice, and Estimation of Individual Saving," *Review of Economic Studies* 53 (1986): 579-601.

Venti, Steven F., and David A. Wise, "IRAs and Saving." In Martin Feldstein, ed., *The Effects of Taxation on Capital Accumulation*. Chicago: University of Chicago Press, 1987.

Venti, Steven F., and David A. Wise, "The Determinants of IRA Contributions and the Effect of Limit Changes," in Zvi Bodie, John B. Shoven, and David A.Wise, eds., *Pensions and the U.S. Economy.* Chicago: University of Chicago Press, 1988.

Venti, Steven F., and David A. Wise, "Have IRAs Increased U.S. Saving: Evidence from Consumer Expenditure Surveys," *Quarterly Journal of Economics* 105 (August 1990): 661-698.

Venti, Steven F., and David A. Wise, "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, 1991.

Venti, Steven F., and David A. Wise, "Government Policy and Personal Retirement Saving," in James M. Poterba, ed., *Tax Policy and the Economy*, vol. 6. Cambridge: MIT Press, 1992.

Venti, Steven F., and David A. Wise, "The Wealth of Cohorts: Retirement Saving and the Changing Assets of Older Americans," Working Paper No. 4600, National Bureau of Economic Research, December 1993.

Endnotes

1. See also Bernheim (1996) and Hubbard and Skinner (1995).

2. The Gale and Scholz (1994)measure of non-IRA wealth also includes the cash value of life insurance, and includes households with heads age 68 and younger.

3. For models along these lines, see, for example, Hubbard and Judd (1987); Engen and Gale (1993); Engen, Gale, and Scholz (1994); and Hubbard, Skinner, and Zeldes (1994, 1995).

4. The presence of annual limits may in fact induce individuals to begin saving in IRAs earlier to allow them to make larger lifetime contributions.

5. Another way to see that IRAs tend to provide marginal incentives in the long term is the insight from Feldstein and Feenberg (1983), that, given the low levels of household holdings of financial assets, it would not take long for most households to exhaust their ability to contribute if contributions were drawn exclusively from existing (taxable) assets.

6. Of course, this result does not prove that the $2,000 contribution represents new saving; it could have been shuffled saving. In a pure life-cycle model, however, taxpayers should not need the prodding of a check due to the IRS to shuffle their saving; they should have done it anyway.

7.One criticism of this "behavioral" approach is that self-control should be a problem among the relatively affluent group of IRA contributors (Gale, 1995). However, a recent focus group conducted by the Federal Reserve Board, and limited to households with net income of more than $250,000 or net worth in excess of $600,000, found frequent mention of "the need to put money 'out of reach' to avoid the temptation to spend it" (Kennickell, Starr-McCluer, and Sunden, 1996, p. 7).

8. Poterba, Venti, and Wise (1996) and Bernheim (1996) argue they are biased downward. Engen, Gale, and Scholz (1996) argue that they are biased upward, in part by *assuming* IRA contributors over 65, excluded from the Joines and Manegold analysis, shuffle all their contributions from existing saving.

9. Our own suspicion is that the saving effect is somewhat larger.

10. Another possible bias pointed out by Engen, Gale, and Scholz (1996) is that the 401(k) balances are pre-tax, in that they

reflect the tax break but not the future tax liabilities. A bias may result from using changes in 401(k) balances to infer something about individual saving.

11.   This point was first made by Leslie Papke. See the other papers in this symposium for a discussion of these issues.

12.   This is why it is quite possible for Gale and Scholz (1994) to estimate a negative coefficient of IRAs on saving if, for example the individual shuffled taxable saving into the IRA, and spent the tax break as well. Joines and Manegold (1995) were careful to net out the tax break before arriving at their estimates.

13. Assuming the contributor plans to smooth withdrawals between age 65 and 80, age 72 is a midpoint.

14. Because of the higher return on stocks, the share of stocks in the portfolio rises over the life of the IRA.

15.   See Feldstein (1995). For example, if the government funds an IRA program using deficit financing at 4 percent, and the IRA is invested in equity paying 12 percent, the government could *make* money on an IRA program when it taxes the appreciated return on the IRA, even if the IRA is funded entirely out of existing taxable saving.

16. The marginal tax rate at retirement is assumed to be 24 percent; the assumed equity share is 50 percent; and the average tax rate on interest, dividends, and capital gains is 26 percent.

17. See Lazear (1994) and Bernheim (1994c) for a detailed discussion of whether there are social benefits of higher saving rates and a larger capital stock.

18. Hassett and Hubbard (1996) review empirical evidence which suggests that the U.S. capital stock is below its golden rule level. For this conclusion to suggest the need for saving incentives, one would need to argue that increases in domestic saving fund domestic investment.

19.   Details available from the authors.

20. See Hubbard and Skinner (1995) for a detailed discussion of these simulation estimates.