Do IRAs Substitute for 401(k)s?

David Love

Yale University

Department of Economics

New Haven, CT 06520

May 19, 2003

Abstract

1 Introduction

An increasing amount of saving in the U.S. is done inside retirement saving accounts such as the Individual Retirement Account (IRA) and the 401(k). As of 2002, approximately 42 million U.S. households owned an IRA, and even more—45 million—participated in a 401(k) plan.[1] The impressive growth of these programs has fueled a debate among economists about the effectiveness of these plans in increasing national saving. Eric Engen, William Gale and John Karl Scholz (1996), for example, argue that the reduction in public saving due to lost tax revenues is not offset by an increase in private saving because individuals fund higher saving in the plan with a reduction in other saving. James Poterba, Steven Venti, and David Wise (1998, 1997) are on the opposite pole of this debate and find that these programs have a large positive effect on total saving. In this chapter, I examine a special case of substitution and assess the degree of substitution between IRAs and 401(k)s. Theoretically, substitution between these accounts should be high, since they share similar characteristics with respect to tax sheltering and illiquidity. If substitution between the accounts is low empirically, then it is less likely that these accounts substitute for other saving.

The conflicting views about the efficacy of saving plans are sustained in part by differences in methodology and incomplete data. While good data are available for IRA and 401(k) contributions, comparable figures for non-IRA/401(k) saving are not available and must generally be constructed. Different assumptions used in constructing these values can lead to wildly variable estimates of the substitution between retirement plans and other assets.[2] Because this parameter flexibility often allows contradictory results to be generated from the same model, it is difficult to assess the accuracy of these experiments. I take advantage of the high quality data on IRAs and 401(k)s to test a more modest measure of substitution: that between the plans themselves.

[1] Investment Company Institute (2002).

[2] Two illustrative examples can be taken from papers written by the above-mentioned authors. In their 1995 paper on IRAs, Gale and Scholz difference asset balances in the IRA between 1986 and 1983 and define contributions as the average yearly amount needed to generate that difference. This approach is quite sensitive to the assumed rate of return and biases the marginal effects they seek to measure. Working in the opposite direction, Venti and Wise (1995) construct a measure of stock balances from asset income, which amplifies the measurement error on that variable.

I use the weaker hypothesis that IRAs and 401(k)s substitute for each other to test the stronger hypothesis that they substitute for other assets. It is important to note, however, that I address only the question of marginal substitutability (i.e., saving changes with respect to a limit increase) and not whether 401(k)s have displaced IRAs over time.

The chapter is organized as follows. Section 2 provides background details and describes the asset substitution effects for retirement saving programs with contribution limits. With a simplified budget constraint, I show that changes in the program rate of return produce only an income effect for persons who contribute to the limit. This section concludes with a review of the previous literature and a brief description of the data used in the analysis.

In section 3, I present budget constraints that incorporate the plan features of limits, withdrawal penalties, and tax breaks. Withdrawal penalties make the IRA and 401(k) imperfect substitutes for other assets, which provides the rationale for treating retirement plan saving separately from non-retirement saving. Noting that the return on a 401(k) weakly dominates that on an IRA, I argue that the following allocation rule is a reasonable strategy for optimizing individuals. An individual participating in both programs will only contribute to the IRA after the 401(k) limit has been reached. I end the section with a discussion of employer matching limits and their probable effect on the analysis.

In section 4, I use a Tobit specification to estimate a model of saving allocation between IRAs and 401(k)s and find little evidence for marginal substitution. While the coefficient measuring substitution is significant, I show that the estimation is biased and inconsistent. This occurs because an estimated error term from the first stage of the analysis enters, but is unaccounted for, in the second stage. If the error terms on the two equations are uncorrelated, however, it is possible to reduce the bias by including an estimate of the first error in the second equation. I do this and find that while the coefficient of substitution is larger, there is still very little substitution between plans. The last section summarizes the paper and offers some direction for future research.

2 Background

2.1 IRA and 401(k)

Initiated in 1974, the Individual Retirement Account (IRA) originally limited participation to individuals who were not covered under a pension plan. Participation in the program was extremely limited until the Economic Recovery Act of 1981 expanded the program to include all employees and raised the annual individual contribution limit from $1,500 to $2,000.[3] Before the Tax Reform Act of 1986, the IRA provided two saving incentives. Contributions were deductible, and interest income accumulated free of tax. Also, early withdrawals were penalized by 10 percent before the age of 59. After the Tax Reform Act of 1986, the tax deductibility of contributions for higher income groups (above $40,000) was removed. Because the majority of IRA participants were members of high-income households, the elimination of up-front deductibility significantly discouraged contributions. Annual contributions dropped from $37 billion in 1986 to $6.5 billion in 1992. Part of this decline, however, might be explained by the introduction of another savings program, the 401(k).

The Revenue Act of 1978 introduced the 401(k) program, but it did not gain popularity until 1981 when the IRS released explanatory regulations. While the 401(k) shares the tax benefits of the original IRA, there are several key differences between 401(k) plans and IRA. First, 401(k) plans are firm-specific. One must be an employee at a participating firm in order to contribute. This contrasts with the IRA, where all wage earners may contribute up to the limit, regardless of employer. Second, the elective contribution limit on 401(k)s is four times larger than that on IRAs. In addition to a more generous contribution cap, 60 percent of 401(k) plans also include employer matched contributions.[4] Assuming a positive employer match rate, the 401(k) returns clearly dominate those on an IRA.[5]

Together, these retirement programs represent a large fraction of total retirement savings

[3]The limits for 2003 are $3,000 for a single individual and $6,000 for a married couple filing jointly.

[4]Olivia Mitchell (1997) reports the modal plan to have a match equal to 50 percent up to 6 percent of salary.

[5]Consider a dollar withdrawn after T years when the investor is older than 59. If m is the contribution rate, r is the interest rate, and τ is the marginal tax rate, then the value of a pre-tax dollar invested in a 401(k) is $V_{401k} = (1+m)(1-\tau)e^{rT}$. The value of a deductible IRA is $V_{ira}^{d} = (1-\tau)e^{rT}$. For a non-deductible IRA, the value is $V_{ira}^{nd} = (1-\tau)e^{r(1-\tau)T}$. Compare these to the value of a dollar invested in a conventional asset.

in the US; by 1992, IRAs and 401(k)s alone accounted for more than 55 percent of retirement savings.[6] Of course, saving for retirement constitutes only one of several motives for saving. Other categories include precautionary saving, saving for large purchases such as a home, and saving for college education. The withdrawal penalties on 401(k)s and IRAs obviously make these programs poor substitutes for the above type of saving, which requires the flexibility to withdraw before age 59.

2.2 Asset Substitution

Despite the saving incentives these programs offer, it is unclear whether they result in new saving. The answer partially depends on the source of contributions. The condition for overall saving to increase requires private saving to rise more than the tax saving generated by each program. Because tax breaks reduce government revenue, each dollar of new saving in an IRA or 401(k) is partially offset by a reduction in public saving. For example, if the present value of an increase in private saving equals the present value of the tax reduction, overall saving remains unchanged. Similarly, if contributions are financed through borrowing or shifting other assets, saving could actually fall. Asset shifting, which occurs when saving in one asset is financed through a reduction in another, can greatly diminish the effectiveness of retirement savings programs. While borrowing is unlikely, since it requires that the return on the program exceed the interest rate, asset shifting might be expected for high savers. It is also possible that IRA or 401(k) contributions simply replace new saving in other assets. If this is the case, then overall saving should decrease by the amount of the tax deductions, excluding income and substitution effects. In order for these programs to increase saving, contributions must be financed through a drop in current consumption, or equivalently, the substitution effect of the tax breaks must outweigh the income effect.

The limits on IRAs and 401(k)s play an important role in the theory predicting saving responses. To see this, suppose there are no limits. Since the 401(k) offers a higher return than either the IRA or conventional saving, the rational investor would save for retirement entirely in

[6] Author's calculation from the SIPP. Retirement saving includes DB, DC, Keogh, 401(k), and IRA.

that account. To avoid dealing with penalties, assume that funds will not be withdrawn before age 59. The higher return generated by the tax breaks on the 401(k) then changes overall saving ambiguously, since both a price and income effect operate (usually in opposite directions) on the individual's saving decision. Because consumption in the future is cheaper relative to today, there is an incentive to trade-off present for future consumption. But the higher return also makes the individual wealthier, and assuming consumption is a normal good in all periods, the wealth effect increases consumption in all periods. Only if the price effect outweighs the income effect will private saving unambiguously rise. With the presence of a contribution limit, however, the substitution effect matters only for infra-marginal savers.

To see how the contribution limits affect marginal saving decisions, consider the simplified two-period budget constraint of a person who can contribute to an IRA, a 401(k), or other assets. If IRA/401(k) and non-IRA/401(k) saving compose total saving, and the only source of wealth is a unit of labor income in period 1 of W, then the individual chooses first period consumption, IRA saving, and non-IRA saving subject to the following constraints.

$$C_1 = W - S_0 - S_{IRA} - S_{401k} \tag{1}$$

$$C_2 = S_0(1 + r_0) + S_{IRA}(1 + r_{IRA}) + S_{401(k)}(1 + r_{401k}) \tag{2}$$

$$S_{IRA} \leq K_{IRA} \tag{3}$$

$$S_{401k} \leq K_{401k} \tag{4}$$

$$r_0 < r_{IRA} < r_{401k} \tag{5}$$

where C_t is consumption in period t, S_0 is non-IRA saving which earns a return of $(1+r_0)$, S_{IRA} is IRA saving and earns $(1+r_{IRA})$, and S_{401k} is 401(k) saving and earns $(1+r_{401k})$. This formulation assumes that there are no bequests and that all remaining wealth is consumed in period 2. Also, S_{IRA} and S_{401k} must not exceed their contribution limits, K_{IRA} and K_{401k} respectively. It is easy to see that if $K_{IRA} + K_{401k} < W$, there is a kink in the budget constraint at both K_{IRA} and K_{401k}. Now suppose there is an increase in the 401k return due to some incentive policy.

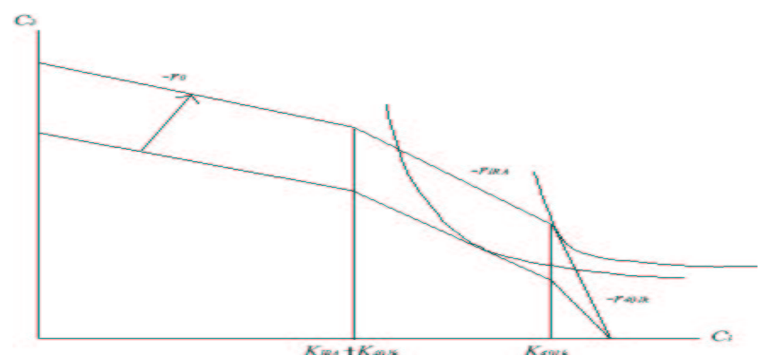

Figure 1: Response to an increase in the 401(k) return

Limit savers ($S_{401k} = K_{401k}$) cannot increase their 401k saving but are richer because of the higher return. Assuming consumption is a normal good, limit savers will increase consumption in both periods. It is therefore possible that savers near the kink (S_{IRA} small) may decrease saving enough to become infra-marginal savers at the new budget constraint. Figure 1 illustrates this result. The consumption change for infra-marginal savers, on the other hand, is ambiguous. The income effect encourages a rise in consumption for both periods, but this interacts with a substitution effect which works in the opposite direction. Whether the absence of a price effect for limit contributors significantly distorts saving incentives depends on the percentage of accounts that contribute the maximum amount. The effects of an increase in the contribution limit K_{401k}

are slightly more complicated. Infra-marginal savers are unaffected since they could have saved more before the limit change but chose not to do so. Savers located between the old kink and the new kink change consumption according to the relative income and substitution effects discussed above, and again, the change is ambiguous. Note that some IRA savers before the change may be induced to contribute nothing in the IRA and at or inside the limit in the 401(k) after the change. Finally, individuals who choose a consumption bundle that is to the left of the new kink before the change will increase consumption in both periods and decrease non-401k saving, possibly so much so, that they become interior contributors after the change. Of course, the results of this discussion apply equally well to the changes expected of non-IRA/401(k) saving with respect to a limit change in either the 401(k) or the IRA.

While the graphs are readily viewable, now is a good time to distinguish between marginal or local substitution between assets and global substitution. Marginal substitution between asset j and asset i is defined as the change in savings in asset j with respect to a small increase in the

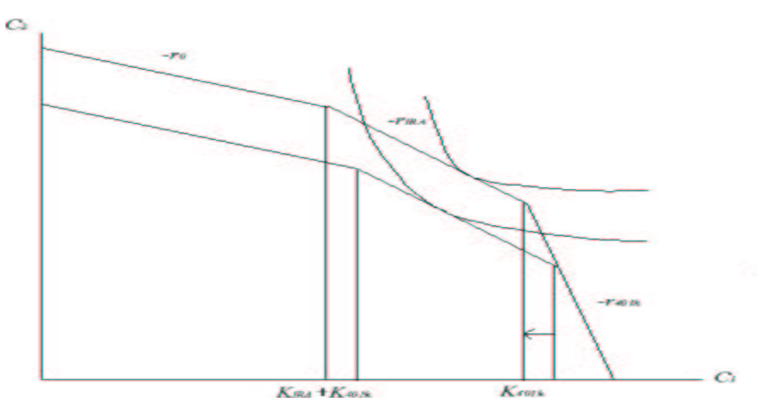

Figure 2: Response to an increase in the 401(k) limit

Table 1: IRA and 401(k) Contributions by category, 1992

Category	% Accounts	% Contributions > 0	% Contribute at Limit
IRA			
Age >59	34.69	6.27	66.67
401(k) contributor	38.60	4.59	33.33
Adjusted income ($1,000)			
<20	16.56	3.13	39.67
20 - 39	28.04	9.74	47.16
40 - 59	33.50	7.50	52.63
60 - 74	46.97	7.20	54.31
75 +	60.41	8.87	60.00
401(k)			
Age >59	3.60	2.13	22.32
Adjusted income ($1,000)			
<20	13.98	6.23	11.96
20 - 39	22.67	17.76	19.42
40 - 59	33.16	32.14	28.43
60 - 74	43.13	26.32	25.85
75 +	38.40	35.76	35.62

Source: Author's calculation from the 1992 SIPP.

contribution limit of asset i. This case is discussed above and is the measure of substitution tested in this paper. Estimates of a marginal substitution parameter can be used to predict the extent to which a limit increase will result in new saving, but they cannot address the question of how much the 401(k) has substituted for the IRA over time, unless it is assumed that the marginal effects are constant for any change in the program limit. Global substitution refers to the total change in asset j that is attributable to a change in asset i over some time period and is better tested using non-parametric analyses along the lines of Venti and Wise (1995).

Table 1 presents information on the number of IRA and 401(k) accounts that contributed to the limit in 1992. For all categories (age, 401(k) status, and income), the percentage of positive IRA accounts that contribute to the limit is high and ranges from about 33 percent to 60 percent, with the percentage increasing monotonically with adjusted income. Investors older than 59 presumably treat IRAs and 401(k)s as perfect substitutes for conventional saving since there is no withdrawal penalty after that age. If only an IRA is held (i.e., no 401(k) or Keogh) then

Table 2: Number of savers by category, 1992

Saver Category	Number of individuals	IRA ($s)		401(k) ($s)	
		Mean	S.D.	Mean	S.D.
Total amount in account					
401(k)>0	3,814	4,400	11,016	18,110	25,313
IRA>0	5,577	14,907	16,757	5,347	17,404
Annual contribution					
401(k) >0	2,635	74	424	2,775	2,923
IRA >0	916	857	507	432	1,762
IRA and 401(k)=0	24,821	-	-	-	-
IRA>0, 401(k)=0	802	1,569	642	-	-
IRA>0, 401(k)>0	114	1,583	621	3,493	3,855
IRA>0, 401(k)=limit	86	1,716	498	8,832	922
IRA=limit, 401(k)=limit	11	1,974	273	8,647	839

Source: Author's calculations using the 1992 Study of Income and Program Participation. Data are censored so that only individuals who are older than 21 are included.

contributions for infra-marginal savers in this age group should represent total saving, since the IRA returns dominate those on conventional assets. In the data, infra-marginal savers hold a significant amount of non-IRA/401(k) financial assets, which probably represents saving that is not targeted for retirement. In practice, it is not always possible to differentiate retirement saving from non-retirement saving. The early withdrawal penalties on IRAs and 401(k)s make them relatively undesirable for pre-retirement saving, so it is unlikely that retirement plan saving would substitute for other forms of saving that are not targeted for retirement. Of particular interest in table 1 is the relatively low percentage of 401(k) contributors that contributed to the limit in an IRA. If 401(k) holders only contribute to their IRA after they have reached the 401(k) limit (a reasonable assumption considering that the return on a 401(k) dominates that on an IRA), then, ceteris paribus, 401(k) limit contributors would be expected to maintain lower IRA balances than individuals contributing only to an IRA.

Since this paper estimates the degree of substitution between IRAs and 401(k)s for limit savers in the 401(k), it is useful to see how many savers contribute infra-marginally in the IRA and at the limit in the 401(k). Table 2 presents the number of individuals in each saver category and the mean and standard deviation of saving in both plans.

The most important figure in Table 2 is the number of IRA contributors who contributed to the limit in their 401(k). As can be seen above, 86 individuals both contributed to the limit in their 401(k) and saved a positive amount in their IRA. Because the measurement of substitution in this paper depends on the characteristics of these individuals, their sample size obviously influences the accuracy of estimates. Relative to the total sample of savers and non-savers, the number of these savers is quite small.

From table 2, it can be seen that the majority of IRA contributors who contribute to a 401(k) are limit contributors in the 401(k) (86 out of 114). This lends support to the allocation strategy hypothesized in this paper: namely, that individuals who have access to a 401(k) only contribute to an IRA after they have hit the limit in their 401(k). Relatively few individuals save to the limit in both plans.

2.3 Previous Research

Studies of targeted retirement programs generally concentrate on the displacement of conventional savings with respect to a change in program contribution. Steven Venti and David Wise published a number of papers that attempt to measure the impact of IRAs and 401(k)s on individual saving.[7] Their early approach conducts a parametric analysis of asset substitution at the beginning of the IRA program in order to directly address the question of whether individuals who save more under the IRA save less in other asset categories.[8] They find that an increase in IRA saving results in virtually no reduction in other forms of financial saving. Criticism of this early approach targets the decision modeling in Venti and Wise (Gravelle, 1989; Burman et al., 1990 Gale and Scholz, 1994). Venti and Wise assume a budget allocation function that fixes consumption to be a fraction of income. As Gravelle (1989) noticed, this allocation function implies an elasticity of the savings rate with respect to the rate of return equal to zero for conventional assets. To get around this problem, the authors assume that IRA saving is not treated as a substitute for conventional saving, and saving and consumption enter the budget constraint separately. This treatment,

[7] See Venti and Wise 1986, 1987.
[8] See for example Venti and Wise (1986, 1987, 1990).

however, is decidedly ad hoc and should be derived rather than assumed. Later approaches by Venti and Wise drop the entire agent optimization framework and instead concentrate on measuring asset changes over time.[9] One experiment they run (Poterba, Venti and Wise, 1995) evaluates the change in other saving when IRA participation status changes. Following the same households over time allows the authors to control for saver heterogeneity, because they can difference out the time-invariant individual effects. This approach does not account, however, for "within-individual" variation over time. They find that IRA saving and conventional saving are independent and that there is little substitution between them. Another method for dealing with saver heterogeneity would be to clump persons together according to the assets they hold or for which they are eligible. Poterba, Venti and Wise (1997) use this technique to study the effect of 401(k)s and IRAs on other financial assets and again find little asset substitution.

The conclusion that IRAs result in new net saving is supported in a cross-country comparison by Caroll and Summers (1987). They argue that the introduction and expansion of RRSPs (an IRA-type contribution plan) in Canada coincided with and partly explained an increase in the savings rate relative to the United States between 1961 and 1985. Unfortunately, their results do not necessarily indicate whether IRAs would lead to a rise in savings within the US: just that RRSP's were a successful experiment in Canada.

Joines and Manegold (1995) estimate the effect of the IRA program by studying the change in IRA contribution limits in 1981 (see section 2.1). Using a panel of individual tax returns, the authors compare the change in total financial assets with respect to the change in the IRA limit for contributing households. Controlling for tax status, they estimate that roughly 26 percent of IRA saving that was induced by the limit change created new saving. A shortcoming of this paper stems from the absence of individual effects such as marital status, income, and age. Without controlling for household characteristics, Joines and Manegold lose the ability to capture heterogeneity and other dimensions of saving.

Gale and Scholz (1994) explicitly model the household savings decision and estimate parameters using data from the 1983-1986 Survey of Consumer Finances. In contrast to other papers

[9]See for example Venti and Wise (1995).

that focus on limit and non-limit contributors, the authors measure substitutability by comparing the decisions of limit contributors to "interior" contributors. They divide saving into IRA and non-IRA, where non-IRA saving consists of non-IRA financial assets short of changes in household debt. Limit contributors have an excess demand for IRAs that can influence saving in other assets, while interior contributors do not. If IRA savers differ from non-IRA savers, their method controls for an important source of agent heterogeneity. Gale and Scholz estimate an empirical specification of their model and find that increasing the contribution limits between 1983 and 1986 would not have resulted in a significant increase in overall saving. Although this paper represents one of the most coherent efforts at explaining the effect of IRAs on saving, it is possible to complain about the assumption that all non-IRA assets are perfectly fungible. For both institutional and psychological reasons, pension assets may be separate from other conventional assets. If "self control" is costly, individuals may keep "mental accounts" that differentiate between retirement saving and other wealth (Thaler, 1990).

Apart from controlling for saver heterogeneity, a difficulty that confronts all of these papers is the unavailability of data on program contributions. For most years of the SIPP, for example, the survey asks only about asset balances. This would not be problematic if there were extended panel data on these individuals, since then contributions could be derived from the yearly change in asset balances, but this is not the case. At best, the SIPP yields two years of data on asset accumulation for individuals and usually no explicit measure of contributions. In the 1990-1992 surveys, however, the SIPP did ask about IRA and 401(k) contributions, which makes it possible to test substitution between these plans without recourse to constructed values of saving. This is the approach taken in this paper.

2.4 Data

This paper uses data from six panels of the 1990-1992 Survey of Income and Program Participation (SIPP). The SIPP collects both cross-sectional and longitudinal data on all persons over 15 years old, as well as any other persons living at the same address. Each year of the SIPP contains topical modules that provide detailed information on household saving, asset accumulation, and income, as well as other non-financial individual characteristics such as age, race, and education. Most importantly, the SIPP offers data on incentive based retirement programs such as IRAs, 401(k)s, and Keoghs. For the IRA and Keogh, survey participants are asked whether they have such a plan, whether they contribute, and if they do contribute, how much.[10] The 401(k) questions ask the following: whether individuals are eligible for such a plan, whether they participate, and if they participate, how much they contribute.[11] This captures the feature that while IRAs are available to all income earners, 401(k)s require the employer to participate.

For each year of interest, I merge two separate topical modules. One contains data for roughly 28,258 individuals on flow variables such as income, spending, and retirement plan contributions. The other provides information for around 31,746 persons on stock variables such as asset accumulation, net wealth, and debt. When merged, the data set includes information on 28,258 individuals, since only those with matching identification numbers are included. Because the paper targets household saving decisions, I censor the data to include only individuals older than 25, which eliminates most dependents. Also, to concentrate on decision-makers, I consider only direct survey respondents and their spouses, if any. Table 3 presents individual characteristics.

Although the SIPP specifically targets asset information for individuals, the reporting of asset levels is not always accurate. For example, although there is a $2,000 limit for individual IRA contributions, a number of sample participants listed contributions in excess of this amount. Also, a large number of 401(k) participants reported zero contributions, hinting at the possibility that some individuals neglected to give their "true" contributions.[12] If this unobserved error is serious, estimates acquired with 401(k) contributions as a variable may be significantly biased.

[10]The actual questions on the SIPP for the IRA are as follows:
"Does ... have an individual retirement account - an IRA - in ...s own name?" (0 = NA, 1 = yes, 2 = no)
"For how many years has ... contributed to IRA accounts?"
"how much did ... contribute?"
"As of (last reference day), what is the total balance or market value of [the IRA]?"
[11]The actual questions on the SIPP for the 401(k) are as follows:
"Does ...'s employer offer a salary reduction plan, sometimes know as a 401(k)?"
"Does ... participate in this plan?"
"how much did ... contribute?"
"As of (last reference day), what is the total amount in this plan?"
[12]Of course, zero contributions may simply indicate that previous contributors no longer choose to participate.

Table 3: Individual characteristics of IRA and 401(k) contributors, 1992

Characteristic	All Individuals	401(k) Contributors	IRA Contributors
Mean Age	48.91	40.79	46.57
Mean High grade	12.90	14.59	14.52
Mean Adjusted income	10,467	27,542	20,458
Median Non-IRA/401k assets	5,337	21,654	64,507
Median Non-liquid assets	49,573	47,062	73,171
Median Debt	30,506	23,350	55,237
Percent holding pension	6.02	33.97	6.26
Number of individuals	25,345	2,635	916

Source: Author's calculations using the 1990-1992 Study of Income and Program Participation. Data are censored so that only individuals who are older than 25 are included.

3 Contribution Limits, the Budget Constraint, and Employer Matching

In order to motivate the empirical analysis in the next section, I will now discuss the inter-temporal budget constraints facing an individual at some time t. Although IRAs and 401(k)s appear to be rather close substitutes for each other, differences in their real returns, withdrawal penalties, and limits require them to be treated separately in the budget constraint. Without explicitly modeling the consumer's optimization problem, it is useful to imagine an agent maximizing his lifetime utility subject to the following budget constraint operating at each time period T.

$$Y_t = C_t + \sum_{j \in J} (1 - I_t \theta_j) S_{jt}, \qquad J = \{401(k),\ IRA,\ Other\} \tag{6}$$

$$W_{jt+1} = R_{jt} W_{jt} + S_{jt} \tag{7}$$

$$W_{IRAt} \geq 0, \quad W_{401kt} \geq 0 \tag{8}$$

$$S_{IRAt} \in [-W_{IRAt}, K_{IRAt}], \qquad S_{401kt} \in [-W_{401kt}, K_{401kt}] \tag{9}$$

Where Y_t is disposable income in period t, W_{jt} is total savings in asset j in period t, R_{jt} is the average gross rate of return on asset j during period t, S_{jt} is saving in asset j during period t, and K_{jt} is the contribution limit on asset j. θ_j $(0 \leq \theta_j \leq 1)$ is the withdrawal penalty for asset j, and

I_t is an indicator variable taking on the value of one if the agent is less than 59 and $S_{jt} < 0$, and equals zero otherwise. These withdrawal penalties become important with the presence of income uncertainty and imply an additional cost to program contributions. Equation (6) requires that income be the sum of consumption and saving, which implies that with dissaving, consumption will exceed income. Equation (7) defines wealth in asset j at time $t+1$ as wealth last period plus new saving. Equation (8) reflects that fact that agents cannot go into debt in their retirement plans. The last equation (9) reflects the contribution limits and non-negativity constraints on savings.[13]

The involvement of these constraints hints at the difficulty of setting up an analytically soluble constrained maximization problem that would include the program features of withdrawal penalties and contribution limits. Namely, many interesting solutions to the problem would involve corner solutions. Several insights may be gained, however, simply through an inspection of the constraints themselves. The first thing to notice is that with income uncertainty, the withdrawal penalties make the IRAs and 401(k)s imperfect substitutes for other assets even if saving is targeted strictly at retirement. This occurs, because there may be a positive probability of needing to withdraw retirement funds before age 59 in reaction to a sufficiently negative income shock. Withdrawal penalties provide sufficient reason for why one would expect to see less than perfect substitution between these programs and other assets.

Apart from the difference in withdrawal penalties on the IRA and the 401(k), and ignoring the costs of joining these plans, the only plausible factor leading to imperfect substitutability between these programs is the difference in their rates of return. The tax deductibility and employer matching in the 401(k) guarantee that the present value of a dollar invested in the 401(k) is at least as large as a dollar put into an IRA. As shown in figures 1 and 2, the presence of contribution limits combined with these differences in rates of return generates kinks in the budget set at the limits.

What is not shown in these diagrams is the effect of matching limits set by the employer.

[13]In reality, 401(k)s can only be withdrawn at retirement or under special conditions such as unemployment. This would imply that saving in 401(k)s must be non-negative during the working life.

Matching limits arise, because employers not only choose the matching rate (if any), they can also set a matching limit. If the employer matching limits are below the program contribution limits, it is easy to see that there will be a new kink in the 401(k) portion of the budget constraint. Unlike the kinks at the contribution limits, however, the kinks due to matching limits will vary according to the policy chosen by the individual's employer. In the SIPP, a substantial proportion (52.14%) of individuals who contribute to a 401(k) report a positive level of employer matching. Unfortunately, the SIPP provides no information on the level of employer matching limits. The question then is whether the variation in unobserved matching limits affects the estimation of substitution in this paper. There are two channels through which matching limits may interfere with the analysis. The first and most important involves the individual's allocation rule. Earlier, I argued that because the 401(k) return dominates the IRA return, a sensible investment strategy would be to max out the 401(k) before investing in an IRA. If the only difference in program return hinges on the presence of matching in the 401(k), then once the matching limit is reached, there is no reason for an investor to prefer the 401(k) to the IRA. Since, however, the return on a 401(k) dominates that on a non-deductible IRA even without matching, the allocation rule posited in this paper still makes sense. This leaves in tact the relevant kink in the budget constraint at the 401(k) limit, which should, in turn, leave the estimates of substitution relatively unchanged. In the case of a deductible IRA, it is true that in the absence of employer matching (or beyond the matching limit) the individual should be indifferent between investing an additional dollar in an IRA and investing it in a 401(k). Here, the allocation story no longer makes sense, since there is no longer a reason for an investor to prefer the 401(k) to the IRA for all levels of the 401(k) beneath the limit. Because all of the 401(k) limit contributors in the SIPP have incomes above the amount for which IRAs become non-deductible, this issue is unimportant.

The second channel through which unobserved matching limits on the 401(k) may affect the estimation involves coefficient estimates on individual saving characteristics. The problem here is analogous to estimating characteristics on a demand equation with an unobserved change in the price of the good being considered. With the 401(k), the price change is also correlated with the amount of the good (saving in a 401(k)) consumed. It is difficult to predict exactly how the unobserved matching biases the coefficient estimates, but the bias is more likely to be quantitative than qualitative. That is, although some estimates may be biased, they are unlikely to be the wrong sign. Most importantly, as long as the allocation rule is unaffected, the estimate of substitution should not be severely biased.

While a rational agent framework suggests that the 401(k) and IRA should be nearly perfect substitutes for retirement saving, this may not be the case empirically. For example, individuals may save according to "rules of thumb" and only alter contributions when a large change in returns or limits occurs. Before attempting to address the question of whether IRAs and 401(k)s substitute for other assets, it therefore seems sensible to test the degree of substitution between the plans themselves.

4 Estimation

4.1 Estimation Model

In this section, I estimate a parametric model that is consistent with the discussion of budget constraints above. Since the return on a 401(k) always dominates that on an IRA, an agent participating in both plans will only contribute to an IRA in proportion to which his desired retirement saving exceeds the 401(k) contribution limit. The same is true regarding non-IRA/401(k) retirement saving. As discussed earlier, this need not hold true with respect to all conventional saving, since withdrawal penalties make both the IRA and the 401(k) undesirable for pre-retirement saving such as precautionary saving, college funds, and saving for a house.

The spirit of the analysis is to first estimate the desired level of 401(k) contributions, controlling for individual characteristics such as age, wealth, income, and education. This desired amount differs from the actual amount for limit contributors and individuals wishing to save negative amounts each period. I then estimate the desired IRA contributions following the same procedure as above but including an independent variable measuring the difference between desired and actual 401(k) contributions. Finally, I interpret the coefficient on this variable as a measure of marginal substitution between IRAs and 401(k)s with respect to a change in the 401(k) limit.

Specifically, the model is as follows:[14]

$$S^d_{401k} = \mathbf{X}\beta + \varepsilon_1 \tag{10}$$

$$S_{401k} = \begin{cases} 0 & \text{if } S^d_{401k} \leq 0 \\ S^d_{401k} & \text{if } 0 < S^d_{401k} \leq K_{401k} \\ K_{401k} & \text{if } S^d_{401k} \geq K_{401k} \end{cases} \tag{11}$$

$$S^d_{IRA} = \mathbf{X}\delta + \lambda\eta(S^d_{401k} - K_{401k}) + \varepsilon_2 \tag{12}$$

$$S_{IRA} = \begin{cases} 0 & \text{if } S^d_{IRA} \leq 0 \\ S^d_{IRA} & \text{if } 0 < S^d_{IRA} \leq K_{IRA} \\ K_{IRA} & \text{if } S^d_{IRA} \geq K_{IRA} \end{cases} \tag{13}$$

$$\lambda = \begin{cases} 1 & \text{if } S^d_{401k} \geq K_{401k} \\ 0 & \text{if } S^d_{401k} < K_{401k} \end{cases} \tag{14}$$

where S^d_j denotes desired saving in plan j (401(k) or IRA), $\mathbf{X}$ is a vector of individual characteristics such as age, income, education, wealth, and debt, S_j is actual saving in plan j, η measures the proportion of excess desired 401(k) saving that is contributed to the IRA, K_j is the contribution limit for plan j, β and δ are vectors of parameters, and ε_1 and ε_2 are normally distributed, mean zero error terms. Following Gale and Scholz (1994), the parameter η measuring substitutability between the plans is assumed to be a linear function of individual characteristics, such that $\eta = \mathbf{X}\gamma$, where γ is a vector of parameters. η can then be evaluated by multiplying the estimated vector γ by the vector of characteristics for limit 401(k) contributors. The lower limit on the 401(k) and IRA equations is zero, because the SIPP provides information on program

[14] This empirical specification is similar to that used in Gale and Scholz (1994). In that paper, they measure substitution between IRAs and other assets.

contributions rather than net saving.

This model does not consider the effects of a limit change on non-IRA/401(k) assets, since this paper seeks only to estimate the substitutability between IRAs and 401(k)s.[15] The basic set-up of the model, in fact, is appealingly simple. An investor who desires to save above the limit in the 401(k) puts some fraction η of the difference between what he would like to save and the limit into an IRA. If this fraction η is close to one, then the plans are close substitutes; if it is close to zero, they are not. I include independent variables based on their theoretical importance in the saving literature.

In order to account for non-linear saving with respect to age predicted in some life-cycle models, I use both age and age squared. This also captures some of the distortions caused by the withdrawal penalties. If future income and consumption are uncertain, withdrawal penalties increase the cost of saving for those younger than 59, because there is a possibility that the assets will be needed before retirement. In contrast, individuals who are older than 59 can withdraw freely and should view the retirement plans as perfect substitutes for ordinary saving. While the lower cost of the program saving for older persons encourages additional saving, the life-cycle pattern of saving implies a drawing down of assets after retirement age. In either case, the impact of IRAs and 401(k)s differs according to age.

Because saving preferences and financial acumen may be a function of education, I include a variable measuring the highest grade completed for each individual (1-18). Even if it only takes a minimal level of financial skill to contribute to an IRA or a 401(k), persons with more education may better understand the importance of saving for retirement than the less well educated.

As can be seen in equation (6), adjusted income enters fundamentally into the agent's budget constraint and is therefore included in the vector of characteristics.[16] For a similar reason, several components of wealth are used as well. Non-IRA/401(k) financial wealth, defined as the total balances in interest bearing accounts, stock and bond holdings, and other asset funds,

[15]It is easy to extend the model to measure substitution with other assets by adding an equation analogous to (8), where saving in other assets responds to the gap between desired IRA contributions and the IRA limit. It is difficult, however, to find an accurate measure of non-IRA/401(k) savings: a point discussed in section 2 of this paper.

[16]Adjusted income is the gross labor income reported by individuals on their tax forms (typically a 1040).

affects retirement saving ambiguously. On the one hand, higher wealth, controlling for saving preferences, should induce individuals to consume more in a given period. On the other hand, it is plausible that wealth is correlated with saving preference. As the latter effect should move in the opposite to the former, the predicted effect of wealth is uncertain. I also add a variable for the amount of non-liquid financial assets (approximated by home equity) held by each person. This controls for the possibility that individuals who hold a lot of non-liquid assets may be reluctant to save in a relatively illiquid retirement saving program. Finally, I include a measure of consumer debt. High consumer debt signals the presence of liquidity constraints, and therefore should move inversely with retirement plan purchases.

4.2 Estimation

In order to accommodate the contribution limits present in the above specification, I estimate equations (10) and (12) with a two-limit Tobit model using maximum likelihood. In the model presented above, zeros in the dependent variable correspond to no contributions, which may be due either to non-participation in the program or a decision not to contribute in the observation year. Table 1 presents data on program participation by income category. Participation in the 401(k) ranges from a minimum of 13.98 percent in the lowest income bracket to 38.4 percent in the highest. For IRAs, participation ranges from 16.56 percent to 60.4 percent. In both cases, the high number of zeros in the dependent variable rules out standard multiple regression models that assume normally and independently distributed error terms. That is, the model specification with limits at zero and K implies the existence of heteroscedasticity in the disturbance.

A natural choice for estimation is the Tobit model. The Tobit is a censored dependent variable model, which combines a Probit model for limit observations with a classical model for non-limit observations. I obtain estimates using maximum likelihood. It is important to keep in mind that the Tobit specification assumes that disturbances are normally distributed. If in reality this is not the case, then the estimator will be inconsistent.[17] The log-likelihood that is maximized

[17] See Amemiya (1973).

is:

$$L = -\frac{1}{2}\sum_{i\in C}\left[\left(\frac{S_{ij}-\widehat{S_j}}{\sigma_j^*}\right)^2 + \log(2\pi\sigma_j^*)\right] + \sum_{i\in L}\log\Phi\left(\frac{S_{ij}-\widehat{S_j}}{\sigma_j^*}\right) + \sum_{i\in R}\log\left[1-\Phi\left(\frac{S_{ij}-\widehat{S_j}}{\sigma_j^*}\right)\right], \quad (15)$$

where $i \in C$ are uncensored observations, $i \in L$ are left censored observations ($S^d \leq 0$), $i \in R$ are right censored observations ($S^d \geq K$), and Φ is the standard cumulative normal with $\varepsilon\tilde{} N(0, \sigma^2 I)$. S_{ij} is the observed saving in retirement plan j for individual i, and $\widehat{S_j}$ is the estimated value of saving.

The error terms in the model specification require some attention. In particular, in equation (12), if $S^d_{401k} \geq K_{401k}$, the actual error is $\varepsilon^* = \lambda\eta\widehat{\varepsilon}_1 + \varepsilon_2$. This occurs because the error in equation (10) enters equation (12) whenever desired 401(k) contributions exceed the limit. To see this, note that when $S^d_{401k} \geq K_{401k}$, $\widehat{\varepsilon}_1 = E[\varepsilon_1 \mid \varepsilon_1 \geq K_{401k} - X\beta] > 0$. Because the error is distributed normally around a zero mean, truncating part of the lower tail necessarily yields a positive expected error. Also, because the truncation point varies negatively with the size of desired 401(k) contributions, the expected error is negatively correlated to the amount of desired saving. The model assumes that $E[\varepsilon_2] = 0$, but because the estimated error term from the first equation $\widehat{\varepsilon}_1$ enters the second, the estimator (15) is biased. This can be seen by noting that ignoring the first term of ε^* is equivalent to an omitted variable, with the usual results following. With a multiple variable regression, it is difficult to predict the direction of bias, but the problem can be severe. One sure way to eliminate the inconsistency generated by the omitted error term would be to define and maximize a single log likelihood function for equations (10)-(14), making explicit the structure of the covariance matrix. While this solution would be optimal, the limiting constraints on the model make the estimation numerically infeasible.

As a second-best alternative to maximizing a single likelihood function, I estimate the error term in the first equation and include it in the estimation of equation (12). This method is a reasonable approximation as long as the correlation between the two errors is close to zero. The estimation model assumes that both ε_1 and ε_2 are distributed normally with a mean of zero, but it does not require the errors to be orthogonal to one another . Using the assumption that

$\varepsilon_1 \tilde{} N(0, \sigma^2 I)$, the conditional expected error term is

$$\widehat{\varepsilon}_1 = E\left[\varepsilon_1 \mid \varepsilon_1 \geq K_{401k} - X\beta\right] = \sigma_1 \frac{\phi\left[(K_{401k} - X\beta)/\sigma_1\right]}{1 - \Phi\left[(K_{401k} - X\beta)/\sigma_1\right]} \tag{16}$$

where σ_1 is the standard deviation, $\phi(.)$ is the standard normal pdf, and $\Phi(.)$ is the standard normal cdf.[18] From the first regression on equations (10)-(11), I estimate the standard deviation by taking the square root of the squared error to get $\widehat{\sigma}_1$. I then use $\widehat{\sigma}_1$ to estimate equation (11) for each individual contributing at the limit in the 401(k). Finally, I estimate a revised version of equations (12)-(13), where (12) is replaced with the following:

$$S^d_{IRA} = \mathbf{X}\delta + \lambda\eta(\mathbf{X}\widehat{\beta} - K_{401k} + \widehat{\varepsilon}_1) + \varepsilon_2 \tag{17}$$

where $\widehat{\beta}$ is the estimated vector of coefficients from the first regression. The appropriate fitted value for desired 401(k) saving now enters into the IRA equation, and the omitted variable bias is removed.

4.3 Results

Table 4 reports the estimates of equations (10) and (17) for 1992. The upper contribution limits are taken to be \$2,000 for the IRA and \$7,000 for the 401(k).[19]

The estimated coefficients for equations (7) and (8) seem to be consistent with research in the saving literature. Maximizing $X\beta$ with respect to age and plugging in the estimated coefficients on age and age squared yields the result that 401(k) saving reaches a maximum at age 41.23. For comparison, a similar calculation yields an optimal age for IRA saving of 50.41. This age discrepancy matches the finding that IRA holders tend to be older and wealthier than 401(k) holders.

Education appears to be an important component of saving in both plans. An additional year of education is associated with a \$318 increase in 401(k) contributions and a \$265 increase

[18]See Johnson and Kotz (1970, p. 81).

[19]Because of inflation adjustments, this undershoots the actual limit on the 401(k).

Table 4: Number of savers by category, 1992

Variable	Coefficient	Standard Error	t statistic
β Constant	-22.614	.9928942	-22.776
Age	.474537	.0385232	12.318
Age^2	-5.7535	.4103632	-14.021
High grade	.318978	.0253069	12.604
Adjusted income	7.21119	.3580271	20.141
Non-IRA/401k assets	.037543	.0398536	0.942
Non-liquid assets	-.17398	.1003823	-1.733
Debt	.328220	.0796591	4.120
IRA contributor dummy	-.062090	.2690775	-0.231
δ Constant	-21.8067	1.178026	-18.510
Age	.4166011	.0392620	10.611
Age^2	-4.13140	.3852716	-10.723
High grade	.2652451	.0284313	9.329
Adjusted income	1.171385	.4151062	2.822
Non-IRA/401k assets	.1323634	.0380587	3.478
Non-liquid assets	.6606268	.0958826	6.890
Debt	.2460981	.0802970	3.065
401k dummy	-.133162	.2100685	-0.634
γ Age	-.065967	.0482271	-1.368
Age^2	.9737634	.5532551	1.760
High grade	.0801339	.0632899	1.266
Adjusted income	-.4568041	.4136862	-1.104
Non-IRA/401k assets	-.0177288	.0132665	-1.336
Non-liquid assets	-.3112146	.1756687	-1.772
Debt	.0730646	.1168250	0.625
η as a constant	-.1310776	.1023080	-1.281
η with no error correction	-.2021056	.1094070	-1.847

Notes: *Age* is censored below 25. *Age*^{2} is in 1,000 years. *IRA* and *401(k) contributions* are in $1,000. *Adjusted income* is in $100,000's. *Non-IRA/401(k) assets* is in $100,000's. *Debt* is in $100,000's. *Non-liquid assets* is in $100,000's. *High grade* takes on values of 1-18. *IRA dummy* = 1 if individual contributed to an IRA. 401*k dummy* = 1 if individual contributed a positive amount to a 401(k). Data for adjusted income include values up to $75,000 and a single value if income exceeds this amount. For these observations, I choose $100,000 for the representative income in order to adjust for unobserved higher incomes in the distribution.

in IRA contributions for individuals saving positive amounts in the respective plans.

Income is positively and significantly correlated with both IRA and 401(k) saving, but the income coefficient for IRA saving is smaller and less significant than for the 401(k). This, however, is not unexpected. While 401(k) participation is relatively evenly distributed across income levels, the distribution of IRA holders tends to be skewed towards the highest income brackets. It may therefore be difficult to distinguish income effects in the IRA if the marginal income effect decreases with income.

A similar problem occurs with wealth. In the 401(k) equation, the coefficient on non-IRA/401(k) assets is insignificant, but in the IRA equation, the value is positive and significant. An explanation for this is that in contrast to IRA savers, who typically maintain high asset balances, the majority of 401(k) contributors hold virtually no financial wealth outside these retirement plans. A similar sign discrepancy occurs with non-liquid wealth, which varies negatively with 401(k) saving but positively with saving in the IRA. This difference may be attributable to the age distribution mentioned above. 401(k) savers are younger and hold less non-IRA/401k wealth on average than IRA savers. Since non-liquid wealth is mostly home equity, it is plausible that 401(k) savers are more likely than IRA holders to save for a house rather than retirement.

The effect of debt is positive for both plans. While this is not a very appealing result for theory, it is possible that debt holders are more sophisticated financially and therefore more likely to contribute to a retirement plan.

The coefficient estimates on the IRA contributor dummy in the 401(k) regression and the 401(k) dummy in the IRA regression are both slightly negative and insignificant. The insignificant estimate of the IRA contributor dummy casts some doubt on the allocation mechanism assumed in this paper. If individuals who are eligible for the 401(k) only contribute to an IRA after they have hit the limit in the 401(k), then one would expect a strongly positive relationship between IRA contributor status and 401(k) contributions. This is shown in figures 1 and 2 above, where individuals only contribute to the IRA after they have reached the 401(k) limit.

The variable of interest in this paper is η, which measures the fraction of excess desired 401(k) saving that is put into an IRA. Estimated coefficients for η reveal the effects of individual characteristics on the substitutability between the plans. Given the low significance of some of these numbers, any interpretation needs to be evaluated with caution. One of the reasons why these estimated coefficients are small and generally insignificant could be that the sample for this variable includes only IRA holders who contributed to the limit in the 401(k). IRA holders already represent a special group of savers who are typically high-income earners, older, and wealthy, and IRA holders with limit contributions in the 401(k) are even more elite. Thus, the sample is not only smaller for this group, which contributes to higher standard errors, their characteristics are more similar. This might explain some oddities like the reversed signs on age and age squared. Here, a negative sign on age combined with the fact that the sample tends to be older supports the hypothesis that individuals in this group have reached the point in their life-cycle where they begin to draw down saving.

When treated as a constant, the estimated coefficient η (-0.131) is negative and somewhat significant, indicating that marginal substitution between these plans is slight. When individual characteristics are controlled for, the predicted value of η is close to zero (0.0492). The results suggest that a small increase in the contribution limit on a 401(k) would not be predictably offset by a decrease in IRA contributions. As discussed above, however, this conclusion must be taken with a grain of salt, since the Tobit estimation I use is neither consistent nor unbiased.

To show the effect of the error correction in equation (12), I estimate η without including the estimated error term from the first regression. The last row of table 4 reports the value of η (-0.202) for the uncorrected estimation. Omitting the estimated error term appears to bias downwards the estimate of substitution.

Taking seriously the small measure of substitutability between 401(k)s and IRAs, considerable doubt is cast upon the hypothesis that these plans substitute marginally for other assets. This would have important implications for policy considerations. For example, suppose the government decides to increase the limit on the 401(k). Limit savers will predictably place the difference between the limits (or the gap between desired 401(k) contributions and the limit if the new limit exceeds desired contributions) into the 401(k). For 401(k) limit savers, total retirement saving would therefore increase by approximately (exactly if the marginal effects hold for any limit change):

$$(1-\eta)\min[K_{401k}^{new}-K_{401k}^{old}, S_{401k}^{old}-K_{401k}^{old}] \geq 0,\ \eta \in [0,1] \tag{18}$$

Thus, the smaller η is, the greater the increase in total retirement saving. While this experiment ignores extra-marginal saving effects, the general result is that an increase in the 401(k) contribution limit would result in new retirement saving for limit contributors.

A weakness of the analysis in this paper is that it does not account for the behavior of limit savers in both the 401(k) and IRA. All that is argued is that low substitution between the plans suggests low substitution between these plans and other assets. Limit savers in both the IRA and 401(k) might desire to save in IRAs and 401(k)s far in excess of their respective limits. To the extent that individuals in this group are more financially sophisticated, they may be supposed to be the most likely to fund increased contributions to the 401(k) with dissaving in other assets in response to the limit increase. This would be expected if instead of following simple rules of thumb for retirement saving, they target an optimal amount. If this is the case, then the substitution effect predicted by η underestimates the actual asset displacement that would occur in response to a change in the 401(k) limit.

5 Conclusion

In this paper, I have sought to measure the marginal substitution between IRAs and 401(k)s. Using a two-stage Tobit analysis, I have estimated a coefficient on excess desired 401(k) contributions in the IRA equation and found it to be close to zero, indicating low substitutability between the plans. The advantage of my analysis is that it does not rely upon the constructed measures of asset saving used in other papers, but instead concentrates on the data on retirement plan contributions. My approach has been to measure the substitution between retirement plans as accurately as possible before considering the more difficult question of whether these plans substitute for other saving. I have then interpreted a low estimated level of substitution between the IRA and the 401(k) as implying a low probability of substitution between retirement programs and other assets.

This approach can be improved upon in at least two directions. First, the saving allocation function used in this paper should be derived from some sort of agent-optimizing model rather than assumed. Solving such a model would reveal more about individual saving decisions, as well as provide a baseline model into which more realistic features could be built. Although this paper concluded that substitution between programs is slight, it did not explain why. Secondly, research should focus on the degree of substitution between IRAs and 401(k)s over a longer period of time. As mentioned earlier in this paper, the greatest barrier to this avenue of research is data quality and availability. This, however, is changing. From 1990, the SIPP has included detailed questions about individuals' 401(k) and IRA contributions and asset levels. Later this year, the 1996 SIPP panel which covers retirement saving will be released. It will then be possible to construct a panel data set that covers at least six years,[20] which would make it possible to control for the individual heterogeneity that frequently confounds estimation of saving behavior.

A low level of substitution between savings plans that seem, for most purposes, to be close substitutes, indicates that saving decisions may not be consistent with an agent-optimizing framework. The data suggest that this may be the case. For example, a substantial number of households that held IRAs and 401(k)s did not contribute to the limit in the 401(k). Because the 401(k) generates a higher real return than the IRA, rational households should only contribute to the IRA if they already contribute to the limit in a 401(k). Applying the traditional analysis, infra-marginal households contributing to both plans save sub-optimally. A flexible model of saving that includes the possibility of "rules of thumb"[21] type decisions may account for this behavior and provide a useful extension of the conventional specification.

References

[1] Amemiya, T. (1973), "Regression Analysis When the Dependent Variable Is Truncated Normal," *Econometrica*, 41, 997-1025.

[20] Currently, SIPP data on 401(k) and IRA contributions is available for only three years (1990-2).
[21] See Thaler (1990).

[2] Bernheim, B. D. (1993), "Private Saving and Public Policy," in James M. Poterba, ed., Tax policy and the economy, Vol. 7. Cambridge, MA: MIT Press, 73-110.

[3] Burman, L.; J. Cordes and L. Ozanne (1990), "IRAs and National Saving." *National Tax Journal*, 43(3),.259-83.

[4] Carroll, C.D., and L. H. Summers (1987) "Why Have Private Savings Rates in the United States and Canada Diverged?" *Journal of Monetary Economics*, 20(2), 249-79.

[5] Engen, Eric M., William G. Gale, and John Karl Scholz (1996), "The Effects of Tax-Based Saving Incentives on Saving and Wealth," *NBER Working Paper Number W5759.*

[6] Gale, W. G. and J. K. Scholz (1994), "IRAs and Household Saving." *The American Economic Review*, 84(5), 1233-1260.

[7] Gravelle, J. G. (1991), "Do Individual Retirement Accounts Increase Savings?" *Journal of Economic Perspectives*, 5(2), 133-148.

[8] Investment Company Institute (2002), "IRA Ownership in 2002" *ICI Brief.*

[9] ———, (2000) "401(k) Plan Participants: Characteristics, Contributions, and Account Activity" *ICI Research Series.*

[10] Johnson, N., and S. Kotz (1970), Distributions in Statistics-Continuous Univariate Distributions-2. New York: John Wiley and Sons.

[11] Joines, D. H. and J. G. Manegold (1991), "IRA and Saving: Evidence from a Panel of Taxpayers." *Federal Reserve Bank of Kansas City Research Working Paper* No. 91-05.

[12] Mitchell, O., "New Trends in Pension Benefit and Retirement Provisions," *NBER Working Paper No. 7381, (1999).*

[13] Poterba, J. M., S. F. Venti,. and D. A. Wise (1997), "The Effects of Special Saving Programs on Saving and Wealth." in M. Hurd and N. Yashiro (ed.), The Economic Effects of Aging in the United States and Japan. University of Chicago Press.

[14] ——— (1995), "Do 401(k) Contributions Crowd Out Other Personal Saving?" *Journal of Public Economics*, 58(1), 1-32.

[15] Thaler, R. H. (1990), "Saving, Fungibility, and Mental Accounts." *Journal of Economic Perspectives*, 4(1), 193-205.

[16] Venti, S. F. and D. A. Wise (1986), "Tax-Deferred Accounts, Constrained Choice and Estimation of Individual Saving." *Review of Economic Studies*, 53(4), 579-601.

[17] ——— (1987), "IRAs and Saving," in Martin Feldstein, ed., The Effects of Taxation on Capital Accumulation. A National Bureau of Economic Research Project Report, University of Chicago Press, 7-48.

[18] ——— (1990), "Have IRAs Increased U.S. Saving? Evidence from Consumer Expenditure Surveys." *Quarterly Journal of Economics*, 105(3), 661-98.

[19] ———, (1995) "Individual Response to a Retirement Saving Program: Results from U.S. Panel Data." *Ricerche-Economiche*, 49(3), 235-54.