

NBER WORKING PAPER SERIES

PERSONAL RETIREMENT SAVING
PROGRAMS AND ASSET ACCUMULATION:
RECONCILING THE EVIDENCE

James M. Poterba
Steven F. Venti
David A. Wise

Working Paper 5599

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

This research was supported primarily by a series of grants from the National Institute on Aging. We also acknowledge the support of the Hoover Institution (Wise) and the National Science Foundation (Poterba). We are grateful to Bill Gale, Jon Gruber, Jon Skinner, and Richard Thaler for comments on an earlier draft of the paper. This paper is part of NBER's research programs in Aging and Public Economics. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

© 1996 by James M. Poterba, Steven F. Venti and David A. Wise. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

PERSONAL RETIREMENT SAVING
PROGRAMS AND ASSET ACCUMULATION:
RECONCILING THE EVIDENCE

ABSTRACT

Over the past several years, we have undertaken a series of analyses of the effect of IRA and 401(k) contributions on net personal saving. We have summarized this research here, together with additional results. Saver heterogeneity is the key impediment to determining the saving effect of these plans. We have used different methods to address this issue and we have organized the discussion according to the method used to correct for heterogeneity. We emphasize that no single method can provide sure control for all forms of heterogeneity. Taken together, however, we believe that the analyses address the key complications presented by heterogeneity. In our view, the weight of the evidence, based on the many non-parametric approaches discussed here, provides strong support for the view that contributions to both IRA and 401(k) plans represent largely new saving. Some of the evidence is directed to the IRA program, other evidence to the 401(k) plan, and some of the evidence to both plans jointly. We believe that the evidence is strong in all cases.

Several other investigators have used different methods to consider the effect of these retirement saving programs on personal saving and in some cases have reached very different conclusions from ours. Thus we have devoted particular effort to trying to reconcile the results, explaining why different approaches, sometimes based on the same data, have led to different conclusions. In some instances, we believe that the limitations of the methods used by others have undermined the reliability of the results. Particular attention is devoted to a recent paper by Gale and Scholz [1994] that is widely cited as demonstrating that IRAs have no saving effect. Based on our analysis of the data used by Gale and Scholz, including calculations based on a replication of their model, we find that their conclusions are inconsistent with the raw data and their formal model does not provide reliable information on the extent of substitution.

James M. Poterba
Department of Economics
E52-350
M.I.T.
Cambridge, MA 02139-4307
and NBER

Steven F. Venti
Department of Economics
6106 Rockefeller Center
Dartmouth College
Hanover, NH 03755
and NBER

David A. Wise
JFK School of Government
Harvard University
79 JFK Street
Cambridge, MA 02138
and NBER

A large fraction of American families reach retirement age with virtually no personal financial assets. The median level of all personal financial assets of families with heads 55 to 64 was only \$8,300 in 1991; excluding Individual Retirement Accounts and 401(k) balances the median was only \$3,000. Mean values are substantially higher. Almost 20 percent of families had no financial assets at all. In 1991, the median value of the future Social Security benefits of retired families with heads 65 to 70 was about \$100,000, the median value of housing was about \$50,000, and the median value of future employer-provided pension benefits was about \$16,000. But other than Social Security and pension benefits, and illiquid housing wealth, the typical family has very limited resources to meet unforeseen expenses.

Two saving programs introduced in the early 1980s were intended to encourage individual saving. Individual Retirement Accounts (IRAs) rapidly became a very popular form of saving in the United States after they became available to all employees in 1982. Any employee could contribute \$2000 per year to an IRA and a non-working spouse could contribute \$250. The contribution was tax-deductible. Annual contributions grew from about \$5 billion in 1981 to about \$38 billion in 1986, approximately 30 percent of total personal saving. Contributions declined precipitously after the Tax Reform Act of 1986, even though the legislation limited the tax deductibility of contributions only for families who had annual incomes over \$40,000 and who were covered by an employer-provided pension plan. By 1994, only \$7.7 billion was contributed to IRAs and while over 15 percent of tax filers contributed in 1986, less than 4 percent contributed in 1994.

The other program, the 401(k) plan, grew continuously and almost unnoticed, with contributions increasing from virtually zero at the beginning of the decade to over \$51

billion by 1991, when almost 25 percent of families contributed to a 401(k). Deposits in 401(k) accounts are also tax-deductible and the return on the contributions accrues tax free; taxes are paid upon withdrawal. But these plans are available only to employees of firms that offer such plans. Prior to 1987 the employee contribution limit was \$30,000, but the Tax Reform Act of 1986 reduced the limit to \$7,000 and indexed this limit for inflation in subsequent years. The contribution limit is \$9,235 for both the 1994 and 1995 tax years.

Although very small at the beginning of the decade, by 1989 contributions to all personal retirement saving plans exceeded contributions to traditional employer-provided pension plans, as shown in Figure 1. It seems evident that if it were not for the Tax Reform Act of 1986, personal retirement plan saving would have been much larger. Whether these programs increase net saving can be of critical importance to future generations of older Americans and to the health of the economy in general. The issue remains an important question of economic debate. In a series of papers based on very different methods of analysis we have concluded that a large fraction of the contributions to these accounts represent new saving. Our previous research is summarized here, along with several new results.

As interest in the saving effect of these programs evolved, several other investigators also directed attention to the issue. In some instances, alternative analyses came to conclusions that differed dramatically from ours. Thus in describing our results we have tried to point out the differences between our methods and alternative approaches that have been used to address the same questions. We have not, however, attempted to comment on all analyses of the relationship between retirement plan saving and total personal saving.

The key impediment to determining the saving effect of IRAs and 401(k)s is saver heterogeneity. Some people save and others don't, and the savers tend to save more in all forms. For example, families with IRAs also have more conventional savings than families without IRAs. Thus a continuing goal of our analyses has been to consider different methods of controlling for heterogeneity. The methods that could be used when each analysis was conducted were largely dependent on the available data. As new data became available we used alternative and possibly more robust methods to control for heterogeneity.

There are several sections of the paper. The first five sections present our results and are organized by the method used to control for heterogeneity. In each case the question is whether IRA and 401(k) contributions substitute for conventional financial asset saving. These sections also contain some discussion of closely related results reported by others. While early work in this area focused on the potential substitution between IRA assets and liquid financial assets, subsequent analyses considered the potential substitution between personal retirement saving plan assets and employer-provided pension assets and housing equity. The sixth section considers others margins of substitution, particularly the possibility that saving in these programs is financed by drawing down home equity. The seventh section addresses the divergence between our conclusions based on Survey of Consumer Finances 1983-1986 summary data -- introduced in section 2 -- and the parametric analysis of the same data by Gale and Scholz [1994]. Conclusions are presented in the eighth section.

1. EARLY PARAMETRIC ANALYSIS OF SUBSTITUTION AT THE OUTSET OF THE IRA PROGRAM

When Venti and Wise began work on the saving effect of IRAs in the mid 1980s, data on asset holdings were available for a limited time period. Assets could typically be measured at only two points in time, one year apart. To use these data, Venti and Wise developed an econometric model that could be used to estimate the relationship between IRA saving and other saving. Within a framework that allowed for any degree of substitution between IRA and non-IRA saving, the analysis asked whether persons who save more in IRAs in a particular year save less in other financial asset forms, controlling for age, income, other personal characteristics, and accumulated housing and financial assets. Given age and income, this approach used accumulated financial assets to control for "individual-specific" saving effects. The analysis accounted for the explicit limit on IRA contributions and placed substantial emphasis on the change in non-IRA saving after the IRA limit is reached.

The first results using this approach were based on data from the 1983 Survey of Consumer Finances (Venti and Wise [1986, 1987]; Wise [1987]). Subsequent analysis was based on the 1980-85 Consumer Expenditure Surveys (Venti and Wise [1990]) and the 1984 panel of the Survey of Income and Program Participation (Venti and Wise [1991]).¹

The results suggested that the majority of IRA saving, even at the outset of the program, represented net new saving, and was not accompanied by substantial reduction

¹In an earlier study, Hubbard [1984] found that the ratio of assets to income was higher for IRA participants, controlling for individual attributes and eligibility. He concluded that the results "provide strong evidence that contributions to IRAs and Keogh plans do increase individual saving." Feenberg and Skinner [1989] show that IRA participants save more than non-participants, controlling for initial wealth.

in other financial asset saving. These findings imply that increasing the IRA limit would lead to substantial increases in IRA saving and very little reduction in other saving. If the IRA limit were raised, one-half to two-thirds of the increase in IRA saving would be funded by a decrease in current consumption and about one-third by reduced taxes; only a very small proportion — at most 20 percent — would come from other saving.

The widely-cited study by Gale and Scholz [1994], based on the 1983-1986 Survey of Consumer Finances, was in some respects in the same spirit as these analyses but their conclusions were radically different, however, suggesting that raising the IRA limit would have virtually no effect on total personal saving. A detailed analysis of the findings in this study is presented in section seven of this paper.

Our subsequent analyses have taken a very different turn, using better data and more robust methods to control for heterogeneity. Our findings based on these approaches are discussed in the next four sections.

2. FOLLOWING INDIVIDUALS OVER TIME AT THE OUTSET OF THE IRA PROGRAM

To frame the discussion in this and subsequent sections, we give a simple algebraic description of the key features of each method that we use to control for heterogeneity. We establish some notation and key ideas at the outset. Consider the flow of saving S_{it} of person i in year t . To capture saving heterogeneity among families suppose that saving of person i depends on an unobserved individual-specific saving effect m_i . This effect is large for more committed savers and small for less eager savers. Saving may also depend on a program effect which is denoted by p_{it} . For program participants, p_{it} is the component of saving that is due to the program; for non-participants, p_{it} is zero. If person i is a saving

program participant and person j is a non-participant: $S_{it} = m_i + p_{it}$ and $S_{jt} = m_j$. The difference in saving between these individuals is

$$(1) \quad S_{it} - S_{jt} = (m_i - m_j) + p_{it}.$$

This difference confounds the program effect with the difference in the taste for saving.

In this simple example the difference in saving rates between participants and nonparticipants does not provide an unbiased estimate of the program effect because the unobserved taste for saving is correlated with program participation. This form of heterogeneity is probably the most important source of potential bias, but there are others as well. In the following sections we present several methods to control for heterogeneity. Each method controls for important sources of heterogeneity, but no single method -- other than a randomized controlled trial -- can control for all possible sources. Each of the methods in this and the following sections is described in a consistent way, trying to highlight both the way that heterogeneity is addressed as well as the potential types of heterogeneity that each method may not address. A specific form of heterogeneity that may confound an estimate obtained by one method may not present the same problem within the context of another method. Thus there is an important advantage to using several methods to address potential heterogeneity.

In practice, each estimate is a difference obtained in one of three ways: (1) By comparing the saving or assets of a "treatment" group in a later period with assets or saving of the same group in an earlier period, relying on within group changes. (2) By comparing saving or assets of two different groups in the same period, relying on between group comparisons. (3) By comparing the assets at a given age of persons who attain that age in different calendar years, using "cohort analysis."

2.1. Change in Other Saving When IRA Status Changes.

2.1.1. The Method: The most direct way to control for heterogeneity is to follow the same household over time, observing the change in S_{it} when program participation changes. Saving in periods t and $t+1$ for household i can be described by

$$S_{it} = m_i + p_{it}$$

$$S_{i,t+1} = m_i + p_{i,t+1}$$

The "within" household change in saving is therefore

$$(2) \quad S_{i,t+1} - S_{it} = (m_i - m_i) + (p_{i,t+1} - p_{it}) = p_{i,t+1} - p_{it}$$

which yields an estimate of the program effect for households that participate in one period but not in the other (p_{it} is zero in the nonparticipating period). The unobserved individual specific saving effects are "differenced" out.

If the heterogeneity is limited to differences in saving commitment among households, and the problem is simply that more committed savers are more likely to be program participants, then this "within" household change in saving provides a clean estimate of the program effect. But this estimate can be confounded by another possible source of heterogeneity: differences in saving commitment over time within the same household. If individual saving commitment changes at the same time that participation status changes, this estimate will capture the effect of a change in the taste for saving as well as the participation effect. With this coincidence, the difference would be

$$(3) \quad S_{i,t+1} - S_{it} = (m_{i,t+1} - m_{it}) + (p_{i,t+1} - p_{it})$$

If the household began to participate in period $t+1$ at the same time that the unobserved propensity to save increased, the difference in saving would overestimate the program effect.

2.1.2 The Results: Venti and Wise [1995a] used this method in analyzing data from consecutive waves of the 1984 Survey of Income and Program Participation (SIPP). The SIPP panel data allow calculation of the change in non-IRA saving when IRA contributor status changes, although non-IRA saving must be inferred from asset income. They considered the change in non-IRA saving between 1984 and 1985 by IRA contributor status.² If non-IRA saving is reduced when IRA saving is increased, then when a household that was not contributing begins to contribute, that household should reduce its non-IRA saving. Likewise, when a household that was contributing stops contributing, non-IRA saving should increase. Venti and Wise find, however, that when the same families are tracked over time, there is little change in other financial asset saving when families begin to contribute to an IRA, or when they stop contributing. Illustrative results are shown in Table 1.

These data reveal little substitution. The key information in this approach is the change in other financial assets when families began to contribute to an IRA. In particular, the non-IRA financial asset saving of families that did not contribute in 1984 but did contribute in 1985 declined by only \$193 between 1984 and 1985.³ This decline in other saving is only a small fraction of the increase in saving from the typical family IRA contribution, \$2,300. These data suggest that even near the outset of the IRA program there was only a small reduction in non-IRA saving when IRA contributions began.⁴

²Non-IRA saving is inferred from capitalized asset income at three points in time, measured in current year dollars. Non-IRA assets include all interest-bearing financial assets including stocks and bonds.

³If the underlying assets are measured in constant dollars, instead of year current dollars, the change is \$186.

⁴The increase of 630.3 when contributions are curtailed also suggests some substitution as well, although the estimate is not significantly different from zero. Again, the amount is much less

Table 1. Change in Non-IRA Saving When IRA Contributor Status Changed Between 1984 and 1985. (Standard errors in parenthesis.)		
	1985 Non-Contributor	1985 Contributor
1984 Non-Contributor	89.4 (102.1)	-193.5 (413.6)
1984 Contributor	630.3 (527.2)	186.2 (303.9)
		F = 0.698

Source: Venti and Wise [1995a].

As emphasized above, this procedure does not correct for "within-individual" change in saving behavior. For example, suppose that the saving behavior of persons who began to contribute in 1985 changed between 1984 and 1985 and that this change happened to coincide with the newly available IRA option. If the IRA option had not been available, it could be argued, the person would have saved in the non-IRA form, but since it was available, the newly awakened saver stored assets in the more advantageous IRA instead. The alternative, of course, is that the IRA option induced them to save in that form, and the new IRA contribution would not have been saved in another form. To us, the results seem more consistent with the conclusion that the two forms of saving are largely independent, with changes in IRA saving having little effect on other saving. It is clear that those who began to contribute in 1985 had not been saving \$2,300 annually prior to 1984. Indeed, their estimated asset balance in 1984 was only \$3,362. The same is true for persons who had contributed to an IRA in 1984 but quit contributing in 1985; their 1984 balance was \$4,816.

It is nonetheless possible that individual behavior could have changed over time, and the method used here cannot formally correct for this. The cohort analysis discussed

than the typical family IRA contribution. Asset balances are measured in May of each year because IRA contributions can be made through the April 15 tax filing deadline.

below, however, accounts for this possibility, and the results are consistent with the conclusions drawn here.

2.2. Attanasio and De Leire's Study of "Old" versus "New" Contributors.

Attanasio and De Leire [1994] (hereafter AD) analyze Consumer Expenditure Survey (CES) data to evaluate the substitution between IRA and other financial assets. The CES data essentially provide a series of independent cross sections, but each cross section is in fact a short panel, providing asset balances at two points in time, one year apart. AD consider the difference between the annual non-IRA saving -- measured by the change in asset balances -- of "old" and "new" contributors. In a given year, old contributors are families that contributed in the previous year (and possibly earlier years as well); new contributors are those who did not contribute in the previous year. Within a regression framework, they find that old contributors save \$1,740 more than new contributors.

At first blush, the results may appear to contradict the evidence just presented. Upon closer inspection, however, the method used by AD -- the comparison of new contributors with old contributors -- can say very little about the saving effect of the IRA program. Indeed, if taken at face value, the AD result suggests that in the first year that an IRA contribution is made, there is a drop in non-IRA financial assets, but in the next year and in future years when the household continues to contribute, other saving reverts to its pre-IRA level. So in the long run there is essentially no offset of IRA saving by a reduction in other saving.⁵

⁵A similar point is made by Hubbard and Skinner [1995].

The key features of the AD method are shown in Figures 2a through 2c. These illustrative figures assume that “old” contributors start to make IRA contributions in the 1982-1983 period and that “new” contributors start to contribute in the 1983-1984 period, that both old and new contributors save \$2,000 per year in the absence of the IRA, and that an IRA contribution is \$2,000. The figures compare the other saving of old and new contributors in the 1983-1984 period.

Figure 2a illustrates the situation when there is no substitution at all. Other saving remains at \$2,000 when IRA contributions begin. When the new contributors begin to contribute, there is no difference in the other saving of old and new contributors. Figure 2b illustrates the situation when there is complete substitution. When the old contributors begin to contribute their other saving falls from \$2,000 to zero and remains at zero thereafter. When the new contributors begin to contribute one year later, their other saving also falls from \$2,000 to zero. When the two groups are compared in 1983-1984, saving is zero for both groups. Like the no substitution case, there is no difference in the other saving of new and old contributors. Thus the value computed by AD cannot distinguish between the two polar cases of no substitution and complete substitution.

Neither of the processes illustrated in Figures 2a or 2b are consistent with the difference between the saving of old and new contributors that AD find. Figure 2c illustrates a process that is consistent with their finding. When old contributors first contribute, their other saving falls from \$2,000 to zero. But in the next year other saving reverts back to the previous level of \$2,000. In this year, new contributors begin to contribute and their other saving falls from \$2,000 to zero. The difference between the other saving of old and new contributors is \$2,000 in this year, approximately the finding of Attanasio and De Leire. The next year, however, “new” contributors become old

contributors and their other saving reverts to its previous level of \$2,000. There is a one-year reduction in other saving but no offset thereafter, and thus little substitution in the long run.

More detail on the non-IRA saving of households who do and do not change IRA status helps to show the limitations of the AD method. The asset data from which the figures in Table 1 were derived are shown in Table 2.⁶ For example, the change in the other saving of persons who began to contribute to an IRA in 1985 is derived from the three asset levels in the upper left of Table 2 [(6546 - 5051) - (5051 - 3362) = -193.5]. Two things are clear from these data: First, the assets of "old" contributors -- in the lower right corner -- are substantially greater than the assets of "new" contributors -- in the upper right. Their annual saving is different as well -- about 1600 for new contributors, about 2300 for old contributors. Second, there was little change in the other saving behavior of either group over the two year period -- -193.5 for new contributors and +186.2 for old contributors. Looking at these data, however, the method used by Attanasio and De Leire would show a difference between the non-IRA financial assets of old and new contributors of -\$894 in 1985 [(6546 - 5051) - (12,048 - 9659) = -894], which bears no particular relationship to substitution.

⁶These asset levels are inferred from asset earnings reported in SIPP at three points in time, approximately one year apart. Inferred assets based on alternative methods are presented in Venti and Wise [1995a]. Although the asset levels differ by method and the growth in assets differs as well, the basic difference-in-difference results, as in Table 1, are very similar.

Table 2. Inferred non-IRA financial asset balances in 1984, 1985, and 1986, by asset and by change in IRA contributor status between 1984 and 1985. (In current dollars)						
1984	1985					
	IRA = 0			IRA > 0		
	May 1984	May 1985	May 1986	May 1984	May 1985	May 1986
IRA = 0	1210	1587	2053	3362	5051	6546
	(98)	(98)	(98)	(399)	(399)	(399)
IRA > 0	4816	5896	7606	7457	9659	12048
	(509)	(509)	(509)	(293)	(293)	(293)
	$R^2 = 0.212$			$F = 434.60$		
Source: Venti and Wise [1995a].						

2.3. Joines and Manegold's Analysis of Saving Change When the IRA Limit Was Increased.

Joines and Manegold [1995] use the "change" in IRA contribution limits determined by the Economic Recovery Tax Act (ERTA) of 1981 to estimate the saving effect of the IRA program. Prior to ERTA, only wage earners without an employer-provided pension plan were eligible to contribute to an IRA. The contribution limit for wage earners in this group was \$1,500. ERTA extended eligibility to all wage earners, beginning in 1982, and increased the limit for each wage earner to \$2,000. In addition, non-working spouses of wage earners could contribute \$250.

Joines and Manegold consider the change in the total annual financial asset saving of contributor households between the 1979-1981 and the 1982-1985 periods as a function of the change in the IRA limit between these time periods. Their analysis is based on a panel of individual tax returns. Saving in each period is determined by the change during the period in total financial assets, which are estimated by capitalizing reported asset income. Two groups of households are considered: "new" contributor households first

contributed at some time in the later (1982-1985) period and “continuing” contributor households contributed in both the earlier and later periods. The “change” in the limit for new contributors was from zero to \$2,000, \$2,250, or \$4,000 for single wage earner families, couples with a single wage earner, and two wage earner families, respectively. The limit change for continuing contributors was from \$1,500 to \$2,000 for single wage earners, from \$1,500 to \$2,250 for couples with a single wage earner, and from \$3,000 to \$4,000 for two wage earner couples.

Joines and Manegold estimate a relationship of the form

$$(4) \quad \text{Saving}_{82-85} - \text{Saving}_{79-81} = \beta(X) * (\text{Limit Change}) + \gamma X ,$$

where the key parameter β is the relationship between the limit change and the change in saving. In some specifications β is estimated as a single parameter; in others it is a function of a vector of covariates X , describing primarily household tax status.⁷ The covariates also enter separately, with coefficient γ . The parameter β is not the saving effect of IRA contributions. It does not represent the relationship between IRA contributions and saving but rather the relationship between saving and the change in the IRA limit. Thus it is not comparable to most other estimates discussed in this paper, which consider the proportion of contributions that represent new saving. For example, suppose that a single new contributor deposited \$1,000 dollars in an IRA. If the \$1,000 were entirely new saving, β would be 0.50 (1000/2000), not 1.00. If the \$1,000 were the only

⁷The variables include a mortgage deduction indicator, number of exemptions, marital status, the first-dollar marginal tax rate, gender, and transitory income (measured by the mean deviation of income from average income over the 1979-1986 period)

deposit of a single wage earner couple, β would be 0.44 (1000/2250). If the \$1,000 were the only deposit of a two wage earner family β would be 0.25 (1000/4000).⁸

When estimated as a single parameter, Joines and Manegold obtain values of β ranging from 0.17 to 0.73, depending on the method of estimation. The largest estimate is obtained by ordinary least squares and is very imprecise. Joines and Manegold favor robust least squares estimates, with β parameterized as a function of X . Their “best guess” estimate of β is 0.26. This implies that substantially more than 26 percent of the IRA contributions associated with an increase in the limit would represent new saving.

Joines and Manegold find that the median of estimated total financial assets of new contributors was only \$4,396 in the 1979-1981 period. The typical IRA contribution in the 1982-1985 period was about \$2,300. Since most contributors in this period were new contributors, the typical contribution of new contributors was clearly much greater than these new contributors had been accustomed to saving prior to the advent of the IRA program. Thus, although the analysis does not purport to estimate the net saving effect of IRA contributions, the Joines and Manegold summary data suggest that the saving effect is likely to have been substantial.

2.4. Analyzing the Change in the Assets of IRA Contributors.

2.4.1. The Method: We now consider the change in non-IRA saving of IRA contributors as their IRA savings accumulate. The specification above relates saving in year t to an individual specific taste effect m_i and a program effect p_i . Most surveys do not

⁸Although Joines and Manegold compare their estimates to those of Venti and Wise [1986, 1987, 1990, and 1991] and Gale and Scholz [1994], their comparisons are inappropriate. Venti and Wise consider the proportion of IRA contributions that represented new saving, and, to indicate the implications of the results, they simulate the proportion of the increase in IRA contributions resulting from an increase in the limit that would be new saving. Gale and Scholz direct their analysis to this proportion as well, as discussed in section 7 below.

obtain direct measures of saving, however, but instead collect information on asset balances; saving must be estimated from changes in the balances. Since asset balances reflect the accumulation of past saving decisions, they also reflect individual-specific saving effects. Suppose an IRA contributor with individual specific component m_i has been saving for s years and has participated in a retirement saving program for n of these years. Then the household's asset balance, after s years of saving and n years of program saving, is

$$(5) \quad A_{si}(n) = (1+r)^s A_0 + \mathbf{m}_i[(1+r)^s - 1]/r + \mathbf{p}_i[(1+r)^n - 1]/r$$

$$= h(s) + \mathbf{m}_i f(s) + \mathbf{p}_i g(n)$$

where A_0 is the level of assets when saving began (possibly zero), r is the rate of return, and $h(\cdot)$, $f(\cdot)$, and $g(\cdot)$ are defined by context. This very stylized formulation need not be interpreted literally and indeed does not reflect saving behavior that might limit accumulation to some precautionary level, for example. Here the formulation simply serves to emphasize that the program effect as well as the individual saving effect are magnified by the number of years over which saving occurs. To illustrate the key features of this and other methods, we write the relationship in the simplified form, highlighting the key parameters in bold.

Now assume that we observe contributors after n years of exposure and then again after $n+k$ years of exposure. After $n+k$ years of program saving, and $s+k$ years of non-program saving, the assets of participants will be

$$A_{si}(n+k) = h(s+k) + \mathbf{m}_i f(s+k) + \mathbf{p}_i g(n+k) .$$

The change in assets over the k years is given by

$$A_{si}(n+k) - A_{si}(n) = h(s+k) - h(s) + \mathbf{m}_i [f(s+k) - f(s)] + \mathbf{p}_i [g(n+k) - g(n)] .$$

If $A_0 = 0$, this expression becomes

$$(6) \quad A_{si}(n+k) - A_{si}(n) = m_i [f(s+k) - f(s)] + p_i [g(n+k) - g(n)] .$$

The change in assets reflects the program effect plus the saving that m_i type families would have done over k years in the absence of the program. To isolate the program effect, we use cross-section data at the earliest observation date to approximate $m_i [f(s+k) - f(s)]$, the expected change in saving over the next k years in the absence of the IRA program. We then compare this estimate with the actual change in assets for IRA contributors.

2.4.2. The Results: Using 1983 and 1986 Survey of Consumer Finances data it is possible to compare the asset balances of the same households over time. Venti and Wise [1992] considered how the assets of IRA contributors changed over this time period. The results are reported in Table 3. Households that made IRA contributions over this period began the period with a median of \$9,400 in other financial assets in 1983. Between 1983 and 1986, the IRA assets of these families increased from \$1,000 to \$7,000. Other financial assets increased from \$9,400 to \$13,500. These families ended the period with total financial assets, including IRAs, of \$24,000, an increase of 100% over assets in 1983. Venti and Wise determined that an increase of this magnitude could not be accounted for by change in age, income, or rate of return between 1983 and 1986.⁹ In particular, they find that the increase in other financial assets is no less than would have been expected in the absence of the IRA program. Thus they conclude that it is unlikely that the IRA contributions simply substituted for saving that would have occurred anyway.

Once again, however, it is possible that at least some IRA contributors experienced a shift in saving commitment that happened to coincide with the emergence of the IRA

⁹Prediction of the expected increase in non-IRA saving in the absence of the program is discussed in section 7 below.

option and that the new commitment to saving was realized through contributions to an IRA instead of contributions to conventional saving accounts. But for the results to be explained by a within-household change in saving behavior would require that most IRA contributors over the 1983-1986 period had not been committed savers prior to this period (to be consistent with the low 1983 asset balances), but became committed savers just as the IRA program became available, and would have become committed savers in the absence of the program. This seems to us an unlikely coincidence of events.¹⁰

Table 3. Survey of Consumer Finances Data Summary			
Contributor Status and Asset	Year:		
	1983	1986	% Change
Contributors in 1986:			
Non-IRA Assets	9400	13500	43.6
IRA Assets	1000	7000	600.0
Total Assets	12075	24000	98.8
Non-Contributors in 1986:			
Total Assets	729	1000	37.2

The numbers in Table 3 come from the same data used by Gale and Scholz [1994] in their analysis of the saving effect of the IRA program. In section 7 below we return to consideration of their methodology and how the conclusions of their formal analysis could be so different from what we believe these simple data suggest.

¹⁰The test reported in Venti and Wise [1990] provides more formal evidence against the coincidence hypothesis. Unlike the SCF data that pertain to the same households in 1983 and 1986, the Consumer Expenditure Survey data used in the Venti and Wise [1990] analysis is based on random samples of similar households for the period 1980 through 1985. For example, the 1980 survey respondents were about the same age as the 1985 respondents. If the saving behavior of contributors changed just as the IRA program was introduced, estimates of saving based on post-1982 data should predict pre-1982 saving poorly. But the formal model estimated on post-1982 data predicts well the pattern of saving by income in the pre-1982 period, prior to the advent of IRAs. If the saving behavior of contributors had changed dramatically over this time period, one would expect a poor match between actual and predicted pre-1982 saving.

3. COMPARING THE ASSETS OF “LIKE” SAVER GROUPS OVER TIME

3.1. Within Group Comparisons.

3.1.1 The Method: Each of the foregoing methods rests on comparing the same individuals over time, so that similar saving propensities can be “differenced out.” Another way to eliminate the unobserved saving effect is to group households with similar saving propensities and then to estimate the program effect by using the within group difference in exposure to retirement saving programs. Poterba, Venti, and Wise [1994a, 1995] use saving program participation itself as a signal of taste for saving. “Like saver” groups are determined by observed saving behavior: families participating in an IRA only are one group, families who participate in both an IRA and a 401(k) are another, and so forth.

We consider the within group difference between the assets of a like group at two points in time, but we do not compare the same households in two periods. Rather the groups are obtained from random cross-sections of households surveyed in different calendar years. Because the cross-section surveys are representative, the demographic attributes of the cross-sections are approximately the same each year. The “like saver” groups will be the same if households that save in a given way in one year are like the households that save in that way in another year. The hope is that two randomly chosen cross-sections from the same “like” group share the same unobserved saving propensities and thus would have the same asset balances, except for differential exposure to the special saving programs which identifies the program effect. Families observed in 1984 had had about two years of exposure to the IRA and 401(k) programs, families observed in 1987 about five years of exposure, and families in 1991 about nine years.

Two factors may complicate this analysis of the effect of program exposure. First, although the IRA program expanded rapidly between 1982 and 1986, the Tax Reform Act

of 1986 reduced the attraction of IRAs for households with incomes above \$30,000 and led to a massive reduction in IRA participation by households at all income levels, even those who were unaffected by the legislation. There were few new contributors after 1986. Second, the 401(k) program grew rapidly throughout the 1980s, with more and more firms offering such plans. In both cases, but especially with respect to IRAs, the characteristics -- and thus the saving commitment -- of participants may have changed over time. In principle there could also be year-specific macro effects that might affect saving of both program participants and non-participants. (The results below show no effects for non-participants.)

Now consider explicitly the assets of a like saver group surveyed in two different randomly selected cross-sections, conducted k years apart. The two random samples of a particular group have been saving in any form for approximately the same number of years s but the sample surveyed in the earlier year has had n years of exposure to the program and the sample surveyed in the later year has had $n+k$ years of exposure. (It is not important that s be known, but only that s be the same for each random cross-section.) Assume that m_l is the typical saving propensity of the sample surveyed at the later date and that m_r is the typical saving propensity of the sample surveyed at the earlier date.

Then

$$A_{si}(n) = h(s) + m_r f(s) + p_i g(n) , \text{ and}$$

$$A_{si}(n+k) = h(s) + m_l f(s) + p_i g(n+k) .$$

This implies

$$(7) \quad A_{si}(n+k) - A_{si}(n) = (m_l - m_r) f(s) + p_i [g(n+k) - g(n)]$$

If $m_i = m_p$, then the difference in assets of the two random samples represents the program effect. If $m_i \neq m_p$, then the difference represents a combination of the program effect and the different saving propensities of the two samples. If, as seems likely, less committed savers are drawn into the program as it matures, then $m_i < m_p$, and the first term is negative. In this case, the direction of the “bias” is clear; the difference in assets underestimates the program effect.¹¹

3.1.2. Results: Poterba, Venti, and Wise [1994a, 1995] used several saving choices to identify “like” saver groups. We grouped families in two ways: first according to whether they contributed to an IRA, a 401(k), or both; and, second, according to whether they were eligible for a 401(k) plan and whether they had an IRA. Altogether, we considered six different groups of “saver types,” not counting those without IRA or 401(k) saving. We focused on the within group change in the other saving of families in these groups using data from the SIPP for 1984, 1987, and 1991. Random samples of saver types are similar in each of these years but the 1984 sample had had only about two years (1982 to 1984) to accumulate 401(k) and IRA balances, the 1987 sample had had about five years, and the 1991 sample about nine years. The central question is whether longer exposure to these plans results in higher levels of saving by families who participate in the programs.

The key test for substitution is whether non-IRA-401(k) assets are lower for the random samples which had been exposed to the IRA and 401(k) programs for longer periods of time and which had accumulated more IRA and 401(k) assets. The answer is typically no. The data for six saver groups are shown in Table 4 (abstracted from Poterba,

¹¹This is what Bernheim [1994] refers to as the dilution effect.

Venti, and Wise [1995]). The key finding is that, with one partial exception, within each saver group the level of other financial assets for the 1991 sample is not noticeably lower than the level of other financial assets for the 1987 and 1984 samples. Indeed, within each saver group, the level of total financial assets for the 1991 sample exceeds the level for the 1987 sample (the total is not available for 1984 because 401(k) assets were not obtained in that year). The only apparent aberration is a decline in the median of other financial assets of 401(k) only savers between 1984 and 1987. For this group, there was a noticeable increase in total financial assets, but little change in non-401(k) assets, between 1987 and 1991. But there was a noticeable increase in the total assets of families who made IRA and or 401(k) contributions, or were eligible for a 401(k) (whether or not they had an IRA). Since there is no evidence of a reduction in other assets for any of these groups, we conclude that the increase in retirement plan assets was funded by a reduction in other financial assets.

Consider, for example, families with an IRA only (group 2a). A comparison of the 1984 and 1991 samples reveals that the median total financial assets of such families increased from \$19,068 to \$23,892. But there was little change in other financial assets, which declined from \$11,595 to \$10,717. Or, consider families with an IRA who were eligible for a 401(k) (group 5a). Because 401(k) asset balances were not reported in 1984, total financial assets are not available in that year, but between 1987 and 1991, total financial assets of this group increased from \$37,882 to \$44,432. Yet there was no decline in other financial assets, which increased slightly from \$16,881 to \$17,212.

Although the key comparison here is the within group change over time in the other financial assets of persons who participated (or were eligible for) the IRA and 401(k) programs, we also show data for families that did not participate in one or both of these

Table 4. Conditional Median Assets (in 1987 \$s), by Saver Group, 1984, 1987, & 1991			
Saver Group and Asset Category	1984	1987	1991
By IRA-401(k) Saver Group			
IRA and 401(k)			
1a. Families With IRA and 401(k)			
Total Financial Assets	--	42655	45724
Other Than IRA or 401(k)	15653	16795	16253
1b. Families With Neither IRA nor 401(k)			
Total Financial Assets	1060	972	939
IRA Only			
2a. Families With IRA Only			
Total Financial Assets	19068	20969	23892
Other Than IRA	11595	10818	10717
2b. Families Without IRA			
Total Financial Assets	1274	1274	1509
Other Than 401(k)	1180	1091	1089
401(k) Only			
3a. Families With 401(k) Only:			
Total Financial Assets	--	8566	9808
Other Than 401(k)	3723	2587	2498
3b. Families Without 401(k)			
Total Financial Assets	3570	3602	3312
Other Than IRA	2472	2339	2145
By 401(k) Eligibility and IRA Saver Group			
All Families			
4a. Eligible for a 401(k)			
Total Financial Assets	--	16763	19608
Other Than IRA or 401(k)	6924	6796	7037
4b. Not Eligible for a 401(k)			
Total Financial Assets	4516	4607	4573
Other Than IRA or 401(k)	3075	3010	3025
Families With an IRA			
5a. Eligible for a 401(k)			
Total Financial Assets	--	37882	44432
Other Than IRA or 401(k)	16881	16032	17212
5b. Not Eligible for a 401(k)			
Total Financial Assets	20686	23537	27094
Other Than IRA or 401(k)	13098	13269	13355
Families Without an IRA			
6a. Eligible for a 401(k)			
Total Financial Assets	--	5748	7013
Other Than IRA or 401(k)	2992	2737	2757
6b. Not Eligible for a 401(k)			
Total Financial Assets	1261	1202	1210
*The estimates are conditional on age, income, education, and marital status. The medians are evaluated at the means of these variables. Source: Poterba, Venti, and Wise [1995].			

programs. In each of these groups, except 5b for which there was a slight increase, there was a decline in other financial assets between 1984 and 1991. For example, the median assets of persons with neither an IRA nor a 401(k) declined from \$1,060 to \$939. The assets of families without an IRA and who were not eligible for a 401(k) declined from \$1,261 to \$1,210. Because the assets of program participants and non-participants are typically very different, however, we avoid between group comparisons of these very dissimilar saver groups. Because their saving propensities are apparently so different, there seems little reason to believe that they would experience similar changes in asset balances in the absence of the saving programs. We return to this issue below.

It is sometimes suggested that these programs may affect households with limited assets but have little effect on wealthier households. We have addressed this issue by comparing the distribution of assets in 1984 and 1991. Again we rely on the fact that households in the 1991 survey had had much more time than their counterparts in earlier years to contribute to the saving programs. As shown in Poterba, Venti, and Wise [1994b], the higher levels of total financial assets held by IRA and 401(k) participant families in 1991 was not limited to families with large or small asset balances. Rather the effect was evident across the entire distribution of households, from those with the least to those with the greatest assets. On the other hand, across the entire distribution, there was almost no change between 1984 and 1991 in the non-IRA-401(k) assets of contributors. At all points in the distribution there was a fall over time in the assets of non-contributors.

For these estimates to control for heterogeneity, it is important that the typical person within a "like saver" group not change substantially over time -- that is, that the unobserved difference in saving propensity $m_i - m_j$ be close to zero. To help to assure that this is true, we have controlled for age, income, education, and marital status in calculating

all the numbers presented in Table 4. Nonetheless, it is possible that there were changes not accounted for by these covariates.

3.2. EGS Between Group Comparisons.

3.2.1. The Method: The critical feature of the PVW “like” group comparison is the within group change in other assets as the retirement assets of a group accumulate with increasing program exposure. This is the technique used to “difference out” the group-specific saving effect. Engen, Gale, and Scholz [1994] (hereafter EGS) follow a very different between group approach and present an alternative comparison as evidence of substitution between retirement saving program assets and other saving. EGS combine two of the PVW “like” groups, and they compare the assets of the combined group to the assets of another of the PVW “like” saver groups. Using a “difference-in-difference” approach, EGS compare the change in the assets of two very different saver groups. The first group -- call it the “treatment” group -- is composed of all 401(k) participants. This is a composite group, some of whom participate only in the 401(k) program and some of whom participate in both the 401(k) and the IRA programs. Assume that this composite program effect is cp_i and that members of this group have saving commitment m_i in the most recent year and $m_{i'}$ in the earlier year. The difference in the assets of two random samples of group i surveyed k years apart is given by

$$A_{si}(n+k) - A_{si}(n) = (m_i - m_{i'})f(s) + cp_i [g(n+k) - g(n)]$$

The EGS second group is composed of IRA participants not eligible for a 401(k) and is thus exposed only to the IRA program, with an IRA program effect denoted by b_j . Let m_j represent the saving commitment of the second group in the most recent period and $m_{j'}$

the saving commitment of this group in the earlier period. Then the difference in assets of two random samples of group j surveyed k years apart is

$$A_{sj}(n+k) - A_{sj}(n) = (m_j - m_{j'})f(s) + b_j [g(n+k) - g(n)]$$

The difference-in-difference estimate used by EGS is

$$(8) \quad [A_{si}(n+k) - A_{si}(n)] - [A_{sj}(n+k) - A_{sj}(n)] = [(m_i - m_{i'}) - (m_j - m_{j'})]f(s) + (cp_i - b_j)[g(n+k) - g(n)]$$

In principle, this method will estimate the program effect if $(m_i = m_{i'})$ and $(m_j = m_{j'})$, or if $m_i - m_{i'} = m_j - m_{j'}$, but there are several confounding effects. First, the estimate is the difference $(cp_i - b_j)$ between the composite group (IRA and 401(k)) program effect and the IRA program effect. Second, the program effect is assumed to be the same for both groups, but it is likely that the same program will have different effects on very dissimilar saver groups, even without a within group change in saving propensity. Third, because IRA participants are a much more select group of savers than 401(k) participants, the between group difference in the saving commitments of the composite group and the IRA group at a point in time $(m_i - m_j$ or $m_{i'} - m_{j'})$ may be very large.¹² Thus there is reason to question whether the within group change in saving commitments will be the same for both groups.

In the EGS case, the within group change in saving commitment $(m_i - m_{i'})$ of the composite group is especially large. The two sub-groups of the composite group have very different saving commitments. And, the subgroup proportions in the composite group change over time. There is a substantially smaller proportion of committed savers and a

¹²The IRA participation rate never exceeded 16 percent; the 401(k) participation rate has been at least 60 percent among eligible households.

larger proportion of less committed savers in the most recent year. Thus m_t is much less than m_p .

3.2.2. Comparison of Results: Our reproduction of the central EGS results is reported in the top panel of Table 5. They compare all participants in a 401(k) plan -- combining PVW groups 1a and 3a in Table 4 -- with IRA contributors who were not eligible for a 401(k) -- PVW group 5b in Table 4. EGS interpret the fall in the assets of 401(k) participants compared to the increase in the assets of the "control" group as evidence that 401(k) contributions did not lead to an increase in financial assets between 1987 and 1991.

The critical feature of our approach to controlling for heterogeneity is comparison of the within group change in non-IRA-401(k) assets as IRA and or 401(k) assets grow, for each "like" saver group. Based on this reasoning, we find that the EGS comparison has two important shortcomings. First, the two groups EGS compare surely exhibited very different saving behavior before the advent of the IRA and 401(k) programs. In 1987, the non-IRA-401(k) assets of the EGS "control" group were almost twice as large as those of the "treatment" group of 401(k) participants (\$11,823 vs \$6,635). From Table 4, it can be seen that similar differences existed in 1984. Thus in our view, the two groups should not be treated as like saver groups and the comparison between them is not meaningful.

Furthermore, the increase in the total financial assets of the "control" group is entirely due to the increase in the IRA assets of this group. There was virtually no change in the non-IRA assets of this group. This can also be seen in the PVW data for group 5b in Table 4. The non-IRA assets of this group remained almost constant, increasing from \$13,098 to \$13,355 between 1984 and 1987. In arguing that IRA and 401(k) plans have no effect on personal saving, it seems awkward to use evidence that suggests a substantial effect of IRAs on saving to show that the 401(k) plan had no effect.

Table 5. EGS Comparison and the Composition Fallacy.				
Saver Group and Asset Category	1987	1991	Change	%Change
EGS Comparison: All 401(k) Participants versus "Control" Group -- IRA Participants Not Eligible for a 401(k).				
1. All 401(k) Participants*				
Total Financial Assets	20,630	19,300	-1,330	-6.4
Net Total Financial Assets**	17,710	15,999	-1,711	-9.7
Other Than IRA & 401(k)	6,635	4,747	-1,888	-28.5
Number	1,489	2,773		
% of Total	100.0	100.0		
2. "Control" Group: IRA Participants not Eligible for a 401(k)				
Total Financial Assets	24,129	28,974	4,845	20.1
Net Total Financial Assets	21,052	26,100	5,048	24.0
Other Than IRA & 401(k)	11,823	11,000	-823	-7.0
Decomposition of EGS All 401(k) Participant Group.				
3a. 401(k) Participants Without an IRA				
Total Financial Assets	8,686	10,000	1,314	15.1
Net Total Financial Assets	5,550	7,149	1,599	28.8
Other Than 401(k)	2,774	2,400	-374	-13.5
Number	780	1,744		
% of Total	52.4	62.9		
3b. 401(k) Participants With an IRA				
Total Financial Assets	44,638	50,275	5,637	12.6
Net Total Financial Assets	41,622	46,099	4,477	10.8
Other Than IRA or 401(k)	29,844	30,000	156	0.6
Number	709	850		
% of Total	47.6	37.1		
*Group 1 is our reproduction of the EGS numbers. Although the match is not exact, it is very close and the qualitative relationships are the same.				
**Net of non-mortgage debt.				
Source: Authors' tabulations from 1987 and 1991 SIPP.				

The second problem with the EGS comparison is also fundamental and leads to an incorrect interpretation of the fall in the assets of the composite 401(k) participant group. This group is in fact composed of two very different groups: 401(k) participants without an IRA, and 401(k) participants with an IRA. The misleading interpretation created by combining these very different groups can be explained with reference to the bottom panel of Table 5. The 1987 non-IRA-401(k) assets of the second group -- 401(k) participants

with an IRA -- are about ten times as large as those of the first group -- 401(k) participants without an IRA.¹³ It is clear that the past saving behavior of the two groups was very different. This confounds inferences made from changes in the assets of the combined group, particularly if the proportions of the two subgroups in the composite group change over time, as they do between 1987 and 1991. The proportion of the second (high saver) group declined from 47.6% to 37.1% of the total combined groups, leading to a fall in the assets of the composite group. The proportion of the low saver group increased from 52.4% to 62.9%. Thus the non-IRA-401(k) assets of the composite group declined. In fact, the total assets of each group separately increased. Total financial assets of 401(k) participants without an IRA increased by \$1,314 and total financial assets of 401(k) participants with an IRA increased by \$5,637. In neither case was there an important change in non-IRA-401(k) assets. This is exactly the result shown by PVW for groups 3a and 1a in Table 4. Thus the composition problem inherent in the EGS comparison creates the illusion of substitution when in fact the data do not show that.¹⁴ The regressions run by EGS and reported in Table 5 of their paper [1994] suffer from the same problem.

In all three of the groups used by EGS, there was an increase in total financial asset saving as IRA and or 401(k) assets grew. In none of the three groups was there a substantial change in non-IRA-401(k) financial assets. The increase in the financial assets

¹³ The other financial assets of the "control" group are three or four times as large as the other assets of group (1) and only a third as large as the other assets of group (2). Thus the past saving behavior of the EGS "control" group doesn't approximate the saving behavior of either of the component groups of the combined group.

¹⁴This composition fallacy is a classic error in empirical analysis made clear by Bickel et al [1975] in their analysis of graduate student admissions at the University of California at Berkeley. While grouping departments made it seem as though there was discrimination against women, looking at individual departments made it clear that in no single department was there discrimination. It was just that women were applying to departments where the admission rate was low and men to departments where the admission rate was high.

of the EGS “control” group was due to an increase in IRA assets, which -- when compared to the fall in the assets of the composite group -- lead EGS to conclude that the 401(k) plan had no effect. But the fall in the assets of the composite group is an illusion created by the changing composition of the group.

4. THE 401(K) ELIGIBILITY “EXPERIMENT”

Another approach relies on the “experiment” that is provided by the largely exogenous determination of 401(k) eligibility, given income. It considers whether eligibility is associated with higher levels of total saving, holding income and other demographic characteristics constant. In this case the key question is whether families who were eligible for a 401(k) in a given year had larger total financial asset balances than families who were not eligible, or, equivalently, did non-401(k) financial assets decline enough to offset the 401(k) contributions of eligible families? This approach is use in Poterba, Venti, and Wise [1994a, 1995].

4.1. The Method.

Unlike the IRA program, only persons whose employers establish a 401(k) plan are eligible to contribute to such a plan. This creates a natural opportunity to compare the saving of eligible and non-eligible households. In this case, we make a between group comparison of assets at a point in time. The data for each year represents a random cross-section of respondents in that year. Thus the random samples in different years have essentially the same demographic characteristics. But samples drawn in later years have had longer exposure to retirement saving plans. This means that s -- the number of years of saving in any form -- remains the same even though retirement program exposure is greater for more recent samples.

Suppose that the saving commitment of the typical eligible household in the most recent period is represented by m_i and the commitment of non-eligible households by m_j . In an earlier period, the commitment of the typical eligible household is represented by $m_{i'}$ and the commitment of non-eligible households by $m_{j'}$. After exposure to the program for n and $n+k$ years respectively, the assets of eligible households are given by

$$A_{si}(n) = h(s) + m_{i'}f(s) + p_i g(n) , \text{ and}$$

$$A_{si}(n+k) = h(s) + m_i f(s) + p_i g(n+k) .$$

The assets of non-eligible households are given by

$$A_{sj}(n) = h(s) + m_{j'}f(s) , \text{ and}$$

$$A_{sj}(n+k) = h(s) + m_j f(s) .$$

To determine the program effect, p_i , we consider the between group difference in the assets of eligible and non-eligible households at a point in time, which after $n+k$ years of exposure is given by

$$(9) \quad A_{si}(n+k) - A_{sj}(n+k) = (m_i - m_j)f(s) + p_i g(n+k)$$

If $m_i = m_j$, the difference represents the program effect. Thus a critical question is whether the saving propensities of the two groups are in fact equal. At the outset of the program, when $n=0$, the assets of the two groups will differ only if $m_{i'}$ and $m_{j'}$ differ, with

$$A_{si}(0) - A_{sj}(0) = (m_{i'} - m_{j'})f(s) .$$

Thus if the two groups have equal assets at the outset of the program the implication is that the saving commitments of the two groups are equal, and vice versa. We use this test to establish approximate equality of taste for saving near the outset of the program.

The estimate presented in equation (9), however, depends not on the equality of saving commitments at the outset of the program, but on equality of saving commitments

at a later point in time (1991 in our case). What would assure that this equality is maintained? Suppose that the two groups are composed of equally committed savers at the outset of the program, with $m_i = m_j$. Over time, more households became eligible for a 401(k). As long as newly eligible households are a representative sample of the former non-eligible households, the two groups will continue to be composed of equally committed savers and the difference in assets at a point in time will represent the program effect.

4.2. The Results.

Inferences about the net saving effect of 401(k) contributions depend on the similarity of the saving behavior of families who are and are not eligible for a 401(k), m_i versus m_j , controlling for income. It is important, for example, that the eligible group not be composed disproportionately of savers. The data show little evidence of this type of difference in saving behavior. The most compelling evidence is for 1984. In that year eligibles and non-eligibles had about the same level of other financial assets, controlling for income. Thus, these data suggest that near the outset of the 401(k) program families that were newly eligible for a 401(k) exhibited about the same previous saving behavior as families that did not become eligible -- m_i and m_j were about the same.

Data for families with incomes between \$40,000 and \$50,000, presented in Table 6, illustrate the findings. In 1984, newly eligible and ineligible 401(k) families had almost identical non-401(k)-IRA assets -- \$5,027 and \$5,082 respectively. By 1991, however, the median of total financial assets of eligible families was \$14,470, compared to \$6,206 for ineligible families. But in 1991, the non-IRA-401(k) assets of the two groups were still about the same, \$4,724 for eligible and \$4,250 for the ineligible group.¹⁵ If families

¹⁵The apparent reduction in the non-401(k)-IRA assets of both groups between 1984 and 1991 is due largely to earnings growth. The income intervals are not indexed and thus families in

reduced saving in other forms when they became eligible for a 401(k) plan, the typical eligible family in 1991 would have accumulated less wealth in other financial assets than the typical non-eligible family. This was not the case.

Table 6. Conditional Median* Asset Balances by 401(k) Eligibility : Families with Income \$40,000 to \$50,000		
Asset Category and Eligibility Status	Results for 1991 (1991 \$s)	Results for 1984 (1984 \$s)
Total Financial Assets		
Eligible for a 401(k)	14470**	--
Not eligible for a 401(k)	6206	--
Non-IRA-401(k) Assets		
Eligible for a 401(k)	4724	5027
Not eligible for a 401(k)	4250	5082
*These are medians controlling for age, marital status, and education.		
**Difference between eligibles and non-eligibles is statistically significant at the 95% confidence level.		
Source: Poterba, Venti, and Wise [1995].		

Similar comparisons are reported in Appendix Table 1 for all income groups. In 1984, the ratio of median non-IRA-401(k) assets of eligibles to non-eligibles, weighted by the number of observations within income intervals, was exactly 1. The ratio of means was 0.87, indicating that the mean of non-IRA-401(k) assets of the eligible group was lower than the mean of non-eligibles. By 1987 the ratio of total financial assets of eligible to non-eligible families was 1.62 and by 1991 this ratio was 2.22. This evidence suggests a sizeable effect of 401(k) saving on the accumulation of financial assets and shows little if any substitution of 401(k) contributions for other financial asset saving.

Indeed, for all income groups, eligible households have greater total financial assets than non-eligible households at virtually all points across the entire distribution of financial

a given interval in 1984 will tend to have greater wealth than families in that same interval in 1991. Comparable calculations with the intervals indexed to 1987 dollars are discussed in section 6.2.2 and reported in Tables 10 and 11 and in Appendix Tables 3 and 4.

assets. But there is virtually no difference across the entire distribution of the other financial assets of eligible and non-eligible households, as shown in Poterba, Venti, and Wise [1995].

For comparisons between eligible and non-eligible households to shed light on the net saving effect of 401(k) plans, it is important that the saving behavior of the eligible and non-eligible groups be comparable. As noted above, after controlling for income, the accumulated assets of the two groups were very close at the outset of the program. Nonetheless, there could have been some change in the composition of the two groups over time, even if eligibles and non-eligibles were very similar in 1984. Data on measured household attributes, however, suggest that there was little composition change.

Many studies of saving behavior have shown that saving commitment is related to household demographic attributes, like age and education. As Appendix Table 5 shows, these characteristics did not change substantially over time. The average age of the head of eligible households was 41.8 in 1984 and 41.4 in 1991. The average years of education of the head of eligible households was 13.6 in 1984 and 13.7 in 1991. Within income interval, there was also very little change in the average age or education of eligible families. Similarly, there was little change in the age or education of non-eligible households. The proportion of households with husband and wife present, which is typically found to be positively related to saving behavior, declined by 7 percentage points, on average, for both eligible and non-eligible households. Much of saving commitment, however, cannot be explained by observed household attributes and we rely on the cohort approach discussed below to provide a check on the eligibility “experiment” results. The cohort analysis is not confounded by the potential difference between the saving commitment of eligible and non-eligible households.

Engen, Gale, and Scholz [1994] question the validity of our comparisons between 401(k) eligible and non-eligible households. They argue that “401(k) eligible families save more in non-401(k) assets than observationally equivalent non-eligible families, even after controlling for other factors.” In our view, however, their numbers differ little from ours. They say, for example, that the two groups had different asset levels in 1984. But they estimate a (statistically insignificant) difference in median financial assets of only \$173. They estimate a difference in median net financial assets of only \$346. Given that our analysis controls for demographic characteristics and compares households within income intervals, while the EGS approach simply includes income as a single variable in a regression equation, their estimates seem hardly different from our findings.

EGS find a difference of \$2,500 in 1984 in the net worth of eligible and non-eligible households. The median for the entire sample is about \$30,000 so the estimated difference represents a percent difference of under 9%. In section 6 below, we repeat our analysis including housing equity, and controlling for income interval as above. We find essentially no difference in the housing equity of eligible and non-eligible families in 1984.

EGS also find that eligible families are more likely than non-eligible families to have a traditional defined benefit employer-provided pension plan. Whether this reflects a difference in saving propensity is questionable. That depends first on whether people choose jobs based on the pension plan. And, if they do, it isn't at all clear that wanting a good pension plan means a stronger preference for saving. It could mean just the opposite. It may well be that choosing a job where the employer saves for you is a means of self-control. If a person is unlikely to save and would not do so if it were not for the employer-guaranteed retirement income, then the people who choose jobs with pensions may be non-savers, not savers.

Assuming further that persons with defined benefit pensions save less in other forms, Engen and Gale [1995] seem to argue that they should have lower financial assets in 1991 than non-eligibles, if eligibility is independent of the taste for saving. In this case, according to their reasoning, a finding that eligible families have about the same, or even more, assets than non-eligible households in 1991, confirms that they have a stronger commitment to save than non-eligibles. Engen and Gale [1995] conclude that even a finding that eligible families save the same amount in non-401(k) assets as ineligible families can be interpreted as evidence that 401(k) eligibility is not exogenous, since eligibles should be saving less, if eligibility is exogenous with respect to saving. This reasoning seems to us self-fulfilling, assuming substitution to demonstrate substitution. It is a remarkable change from all earlier studies of IRAs, in which the central hypothesis was that IRA participants should have less non-IRA financial wealth than non-participants, if IRAs and other financial assets are substitutable. Here, the possibility that 401(k) eligibles may have more non-401(k) assets than non-eligibles is used as evidence for substitution.

The weight of the evidence, however, is that persons with pensions do not reduce saving much, if at all, relative to persons without pensions. Gale [1995] argues that the methods that have been used by others to address this question are plagued by a series of biases that lead to an underestimate of the reduction in other saving for persons with employer-provided pensions. The key method that he proposes to avoid bias in his empirical analysis is a derived adjustment factor that multiplies pension wealth in regression equations relating other saving to pension wealth.

While Gale [1995] has pointed to a number of possible difficulties with prior estimates, his methodology, particularly his derived adjustment factor, raises new questions of interpretation. He derives the adjustment factor in a stylized life cycle model

that assumes that a household's consumption and saving, even before retirement, are proportional to the present discounted value of lifetime wage and pension income. There are three difficulties with this approach. First, the model assumes that persons view pension wealth and other financial assets as perfect substitutes, that there is a "complete offset between pensions and other wealth."¹⁶ Because the adjustment factor used in the empirical analysis is derived assuming complete substitution, the analysis cannot provide an unambiguous test of the extent of substitution. Much of the empirical research on saving, including the analysis of tax-deferred saving summarized here, suggests that the assumption is inappropriate. More conceptual "behavioral" explanations of saving behavior, in particular the work that emphasizes the "mental accounts" approach to saving, also bring into question this assumption. Second, the approach abstracts from the likely possibility that many households face liquidity constraints that make it difficult to consume out of pension wealth before retirement. Third, only a small fraction of persons who take a job with a pension early in the life cycle will retire from that job and acquire the rights to the defined pension benefit at age 65. Kotlikoff and Wise [1985, 1988, 1989], explain that a person who leaves a firm at 40, for example, will have accrued only a small fraction of the age 65 pension. Such considerations imply that Gale's [1995] findings are subject to new biases that cloud interpretation of his results.

4.3. 401(k) Eligibility and Other Pension Plans.

Our analysis has focused on the substitution between 401(k) assets and other financial assets. Another potential tradeoff is employer substitution of 401(k) plans for other employer-provided pension plans. The possibility of such substitution arises more directly with 401(k) plans than with IRAs, because 401(k) plans are part of the

¹⁶Gale [1995], pages 16-17.

workplace benefits package and their availability, like the availability of defined benefit or defined contribution pension plans, is subject to employer choice.

Although substitution between traditional pensions and 401(k) plans is a theoretical possibility, existing empirical evidence provides little support for such substitution in practice. As discussed above, Engen, Gale and Scholz [1994] present evidence that workers who are eligible for 401(k) plans are more likely to be covered by a defined benefit pension plan than are workers without 401(k) eligibility. Although they interpret this as evidence of saver heterogeneity, it is prima facie evidence against the pension substitution hypothesis. Papke (1995) uses data from the 1985 and 1991 IRS Form 5500 filings, and Papke, Petersen, and Poterba (1996) use data from a survey of 401(k) providers, to provide further evidence on this question. There is essentially no evidence that large firms offering 401(k) plans substituted these plans for other pension plans; the first 401(k)s were typically offered as "add-on" plans in large firms with pre-existing defined benefit pension plans. Papke (1995) finds some evidence of substitution at smaller firms that have introduced 401(k) plans in recent years.

5. COHORTS AND THE EFFECTS OF RETIREMENT SAVING PROGRAMS

This method compares the assets of persons who are the same except that they reached a given age in different calendar years. Hence some cohorts had longer than others to contribute to special saving programs. For example, families that reached age 65 in 1984 had had only two years to contribute to an IRA or to a 401(k) plan, but families who attained age 65 in 1991 had had nine years to contribute. If these programs affect personal saving, they should lead to differences in asset accumulation by cohort.

5.1. The Cohort Method.

Consider a random sample of all families, and assume for the moment that the typical family i has saving commitment m_i . Cohorts are distinguished by age (c) in 1984. We assume that cohort c has had s years to save. Suppose that in 1984 cohort c has had s years to save and during n of these s years was able to contribute to special retirement saving programs. We follow each cohort from 1984 to 1987 to 1991. For simplicity, we can assume that the cohort c began the period having saved for s years and ended the period having saved for $s+k$ years. By 1991, each cohort has had $n+k$ years to contribute to special retirement saving programs. (In 1984, n is about 2; in 1991, $n+k$ is about 9.) For cohort c , assets after n and $n+k$ years of exposure are given by

$$A_{c_i}(n) = h(s) + m_i f(s) + p_i g(n)$$

$$A_{c_i}(n+k) = h(s+k) + m_i f(s+k) + p_i g(n+k)$$

Consider now an older cohort ($c+k$) that began the period having saved for $s+k$ years. For this cohort assets after n and $n+k$ years of exposure are given by

$$A_{c+k,i}(n) = h(s+k) + m_i f(s+k) + p_i g(n)$$

$$A_{c+k,i}(n+k) = h(s+k+k) + m_i f(s+k+k) + p_i g(n+k)$$

This second cohort is k years older than the first cohort in the same calendar year. Thus the difference in the assets of the two cohorts when both have saved for $s+k$ years is given by

$$(10) \quad A_{c_i}(n+k) - A_{c+k,i}(n) = (m_i - m_i) f(s+k) + p_i [g(n+k) - g(n)] = p_i [g(n+k) - g(n)]$$

Thus the difference in the assets of families who reached the same age in calendar years n and $n+k$ is due to the program effect.

If the different cohorts had different saving commitments, however, the term $(m_i - m_i)$ in equation (10) would not be zero and the difference would reflect this, as well as the program effect. Judgments about the likely importance of such differences may be based on several features of the analysis. Cohort effects are obtained for a succession of cohorts ranging in age from 42 to 70 in 1984. These cohort effects are obtained for several asset categories: special retirement saving program assets, total financial assets, and conventional financial assets. Differences in the cohort effects for different assets can be used to judge whether there was a systematic change in taste for saving over time. It is also possible to compare cohort effects for participants and non-participants in retirement saving programs.

5.2. The Results.

The cohort method was used by Venti and Wise [1996]. They find that households who attained a given age in 1991 had consistently larger total real financial assets than households who reached that age in 1984. The larger assets of the younger cohorts is accounted for almost entirely by more assets in IRA and 401(k) plans. There is on average no difference between the other financial assets of the older and younger cohorts. The results can be illustrated by comparing the assets of families who reached ages 60 to 64 in 1984 with the assets of families that attained that age in 1991, as shown in Table 7.

To control for heterogeneity, the data for all families -- both contributors and non-contributors combined -- are the most compelling. In this case it is the typical saving propensity m_i over all families that is important and the possible effect of the changing composition of participant and non-participant families is avoided. (Because fewer than half of all families participate in these programs, the median of program assets for all

families is zero and thus not informative.) The mean of total financial assets of all families that attained age 60 to 64 in 1984 was \$42,250; the mean of those who attained this age in 1991 was \$50,419 (both values are in 1991 dollars and control for income, age, education, and marital status). The increase was accounted for almost entirely by personal retirement saving -- \$5,118 for the cohort that attained ages 60 to 64 in 1984 compared to \$14,156 for the cohort that attained this age in 1991. There was essentially no cohort difference in other financial assets (\$37,132 for the older cohort and \$36,263 for the younger cohort). Thus there is little evidence of substitution of personal retirement saving for other financial assets.

Table 7. Summary of Cohort Effects at Age 60-64 (1991 \$s)		
	1984*	1991
Contributors and Non-Contributors Combined		
Mean:		
Personal Retirement Assets	5118	14156
Other Personal Financial Assets	37132	36263
Total Personal Financial Assets	42250	50419
Contributors		
Percent	38	42
Median:		
Personal Retirement Assets	8171	22148
Other Personal Financial Assets	22983	21528
Total Personal Financial Assets	34975	50182
Non-Contributors		
Percent	62	58
Median:		
Total Personal Financial Assets	2687	2134
*The means and medians reported in this table are controlling for age, income, marital status and education. The 1984 totals exclude 401(k) assets, which were small at that time. Thus the data for personal retirement and for total personal financial assets are affected to some degree by this omission. But the data on other personal financial assets are unaffected.		
Source: From Venti and Wise [1996], converted to 1991 dollars.		

The data for families who participated in personal retirement saving plans provides a better measure of the potential of the plans to augment the financial assets of retirees.

The median level of total personal financial assets of contributor families that attained age 60 to 64 in 1984 was \$34,975, compared with \$50,182 for families who attained that age in 1991. The median level of personal retirement plan assets of the families that reached this age in 1984 was \$8,171, compared with \$22,148 for families who reached age 60 to 64 in 1991. On the other hand, the other financial assets of these families were about the same in 1984 and 1991 (\$22,983 and \$21,528 respectively). Although not as compelling as the data for both groups combined, these data also provide little evidence of substitution. In contrast, the financial assets of families that attained age 60 to 64 in 1991 but did not participate in personal retirement plans were somewhat lower than the assets of similar families who reached this age in 1984.

The results for other age groups are summarized in Figure 3. To understand the figure, consider age 66: The cohort that reached this age in 1984 had about \$5,000 less in personal retirement assets (the heavy lines) than the next younger cohort that reached that age about four years later. The difference in the total financial assets (the lighter lines) of these two cohorts is also about \$5,000. But there is very little difference in the other financial assets of these two cohorts.

Results of more formal estimation of cohort effects are shown in Table 8. The estimates are obtained by fitting a cubic function in age to the cohort means, allowing for cohort shifts -- the cohort effects -- in the relationship between age and assets.¹⁷ The

¹⁷We fit the actual cohort means with a specification of the form

$$A_{ic} = \alpha + \beta_c + \gamma_1(\text{Age}_i) + \gamma_2(\text{Age}_i)^2 + \gamma_3(\text{Age}_i)^3 + \varepsilon_{ic} \tag{1}$$

where A represents an asset category -- personal retirement assets, other personal financial assets, total personal financial assets -- c indexes cohort and i the ith cohort mean. The β_c are cohort effects with $\sum \beta_c = 0$. Thus the individual estimates represent deviations from the mean effect, which is set to zero. The specification is intended to fit the age-asset accumulation pattern, allowing the differences in the levels of the assets between successive cohorts to be maintained and to

estimate of the youngest cohort effect for personal retirement assets (column 1) is \$14,076 above the mean while the estimate for the oldest cohort is \$13,105 below the mean, a difference of \$27,181. If there were no counterbalancing cohort effects with respect to other personal financial assets, the total personal financial asset cohort effects should approximately parallel the retirement asset cohort effects. The estimates show that the other personal financial asset cohort effect (column 2) for the youngest cohort is \$16,003 above the mean and the cohort effect for the oldest cohort is \$14,083 below the mean, a difference of \$30,086. The other personal financial asset cohort effects are typically small and not statistically different from zero. An F-test does not reject the hypothesis that all the cohort effects with respect to other personal financial assets are zero.

The analysis suggests that if current patterns persist families who reach retirement age 25 or 30 years from now will have much more in financial assets than families currently attaining retirement age, and the difference will be due solely to assets in personal retirement accounts.

We believe that the cohort approach provides the surest way of controlling for heterogeneity. When both contributors and non-contributors are considered jointly, the overall saving effects are not contaminated by the potential changing composition of the two groups. Nor are the cohort estimates confounded by the “coincidence” possibility that may affect the difference-in-difference estimates discussed in section 2. In principle, the

cumulate as the cohorts age. It is assumed, for example, that the estimated difference between the assets of the two youngest cohorts, C42 and C46, will be maintained as the cohorts age. It is likely that this assumption implies a conservative estimate of the projected cohort differences. Constant percentage differences as the cohorts age, for example, imply much larger absolute differences at advanced ages than this model does.

cohort analysis compares families who differ only in the calendar year in which they reached a given age and therefore in their exposure to retirement saving programs.

Cohort	Personal Retirement Assets		Total Personal Financial Assets		Other Personal Financial Assets	
	Coefficient	t-Stat	Coefficient	t-Stat	Coefficient	t-Stat
C42	14076	19.0	16002	8.2	1927	1.0
C44	11085	17.9	12024	7.3	939	0.6
C46	9997	17.3	9568	6.3	-428	-0.3
C48	7821	14.8	6556	4.7	-1264	-0.9
C50	5759	11.9	4132	3.2	-1626	-1.3
C52	3814	8.6	1459	1.2	-2354	-2.1
C54	1944	4.7	452	0.4	-1492	-1.4
C56	363	0.9	734	0.7	370	0.4
C58	-1604	-3.9	-1682	-1.6	-78	-0.1
C60	-3815	-8.7	-5165	-4.5	-1349	-1.2
C62	-5813	-12.1	-3796	-3.0	2017	1.7
C64	-8130	-15.4	-5234	-3.7	2895	2.2
C66	-10345	-18.0	-8766	-5.8	1578	1.1
C68	-12049	-19.2	-12203	-7.3	-154	-0.1
C70	-13,103	-17.8	-14,081	-6.1	-981	0.4

Source: Venti and Wise [1996].

A potential, although we believe unlikely, confounding influence would be an overall change in saving behavior, with each successively younger cohort wanting to save more than its older cohorts. The evidence suggests that such a systematic increasing taste for saving must have been realized only in contributions to the special retirement saving programs. We find this an unlikely possibility for two reasons. There are no cohort effects in other financial assets, as we would expect if there were an underlying change in taste for saving. Nor are there cohort effects for non-participants, as we would also expect if there were an overall change in the taste for saving. Therefore, we interpret the cohort results as supporting the results of the other methods of correcting for heterogeneity.

5.3. Further Results.

Registered Retirement Saving Plans were first introduced in Canada in 1957. Like the IRA in the United States, an individual can make contributions to an RRSP and deduct the contributions from income for tax purposes. Interest accrues tax-free until withdrawal, when taxes are paid. The contribution limits were increased substantially in the early 1970s and RRSPs were widely promoted. Since then, they have become a very prominent form of saving. Annual contributions grew from \$225 million in 1970 to almost \$3.7 billion in 1980 to \$16 billion by 1992, when they accounted for about one-third of aggregate personal saving. In 1992 about 33% of families contributed with an average contribution of \$4,180. Now RRSP contributions exceed the total of employee and employer contributions to employer-provided pension plans.

Based largely on "cohort" analysis like the procedure described above, Venti and Wise [1995b] conclude that RRSPs have contributed substantially to personal saving in Canada. In virtually no case do the data suggest substitution of RRSP for other forms of retirement saving. In the two decades prior to the growth in RRSP popularity, the personal saving rate in Canada was typically below the U.S. personal saving rate. Since that time, the personal saving rate in Canada has become much higher than in the United States. Although it is difficult to make judgements about the RRSP saving effect based only on the trends in U.S. and Canadian aggregate saving rates, the cohort analysis suggests that a large fraction of the current difference can be accounted for by RRSP saving. Engelhardt [1995a] analyzed the similarly tax-advantaged Registered Home Ownership Saving Program (RHOSP), designed to encourage saving for home purchase. He finds that the RHOSP program also increased total personal saving.

6. OTHER MARGINS OF SUBSTITUTION: HOME EQUITY

The foregoing discussion focuses on the substitution between contributions to special retirement saving plans and other financial assets. There are at least two other potential margins of substitution: employer-provided pension assets and home equity. As mentioned above, many analysts have considered the substitution between employer-provided pension assets and personal financial assets. The results are mixed but the weight of the findings suggests little substitution.¹⁸ Venti and Wise [1993 and 1995b] have also addressed this question, considering the assets of retired persons for whom pension assets are known, and using Social Security benefit percentiles to control for lifetime income. They conclude that there is essentially no relationship between employer-provided pension assets and either personal retirement saving plan assets or other financial assets. They considered the same question for retired persons in Canada, where the RRSP program has been widely used for several decades, again finding essentially no relationship between employer-provided pension assets and personal financial assets. We will not address that question further here.

We will, however, consider the potential substitution between housing equity and retirement saving plan assets. Our focus on the relationship between retirement saving assets and other financial assets neglects the possible interaction between these retirement plan assets and home equity, which is the largest asset of a large fraction of households. While many of the factors that are likely to determine whether to purchase a home, the value of the home, and how to finance a home purchase may be unrelated to the accumulation of retirement saving assets, it is possible that some

¹⁸Gale [1995] has summarized the results in table 1 of his paper.

of the build-up in these accounts has been financed through reduced accumulation of housing equity.

Several studies have considered the relationship between housing prices and financial assets. In his review article, Skinner [1994] finds little relationship between exogenous shocks to housing value and personal financial assets. Several other studies are based on the Panel Survey of Income Dynamics: Skinner [1993] finds a small relationship for younger households and no relationship for older households. Hoynes and McFadden [1994] find little relationship between exogenous changes in home values and changes in financial assets. Engelhardt [1995b] finds no decrease in financial asset saving among households with an increase in home values but finds a small increase among households with falling home values. Engen and Gale [1995] have considered the relationship between home equity and 401(k) assets based on SIPP data. While their results largely confirm our findings on the relationship between 401(k) and other financial assets, they conclude that the increase in the financial asset saving of 401(k) participants (or eligibles) between 1987 and 1991 was offset by a reduction in home equity.

We consider the relationship between retirement saving plan contributions and home equity using cohort analysis as in section 5 and comparison of 401(k) eligible and non-eligible families as in section 4. The most important conclusion from the cohort analysis is that the timing of changes in mortgage debt and net home equity is inconsistent with a causal relationship between personal retirement plan contributions and mortgage debt. With respect to 401(k) contributions in particular, we conclude from the eligibility comparison that there was no apparent offset to 401(k) contributions through a reduction in home equity. We consider briefly one possible reason for the difference between our results and those reported by Engen and Gale [1995].

6.1. Cohort Analysis.

Cohort data make it easy to compare the trends in personal retirement saving and housing assets. As in the analysis above, we consider IRA and 401(k) participants and non-participants together.¹⁹ The interpretation of financial asset versus housing equity trends must be tempered by at least two factors. First, market trends in housing values and financing practices that are unlikely to be induced by IRA and 401(k) contributions can have substantial effects on housing equity. There was probably little relationship between retirement saving plan contributions and the concerns that led to elimination of tax-deductibility of non-mortgage interest as part of the Tax Reform Act of 1986. But this provision may have had a substantial effect on home mortgage debt.²⁰ Thus the home equity data may be subject to very important time effects.

Second, unlike other consumer debt, mortgage debt may, in the long run, increase future saving. Many financial planners tout mortgage debt repayment as the surest way for a household to commit to a long-term saving strategy. Regularly scheduled mortgage payments can thus be viewed as means of self control as stressed by Thaler and Shefrin [1981], Shefrin and Thaler [1988], and Thaler [1990]. While increased mortgage debt may appear as a reduction in wealth today, it may assure greater rather than reduced wealth at retirement. Similarly, a home equity loan that is repaid before retirement may not affect wealth at retirement.

¹⁹To maintain comparability with the cohort analysis discussed above, we have converted current dollar amounts to 1991 values using the BLS earnings index. Using a price index instead has little effect on the trends.

²⁰Skinner and Feenberg [1990] find that each dollar of reduced consumer debt following the TRA 86 was offset by a 67 cent increase in mortgage debt.

The central results of the cohort analysis are presented in Figure 4, which shows the relationships between contributions to personal retirement saving plans and housing market data. Like the results presented in section 5, the analysis here is based on 1984, 1987, and 1991 data for 15 cohorts: the youngest was 42 and the oldest was 70 in 1984.²¹ The cohort data on housing value, mortgage debt, and home equity are shown in Appendix Tables 2a through 2c. Data for selected cohorts are graphed in Figures 4a through 4d. Figure 4a shows data for mean personal retirement assets (including IRA, 401(k), and Keogh saving balances). Figure 4b pertains to home value, Figure 4c to mortgage debt, and Figure 4d to net home equity. All values are in 1991 dollars. The figures can be explained with reference to Figure 4a. For each of the cohorts, mean retirement assets are shown for 1984, 1987, and 1991. For example, cohort "C42" was age 42 in 1984, 44 in 1987²², and 48 in 1991. By 1991, this cohort had had nine years to contribute to the retirement saving program and had mean assets of \$8,000 in these accounts at age 48. In contrast, cohort "C48" had had only about two years to contribute to such accounts when first observed in 1984 and had only about \$2,000 in these retirement assets at age 48. Similar comparisons can be made at ages 54, 60, and 66. The cohort that attained the given age later had much large amount in these retirement assets at that age. The cohort data in Figures 4b through 4d present housing data for the same cohorts and the trends can be compared to the cohort trends for retirement assets.

²¹In principle, we would like to consider younger cohorts as well, but we wanted these data to be comparable to our earlier analysis of financial asset data that was directed to families approaching and entering retirement.

²²The 1984 survey was administered between September and December 1984, and the 1987 survey between January and April of 1987, a difference of approximately 28 months.

Figure 4b shows a substantial fall in real home value between 1984 and 1987 for younger cohorts but an increase for older cohorts. For all cohorts, but especially for the younger cohorts, there was a large increase in home values between 1987 and 1991. Given that housing values were falling during the rapid rise in retirement saving plan assets -- and only rising later on -- these trends apparently reflect housing market effects that are unrelated to 401(k) and IRA contributions. It is clear, however, that at ages where direct comparisons can be made, the home values of younger cohorts are much greater than those of older cohorts. For example, the cohort that reached age 48 in 1991 had a real mean home value of about \$80,000. The cohort that attained age 48 in 1984 had a mean home value at that age of about \$72,000, in 1991 dollars.

Figure 4c shows a fall in mortgage debt between 1984 and 1987 for all cohorts. This pattern persists even for older cohorts that experienced an increase in home values between 1984 and 1987. Yet over this period there was a sharp increase in the IRA and 401(k) assets of these cohorts, as shown in Figure 4a. Between the early 1980s and 1986 contributions to these programs grew from about \$3 billion to almost \$74 billion. Contributions to 401(k) plans almost doubled between 1984 and 1986. Yet it is clear that over this period when contributions to special retirement saving plans were growing dramatically there was no countervailing increase in home mortgage debt.

There was an enormous increase in home mortgage debt between 1987 and 1991 for all cohorts. Although assets in personal retirement saving plans continued to grow over this period, the increase was not as rapid as over the earlier period, when mortgage debt was declining. Indeed, new contributions to special retirement saving programs declined between 1986 and 1991. Because of the 1986 cutback in the IRA program, contributions to that program fell from almost \$40 billion in 1986 to less than \$10 billion by 1991.

Contributions to all special retirement programs decreased from about \$74 billion 1986 to about \$68 billion in 1991, a decline of about 9%. Thus when contributions to these programs were growing dramatically there was a fall in mortgage debt, and when contributions to the retirement saving programs were declining there was a dramatic increase in mortgage debt. This pattern does not appear to be consistent with substitution of IRA and 401(k) assets for housing equity.

The cohort data confirm that changes in mortgage debt, as well as changes in home value, were not induced by contributions to retirement saving plans. It seems likely that the increase in mortgage debt for all cohorts after 1987 was prompted by the provisions of the Tax Reform Act of 1986 that eliminated the tax deductibility of non-mortgage debt. We consider, though, whether there was a difference in the behavior of younger and older cohorts over this period.

Figure 4d summarizes the cohort data for home equity, which of course is the difference between housing value and mortgage debt. There is a change in the cohort relationships, starting with the cohort that attained age 52 in 1984. The youngest cohorts have lower home equity than successively older cohorts up to the age 52 cohort. For example, the younger cohorts that reached ages 48 and 54 in 1991 had lower mean values of home equity than the older cohorts that attained those ages in 1984. But for older cohorts, the reverse is true, younger cohorts have greater housing equity than successively older cohorts -- at ages 60 and 66 for example. The cohort effects in home equity are very dissimilar from the cohort effects readily apparent in retirement saving assets, and thus we judge were not prompted by contributions to special retirement saving programs.

The time effects in home value and mortgage debt complicate the identification of cohort effects.²³ Nonetheless, to provide some indication of the housing equity of successively older cohorts, we have estimated cohort effects (as above) by fitting the cohort means with a function cubic in age. The results are shown in Figure 5. The first series shows estimated home value cohort effects, the second series shows cohort mortgage debt effects, and the third series shows home equity effects. The home value effects range from +25,667 for the youngest cohort to -36,407 for the oldest cohort, a difference of 62,074. The mean home value of each successively older cohort is lower than the mean for the immediately younger cohort. Interpreted literally, if there were no changes in the housing market, these data would suggest that when the current youngest cohort attains the age of the oldest cohort, the mean home value of the current young cohort will be \$62,074 more than the mean of the oldest cohort.

The home mortgage cohort effects show a similar pattern, ranging from a high of +26,180 to a low of -20,951. Again, interpreted literally, these estimates would suggest that when the youngest cohort attains the age the oldest cohort, that (future old) cohort will have \$47,131 more in mortgage debt than the current old cohort. But here it becomes clear that the projections are likely to be exaggerated. Mortgage debt is likely to be paid down. If it were completely paid off by age 72 say, then the current young cohort would be wealthier than the current old cohort -- by an amount given by the difference in their home value cohort effects (\$62,074). The key question, which is not addressed by these data, is how much the current mortgage debt of the younger cohorts will be paid down.

²³This issue is discussed further in Appendix A.

The home equity cohort effects mirror the pattern shown in Figure 4d. The estimated effect for the youngest cohort is -513, while for the oldest cohort the effect is -15,456. If the home mortgage were not reduced, the difference of \$14,943 indicates that when the youngest cohort attains the age of the oldest cohort, the young cohort would have \$14,943 more in home equity than the current oldest cohort.²⁴ On balance, the home equity cohort effects magnify the financial asset cohort differences, showing successively greater financial assets with each younger cohort. (See Figure 6.) But if, as emphasized above, all mortgage debt is reduced with age and the trend in housing value persists, the difference in the assets of the younger and the older cohorts at retirement would be more closely indicated by the difference in home value. Since most retirement assets are likely to be accumulated until retirement, if mortgage debt is paid off by retirement age, wealth at that time will include retirement saving balances plus home value.

6.2. The 401(k) Eligible Non-Eligible Comparison: Evidence on Housing Equity.

An approach that is not complicated by a coincidental growth in retirement saving and mortgage debt is to compare the assets of 401(k) eligible and non-eligible families in a given year, based on a random cross-section of respondents of all ages. As described in section 4, for this comparison to be compelling, it is important that the eligible and non-eligible households be similar with respect to saving propensity, controlling for income interval. As above, we use 1984 data, near the outset of the 401(k) program, to demonstrate the similarity of the saving propensities of eligible and non-eligible families.

²⁴The difference between the home equity cohort effects varies with age, however. The effects increase from the youngest cohort to the cohort that is age 54 in 1984. Thereafter the cohort effects decline, with successively older cohorts having less home equity. The difference between the youngest and the "C54" cohort is \$7,193. The difference between the C4 and the oldest cohort is \$22,136.

We first discuss data on trends in the housing equity of eligible and non-eligible households.

6.2.1. Trends in the Housing Equity of Eligible and non-Eligible Households.

The cohort data described above show changes in home mortgage debt and home values from 1984 to 1987 and from 1987 to 1991. Here we consider changes in home equity for families that are and are not eligible to contribute to a 401(k) plan. Table 9 shows that the trend was essentially the same for both groups. (These data show differences in the mean levels of home equity of eligible and non-eligible households, without controlling for income interval. Within income interval, the differences are typically small, as discussed below.) There was very little change between 1984 and 1987 in mean home equity for eligible or for non-eligible households. This is consistent with the cohort data, which shown a decrease for some cohorts and an increase for others. Between 1987 and 1991 there was a substantial decline in home equity for both eligible and non-eligible households. The absolute decline is larger for eligible than for non-eligible households, reflecting their larger absolute level of housing equity at the beginning of the period. The percent declines were approximately the same for both groups, about 17% for non-eligible families and about 19% for eligible families.²⁵ Given that the absolute effect of both market determined housing price changes, and the availability of home equity loans, are a function of initial housing equity, it is not surprising that the changes are roughly proportional to initial equity.

Although percent changes in mean values were about the same for both groups, the proportionate decline in medians, and other quantiles, was much greater for non-eligible

²⁵Home ownership declined 10% for non-eligible and 4% for eligible families, and mean home value of home owners declined 12% for non-eligible and 16% for eligible families.

than for eligible households. Quantile values (50th, 75th, and 90th) for eligible and non-eligible households are shown in Figure 7. Because a large fraction of households don't own a home, medians can be substantially affected by small changes in mean values. Like the means, the quantile changes between 1984 and 1987 were much smaller than the changes between 1987 and 1991. Between 1987 and 1991, median home equity for eligible households declined by 40%, the 75th percentile by 18%, and the 90th percentile by 5%; the declines for non-eligible households were 71%, 25%, and 10% respectively. These tabulations suggest that the forces that induced changes in home equity applied more or less equally to eligible and non-eligible households during the 1987-1991 period.

Table 9. Trends in Home Equity by 401(k) Eligibility, 1984, 1987, and 1991 (In 1991 Dollars).			
	Year		
	1984	1987	1991
Percent Own			
Eligible	0.78	0.75	0.75
Not Eligible	0.63	0.60	0.57
Percent Own -- Relative to 1984			
Eligible	1.00	0.96	0.96
Not Eligible	1.00	0.96	0.90
Mean Home Equity Given Own			
Eligible	70723	71189	59880
Not Eligible	61197	61688	54629
Mean Home Equity Given Own -- Relative to 1984			
Eligible	1.00	1.01	0.85
Not Eligible	1.00	1.01	0.89
Mean Home Equity			
Eligible	49747	47685	40425
Not Eligible	34073	33088	28273
Mean Home Equity -- Relative to 1984			
Eligible	1.00	0.96	0.81
Not Eligible	1.00	0.97	0.83

6.2.2. The Eligibility Comparison.

We expanded our comparison of the assets of eligible and non-eligible households at a point in time to include net housing equity. As emphasized above, the validity of this difference as an estimate of the eligibility effect depends on the similarity of the underlying taste for saving of the two groups. Again, we rely on comparison of the assets of eligible and non-eligible households near the outset of the program, in 1984, to establish the extent of similarity.

Median 1984 asset balances are shown in Table 10, by income interval.²⁶ Within an income interval, the medians control for age, marital status, and education. The assets of eligible and non-eligible families were roughly the same at the outset of the 401(k) program, regardless of whether the asset measure includes or excludes net housing equity.²⁷ There is, however, a noticeable difference for the 75,000+ income interval. There are only 83 families in the 401(k) eligible group in this interval. Because this top interval is open-ended, the incomes of eligible and non-eligible households in this interval may be quite different. For most families, net non-IRA-401(k) assets were negative or very small in 1984. Thus, any significant contributions to a saving plan, from which assets are not withdrawn, would therefore represent a net increase in the financial asset saving of most families.

²⁶The estimates are evaluated at the median of sample values for age, marital status, and education. Thus they differ from similar calculations reported in Poterba, Venti, and Wise [1995], which are evaluated at the means of control variables.

²⁷Appendix Table 3 shows that eligibles and non-eligibles had similar levels of other asset measures at the outset of the 401(k) program in 1984. We were unable to calculate conditional medians for home value and home mortgage.

Table 10. Conditional median asset balances by 401(k) eligibility and income interval, 1984. (In 1984 dollars.)							
Asset Category and Eligibility Status	Income Interval**						
	<10	10-20	20-30	30-40	40-50	50-75	75+
Net non-IRA-401(k) financial assets							
Eligible	-1288	-651	302	716	2815	6241	22068
Not eligible	-607	-348	130	775	2080	5208	17802
Difference	-681	-304	172	-60	735*	1034*	4267*
Net non-IRA-401(k) financial assets plus home equity							
Eligible	11594	16616	21371	28136	38799	53060	104748
Not eligible	11293	14398	18632	28461	36327	44462	83338
Difference	301	2218	2739	-325	2472	8598*	21410*

**Income intervals indexed to 1987 dollars. *Statistically significant at the 5% level.
Source: Authors' tabulations from the 1984 SIPP.

Median assets balances in 1991 are shown in Table 11. Although at the outset of the program, eligible and non-eligible families had approximately the same level of net financial assets, by 1991 eligible families had substantially greater median levels of net total financial asset balances, and greater levels of financial asset plus home equity balances, than non-eligible families.

The first two panels of Table 11 show net total financial assets, and net total financial assets plus home equity, for eligible and non-eligible households. For the most part, the difference between the assets of eligible and non-eligible families remains about the same when home equity is added to net total financial assets. At the outset of the program, the financial assets and the home equity of eligible and non-eligible families were about the same. They were also about the same in 1991.²⁸ The difference between median levels of net non-IRA-401(k) financial assets plus housing equity of eligible and non-eligible families is about the same as the difference in net non-IRA-401(k) financial

²⁸More detail is shown in Appendix Table 4.

assets. Thus these data suggest that the greater financial assets of 401(k) eligible families were not offset by a disproportionate reduction in the housing equity of eligible families. This is consistent with the data that show approximately equal proportional changes in the housing equity of eligible and non-eligible families between 1984 and 1991.

Table 11. Conditional median asset balances by 401(k) eligibility and income interval, 1991 (In 1991 dollars).							
Asset Category and Eligibility Status	Income Interval**						
	<10	10-20	20-30	30-40	40-50	50-75	75+
Net Total financial assets							
Eligible	1102	1073	2464	7554	17022	34726	67878
Not eligible	-483	-57	370	2307	3652	11597	39218
Difference	1585*	1130*	2094*	5247*	13370*	23129*	28660*
Net total financial assets plus housing equity							
Eligible	14509	14150	20538	32875	49361	84511	151834
Not eligible	9185	13121	15106	28502	38139	60945	122341
Difference	5324	1029	5432*	4373*	11222*	23566*	29499*
Net non-IRA-401(k) financial assets							
Eligible	-491	-262	-95	1089	3094	8838	18925
Not eligible	-327	-142	116	907	1968	5667	26909
Difference	-164	-120	-211	182	1126*	3171*	-7984*
Net non-IRA-401(k) financial assets plus home equity							
Eligible	9030	10361	14017	24168	34682	61358	108290
Not eligible	8059	11557	13522	25468	35275	56360	105924
Difference	971	-1196	495	-1300	-593	4998*	2366
**Income intervals indexed to 1987 dollars. *Statistically significant at the 5% level. Source: Authors' tabulations from the 1991 SIPP.							

6.3. The EG Between Group Results.

Using a different approach, Engen and Gale [1995] conclude that the increase in the financial assets of eligible families was offset by a reduction in home equity. We do not explore the differences between their results and ours in detail here but we do describe

the key elements of their method and provide some conjectures about possible reasons for the differences.

6.3.1. The Method.

To study the relationship between 401(k) saving and home equity, EG use an approach similar to the one described in section 3.2.1 above. They consider several between group comparisons including 401(k) eligible versus non-eligible families and 401(k) eligibles who have an IRA versus 401(k) non-eligibles who have an IRA. For illustration, we consider first the former comparison, denoting the first group by i and the second group by j . We treat the Tax Reform Act (TRA) of 1986 as a “program,” with an effect on both groups. Using the same terminology as in equation (7), the “treatment” group (i) is subject to both the 401(k) and the TRA 86 “program” effects, p_i and r_i , respectively, while the comparison group (j) is subject only to the TRA 86 effect. The difference-in-difference estimate including both the saving program and the TRA effects would be

$$(11) \quad [A_{si}(n+k) - A_{si}(n)] - [A_{sj}(n+k) - A_{sj}(n)] = \\ \quad [(m_i - m_r) - (m_j - m_r)]f(s) + \{p_i[g(n+k) - g(n)] + (r_i - r_j)[q(u+k) - q(u)]\}.$$

To make clear that years of exposure to the TRA may differ from exposure to the saving program, we let u indicate the number of years of exposure to the TRA at the first observation (e.g. 1987) and $u+k$ the number of years of exposure to the TRA at the second observation (e.g. 1991). This method will estimate the 401(k) program effect p_i if two conditions are met: $m_i = m_r$ and $m_j = m_r$, or if $(m_i - m_r) = (m_j - m_r)$, as discussed above, and $r_i = r_j$. But if the two groups have very different levels of home equity at the outset, it is unlikely that the effects of the TRA, r_i and r_j , will be equal, at least in levels.

If the comparison is between 401(k) eligibles who have an IRA versus 401(k) non-eligibles who have an IRA, the “treatment” group (i) is subject to three program effects: 401(k), IRA, and the TRA 86 -- p_i , b_i , and r_i , respectively. The comparison group (j) is subject to two program effects: IRA and TRA 86, indicated by b_j , and r_j . In this case, the difference-in-difference estimate is

$$(12) \quad [A_{si}(n+k) - A_{si}(n)] - [A_{sj}(n+k) - A_{sj}(n)] =$$

$$[(m_i - m_{i'}) - (m_j - m_{j'})]f(s) + [(p_i + b_i) - b_j][g(n+k) - g(n)] + (r_i - r_j)[q(u+k) - q(u)] .$$

The program effect p_i will be isolated if three conditions are met: $m_i = m_{i'}$ and $m_j = m_{j'}$, or if $(m_i - m_{i'}) = (m_j - m_{j'})$, $r_i = r_j$, and if $b_i = b_j$. Again, whether these conditions are approximately met is likely to depend on the initial conditions of the two groups. As emphasized above, if the initial conditions of the two groups are very different it is more likely that the program effects on the two groups will differ as well. The effects of the TRA are likely to depend on the initial home equity levels and the potential effects of the saving programs may vary in non-systematic ways with the initial financial assets of the two groups. Some committed non-savers may be completely unaffected by the programs, for example. Thus it is problematic whether any differencing procedure will adequately account for differences in the potential program responses of very dissimilar saver groups.

6.3.2. Results versus Method: Some Illustrations.

A complete understanding of why our results differ from those obtained by Engen and Gale [1995] will have to await further analysis and discussion, but we believe one explanation is their use of dissimilar groups in computing a difference-in-difference

estimator. Recall that PVW [1994a, 1995] emphasize within group estimates in the like saver group comparisons discussed in section 3. EG use a between group approach.²⁹

The within group approach that we used to evaluate the effect of the saving programs on financial asset saving may not extend satisfactorily to include housing equity. Although housing equity may be affected by 401(k) eligibility, it is also likely to have been affected by the TRA of 1986. A within group estimator cannot distinguish the separate effects of the two programs. Thus it is natural to seek a saver group affected by the TRA of 1986 but not by the 401(k) plan, with which the 401(k) group can be compared. This is what the between group estimate that EG use is intended to do. But typically their comparisons are between dissimilar saver groups with very different saving behavior. The question then is how to obtain reliable estimates from between group comparisons when the groups are so different. There may be no completely satisfactory way to do this -- other than a randomized controlled trial -- and we do not try to solve the problem.

We do, however, illustrate the issue using data for the saver groups defined by 401(k) eligibility and IRA participation status. These groups were considered by PVW in the "like" group analysis discussed in section 3. (The financial asset data for the groups are shown in the bottom panel of Table 4.) In their within like group analysis, PVW emphasized that there was virtually no change between 1984 and 1991 in non-IRA-401(k) assets of 401(k) eligible households, 401(k) eligible households with an IRA, or 401(k) eligible households without an IRA (or in the non-IRA-401(k) assets of 401(k) non-eligible households with an IRA). Yet for each of these saver groups there was a large increase

²⁹Engen and Gale [1995] cite Poterba, Venti, and Wise [1995] as the source for their method. This is a misunderstanding of our "like" group method which relies on within group comparisons.

in total financial assets. Based on between group comparisons, EG argue that the increase in financial assets between 1987 and 1991 was offset by a reduction in home equity.

But this conclusion depends critically on whether the groups compared are similar or dissimilar, as the data in Table 12 show. The question is whether there was a differential effect of the TRA of 1986 on the 401(k) eligible households, compared to the “comparison” households. Consider first 401(k) eligibles with an IRA compared to 401(k) non-eligibles with an IRA. These two groups had similar housing assets in 1984 at the outset of the 401(k) program and experienced similar declines in housing equity (10.2% for the eligible and 10.6% for the non-eligible group). In this case, the dollar declines were about the same as well. The decline for 401(k) eligible households was \$379 greater than the decline for non-eligible households. Were one to assume that this is the decline due to 401(k) eligibility -- which we would not -- this amount would offset very little of the increase in the total financial assets of 401(k) eligible households between 1987 and 1991.

Saver Group	1984	1987	1991	Difference 1984 to 1991		Difference 1987 to 1991	
				Percent	Level	Percent	Level
All Families:							
401(k) eligible	32658	28743	24230	-25.8%	-8428	-15.7%	-4513
401 (k) not eligible	18699	16469	14215	-24.0%	-4484	-13.7%	-2254
Difference-in-difference				-1.8%	-3944	-2.0%	-2259
Equal % \$ equivalent					-597		-579
Families with an IRA:							
401(k) eligible	52621	48451	43531	-17.3%	-9090	-10.2%	-4920
401 (k) not eligible	46385	42913	38372	-17.3%	-8013	-10.6%	-4541
Difference-in-difference				0.0%	-1077	0.4%	-379
Equal % \$ equivalent					0		207
Families without an IRA:							
401(k) eligible	22905	19704	15578	-32.0%	-7327	-20.9%	-4126
401 (k) not eligible	12399	10575	8696	-29.9%	-3703	-17.8%	-1879
Difference-in-difference				-2.1%	-3624	-3.2%	-2247
Equal % \$ equivalent					-486		-625

When very dissimilar groups are compared, however, this approach can yield misleading results. For example, even though all 401(k) eligible and all 401(k) non-eligible households experienced similar proportional declines in housing equity (15.7% versus 13.7%), the dollar declines were very different (\$4,513 versus \$2,254) because the two groups had very different levels of housing equity at the outset of the program. Thus it is misleading to ascribe the greater decline in the housing equity of the 401(k) eligible group to 401(k) eligibility per se, as the between group comparison does. The greater decline for eligibles may simply reflect their larger initial housing equity. The groups also had very different levels of financial assets. This is why we emphasize within group comparisons, and avoid inferences based on between group comparisons, in our “like” group analysis.

One way to estimate the reduction in housing equity attributable to the 401(k) program would be as the difference between the actual reduction in home equity (28,743 - 24,230) and the reduction had the treatment group experienced the same percent reduction as the non-eligible group. This yields an estimate of \$576, which is small compared to the increase in the financial assets of this group. This approach seems plausible in this case because programs that affect housing values and mortgages are likely to have effects proportional to initial housing values.

7. SCF: SUMMARY DATA AND GALE AND SCHOLZ PARAMETRIC ANALYSIS

In section 2.4. above, we discussed the change in the financial assets of IRA contributors between 1983 and 1986, as their IRA assets accumulated. We concluded that these data, from the Survey of Consumer Finances, showed no substitution of IRA contributions for other forms of saving, and that the IRA contributions between 1983 and 1986 represented largely new saving. Based on parametric analysis of the same data,

Gale and Scholz [1994] (hereafter GS) concluded that virtually none of the additional IRA saving resulting from an IRA limit increase would be new saving. This result has often been interpreted to imply that none of the IRA saving undertaken during the 1983 to 1986 period represented new saving, although GS are careful to emphasize that their analysis pertains to a limit increase. They conclude that, of the increase in IRA contributions resulting from an increase in the limit, 31% would be financed by lower taxes, 2% would be funded by a decrease in consumption, and 67% would come from a reduction in other saving. Our conclusion and that of Gale and Scholz are not necessarily inconsistent although it seems unlikely that they could both be true. Thus we now consider what lies behind our different conclusions.

We first discuss the data upon which our results in section 2.4 and the GS results are based. We consider the deletion of observations that preceded the GS estimation and we draw attention to the potentially important effect of sample selection on the GS results. Then we discuss more carefully the change in the non-IRA saving of contributors that would have been expected in the absence of the IRA program. We conclude that it is virtually impossible that IRA contributions between 1983 and 1986 came entirely from a reduction in other saving. We consider whether this conclusion could be consistent with the possibility that an increase in the IRA limit would result in no new saving. Finally, we explore the GS estimation procedure in detail, and find that their results are not a robust reflection of the SCF data, but rather are an artifact of their specification and estimation procedure.

7.1. The Data.

The data and the GS estimation sample are described with reference to Table 13. For background, several features of the SCF data are important: (1) Only households who

were surveyed in both the 1983 and 1986 waves of the SCF can be used in the analysis. (2) Some households are deleted because they did not meet the criteria for IRA participation or for other reasons were unlikely to contribute to an IRA. (3) The SCF data are comprised of two samples: an “area probability” sample and a “high income” sample that over-sampled high income households. (4) Whether a family contributed to an IRA during the period is not reported in the SCF, so contributor status must be inferred. (5) Flow saving in other conventional forms must also be inferred from the reported levels of assets in 1983 and 1986. Although we have been unable to match exactly the GS estimation sample, we believe that the differences do not materially affect the conclusions that we draw below.

IRA Status and Area Probability versus High Income Samples	Using 1986 IRA Balance		Using GS Contributor Assignment		
	Without Savings GS Deletions	With GS Savings Deletions	Without GS Savings Deletions	With GS Savings Deletions	
				PVW Replication	GS Estimation Sample
All	1670	1486	1670	1486	1483
Area Probability	1489	1445	1489	1445	
High Income	181	41	181	41	
Non-Contributors	1021	988	1099	1038	1035
Area Probability	996	982	1042	1025	
High Income	25	6	57	13	
Contributors	649	498	571	448	448
Area Probability	493	463	447	420	
High Income	156	35	124	28	
Within Limits			403	341	331
Area Probability			340	324	
High Income			63	17	
At Limit			168	107	117
Area Probability			107	95	
High Income			61	11	

The 1983 SCF sample included 4262 respondents, of whom 3824 were in the area probability sample and 438 in the high income sample. Of these, 2791 were surveyed in

1986 as well.³⁰ Excluding families in which either the respondent or the spouse were self-employed, the age of the head was less than 25 in 1986, the age of the head was greater than 65 in 1983, or there was a change in marital status between 1983 and 1986, leaves a total of 1670 households. The first column of Table 13 gives a breakdown of this sample by area probability versus high income sample status and by IRA contributor status.

Because the SCF reports IRA balances but not annual contributions, to determine whether a household contributed to an IRA between 1983 and 1986 requires assumptions about the return on assets, as well as other conventions. The GS assumptions are explained in appendix A to their paper. The first column in Table 13 shows the number of observations using whether or not the family had an IRA balance in 1986 to indicate IRA contributor status. Of course some of these households could have contributed in 1982 but not thereafter and thus were not active contributors in the 1983-1986 period. There were 1021 households without an IRA balance in 1986 and 649 with a positive IRA balance. Of the households with an IRA balance in 1986, 24% of the respondents (156 of 649) were from the high income sample. Appropriately weighted, only 4.5% of all contributors would be from households in this high income group.

The third column shows the breakdown of "contributors" and "non-contributors" based on the GS contributor status assignment conventions. They use balances in 1983 and 1986 together with an assumed rate of return on 1983 balances to infer new contributions between the two years. These assignments yield fewer contributors than the number of households with a positive 1986 IRA balance (571 versus 649), as expected.

³⁰The 1983 survey was conducted between February and August 1983, with the majority of interviews in March and April. Thus, 1983 IRA balances represent 1982 contributions for the most part. The 1986 survey was conducted in June through September 1986, so 1986 IRA balances represent contributions through the 1985 tax year.

GS also use their assumptions to allocate households to limit contributor status (those with estimated three-year contributions greater than the estimated three-year limit) and nonlimit contributor status.

To estimate their model, GS eliminate a large number of additional observations, those with 1983 to 1986 estimated saving less than -\$100,000 or greater than +\$100,000. The resulting sample is labeled "With GS Savings Deletions" in Table 13. Their procedure removes 184 of 1670 households, 61 non-contributors and 123 contributors. Of the 168 limit contributors, 61 are removed, including 82% (50 of 61) of the high income sample limit contributors. Estimates based on the remaining 107 limit contributors determine the results of the GS estimation procedure. We were unable to replicate exactly the 117 observations used by GS. The last column shows the number of observations used in the GS estimation procedure, as reported in their paper.

These deletions have an enormous effect on the distribution of saving and assets in the estimation sample and on formal parameter estimates, as GS show. Mean and median estimated saving between 1983 and 1986 with and without these deletions, as well as non-IRA financial assets with and without the deletions, are shown in Table 14. As the table shows, the sample deletions have enormous effects on the sample means. The GS estimation procedure is based on means, and their results are essentially determined by the few limit contributors in the sample, so sample deletions can have an enormous effect on the results. While in principle there is nothing wrong with "trimming" the data, we show below that the key parameter estimate is extremely sensitive to exactly which observations are deleted and the sample deletions that are made essentially determine the conclusions that GS report.

Table 14. Change in mean and median saving and 1983 non-IRA financial assets with and without sample deletions.		
Measure and Deletion Choice	Mean	Median
Saving:		
Without GS Deletions	-13,303	1,132
With GS Deletions	2,378	1,044
1983 Non-IRA Financial Assets		
Without GS Deletions	217,668	5,700
With GS Deletions	22,244	4,222

7.2. A Simple Reality Check.

In section 2.2. above, we discussed summary data (Table 3) based on these same SCF surveys. We return to a similar discussion here, based on data reported in Table 15. As explained above, using the 1983 and 1986 waves of the SCF it is possible to compare the asset balances of the same households over time. We begin with respondents to the 1986 survey. We exclude households with self-employed members and households with a change in marital status between 1983 and 1986. There are two reasons why the values in Table 15 may differ from “comparable” values reported by GS: First, we restrict the sample to all households between ages 25 and 65 in 1986. Because GS limit their sample to households with heads 65 and under in 1983, some heads are as old as 68 in 1986. We believe that our sample is a better representation of the pool of potential contributors, i.e. non-retirees. Second, we also use a narrower definition of financial assets, including only those assets that we believe are most likely to be substituted for IRAs. Our measure includes checking accounts, statement, passbook, share draft, and other saving accounts; stocks and mutual funds, saving bonds, corporate, municipal and all other bonds; money market accounts, and CDs. The GS measure includes in addition the cash value of life insurance, trusts, managed investment accounts, and notes and land contracts owed to

the household.³¹ In addition, the summary data reported by GS in their tables 1 to 3 are based on different age criteria than their estimation sample, including all households over 25, even those over 65.³²

Table 15. Survey of Consumer Finances summary data, using GS estimation sample definitions.			
Contributor Status and Asset	Year		% Change
	1983	1986	
Medians			
Contributors in 1986:			
Non-IRA Assets	8,800	13,400	52.3
IRA Assets	600	6,857	1043.0
Total Financial Assets	11,800	23,000	94.9
Non-Contributors:			
Total Financial Assets	750	1050	40.0
Median of Natural Logarithms*			
Contributors in 1986:			
Non-IRA Assets	9.083	9.503	42.0
IRA Assets	6.397	8.833	243.6
Total Financial Assets	9.376	10.043	66.7
Non-Contributors:			
Total Financial Assets	6.620	6.957	33.7
*The percent change is approximated by the difference in the logarithms. Source: Authors tabulations using the 1983 and the 1986 SCF.			

Median IRA and non-IRA financial asset balances in 1983 and 1986, and the change in balances between these years, are shown in the first panel of Table 15, by whether the respondent had a positive 1986 IRA balance. The table also shows total assets of contributors, including both IRA and non-IRA balances. (In anticipation of estimation results discussed below, the table also includes the median of the logarithm of

³¹These additions change the magnitude but not the pattern of the data reported in Table 15. Data based on the GS definitions are shown in Table 17 below.

³²In addition, although the pattern revealed by the data is the same in both cases, the values reported here differ slightly from the numbers reported in Table 3 for three reasons: (1) The 1986 SCF combines IRA and Keogh balances. GS present a method for inferring the 1986 IRA balance based on the 1983 response and we use the GS method here. (2) We use the newer set of sample weights here. (3) Here we use the GS definition of a change in marital status.

assets.³³) Several features of the data are important: (1) The median 1983 non-IRA asset balance of households with IRA accounts in 1986 was \$8,800. Clearly, prior to 1983, this group had not been accumulating assets at the rate of the typical household IRA contribution, about \$2,300 per year. (2) The \$6,257 increase in IRA balances (from \$600 in 1983 to \$6,857 in 1986) clearly was not funded by transferring funds from the 1983 balance in non-IRA accounts, which was only \$8,800 at the beginning of the period. (3) Indeed, the non-IRA assets of contributors did not decline at all as IRA assets increased between 1983 and 1986. On the contrary, they increased over 52%, from \$8,800 to \$13,400.

Without the IRA program, what increase in the 1983 non-IRA asset balance would have occurred over the next 3 years? The observed 52.3% increase was equivalent to an annual asset growth rate of 15 percent. If IRA contributions were funded either by withdrawing funds from non-IRA balances, or by reducing new saving in non-IRA assets, then the increase in non-IRA assets between 1983 and 1986 should have been much less than would have been expected in the absence of the IRA program. That is, the expected increase in non-IRA assets should have been much more than the observed increase -- from \$8,800 to \$13,400. We consider a simple prediction of asset growth in the absence of IRAs.

Assets tend to increase with age and income. A simple way to estimate the expected increase in non-IRA assets between 1983 and 1986 is to predict the increase based on the 1983 relationship between age and income on the one hand and assets on the other, with some allowance for change in the rate of return on assets.

³³When logarithms are used, assets of zero are set to 1 and assigned a logarithm of 0.

Even simple estimates of the income-asset profile are confounded by the nature of the data. There is enormous “residual” variance with respect to assets. For example, a linear regression of assets on age and income yields an R^2 value of about 0.06 with a residual standard deviation of \$2,400,000. If the data are weighted by the appropriate sampling probabilities, the R^2 is about .07 and the residual standard deviation is about \$540,000. (This portends the finding that sample deletions can have an enormous influence on the results.) In addition, the data exhibit enormous heteroskedasticity, which we attempt to correct for by using a semi-log specification of the form:

$$(13) \quad \text{Log } A_{83} = a + b(\text{Income}_{83}) + c(\text{Age}_{83}) + e .$$

The predictions for 1986 are based on

$$(14) \quad \text{Log } A_{86} = \text{Log } A_{83} + b(\text{Income}_{86} - \text{Income}_{83}) + c(\text{Age}_{86} - \text{Age}_{83})$$

Thus the predictions account for the change in assets associated with an increase in age between 1983 and 1986, and the change associated with the change in income.

Predictions based on equation (14) are shown in Table 16. The predicted increase in median assets based on a weighted median regression is about 23%, which is less than the actual increase (based on the difference in logs) of about 45%. The predicted increase in mean assets is about 29%, compared to an actual increase of about 56%.

Table 16. Logarithm of 1983 Median and Mean Actual Assets and Predicted Median and Median 1986 Assets			
Measure	1983	1986	
		Predicted	Actual
Weighted			
Median	9.048	9.280	9.503
Mean	8.786	9.072	9.344
Unweighted			
Median	9.598	9.817	9.957
Mean	9.594	9.839	10.555

Thus, we predict a 1986 non-IRA asset level that is lower than the observed level, not higher than the observed level as would be expected if IRA contributions simply substituted for saving that would have occurred anyway.

The 1983 cross-section regression implies a difference in the assets of families by age. We want to predict the increase for families who age three years and whose earnings change over these three years. The 1986 prediction based on the 1983 cross-section estimates accounts for the increase in age between 1983 and 1986, and it accounts for the change in earnings by using in the prediction the 1986 earnings of the respondents. But the prediction does not account for any change due to the return on initial asset holding. Inherent in the 1983 regression estimates of the difference in assets of people differing in age by three years is also a rate of return, but for an earlier period. If the prior return differs from the 1983-1986 return, the projected asset increase may not apply to this later period. The magnitude depends on the difference between the prior and ex post rates of return. Consider, for example, the AAA bond rate in 1980 through 1982 versus 1983 versus 1985. The average during the first period was 13.30 and during the second period 12.04. So correction for the rate of return would reduce the estimated increase. The return on other assets may give a different sign; more detail on this issue is presented in Appendix B.

7.3. The Saving Effect of Program Contributions versus the Saving Effect of a Limit Increase.

The foregoing analysis suggests to us that it is very unlikely that the bulk of IRA contributions were financed at the expense of withdrawals from non-IRA accounts, or from a reduction in new saving in non-IRA accounts. Indeed, if anything, the data taken at face value suggest that other saving increased as IRA contributions increased. Thus from these

data alone we would argue that the contributions under the existing program represented largely new saving. Yet GS conclude, and they say explicitly in their paper, that an increase in the IRA limit would not increase saving. Here we consider the summary data that GS highlight in foretelling their formal results. In particular, we consider how the inferences that GS draw from the summary data can be so different from our judgements based on the same data.

GS argue that limit contributors in particular, but non-limit contributors as well, had substantial non-IRA financial assets in 1986. The implication is that if the limit were raised, these families could easily fund an IRA by transferring assets from non-IRA to IRA accounts without increasing net saving, and that because they could do that, they would. We emphasize the low level of non-IRA assets in 1983, at the outset of the program, and the increase in these non-IRA assets as IRA contributions were accumulating. We infer from these data that the IRA accumulation could not possibly have been funded by withdrawing funds from non-IRA balances and was unlikely to have been funded by reducing new non-IRA saving that otherwise would have occurred. We want to understand what accounts for the difference between the \$8,800 level that we emphasize and the \$41,269 for limit contributors emphasized by GS. Part of the difference is simply their emphasis on 1986 assets versus our emphasis on 1983 assets. Part of the difference is the definition of non-IRA financial assets. Part of the difference comes from differences in the meaning of limit contributor. We consider the last issue first and then turn to differences in financial asset definitions.

7.3.1. Limit Contributors Versus All Contributors.

We have framed our judgements in terms of the addition to net saving represented by the contributions of all contributors, both limit and non-limit. But because such a large

fraction of contributions were at the limit, we believe that a higher limit would have led to still greater net saving. Based on an analysis of 1983 tax returns, Burman, Cordes, and Ozanne [1990] find that 75.3% of all IRA contributions were at the family limit and that an additional 11.3% were at the limit for one spouse in households filing joint returns. Engen, Gale, and Scholz [1994] report that 63.3% of contributions were at the family limit over the period 1982 to 1986. With such a large fraction of households at the limit, if all limit contributors funded IRA contributions by transferring funds or by reducing other saving, the summary data would show that. But they don't. We infer, therefore, that if the limit had been higher, we would have seen a greater increase in assets by 1986 than actually occurred.

In considering the effect of a higher limit, the number of individual annual contributions made at the limit is the relevant statistic. Presumably each contribution at the limit would have been at least somewhat greater had the limit been higher. Thus we should have in mind that between 60% and 85% of contributions are by families in this category.

GS point to an entirely different measure, suggesting that only 22% of contributions are at the limit. The families that GS call "limit contributors" are those who are assigned limit status in each of three consecutive years (1983 through 1985) based on their assignment criteria. They report 21.8% at the limit based on these criteria. Thus the actual proportion of contributions at the limit is three or four times as large as the proportion assumed by GS. In considering whether persons at the limit would have contributed more, and saved more, if the limit had been higher, recognizing that a much larger share of contributions are at the limit may well alter one's prior expectations about the saving effect of raising the limit.

7.3.2. Assets of All Contributors and Limit Contributors.

To understand the differences in the asset levels that we emphasize and those reported by GS, begin with the non-IRA financial asset values reported in Table 15, which are \$8,800 and \$13,400 for 1983 and 1986 respectively. Following the presentation of the data in their table 3, GS would emphasize the 1986 balance, corresponding to \$13,400. The 1986 balance reflects the increase in non-IRA assets during the time that IRA assets were accumulating.³⁴

The GS asset definition also differs from ours. Based on the GS non-IRA financial asset definition -- but still considering the assets of households with IRA balances in 1986, not the GS assignment procedure -- we find the values reported in Table 17, for households aged 25 to 65 in 1986.³⁵ The pattern is the same as that reported in the top panel of Table 15. In particular non-IRA assets by this definition are \$13,085 and \$19,000 in 1983 and 1986 respectively. But even including asset balances from which we believe IRA contributions are unlikely to be taken, the \$13,085 balance in 1983 suggests that contributors had not previously been accumulating assets at the rate of \$2,300 annually.

The asset balances reported by GS differ in still other respects from those in Table 17. GS report 1986 non-IRA financial assets of \$21,695 for households with inferred contributions between 1983 and 1986. In addition to the inferred contributor definition, this estimate incorporates a broader age range, including all persons over 24, even those who

³⁴Indeed, in the 1990 version of their paper, GS emphasized that the level of 1986 non-IRA assets of contributors was \$13,500, very close to our measure of \$13,400.

³⁵GS [1994] uses an expanded definition on non-IRA financial assets that includes the cash value of life insurance, trusts, managed investment accounts, and note and land contracts owned to the household. This sample also includes all households over 25 in 1986, including those over 65.

are over 65 and unlikely to make IRA contributions. Because older households tend to have greater assets than younger households, expanding the upper age limit may significantly affect the results.

Table 17. SCF Summary Data: Using GS Definitions			
Contributor Status and Asset	Year		% Change
	1983	1986	
Medians			
Contributors in 1986:			
Non-IRA Assets	13,085	19,000	45.2
IRA Assets	600	6,857	1043.0
Total Financial Assets	14,100	30,000	112.8
Non-Contributors:			
Total Financial Assets	1,200	2,269	89.1
Source: Authors' calculations based on the 1983 and 1986 SCF.			

Based on the three-year inferred limit criterion, the more inclusive definition of financial assets, and the all-inclusive age range, GS report median financial assets of their “three-year limit” contributors of \$41,269.³⁶ Because GS so severely underestimate the proportion of contributions at the annual limit, the assets of the much larger number of persons who actually contribute at the limit is probably lower than this. But there is no data-based value to compare with this figure, since the SCF does not report contributions between 1983 and 1986. Using the 1983 and 1986 Consumer Expenditure Surveys, we calculate that the median non-IRA financial assets of limit contributors was \$14,250 in 1983 and \$19,500 in 1986.³⁷

³⁶Based on our approximation to their sample, prior to their saving deletions, their estimation sample included approximately 168 persons at this limit. The number used to obtain this median asset balance will be somewhat larger than this because it included more older households.

³⁷The CES financial asset definition matches approximately the definition used in Table 12, but is less inclusive than the definition used by GS.

In our view, the summary data reported in Table 15 suggest that most contributions between 1983 and 1986 represented a net addition to saving, and, in addition, are inconsistent with the possibility that families who contributed at the IRA limit did not increase net saving. Since a large fraction of contributions are at the limit, most of the increase in non-IRA financial assets as IRA contributions were accumulating must be attributed to limit contributors. Thus it is implausible that if the limit had been higher, these limit contributors would not have increased total saving still more. GS emphasize large financial asset values for contributors by citing assets in 1986 instead of 1983, by using a broader definition of financial assets, and, in the case of limit contributors, by citing the assets of families at their "three-year limit," rather than at the annual limit.

7.4. The GS Model.

GS present estimates of a formal model that they believe supports the view that limit contributors would not increase net saving if the limit were increased. We explain here why we believe that their conclusion is not supported by the data. We first describe their model, and then present and discuss two stage "consistent" estimates of it that may be easier to understand than the full maximum likelihood estimates. Finally we explore directly the properties of the maximum likelihood specification used by GS. In each case, we decompose the specification to identify its critical features. The important features of the specification may be lost when looking at the whole, but are easily discerned if the procedure and results are decomposed and built up step by step. We believe that the key GS results are an artifact of their model specification.

The GS specification can be written as

$$(15) \quad I^* = X\beta + u$$

$$(16) \quad S = X\gamma_1 + \epsilon_1 \quad \text{if } I < 0$$

$$(17a) \quad S = X\gamma_2 + \epsilon_2 \quad \text{if } 0 < I < L$$

$$(17b) \quad \begin{aligned} S &= X\gamma_2 + \eta(I^* - L) + \epsilon_2 \\ &= X\gamma_2 + \eta(X\beta - L) + \eta u + \epsilon_2 & \text{if } I > L \\ &= X\gamma_2 + \eta X\beta - \eta L + \eta u + \epsilon_2 \\ &= X(\gamma_2 + \eta\beta) - \eta L + \eta u + \epsilon_2 \end{aligned}$$

$$(18) \quad \eta = X\delta$$

where I^* represents desired IRA saving and S represents non-IRA financial assets saving. The variables X are a set of household attributes. The specification can be described with reference to Figure 8. The line labeled I^* represents equation (15), which is desired IRA saving and is limited at L . Other financial asset saving of IRA contributors is represented by the line labeled S and then S_L , and has two parts. Up to the kink point -- associated with some X^* -- it reflects equation (17a). We sometimes refer to this component as “underlying” saving. The change in slope after X^* reflects equation (17b) and in particular the value of η . The segment after the kink -- with the steeper slope -- is labeled S_L in the figure. The steeper slope recognizes the possibility that desired IRA saving in excess of the limit may be made up by increasing other saving. If η equals 1, the difference between desired IRA saving and the limit L is the same as the difference between S and S_L . Thus, the key parameter is η .³⁸ Equation (14), which describes the other saving of non-contributors, is essentially irrelevant in this specification and is not represented in the figure. In fact, in the GS specification, the saving of non-IRA contributors provides no information about the substitution between IRA and non-IRA saving.

³⁸As GS point out, the difference in the tax treatment of IRA and conventional saving could lead to values of η greater than 1. For example, if households wanted to reach a given asset goal by retirement age, the amount saved in conventional forms would have to be greater than the amount saved in an IRA.

An identical figure was used by Venti and Wise to describe their specification (e.g. Figure 4.2 in Venti and Wise [1991]). Thus in spirit the two specifications are very similar. But here the similarity ends. The method used to identify the change in the slope of S after the IRA limit is reached is very different in the two approaches.

How is η identified? We show below that in the GS model specification a downward bias in the “underlying” other saving function for limit contributors is balanced by an upward bias in η . One may conclude therefore that the estimation procedure does not identify an η , the key behavioral parameter. Before turning to this matter, however, we consider how η might in principle be identified and what the estimated value might mean.

Assume that β -- the effect of X variables on IRA contributions -- is identified from equation (15). Then η is identified by variation in L. Since most variation in L is due to marital status, η is determined in large part by marital status. If the limit for married couples is \$2,000 higher than the limit for single persons, for example, and if $\eta = 1$, married limit contributors should save \$2,000 less in other financial assets than single persons. If this is not the difference, in principle, η would change accordingly. Of the 107 limit contributors in our estimation sample, 80 are married and 27 are single. This does not provide much evidence on which to base an estimate of η . In addition, the specification assumes that marital status does not influence other financial asset saving, except through the lower IRA spillover effect. This raises a further confounding issue. Marital status is not allowed to affect other saving directly, nor to affect IRA saving. Thus, it only enters the equation for limit contributors. If, as most prior research suggests, marital status should properly be an explanatory variable in the underlying other saving equation, then, for limit contributors, η will pick up this effect, in addition to any spillover effect.

Although marital status is critical in the identification of η , GS add to the specification some complexity, which also influences the estimated value of η . They allow η to be a function of covariates X , with $\eta = X\delta$, but they do not allow a constant term in the relationship. (The absence of a constant is associated with an important error in the GS interpretation of the results, and this issue is taken up below.) Now the specification for limit contributors becomes $X\gamma_2 + (X\delta)(X\beta) - (X\delta)L = X[\gamma_2 + (X\delta)\beta] - (X\delta)L$. In this case, the effect of marital status is allowed to depend on other covariates.³⁹

To demonstrate the critical features of the GS model we begin with a simplified version, using only income as an explanatory variable. Then we proceed to the full GS specification. We use both a two-stage procedure that provides consistent estimates under the GS assumptions, and the maximum likelihood procedure used by GS. All of our estimates are based on our replication of the GS sample, described in Table 13. We use all the GS variable definitions as well as their procedure to assign limit and non-limit contributor status.

7.4.1. A Two-Step Procedure.

Suppose that equation (15) is estimated independently using a Tobit specification to obtain estimates of the β coefficients. The other saving equation is

$$(19) \quad S = X\gamma_2 + \eta(X\hat{\beta} - L) * D + (\eta u * D + \epsilon_2),$$

where D identifies limit contributors. Under the assumptions of the GS model, the expected value of S can be written as

$$E(S) = X\gamma + \eta(X\beta - L) * D + \eta E(u|u > L - X\beta) * D \\ + E(\epsilon_2|u > L - X\beta) * D + E(\epsilon_2| -X\beta < u < L - X\beta) * (1 - D).$$

³⁹The multiplication of terms in X is also likely to make identification tenuous.

Using standard results, we can write

$$(20) \quad E(u|u > L - X\beta) = \sigma_u \frac{\phi\left(\frac{L - X\beta}{\sigma_u}\right)}{1 - \Phi\left(\frac{L - X\beta}{\sigma_u}\right)} = \sigma_u \lambda_1$$

$$(21) \quad E(\epsilon_2|u > L - X\beta) = \rho\sigma_\epsilon \frac{\phi\left(\frac{L - X\beta}{\sigma_u}\right)}{1 - \Phi\left(\frac{L - X\beta}{\sigma_u}\right)} = \rho\sigma_\epsilon \lambda_2$$

$$(22) \quad E(\epsilon_2| -X\beta < u < L - X\beta) = \rho\sigma_\epsilon \frac{\phi\left(\frac{L - X\beta}{\sigma_u}\right) - \phi\left(\frac{-X\beta}{\sigma_u}\right)}{\Phi\left(\frac{-X\beta}{\sigma_u}\right) - \Phi\left(\frac{L - X\beta}{\sigma_u}\right)} = \rho\sigma_\epsilon \lambda_3$$

Under the assumptions of the GS model, consistent estimates can be obtained by estimating

$$\begin{aligned} S &= X\gamma + \eta(X\beta + \sigma_u \lambda_1 - L) * D + D * \rho\sigma_\epsilon \lambda_2 + (1 - D) * \rho\sigma_\epsilon \lambda_3 \\ &= X\gamma + \eta(X\beta + \sigma_u \lambda_1 - L) * D + \rho\sigma_\epsilon [D * \lambda_2 + (1 - D) * \lambda_3] \end{aligned}$$

where β and σ_u are estimated from the first stage Tobit equation and γ , η , and $(\rho\sigma_\epsilon)$ are estimated in the second stage.

We have estimated several specifications using this procedure and the estimated values of η are reported in Table 18. The equations use a variety of covariates X , ranging from income only to the full set of variables used by GS, and use several different specifications of η . There are two important features of these results. First, in the most

inclusive specifications, η is small. The estimate is 0.10 in specification 4 and 0.192 in specification 8. Specification 8 is analogous to the GS specification. If only a constant is included in the set of explanatory variables η , the resulting estimate is $\eta = 0.103$. Second, the inclusion of 1983 non-IRA financial assets (NIFA) produces a large jump in the estimated value of η . This is a feature of both the two-stage procedure and the maximum likelihood procedure discussed next.

#	Variables		Estimated η	Standard Error
	In IRA and Other Saving Equations	In η		
1	Income	Constant	-1.193	1.879
2	Income & NIFA	Constant	0.835	1.828
3	All Xs Excl NIFA	Constant	-1.668	1.762
4	All Xs Incl NIFA	Constant	0.103	1.712
5	Income	Income, No Constant	-1.591*	
6	Income & NIFA	Income & NIFA, No Constant	-0.413*	
7	All Xs Excl NIFA	All Xs Excl NIFA, No Constant	-1.889*	
8	All Xs Incl NIFA	All Xs Incl NIFA, No Constant	0.192*	
9	All Xs Excl NIFA	All Xs Excl NIFA, Plus Constant	-1.668*	
10	All Xs Incl NIFA	All Xs Incl NIFA, Plus Constant	0.103*	

*Evaluated at the mean of predicted η for limit contributors.

Why does the inclusion of the 1983 level of NIFA lead to such large changes in η ? For illustration, we compare specifications 7 and 8. Non-IRA financial assets are included presumably to control for past saving behavior. In specification 8, aside from income, NIFA is the only statistically significant variable of the 11 variables in the non-IRA saving equation for IRA contributors. Its estimated coefficient is negative, -3.447 with a t-statistic of 9.429. Taken literally, this result says that the greater the level of non-IRA assets in 1983 -- controlling for age, income, and other variables -- the lower the level of saving over the subsequent three years. Thus NIFA is clearly not serving as a control for past saving behavior. Instead, it seems apparent that the coefficient reflects enormous error in the

measurement of 1983 non-IRA financial assets. Recall that saving is inferred from 1983 and 1986 NIFA balances, using a variant of $S = NIFA_{1986} - NIFA_{1983}(1+r)^3$. Thus any error in $NIFA_{1983}$ will impart a negative bias to the coefficient on this variable and in this case the measurement error is surely very large -- large enough to more than offset the intended role of NIFA as a control for heterogeneity.

The effect of NIFA on η can be understood by considering the components of the specification that determines η . It is useful to recall that to a first order of approximation, η can be thought of as

$$\eta = [(Actual\ other\ saving) - (Underlying\ other\ saving)] / (I^* - L),$$

where “underlying” other saving refers to other saving as a function of X before the IRA limit is reached. The important aspect of this formula is that, given actual other saving, if the prediction for underlying saving is arbitrarily low, η will be arbitrarily large to compensate for the low underlying saving. In particular, if predicted underlying saving for limit contributors is lower than the actual saving of limit contributors, the shortfall between underlying and actual saving can be bridged by a large value of η . We show that adding NIFA to the specification yields implausible estimates of underlying saving and corresponding large offsetting increases in η . The first column in Table 19 shows the mean value of estimated $X\gamma$. The second column shows $X\gamma + \rho\sigma_\epsilon[D*\lambda_2 + (1-D)*\lambda_3]$ and represents the predicted value of underlying saving for limit and non-limit contributors. The third column shows predicted values of other saving, accounting for the upper slope component of other saving for limit contributors.

It's easy to see from this table why the effect of NIFA is so large. Without NIFA (specification 7) the systematic component of underlying saving ($X\gamma$) for non-limit contributors is predicted to be -2452 and for limit contributors 759. When NIFA is added

the predicted underlying saving for limit contributors is reduced from 759 to -1959 and for non-limit contributors is increased from -2452 to 682. Including the λ_2 and λ_3 terms, without NIFA, predicted underlying saving for limit contributors is substantially higher than for non-limit contributors (10,782 versus 2,084). But when NIFA is added, predicted underlying saving of limit contributors is reduced from 10,782 to 1,287 and underlying saving for non-limit contributors is increased somewhat from 2,084 to 2,149. With reference to Figure 8, the underlying saving function for limit contributors is lowered. In particular, predicted underlying saving is now lower for limit than for non-limit contributors, which is inconsistent with the prevailing heterogeneity concern, i.e. that contributors want to save more in all forms than non-contributors and that limit contributors want to save more than contributors. This means that η , the slope of the portion above the kink point, must be increased to compensate for the low predicted value of underlying saving for limit contributors. Indeed, the difference is made up by the larger η . The average predicted value of total other saving for limit contributors is close to the actual average -- 2920 versus 3089. Taken literally, specification 7 says that limit contributors save more in non-IRA financial assets than non-limit contributors, as might be expected. But specification 8, with NIFA included, says just the reverse, that limit contributors have a lower propensity than non-limit contributors to save in other forms.

This feature of the specification is created by the large negative coefficient on NIFA, which seems clearly to reflect error in measurement and not, as intended, a control for saving propensity. Thus the error in measurement of this variable imparts substantial bias to the results. We show below that this feature of the specification applies equally to the joint maximum likelihood estimation.

Table 19. Predicted underlying other saving, total other saving, and actual other saving, for limit and non-limit contributors, based on two-step estimates.				
Specification	Predicted Underlying Other Saving -- X γ Only	Predicted Underlying Other Saving -- X γ Plus λ_2 and λ_3	Predicted Other Saving -- X γ Plus λ_2 and λ_3 and λ_1	Actual Other Saving
Non-Limit Contributors				
Specification 7	-2452	2084	2084	1989
Specification 8	682	2149	2149	1989
Limit Contributors				
Specification 7	759	10782	2455	3089
Specification 8	-1959	1287	2920	3089

7.4.2. Joint Maximum Likelihood Estimation.

If the GS model specification is a correct representation of the data, then both the two-step procedure and joint maximum likelihood estimation provide consistent estimates of the model parameters.⁴⁰ This, of course, is not true if the specification does not capture the empirical regularities in the data. As with the two-step procedure, we begin by using only one X variable -- income -- and then expand the specification to include the full set of variables used by GS. To emphasize the key features of the model, in most specifications, we fix the β parameters at those obtained in a single equation Tobit estimate of the IRA equation.⁴¹ The results are presented in Table 20.

As with the two-step estimates, the results change dramatically when non-IRA financial assets are included among the X variables. No matter what the model specification, the estimated η jumps wildly when NIFA is added. For example, when NIFA

⁴⁰In Appendix A of their paper, GS describe the components of the likelihood function used in their analysis. The third component for limit contributors is incorrect. The numerator of the first term should include $\hat{S}_1 - L$, but the L is not included. We assume that this is only a typographical error in the paper, but that the likelihood function is in fact programmed correctly.

⁴¹This has very little effect on the results. In specifications I and J in Table 20 β is estimated jointly with all other model parameters.

is added to the specification including all X variables but with η estimated as a constant, the estimated value of η jumps from -1.468 to 4.355 (specification D versus C).

Specification	Variables		Estimate η & Standard Error	
	X Variables (Except in η)	Variables in η	Estimate	Standard Error
A	Income	Constant	-0.790	0.180
B	Income + NIFA	Constant	4.644	0.887
C	All Xs Excl NIFA	Constant	-1.468	1.115
D	All Xs Incl NIFA	Constant	4.355	0.904
E	All Xs Excl NIFA	All Xs, No Constant	-0.209**	
F	All Xs Incl NIFA	All Xs, No Constant	1.170**	
G	All Xs Excl NIFA	All Xs, Plus Constant	-0.006**	
H	All Xs Incl NIFA	All Xs, Plus Constant	1.254**	
I	All Xs Excl NIFA	All Xs, No Constant	-0.222**	
J	All Xs Incl NIFA	All Xs, No Constant	1.116**	

* β is fixed at single equation Tobit estimates, except in specifications I and J.
**Evaluated at the mean of the X values for limit contributors.

Although we were unable to match the GS sample precisely, specification J is the same specification used by GS, and specification I is the GS specification excluding NIFA. Key parameters of these specifications along with the GS estimates are summarized in Table 21, for four key variables: income, debt, non-liquid assets, and NIFA. (None of the estimated coefficients on the six other variables included in the GS analysis is statistically different from zero in any of these relationships.) The model J parameter estimates are very close to the estimates presented by GS, although our estimate of η differs somewhat: from the value obtained by GS -- 1.116 versus 1.85 reported by GS. The difference arises for two reasons. First, our estimation sample differs from the GS sample (107 versus 117 limit contributors). Second, η is parameterized without a constant and estimated values are extremely sensitive to even small changes in the X values that may arise from slight differences in the sample. (This feature of the estimation procedure is documented below.)

The critical features of the maximum likelihood estimates are the same as those of the two-step procedure.

Variable	Non-Contributor Other Saving Equation			Contributor Other Saving Equation			η Equation		
	Without NIFA	With NIFA	GS Estimate	Without NIFA	With NIFA	GS Estimate	Without NIFA	With NIFA	GS Estimate
Income	5.654 (0.952)	5.638 (1.243)	6.012 (1.026)	-0.963 (1.026)	7.765 (1.505)	7.784 (1.391)	-0.322 (0.693)	-2.773 (0.821)	-3.170 (0.836)
Debt	1.175 (0.126)	1.187 (0.142)	1.056 (0.248)	1.412 (0.332)	1.154 (0.356)	1.151 (0.373)	-0.735 (0.248)	-0.486 (0.255)	-0.466 (0.192)
Non-liquid assets	-3.874 (0.330)	-3.866 (0.357)	-3.998 (0.492)	0.419 (0.712)	-0.120 (1.375)	-0.599 (0.577)	1.094 (0.885)	0.802 (0.869)	0.471 (0.541)
NIFA		-0.001 (0.045)	0.077 (0.059)		-3.512 (0.281)	-3.686 (0.402)		0.782 (0.182)	0.985 (0.235)

Estimates for the key variables in Table 21 indicate that the influence of NIFA is enormous . In the contributor saving equation -- the underlying level of saving for IRA contributors that have not reached the limit -- NIFA has a large and negative coefficient (our estimate is -3.512 and the GS estimate is -3.686). This suggests that a one standard deviation increase in the 1983 level of non-IRA financial assets (about \$100,000) is associated with a decrease in saving of over \$35,000 in the 1983-1986 period! Apparently the measurement error in 1983 NIFA, from which the dependent variable is constructed, swamps any role NIFA might play as a control for heterogeneity.

On the other hand, when NIFA is added to the η equation the estimated coefficient is large and positive, offsetting the large negative effect in the underlying saving equation. The size of the coefficient (0.782 in our model J and 0.985 in GS) is implausibly large,

implying that a one standard deviation increase in the 1983 level of non-IRA financial assets (about \$100,000) will increase η by almost 10!

As above, we can understand better the effect of NIFA on η by considering predicted values of underlying saving with and without this variable in the specification. The appropriate predictions are shown in Table 22. Without NIFA, predicted underlying saving of limit contributors is higher than that of non-limit contributors (4533 versus 1651). But when NIFA is added, the underlying saving of limit contributors is lowered from 4533 to -760, leaving a large gap between underlying and actual other saving, which is 3089. The underlying saving of non-limit contributors is increased somewhat, from 1651 to 1713. Once again, to fit the actual saving data, the gap between underlying and actual saving of limit contributors is bridged by the large estimated value of η . In this case, predicted other saving of limit contributors is well above actual other saving -- 4461 versus 3089.⁴² That is, when underlying saving is depressed, the slope η of S_L must be raised. With reference to Figure 8: when S is reduced to S' , S_L must be increased to $S_{L'}$. Or, the measurement error in NIFA biases the underlying saving downward and this is offset by an upward bias in η . Thus, at least in the presence of NIFA, no behavioral interpretation can be ascribed to the estimated value of η .

Although the addition of NIFA has an enormous effect on the other saving equation for contributors, the addition of this variable has virtually no effect on estimates in the other saving equation for families who don't contribute to an IRA, as shown in the first two columns of Table 21.

⁴²GS do not report predicted other saving of limit and non-limit contributors separately, but the predicted value for all contributors, which they do report, is far less than actual other saving of contributors (\$806 versus \$2,184).

Table 22. Predicted underlying other saving, total other saving, and actual other saving, for limit and non-limit contributors, based on maximum likelihood estimates.				
Specification	Predicted Underlying Other Saving -- X γ Only	Predicted Underlying Other Saving -- X γ Plus λ_2 and λ_3	Predicted Other Saving -- X γ Plus λ_2 and λ_3 and λ_1	Actual Other Saving
Non-Limit Contributors				
Specification I -- Excl NIFA	704	1651	1651	1989
Specification J -- Incl NIFA	1082	1713	1713	1989
Limit Contributors				
Specification I -- Excl NIFA	1082	4533	3298	3089
Specification J -- Incl NIFA	-2160	-760	4461	3089

In addition to the artificial increase in η caused by the introduction of NIFA, the GS parameterization of η -- excluding a constant term -- means that predicted changes in η with changes in X are likely to have little meaning. Without a constant term in the specification, the expected value of η , given X, is not captured by $X\delta$ unless the constant is in fact zero.⁴³ Predictions of η vary wildly in response to small changes in X variables. Indeed, by judicious selection of X values, a wide range of results can be obtained. GS have highlighted the estimated values of η based on selected X values. For example, GS show values for a "typical 35 year old," defined by particular X values. The "typical 35 year old" does not have an IRA in 1983 and does not have an employer-provided pension. Their "predicted" value of η is 0.68. But if in addition to the X values GS use, the person is defined to have an IRA and a pension, the value of η is -1.565. Based on other X values, almost any η could have been emphasized. Table 23 makes the possibilities clear. For

⁴³Specifications G and H in Table 20 are estimated including a constant term in η , as well as the X variables. Although the constant in these specifications is not significantly different from zero it is measured extremely imprecisely and identification is tenuous.

each specified change in an X variable the table gives the change in η implied by the GS estimates, reported in their table 5. For each of the continuous variables, the indicated change in the X variable is approximately one standard deviation. Thus the GS estimates themselves yield implausible responses to changes in household characteristics.

Table 23. Change in η for selected changes in X, based on GS estimates.*			
#	Variable	Change*	Change in η
1	Age	Increase by 12 years	0.331
2	Income (3 years)	Increase by \$120,000	-3.170
3	Pension	No to yes	-1.376
4	Education	Increase by 3 years	0.666
5	Family size	Increase by 1.5 persons	1.424
6	1983 NIFA	Increase by \$100,000	9.850
7	Debt	Increase by \$40,000	-1.864
8	Non-liquid assets	Increase by \$200,000	0.942
9	IRA in 1983	No to yes	-0.869

*For each of the continuous variables the indicated change in the X variable is approximately one standard deviation.

As indicated above, estimation results may vary enormously based on the sample used in estimation. Since the results depend critically on the small number of limit contributors -- 117 in the GS estimates -- any selection that changes this number can change the results enormously, especially given the vast variation in assets and saving. The reported estimates are based on a sample deleting households with inferred 1983-86 saving less than or greater than \$100,000. Although they don't report the η values with samples based on different saving thresholds, they do report their estimates of the proportion of an increase in IRA saving, resulting from an increase in the IRA limit, that would be net new saving. Here are their examples which document how sensitive their results are to the choice of a saving threshold:

Sample Saving Deletions	% Net Saving
More or less than +/- \$75,000	-17.5%
More or less than +/- \$100,000	2.1%
More or less than +/- \$200,000	-382.2%

Based on our estimates of the GS model (including NIFA), small differences in this critical sample selection criterion yield very different values of η . Estimates for selection thresholds ranging from plus or minus \$50,000 to plus or minus \$200,000 are shown in Figure 9. Not only is the variation great, but most of the estimates are very small, even including NIFA.

7.5. Summary.

In our judgement, a descriptive summary of the data used by GS suggests very strongly that the contributions of participants to the IRA program between 1983 and 1986 represented largely new saving. It is also clear from these data that the typical IRA contributor had not been saving close to the typical annual IRA contribution of \$2,300 per year. Simple predictions of 1986 non-IRA financial assets based on the 1983 cross-section relationship between age, income, and non-IRA financial assets (and accounting for differences in rate of return), bolster the message of the raw data, suggesting that most IRA contributions must have represented new saving. These "reality checks" would seem to be inconsistent with the GS results.

But GS frame their conclusions in terms of the effect of an increase in the IRA limit, arguing that virtually none of the increase in IRA contributions resulting from a limit increase would represent new saving. The summary data also suggest that this conclusion is inconsistent with the underlying data. A large fraction -- 60% to 85% -- of annual IRA contributions are at a family or an individual IRA limit. Thus most of the increase in net saving that seems evident from the summary data must be attributed to

persons who contributed at the limit. And it seems implausible to conclude that if the limit had been higher these people would not have increased their saving even more.

GS acclimatize the reader to their conclusion by highlighting the 1986 non-IRA financial assets of persons at a constructed IRA limit. GS must infer IRA contributions from 1983 and 1986 IRA balances, and from these inferred contributions, they infer limit status in the three consecutive years between 1983 and 1986. They classify 117 families in this group and conclude that about 22% of contributors are at this “constructed three-year limit.” They point to the rather large non-IRA financial assets of this group, suggesting that because these families did have assets that could be transferred to an IRA, that is what they would do if the IRA limit were reached. Even with this unverified possibility in mind, the number GS highlight is misleading for several reasons: (1) While GS point to the assets of 22% of participants at a constructed limit, in the proportion of contributions at either an individual or a family annual limit is three or four times this large. It is the assets of this much larger group that are relevant. (2) GS further exaggerate relevant non-IRA financial assets by including the assets of all families over 25, including those over 65, who are unlikely to contribute to an IRA. GS also use a very inclusive definition of financial assets, including assets that we believe are unlikely to be substituted for IRA assets. (3) Finally, GS emphasize 1986 asset levels. The 1986 data are relevant if one is drawing attention to what might be expected from a “future” increase in the IRA limit. But the 1986 number is at the same time misleading because there was a substantial increase in non-IRA financial assets of contributors during the 1983 to 1986 period during which IRA contributions were made. In considering the summary data, we emphasize 1983 assets, drawing attention to the fact that contributors had not been saving at the typical IRA annual contribution rate before the advent of the program. Indeed the median level of the non-IRA

financial assets of 1983 IRA contributors was only about 20% of the assets of inferred “three-year limit contributors” emphasized by GS. And the median 1983 non-IRA financial assets of the persons who contributed at an IRA limit during the 1983 to 1986 period was probably less than 40% of the level emphasized by GS.

The formal GS results are not based on summary data, however, but rather on a complex estimation method. Having reproduced their estimation procedure and analyzed it closely, we conclude that the data provide little support for their conclusions. The value of η , the key substitution parameter reported by GS, is estimated with substantial bias because of measurement error in 1983 non-IRA financial assets (NIFA). Furthermore, because the parameterization of η used by GS does not include a constant term, the reported variations in η by family attributes have no behavioral meaning whatsoever. These features of the GS estimation procedure, together with values of η that vary wildly with small changes in the sample used in estimation, mean that judicious choice of sample family attributes at which to evaluate η can produce virtually any result.

But can any specification of the GS model be given credence? In particular, can estimates that exclude NIFA and that estimate a constant η be viewed with confidence? We know that even in this case, the estimates vary widely depending on the sample selected for estimation. This is especially critical given that the key parameter η is essentially determined by limit contributors. Excluding NIFA, all specifications yield a negative value of η , which in many but not all instances is not significantly different from zero. We also find that although inclusion of NIFA changes the estimated value of η , it does not change the residual correlation between IRA and other non-IRA financial asset saving. When all X variables other than NIFA are used, the estimated correlation is essentially zero. And when NIFA is added the correlation remains essentially zero. Thus

NIFA does not seem to serve its intended goal of providing further control for heterogeneity, but does impart substantial bias. It seems evident that the results provide little support for a positive η , and thus little support for substitution of non-IRA financial asset saving for IRA saving. But the results obtained without NIFA may also be so fragile as to provide unreliable evidence of no effect.

8. CONCLUSIONS

Over the past several years we have undertaken a series of analyses of the effect of IRA and 401(k) contributions on net personal saving. We have summarized this research here, together with additional results. Saver heterogeneity is the key impediment to determining the saving effect of these plans and in our studies we have used different methods to address this issue. We have organized the discussion according to the method use to correct for heterogeneity. We emphasize that no single method can provide sure control for all forms of heterogeneity. Taken together, however, we believe that the analyses address the key complications presented by heterogeneity. In our view, the weight of the evidence, based on the many non-parametric approaches discussed here, provides strong support for the view that contributions to both IRA and 401(k) plans represent largely new saving. Some of the evidence is directed to the IRA program, other evidence to the 401(k) plan, and some of the evidence to both plans jointly. We believe that the evidence is strong in all cases.

Several other investigators have used different methods to consider the effect of these retirement saving programs on personal saving and in some cases have reached very different conclusions from ours. Thus we have devoted particular effort to explaining why different approaches, sometimes based on the same data, have led to different conclusions. In some instances, we believe that the limitations of the methods used by

others have undermined the reliability of the results. Particular attention is devoted to a recent paper by Gale and Scholz [1994] that is widely cited as demonstrating that IRAs have no saving effect. Based on our analysis of the data used by Gale and Scholz, including calculations based on a replication of their model, we find that their conclusions are at odds with the patterns of asset holding and saving in the raw data.

REFERENCES

- Attanasio, Orazio and Thomas De Leire. 1994. "IRAs and Household Saving Revisited: Some New Evidence." NBER Working Paper No. 4900, October.
- Bernheim, B. Douglas. 1994. "Comments on: Do Saving Incentives Work?" *Brookings Papers on Economic Activity* 1: 152-166.
- Bickel, Peter J., Eugene A. Hammel, and J. William O'Connell. 1975. "Sex Bias in Graduate Admissions: Data from Berkeley." *Science* 87: 398-404, February 7.
- Burman, Leonard, Joseph Cordes, and Larry Ozanne. 1990. "IRAs and National Savings." *National Tax Journal* 43: 259-283.
- Engelhardt, Gary V. 1995a. "Tax Subsidies and Household Saving: Evidence from Canada." Mimeo.
- _____. 1995b. "House Prices and Homeowner Saving Behavior." *Regional Science and Urban Economics*, forthcoming.
- Engen, Eric M. and William G. Gale. 1995. "Debt, Taxes and the Effects of 401(k) Plans on Household Wealth Accumulation." Unpublished.
- Engen, Eric, William G. Gale, and John Karl Scholz. 1994. "Do Saving Incentives Work?" *Brookings Papers on Economic Activity*. 1:85-151.
- Feenberg, Daniel and Jonathan Skinner. 1989. "Sources of IRA Saving." *Tax Policy and the Economy* 3: 25-46. Cambridge: MIT Press.
- Gale, William G. 1995. "The Effects of Pensions on Wealth: A Re-Evaluation of Theory and Evidence." Mimeo, Brookings Institution.
- Gale, William G. and John Karl Scholz. 1994. "IRAs and Household Saving." *American Economic Review* 84, 5: 1233-1260, December.

- Hoynes, Hilary and Daniel McFadden. 1994. "The Impact of Demographics on Housing and Non-Housing Wealth in the United States." NBER Working Paper No. 4666, March.
- Hubbard, R. Glenn. 1984. "Do IRAs and Keoghs Increase Saving?" *National Tax Journal* 37: 43-54.
- Hubbard, R. Glenn and Jonathan S. Skinner. 1995. "The Effectiveness of Saving Incentives: A Review of the Evidence." Mimeo.
- Joines, Douglas H. and James G. Manegold. 1995. "IRAs and Saving: Evidence from a Panel of Taxpayers." Mimeo, University of California, revised February 1995.
- Kotlikoff, Laurence J. and David A. Wise. 1985. "Labor Compensation and the Structure of Private Pension Plans: Evidence for Contractual versus Spot Labor Markets." In D. Wise (ed.), *Pensions, Labor, and Individual Choice*, D. Wise (ed.), University of Chicago Press.
- _____. 1988. "Pension Backloading, Wage Taxes, and Work Disincentives." In L. Summers (ed.) *Tax Policy and the Economy*, Vol. 2, MIT Press.
- _____. 1989. "Employee Retirement and a Firm's Pension Plan." In D. Wise (ed.), *The Economics of Aging*, University of Chicago Press.
- Papke, Leslie E. 1995. "Does 401(k) Introduction Affect Defined Benefit Plans?," *Proceedings of the National Tax Association-Tax Institute of America*. Columbus, OH: NTA-TIA): 173-177.
- Papke, Leslie E., Mitchell Petersen, and James M. Poterba. 1996. "Did 401(k) Plans Replace Other Employer Provided Pensions?" In D. Wise (ed.), *Papers in the Economics of Aging*. University of Chicago Press.

- Poterba, James M., Steven F. Venti, and David A. Wise. 1994a. "401(k) Plans and Tax-Deferred Saving." In D. Wise (ed.), *Studies in the Economics of Aging*. University of Chicago Press.
- _____. 1994b. "Targeted Retirement Saving and the Net Worth of Elderly Americans." *American Economic Review* 84,2: 180-185.
- _____. 1995. "Do 401(k) Contributions Crowd Out Other Personal Saving?" *Journal of Public Economics* 58: 1-32.
- Shefrin, Hersh M. and Richard H. Thaler. 1988. "The Behavioral Life-Cycle Hypothesis." *Economic Inquiry* 26: 609-643.
- Skinner, Jonathan. 1993. "Is Housing Wealth a Sideshow?" NBER Working Paper No. 4522, November. (Forthcoming in D. Wise (ed.), *Advances in the Economics of Aging*, University of Chicago Press.)
- _____. 1994. "Housing and Saving in the U.S." In Y. Noguchi and J. Poterba (eds.), *Housing Markets in the United States and Japan*. University of Chicago Press.
- Skinner, Jonathan and Daniel Feenberg. 1990. "The Impact of the 1986 Tax Reform Act on Personal Saving." In J. Slemrod (ed.), *Do Taxes Matter? The Economic Impact of the Tax Reform Act of 1986*. Cambridge: MIT Press.
- Thaler, Richard H. 1990. "Saving, Fungibility, and Mental Accounts." *Journal of Economic Perspectives* 4: 193-205.
- Thaler, Richard H. and H.M. Shefrin. 1981. "An Economic Theory of Self-Control." *Journal of Political Economy* 89, 2: 392-406.

- 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.
- _____. 1987. "IRAs and Saving." In M. Feldstein (ed.), *The Effects of Taxation on Capital Accumulation*. University of Chicago Press.
- _____. 1990. "Have IRAs Increased U.S. Saving?: Evidence from Consumer Expenditure Surveys." *Quarterly Journal of Economics*. August: 661-698.
- _____. 1991. "The Saving Effect of Tax-Deferred Retirement Accounts: Evidence from SIPP." In B.D. Bernheim and J.B. Shoven (eds.), *National Saving and Economic Performance*. University of Chicago Press.
- _____. 1992. "Government Policy and Personal Retirement Saving." In J. Poterba (ed.), *Tax Policy and the Economy* 6. Cambridge: MIT Press.
- _____. 1995a. "Individual Response to a Retirement Saving Program: Results from U.S. Panel Data." *Ricerche Economiche* 49: 235-254.
- _____. 1995b. "RRSPs and Saving in Canada." Mimeo.
- _____. 1996. "The Wealth of Cohorts: Retirement Saving and the Changing Assets of Older Americans," forthcoming in S. Schieber and J. Shoven (eds.), *Public Policy Toward Pensions* (tentative title), MIT Press.
- Wise, David A. 1987. "Individual Retirement Accounts and Saving." In M. Feldstein (ed.), *Taxes and Capital Formation*. University of Chicago Press.

Fig 1. Retirement Plan Contributions




Figure 2. Properties of Attanasio De Leire Method

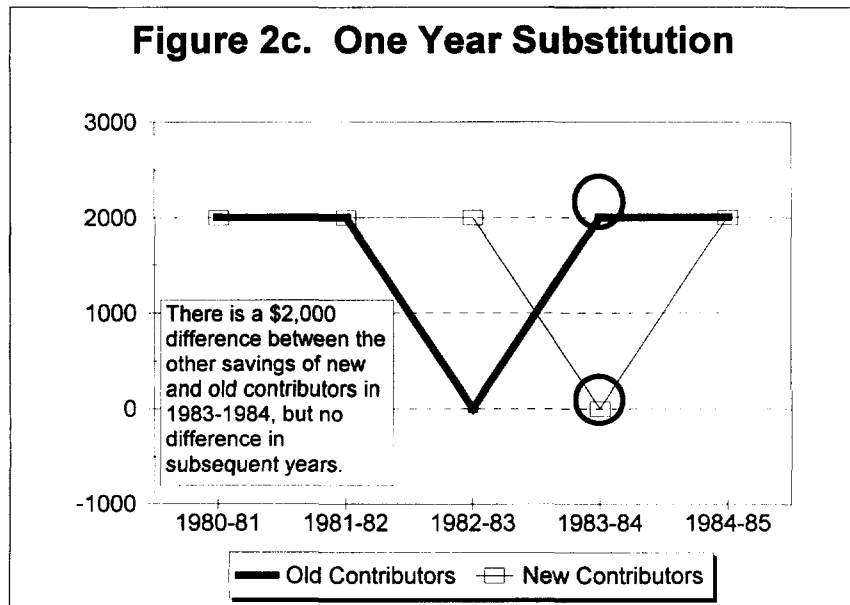




Fig 3. Personal Financial Assets of Selected Cohorts.

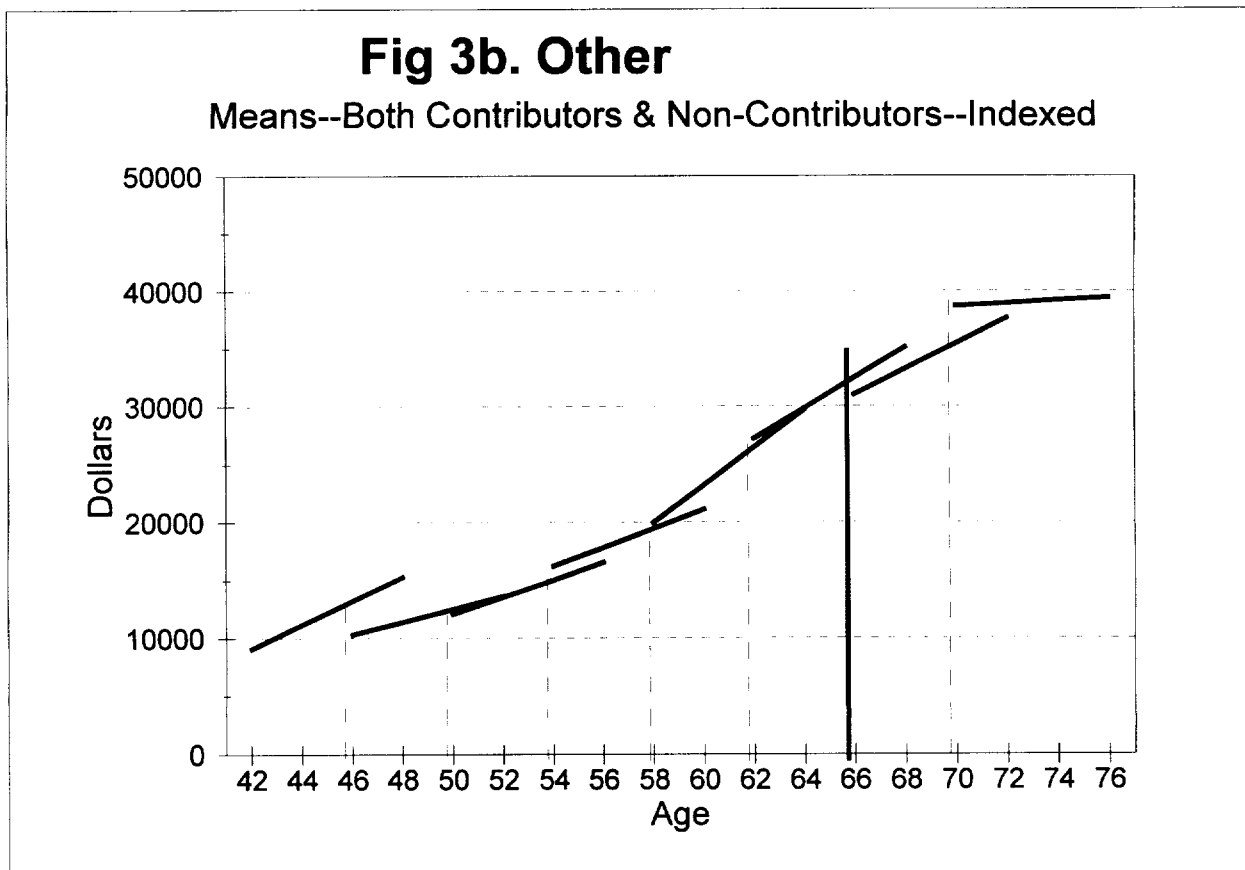



Figure 4. Summary of cohort trends in retirement saving vs home value, mortgage, and equity.

Fig 4a. Retirement Assets Summary




Fig 4b. Home Value Summary




Fig 4c. Home Mortgage Summary




Fig 4d. Home Equity Summary




Fig 5. Estimated Cohort Effects




Fig 6. Cohort Effects

Financial Assets and Home Equity




Fig 7. Home Equity Quartiles By Eligibility and Year

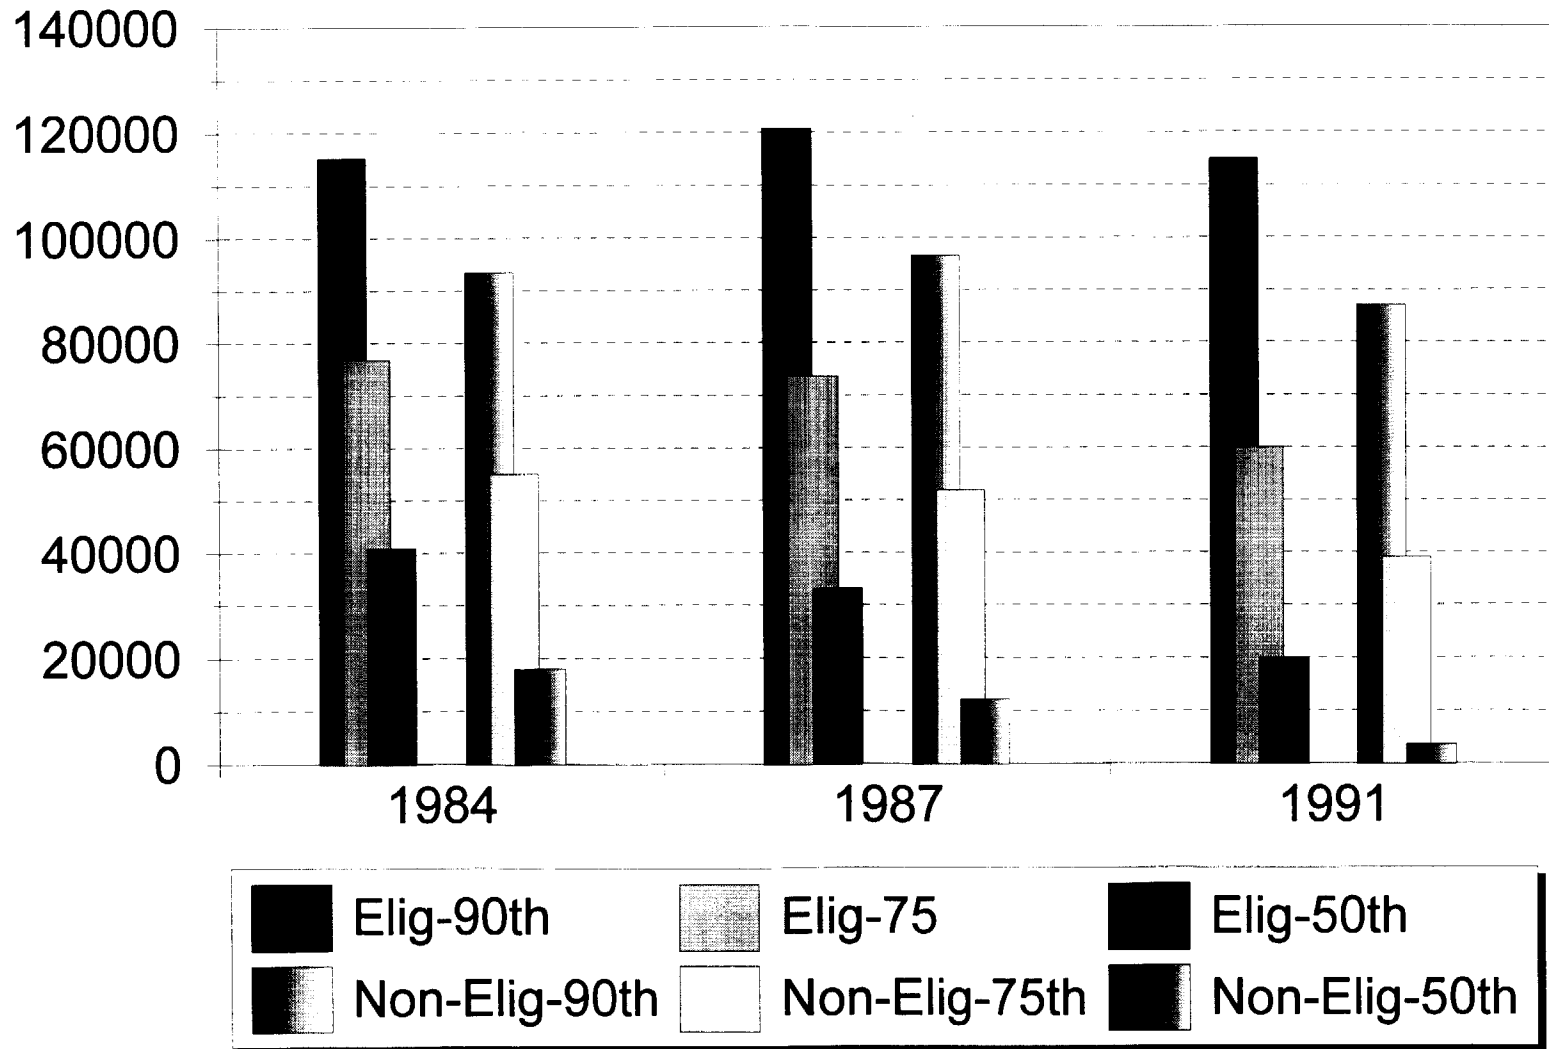

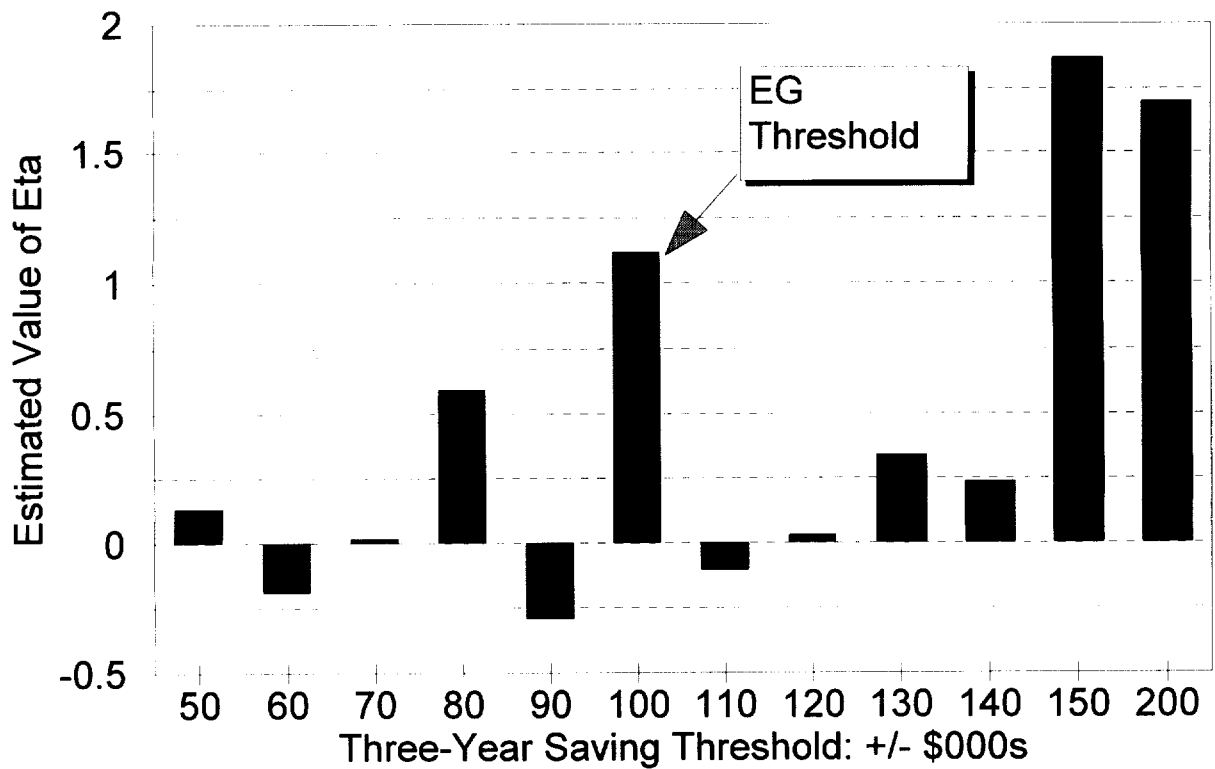


Figure 8. The GS Model



$I = IRA$. $S-SL$ without NIFA. $S'-SL'$ with NIFA.

Fig 9. Estimated Eta Values
By Three-Year Saving Threshold



Appendix A: Cohort and Period Effects

Although the data make clear that the timing of housing market trends and trends in mortgage debt do not coincide, the apparent market, or “period” effects in the housing market complicate the estimation of precise housing data cohort effects. For example, the mortgage data by cohort between 1984 and 1987 show very small cohort effects. But if only the 1987 and 1991 data were considered, it would appear that there were substantial cohort effects. But it is likely that these within cohort changes reflect “period” effects that can show up as cohort effects. This is illustrated in Figure A-1, which shows mortgage debt for the C42 and C48 cohorts. Looking only at the 1984 and 1987 data only, no cohort effect is apparent, as shown by the narrow line. But if the data are fitted by cohort, including the 1991 values, there appears to be a cohort effect. If a 1991 year effect were accounted for, the apparent cohort effect would essentially disappear. Suppose the increase in mortgage debt between 1987 and 1991 resulted from the Tax Reform Act of 1986. Then the data might be interpreted this way: The 1986 legislation induced an increase in the debt of the younger cohorts between ages 44 and 48. For the older cohort, the increase was induced between 50 and 54. This “period” effect raises the debt of the younger cohort at age 48 and the debt of the older cohort at age 54. This creates the illusion of a cohort effect, illustrated by the vertical distance between the two heavy lines. In this case, the apparent cohort effect is really a period effect and should be distinguished from a true cohort effect. But with so few observation per cohort, we have not tried to do that.¹

¹As is well known, it is not possible to distinguish age effects from time effects within the same cohort. But if age effects are assumed not to depend on cohort -- as is assumed when the effect of age is parameterized as in equation (1) in footnote 17 on page 42 -- then, in principle, time effects can be estimated. In effect, shifts that correspond to changes between years are interpreted as year effects rather than within cohort age effects. It is seems evident from the summary graphs that the data is somewhat more complicated than this, because the year effects seem in some cases to have a differential effect on young versus older cohorts.

Appendix B: Rate of Return Effects

The potential magnitude of this effect can be approximated as follows: Consider the predicted (by equation 12) assets of persons age $a+3$ in 1983. The average asset level predicted by the 1983 cross-section regression could -- if the appropriate data were known -- be decomposed this way:

$$(B-1) \quad A_{a+3} = A_a(1+r_{_})^3 + S(A_a, Y; r_{_})$$

Here, A_a represents the assets that persons age $a+3$ had three years earlier. The $r_{_}$ pertains to the rate of return that applied during the three years preceding 1983. New saving, $S(A_a, Y; r_{_})$, is some function of income over the three year period. The $r_{_}$ in this function recognizes that people might save less if r is higher, because the gain from existing assets is greater. Similarly, the assets of persons age a could be decomposed as

$$(B-2) \quad A_a = A_{a-3}(1+r_{_})^3 + S(A_{a-3}, Y; r_{_})$$

The difference between assets at $a+3$ and assets at a can then be described as

$$(B-3) \quad A_{a+3} - A_a = (A_a - A_{a-3})(1+r_{_})^3 + S(A_a, Y; r_{_}) - S(A_{a-3}, Y; r_{_})$$

Accounting for the change in income between 1983 and 1986, the difference predicted by equation (13) could be decomposed this way. But the rate of return that determines the difference in assets with a three-year age difference in 1983 may be different from the rate that obtained in the next three years. To predict over the next three years, we would want to use the rate that applied during those years. In this case, we would have

$$(B-4) \quad A_{a+3} - A_a = (A_a - A_{a-3})(1+r^+)^3 + S(A_a, Y; r^+) - S(A_{a-3}, Y; r^+)$$

where r^+ is the rate of return that applied between 1983 and 1986. The difference between two predictions is

$$(B-5) \quad (A_a - A_{a-3})[(1+r^+)^3 - (1+r_{_})^3] + [S(A_a, Y; r^+) - S(A_a, Y; r_{_})] - [S(A_{a-3}, Y; r^+) - S(A_{a-3}, Y; r_{_})]$$

Assume that the last two terms approximately cancel. Then the difference is given by the first term. Consider, for example, the triple a bond rate in 1980 through 1982 versus 1983 versus 1985. The average during the first period was 13.30 and during the second period 12.04. The second component of this term is thus negative. So correction for the rate of return would reduce the estimated increase. The return on other assets may give a different sign, but it seems evident that differences in the rate of return could not account for much of the difference between 1983 and 1986 non-IRA financial assets.

Appendix Table 1. Conditional Median Asset Balances by 401(k) Eligibility and Income

Asset Category and Eligibility Status	Income						
	<10	10-20	20-30	30-40	40-50	50-75	>75
A1a. Results for 1991 (1991 \$s)							
Total Financial Assets							
Eligible for a 401(k)	2033	4045*	5499*	8683*	14470*	26093*	51080*
Not eligible for a	1378	1997	2558	3256	6206	10080	29842
Non-IRA-401(k) Assets							
Eligible for a 401(k)	538	1138	1500	2835*	4724	8699*	18188*
Not eligible for a	663	1063	1411	2052	4250	5437	17000
401(k) Assets							
Eligible for a 401(k)	1171	1008	1211	2092	3073*	4833*	14300*
Not eligible for a	0	0	0	0	0	0	0
IRA Assets							
Eligible for a 401(k)	0	0	0	0	0	1437	6029*
Not eligible for a	0	0	0	0	0	978	2882
A1b. Results for 1987 (1987 \$s)							
Total Financial Assets							
Eligible for a 401(k)	2061	2404	4206*	9062*	12588*	24384*	57348*
Not eligible for a	1581	1902	2624	4605	6726	14108	30971
Non-IRA-401(k) Assets							
Eligible for a 401(k)	591	1029	1711	3398*	5663*	10776*	24044*
Not eligible for a	799	1004	1554	2904	4246	8462	20383
401(k) Assets							
Eligible for a 401(k)	456	474	607	895	1255*	1755*	8056*
Not eligible for a	0	0	0	0	0	0	0
IRA Assets							
Eligible for a 401(k)	0	0	0	0	0	3564	9064*
Not eligible for a	0	0	0	0	0	2770	4950
A1c. Results for 1984 (1984 \$s)							
Non-IRA-401(k) Assets							
Eligible for a 401(k)	561	1042	1988	3861*	5027	11683*	28824*
Not eligible for a	754	1138	1746	3076	5082	10846	21485
IRA Assets							
Eligible for a 401(k)	0	0	0	0	0	2250	3181
Not eligible for a	0	0	0	0	0	1484	2084

*Difference between eligibles and non-eligibles is statistically significant at the 95% confidence level.
Source: Poterba, Venti, and Wise [1995].

Appendix Table 2. Summary of Cohort Trends in Home Value, Home Mortgage, and Home Equity

Age	Cohort														
	C42	C44	C46	C48	C50	C52	C54	C56	C58	C60	C62	C64	C66	C68	C70
A2c. Mean Home Equity Cohort Data															
42	47003														
44	39802	47672													
46		41720	47064												
48	46050		44430	49436											
50		48182		46305	51872										
52			53853		49996	58784									
54				54169		52107	58827								
56					57808		56554	58914							
58						58580		57974	57113						
60							61551		58572	57735					
62								65675		60742	61984				
64									70784		62141	61516			
66										66502		63273	59059		
68											70815		67867	53334	
70												69640		65204	59087
72													65764		59235
74														73873	
76															78323

Note: All values are in 1991 dollars.

Appendix Table 3. Conditional Median Asset Balances by 401(k) Eligibility and Income
1984

Asset Category and Eligibility Status	Income						
	<10	10-20	20-30	30-40	40-50	50-75	75+
Non-IRA-401(k) Financial Assets							
Eligible	147	550	1454	2404	4732	7901	31485
Not eligible	220	545	1034	2043	3748	7059	21778
Difference	-73	6	420*	361*	985*	842*	9708*
Net Non-IRA-401(k) Financial Assets							
Eligible	-1288	-651	302	716	2815	6241	22068
Not eligible	-607	-348	130	775	2080	5208	17802
Difference	-681	-304	172	-60	735*	1034*	4267*
Housing Value							
Eligible	--	--	--	--	--	--	--
Not eligible	--	--	--	--	--	--	--
Difference	--	--	--	--	--	--	--
Housing Mortgage							
Eligible	--	--	--	--	--	--	--
Not eligible	--	--	--	--	--	--	--
Difference	--	--	--	--	--	--	--
Home Equity							
Eligible	11377	16210	17486	26138	31101	43185	65232
Not eligible	12384	13725	16007	25123	30833	34348	52746
Difference	-1007	2486	1478	1014	268	8837*	12486*
Net Non-IRA-401(k) Financial Assets Less Mortgage Debt							
Eligible	-5329	-5370	-4328	-12021	-16737	-16163	-15560
Not eligible	-4468	-3670	-4386	-7810	-14317	-18665	-10264
Difference	-861	-1700	58	-4212*	-2420*	2503*	-5297*
Net Non-IRA-401(k) Financial Assets Plus Home Equity							
Eligible	11594	16616	21371	28136	38799	53060	104748
Not eligible	11293	14398	18632	28461	36327	44462	83338
Difference	301	2218	2739	-325	2472	8598*	21410*
IRA							
Eligible	0	0	0	0	0	1083	5100
Not eligible	0	0	0	0	0	0	4200
Difference	0	0	0	0	0	1083	900

Note: Evaluated at the medians of age, marital status, and education except for last column which is evaluated at the medians of income, age, marital status and education. Difference between eligibles and noneligibles statistically significant at the 95% confidence level. Income intervals are in 1987 dollars.

Appendix Table 4. Conditional Median Asset Balances by 401(k) Eligibility and Income, 1991

Asset Category and Eligibility Status	Income						
	<10	10-20	20-30	30-40	40-50	50-75	75+
Total Financial Assets							
Eligible	2323	2473	5443	10263	19628	37166	70954
Not eligible	563	1030	1630	3863	5523	15109	42953
Difference	1760*	1443*	3813*	6400*	14105*	22057*	28002*
Net Total Financial Assets							
Eligible	1102	1073	2464	7554	17022	34726	67878
Not eligible	-483	-57	370	2307	3652	11597	39218
Difference	1585	1130*	2094*	5247*	13370*	23129*	28660*
Non-IRA-401(k) Financial Assets							
Eligible	488	616	1325	3128	5692	11487	21414
Not eligible	264	544	1086	2504	3973	8944	27120
Difference	224	72	239*	624*	1719*	2543*	-5709*
Net Non-IRA-401(k) Financial Assets							
Eligible	-491	-262	-95	1089	3094	8838	18925
Not eligible	-327	-142	116	907	1968	5667	26909
Difference	-164	-120	-211	182	1126*	3171*	-7984*
Home Equity							
Eligible	8105	10973	13937	20293	25400	44839	76176
Not eligible	9210	11044	11863	18751	27132	34834	58420
Difference	-1105	-71	2073*	1541	-1732	10005*	17756*
Net Total Financial Assets Plus Housing Equity							
Eligible	14509	14150	20538	32875	49361	84511	151834
Not eligible	9185	13121	15106	28502	38139	60945	122341
Difference	5324	1029	5432*	4373*	11222*	23566*	29499*
Net Non-IRA-401(k) Financial Assets Less Mortgage Debt							
Eligible	-7112	-7534	-10790	-23701	-41359	-39584	-63414
Not eligible	-5800	-5689	-7795	-22799	-29345	-38977	-16320
Difference	-1312	-1845*	-2995*	-902	-12014*	-607	-47094*
Net Non-IRA-401(k) Financial Assets Plus Home Equity							
Eligible	9030	10361	14017	24168	34682	61358	108290
Not eligible	8059	11557	13522	25468	35275	56360	105924
Difference	971	-1196	495	-1300	-593	4998*	2366
IRA							
Eligible	0	0	0	0	0	0	0
Not eligible	0	0	0	0	0	0	0
Difference	0	0	0	0	0	0	0
401(k)							
Eligible	405	431	1164	2072	4053	6942	15832
Not eligible	0	0	0	0	0	0	0
Difference	405	431	1164*	2072*	4053*	6942*	15832*

Note: evaluated at the medians of age, marital status, and education except for last column which is evaluated at the medians of income, age, marital status and education.

Income intervals are in 1987 dollars.

*Difference between eligibles and noneligibles statistically significant at the 95% confidence level.

Appendix Table 5. Mean Demographic Characteristics by Income Interval, Eligibility, and Year

		Income Interval							
		<10	10-20	20-30	30-40	40-50	50-75	75+	ALL
1984									
401(k) Eligible	Income	6641	16276	25418	34733	44772	59170	97115	41393
	Age	40.8	41.5	41.5	40.9	40.5	43.0	46.2	41.8
	Education	12.9	12.1	13.0	13.5	14.1	14.4	15.2	13.6
	Married	0.28	0.42	0.64	0.75	0.86	0.93	0.98	0.76
	Number	31	140	243	306	212	262	96	1290
401(k) Not Eligible	Income	6562	15163	25009	34528	44539	58600	96893	29276
	Age	40.8	40.0	39.4	41.3	41.9	43.2	45.4	40.8
	Education	11.9	12.0	12.6	12.9	13.4	14.1	14.8	12.7
	Married	0.23	0.45	0.66	0.84	0.88	0.90	0.91	0.65
	Number	712	1820	1875	1388	798	645	188	7426
1991									
401(k) Eligible	Income	7415	15594	24876	34573	44695	59565	91499	38895
	Age	41.0	40.2	40.0	41.4	41.9	43.0	43.6	41.4
	Education	12.5	12.2	13.2	13.7	14.2	14.7	15.6	13.7
	Married	0.21	0.37	0.55	0.75	0.82	0.90	0.93	0.69
	Number	84	543	807	784	564	633	266	3681
401(k) Not Eligible	Income	6631	14949	24525	34521	44623	59506	92111	26533
	Age	40.0	40.2	40.4	40.6	41.3	43.3	45.3	40.7
	Education	11.6	12.3	13.1	13.5	14.1	14.8	15.3	12.9
	Married	0.28	0.41	0.63	0.78	0.84	0.90	0.89	0.58
	Number	853	1928	1487	880	507	404	171	6230

Note: The income intervals are indexed to 1987 dollars.

Fig A-1. Period Vs Cohort Effect?

