Volume Title: National Saving and Economic Performance

Volume Author/Editor: B. Douglas Bernheim and John B. Shoven, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-04404-1

Volume URL: http://www.nber.org/books/bern91-2

Conference Date: January 6-7, 1989

Publication Date: January 1991

Chapter Title: The Saving Effect of Tax-deferred Retirement Accounts: Evidence from SIPP

Chapter Author: Steven F. Venti, David A. Wise

Chapter URL: http://www.nber.org/chapters/c5989

Chapter pages in book: (p. 103 - 130)

4 The Saving Effect of Tax-deferred Retirement Accounts: Evidence from SIPP

Steven F. Venti and David A. Wise

Individual retirement accounts (IRAs) rapidly became a very popular form of saving after they became available to all employees in 1982. Annual contributions grew from about $5 billion in 1981 to about $38 billion in 1986. Preliminary data indicate that contributions declined precipitously after the Tax Reform Act of 1986, even though the legislation limited the tax deductibility of contributions only for families who have annual incomes over $40,000 and who are covered by a firm pension plan. Whereas over 15 percent of tax filers made contributions in 1986, only 7 percent contributed in 1987. Two claims received considerable attention in the legislative debate over the tax treatment of IRAs. One was that the accounts were held primarily by the wealthy, a claim that is not supported by the data. Although wealthier households are much more likely than poor households to have IRAs, approximately two-thirds of accounts are held by households with incomes less than $50,000. The second claim was that IRAs produced no net saving; funds were simply transferred from other saving balances, or, if there was new saving, it would have taken place anyway. In earlier papers (Venti and Wise 1986, 1987a, 1987b, 1988b; Wise 1987) we analyzed the relationship between IRA saving and other financial asset saving. Those studies were based on household data from the 1983 Survey of Consumer Finances (SCF) and the 1980–85 Consumer Expenditure Surveys (CESs). At most, the evidence from these studies showed only a very modest substitution of IRA for other forms of

Steven F. Venti is an associate professor of economics at Dartmouth College and a research associate of the National Bureau of Economic Research. David A. Wise is John F. Stambaugh Professor of Political Economy at the John F. Kennedy School of Government, Harvard University, and a research associate of the National Bureau of Economic Research.

Financial support from the National Institute on Aging and the Hoover Institution is gratefully acknowledged. The authors are also grateful for discussion with Michael Rothschild, some of whose initial comments as a discussant have been incorporated in this version of the paper. The authors also thank Douglas Bernheim for his comments.

saving, indicating that the net saving effect was substantial. Recent analysis by Feenberg and Skinner (1989), using a panel of individual tax returns for 1980–84 also finds little evidence of substitution.

The results on IRAs are consistent with analysis of contributions to Registered Retirement Saving Plans (RRSPs) in Canada by Wise (1984 and 1985), and with the comparison of Canadian versus U.S. savings rates over time by Carroll and Summers (1987). A program comparable to the IRA has existed in Canada since 1956. In the early 1970s the contribution limits were increased substantially and the program was widely publicized. The maximum individual limit was $3,500. New limits will be as high as $15,000. Although the program has been in existence in Canada much longer than in the United States, and although the limits are based on income and for some are much higher than in the United States, Wise (1985) shows that the relationship between desired contributions and income is virtually the same in the two countries after accounting for the differences in the limits. Carroll and Summers (1987) show that after moving in tandem for almost 25 years the private savings rates in the two countries diverged dramatically after 1975, following expansion of the RRSP program. Corporate saving in the two countries, they find, has shown no long-term trend since 1954. The increase in the Canadian private saving rate and the decrease in the U.S. rate resulted from changes in the behavior of individuals, not corporations. Whether the increase in Canada was due to the RRSP program can only be judged by the coincidence of the two events and by the apparent lack of other explanations.

Nonetheless, simple forms of theoretical reasoning raise doubts about the net saving effect. Thus the question is reconsidered in this paper, based on data that are, in principle, better than the other data that we have used. The analysis here is based on the Survey of Income and Program Participation (SIPP). A total of almost 20,000 households were covered in the first nine waves—now available—of this panel survey. Each household in the survey is interviewed quarterly for 32 months. In principle, the data provide information on IRA contributions in two consecutive years, allowing statistical correction for individual-specific saving effects. Such effects may have influenced to some extent our prior results. Unfortunately, these data have not been entered on the data tapes released to us to date. Thus the analysis in this paper is based on contributions in a single year only, calculated as the difference between balances reported in the fourth (September–December 1984) and the seventh (September–December 1985) waves of the survey.[1]

We begin with descriptive data on IRAs and other forms of saving. Because the paper is directed to IRA contributions, self-employed persons and those over 65 and under 21 have been excluded.[2] Most of the descriptive data can be compared with information from the 1983 Survey of Consumer Finances and from the Consumer Expenditure Surveys, with no major inconsistencies. The following conclusions may be drawn from these data:

- The typical American family has very little financial asset saving, consistent with evidence from other surveys. The median of financial assets, including stocks, was only $600 in 1985. The majority of the saving of most families is in the form of housing.
- Families who have contributed to IRAs since 1982 had not, prior to that time, accumulated financial assets at a rate even close to the annual IRA limit.
- Comparison of IRA balances with other asset balances, or of the annual change in IRA balances with the change in other asset balances, provides no evidence of substitution of IRAs for other saving, even after controlling for several family attributes like age and income.
- These data apparently reveal individual-specific savings effects; individuals who save in one form are also likely to save in other forms as well.
- The data provide no evidence that IRAs have been funded by borrowing.

The incentive effects of IRA accounts are considered next. Attention is directed to the possibility that retirement saving and saving for other purposes may be treated by individuals as distinct "goods." That is, it may be incorrect to think of the IRA tax deduction as simply a subsidy to the one and only form of saving. To the extent that this is true it invalidates the simple theoretical reasoning that suggests little net saving effect of IRAs.

The formal statistical model that we estimate is summarized next and the estimation results are discussed. The conclusions are summarized by simulating the effect of an increase in the IRA limit.

- If the IRA limit of each family in the sample were increased by $1,000, the annual IRA contributions of families at the current limit would increase by an average of $856.
- About two-thirds of the increase would be financed by reduced consumption and about one-third by reduced taxes. Very little would be financed by reducing other saving or by increasing debt.

The last section contains a discussion and summary of the paper.

4.1 Descriptive Data

The SIPP data are organized by household and by subfamilies within households. Other surveys, like the SCF and the CES collect data only by household. Thus for comparative purposes most of the data presented here are also by household; the family data are also presented in most instances. In principal, the IRA information should be analyzed by family unit; they are most likely to correspond to tax units. In practice, however, the difference may be small. Data on accumulated wealth are presented first, then data on annual saving (change in asset balances). In each case the relationship between IRA and other saving is emphasized.

4.1.1 Accumulated Wealth

Household Assets

The data in table 4.1a confirm that the vast majority of the personal wealth of most households is housing equity.[3] The table shows the median of assets by type of asset and by income and age. The median of total wealth is

Table 4.1a **Median Household Wealth by Type of Asset and by Income and Age in 1985**

Type of Asset and Age	Income (in thousands) <10	10–20	20–30	30–40	40–50	50–75	>75	All
Total wealth:								
<25	0	838	4,425	7,797	16,794	58,469	285,476	2,031
25–35	0	2,053	7,394	12,488	22,535	35,450	53,775	6,325
35–45	500	8,569	26,850	34,108	41,400	60,375	111,488	30,112
45–55	1,500	14,275	36,200	42,242	61,850	88,675	129,236	45,724
55–65	10,175	38,750	58,500	71,284	81,700	99,730	171,715	56,241
All	500	5,822	21,340	30,850	43,329	68,197	120,483	25,067
Housing equity:								
<25	0	0	0	0	3,250	25,350	66,000	0
25–35	0	0	126	6,250	14,250	20,000	20,000	0
35–45	0	930	19,000	27,000	30,400	44,000	60,500	22,400
45–55	0	10,000	27,500	31,000	44,000	56,000	65,000	32,000
55–65	6,000	30,000	44,000	50,000	50,000	62,100	76,000	40,000
All	0	0	14,000	23,000	30,000	45,800	62,500	16,974
Financial assets, including stocks:								
<25	0	188	843	1,500	1,827	6,110	22,975	430
25–35	0	200	820	1,724	3,607	6,500	12,572	840
35–45	0	200	1,100	2,100	3,640	7,650	22,200	1,550
45–55	0	200	1,200	2,734	5,000	9,200	21,198	2,400
55–65	50	2,625	9,498	8,873	9,500	20,419	48,470	6,350
All	0	300	1,300	2,220	4,250	9,236	23,867	1,600
Financial assets, excluding stocks:								
<25	0	188	793	1,298	1,750	2,250	12,508	400
25–35	0	200	600	1,318	3,000	4,058	10,000	720
35–45	0	171	900	1,600	2,500	5,500	14,374	1,200
45–55	0	200	801	2,040	3,503	5,740	14,150	1,830
55–65	42	2,100	6,600	6,002	7,000	13,500	36,470	4,675
All	0	290	1,000	1,800	3,124	6,000	15,550	1,275
Debt:								
<25	0	492	1,170	1,700	2,105	2,290	3,556	800
25–35	0	606	1,000	1,600	1,765	3,000	4,050	1,000
35–45	0	500	1,000	1,400	1,600	1,600	3,050	1,000
45–55	0	213	545	1,025	1,560	2,300	3,025	849
55–65	0	47	195	442	600	1,240	2,000	200
All	0	350	700	1,200	1,540	2,000	3,075	750

Note: Sample is weighted to represent the national population of households with head age 21 to 65 and not self-employed.

$25,100.[4] The median of housing equity is $17,000. Including stocks, the median level of financial assets is only $1,600; excluding stocks it is only $1,275. Thus saving in the form of financial assets is typically very limited.[5] It is even smaller taking the family as the unit of analysis, as shown in table 4.1b.[6] Including stocks, the median of family financial assets is only $600. The median of total family wealth is only $8,100. Consistent with analysis

Table 4.1b **Median Family Wealth by Type of Asset and by Income and Age in 1985**

	Income (in thousands)							
Age	<10	10–20	20–30	30–40	40–50	50–75	>75	All
Total wealth:								
<25	20	11,181	3,800	5,752	9,157	41,305	356,945	500
25–35	0	2,005	7,805	14,515	29,375	44,473	67,450	3,249
35–45	500	6,425	24,600	35,831	52,489	80,248	130,448	23,978
45–55	670	20,950	41,463	57,863	76,700	117,025	163,375	40,025
55–65	6,521	46,909	76,899	89,369	111,200	145,021	274,690	57,436
All	200	3,700	19,400	33,180	55,397	86,770	162,604	8,069
Housing equity:								
<25	0	0	0	0	0	23,000	0	0
25–35	0	0	0	5,300	15,500	18,025	33,000	0
35–45	0	0	13,250	25,000	32,250	46,000	63,500	13,500
45–55	0	10,000	30,000	36,500	48,398	57,000	70,000	25,000
55–65	0	30,000	45,000	54,323	57,000	70,000	86,500	35,000
All	0	0	8,000	20,000	32,000	46,000	69,500	0
Financial assets, including stocks:								
<25	12	300	843	1,200	1,827	7,750	351,570	150
25–35	0	200	950	1,800	3,698	6,349	10,000	400
35–45	0	200	1,196	2,120	4,025	8,200	20,099	1,000
45–55	0	399	1,500	3,309	6,000	11,500	20,698	1,284
55–65	0	2,358	8,048	12,785	17,279	35,400	79,700	3,700
All	0	314	1,325	2,500	5,000	10,453	30,100	600
Financial assets, excluding stocks:								
<25	10	251	750	978	1,300	1,189	1,570	125
25–35	0	200	737	1,400	3,000	4,457	7,499	350
35–45	0	200	955	1,548	2,800	5,600	14,374	800
45–55	0	350	1,200	2,499	4,400	6,800	13,450	1,000
55–65	5	2,000	5,235	9,000	10,500	25,300	47,900	3,000
All	0	300	1,030	1,898	3,510	6,500	15,000	500
Debt:								
<25	0	150	800	1,266	1,000	282	1,500	0
25–35	0	350	785	1,500	1,500	2,000	3,540	400
35–45	0	468	900	1,250	1,475	1,400	2,900	1,650
45–55	0	200	500	1,130	1,200	1,200	1,850	350
55–65	0	50	150	300	80	620	506	24
All	0	240	600	1,100	1,100	1,200	1,575	275

Note: Sample is weighted to represent the national population of households with head age 21 to 65 and not self-employed.

based on the SCF and the CES, these data make clear that the typical family was not, prior to the introduction of IRAs, accustomed to saving even close to the IRA annual limit, $2,000 per year per worker.

The Distribution of IRA Accounts by Age and Income

The percentage of households with IRA accounts and the mean balance in these accounts is shown in table 4.2, panel A, by age and income; comparable data for families is shown in table 4.2, panel B.[7] Overall, 25 percent of households have IRA accounts. The percentage increases with both age and in-

Table 4.2 IRA Accounts by Age and Income

	Income (in thousands)							
Age	<10	10–20	20–30	30–40	40–50	50–75	>75	All
	A. By Household							
Percentage of households with IRA accounts:								
<25	.0	4.1	6.8	7.5	8.7	29.0	29.6	6.2
25–35	1.1	6.2	11.5	17.0	27.5	33.9	54.8	14.1
35–45	1.2	8.1	22.5	21.8	34.8	48.3	70.4	24.4
45–55	8.3	13.8	25.3	29.7	40.3	56.6	68.6	32.6
55–65	9.8	257.1	40.6	43.5	54.4	67.2	76.9	38.6
All	4.7	12.1	22.0	24.5	36.6	51.2	68.3	25.0
Mean IRA balance:								
<25	0	162	80	258	419	3,372	2,346	342
25–35	16	195	391	526	1,378	1,886	3,810	607
35–45	18	367	1,022	1,222	2,240	3,787	6,540	1,588
45–55	551	660	1,574	1,976	2,858	4,924	8,010	2,588
55–65	562	2,028	3,415	3,817	5,314	6,908	8,674	3,495
All	260	691	1,314	1,508	2,593	4,343	7,071	1,818
	B. By Family							
Percentage of families with IRA accounts:								
<25	1.2	3.0	4.9	7.6	3.4	21.6	.0	2.4
25–35	1.5	6.6	12.7	17.9	28.9	33.3	62.5	10.9
35–45	2.4	9.4	23.4	22.4	37.0	53.8	74.4	22.2
45–55	6.2	16.2	27.6	35.2	49.8	66.7	74.2	29.4
55–65	8.6	26.8	42.9	56.5	67.5	74.6	74.1	34.9
All	3.3	10.8	22.4	27.5	40.8	56.1	72.6	18.8
Mean IRA balance:								
<25	22	86	65	166	136	228	0	52
25–35	34	195	445	572	1,441	1,791	4,889	427
35–45	49	374	1,048	1,273	2,313	4,288	7,314	1,359
45–55	360	848	1,680	2,170	3,903	7,124	8,929	2,305
55–65	464	1,854	3,660	5,320	6,961	7,848	8,951	3,081
All	141	551	1,322	1,747	2,979	5,067	7,974	1,290

Note: Sample is Weighted to represent the national population of households with head age 21 to 65 and not self-employed.

come. About 19 percent of families have accounts. Although wealthier households and families are much more likely than poorer households and families to have accounts, most account holders are not wealthy, as shown in table 4.3, panels A and B, for households and families, respectively. About two-thirds of households with at least one account have household income less than \$50,000; these households own about 52 percent of IRA assets. Of families with accounts, about 76 percent have income less than \$50,000 and these families own about 66 percent of IRA assets.

IRAs and Other Financial Asset Saving

IRA account holders also save more in other forms as well, consistent with evidence from other surveys.[8] In addition, IRA holders also have less debt. The data are shown in table 4.4, panels A and B, for households and families, respectively. These data provide no evidence that IRA saving substitutes for other financial asset saving. Nor do the data indicate that IRA accounts are funded by borrowing, as has been suggested by some commentators. Rather, these data apparently reflect individual-specific saving behavior; savers save more than nonsavers in all forms, including IRAs. And, almost by definition, savers borrow less. Even typical IRA holders, however, had not accumulated financial assets at close to the IRA annual limit, as is evident from the median balances.

Table 4.3 Distribution of IRA Accounts and Balances by Income, 1985

Income (in thousands)	Percentage with Accounts	Mean Balance		Cumulative Percentage	
		All	IRA > 0	All Accounts	Total Balances
			A. By Household		
<10	4.7	260	5,754	2.6	2.0
10–20	12.1	691	5,628	12.8	10.0
20–30	22.0	1,314	6,058	31.8	25.6
30–40	24.5	1,508	5,887	48.2	39.4
40–50	36.6	2,593	7,091	63.8	54.6
50–75	51.2	4,343	8,408	87.4	82.1
75+	68.3	7,071	10,460	100.0	100.0
All	25.0	1,818	7,303	. . .	. . .
			B. By Family		
<10	3.3	141	4,325	4.8	3.0
10–20	10.8	551	5,084	20.1	14.3
20–30	22.4	1,322	5,901	43.6	34.6
30–40	27.5	1,747	6,352	61.2	50.8
40–50	40.8	2,979	7,295	75.9	66.4
50–75	56.1	5,067	9,040	93.2	89.2
75+	72.6	7,974	10,982	100.0	100.0
All	18.8	1,290	6,873	. . .	. . .

Table 4.4 **Financial Assets and Debt of IRA Account Holders and Nonholders, by Income, 1985**

Income (in thousands)	Median Financial Assets Including Stocks		Median Financial Assets Excluding Stocks		Median Debt	
	IRA > 0	IRA = 0	IRA > 0	IRA = 0	IRA > 0	IRA = 0
			A. By Household			
<10	7,625	0	6,500	0	0	0
10–20	6,538	200	4,800	200	250	400
20–30	6,365	900	5,000	700	400	800
30–40	6,015	1,692	4,080	1,349	600	1,475
40–50	10,000	2,694	6,800	2,005	800	2,000
50–75	14,516	5,100	9,709	3,367	1,500	2,581
75+	36,085	9,735	21,475	6,687	2,613	4,425
All	10,800	728	7,641	600	900	700
			B. By Family			
<10	2,600	0	2,065	0	0	0
10–20	4,000	250	3,000	200	300	238
20–30	6,000	950	4,998	737	300	700
30–40	6,756	1,800	4,320	1,400	554	1,400
40–50	10,450	3,000	7,000	2,420	650	1,600
50–75	17,900	5,000	10,100	3,862	1,000	2,020
75+	37,700	11,000	19,000	6,877	1,400	3,000
All	8,600	300	5,922	270	500	200

Regression Summary of IRA and Other Saving

The relationship between IRA balances and other assets may be summarized by regressions of other wealth on IRA balances. The results are shown in table 4.5, panels A and B, for households and families, respectively. In addition to IRA balances, the regressions control for current income, age, age × income, education, marital status, and private pension coverage. It is clear that larger IRA balances are associated with greater wealth in all other forms, not less. Again, the data apparently reflect individual-specific saving effects.

4.1.2 Annual IRA Contributions and Other Saving

We next consider the relationship between IRA contributions (change in IRA balances) and the change in other saving balances between 1984 and 1985, first by considering summary tabulations and then by simple descriptive regressions.

Summary Tabulations

The relationship between IRA saving and other financial asset saving and debt is shown for households and families in table 4.6, panels A and B, respectively. The figures in the first two columns are the percentage of house-

Table 4.5 **Regression Parameter Estimates, Other Assets on IRA Balances, 1985**

Other Asset Category	IRA Balance Coefficient	Standard Error
	A. By Household[a]	
Total wealth	2.80	(.26)
Housing equity	1.02	(.09)
Nonhousing wealth	1.78	(.24)
Financial assets, including Stocks	1.25	(.11)
Financial assets, excluding Stocks	1.00	(.05)
Debt	−.07	(.03)
	B. By Family[a]	
Total wealth	2.48	(.17)
Housing equity	1.05	(.07)
Nonhousing wealth	1.44	(.14)
Financial assets, including Stocks	.76	(.09)
Financial assets, excluding Stocks	.82	(.04)
Debt	−.06	(.02)

[a]The regressions also control for current income, age, age × income, education, marital status, and private pension coverage.

Table 4.6 **IRA and Other Financial Asset Saving and Debt, 1984–85**

Income (in thousands)	% Non-IRA Saving >0[a]		% Δ Debt > 0		Median Δ Debt	
	IRA > 0	IRA = 0	IRA > 0	IRA = 0	IRA > 0	IRA = 0
			A. By Household[a]			
<10	49	21	17	22	−10.5	0
10–20	48	39	44	37	0	0
20–30	50	49	44	41	0	0
30–40	60	54	43	49	0	0
40–50	56	55	49	55	0	191
50–75	58	60	46	51	0	65
75+	62	71	55	47	50	0
All	55	43	46	39	0	0
			B. By Family[a]			
<10	42	22	21	19	0	0
10–20	53	40	30	36	0	0
20–30	55	51	42	42	0	0
30–40	60	55	44	50	0	50
40–50	53	56	42	48	0	0
50–75	58	61	49	47	0	0
75+	71	74	53	39	325	0
All	56	40	42	34	0	0

[a]Excluding stocks.

Table 4.7 Regression Parameters, Change in Other Assets on IRA Saving, 1985–84

Change in Other Asset Balances	IRA Saving Coefficient	Standard Error
	A. By Household[a]	
Total wealth	.65	(.24)
Housing equity	.23	(.13)
Nonhousing wealth	.42	(.19)
Financial assets, including stocks	.49	(.12)
Financial assets, excluding stocks	.31	(.08)
Debt	.07	(.07)
	B. By Family[a]	
Total wealth	.85	(.18)
Housing equity	.26	(.12)
Nonhousing wealth	.60	(.14)
Financial assets, including stocks	.33	(.09)
Financial assets, excluding stocks	.21	(.07)
Debt	.05	(.04)

[a]The regressions also control for current income, change in incomes between 1984 and 1985, age, age × income, education, marital status, and private pension coverage. Total wealth and nonhousing wealth exclude IRAs.

holds with positive non-IRA saving, distinguished by whether the family was an IRA contributor (IRA > 0) or a noncontributor (IRA = 0).[9] Controlling for income, it is clear that IRA savers are at least as likely as non-IRA savers to save in other financial asset forms. The next four columns show the change in debt for IRA contributors and noncontributors. There is little relationship between IRA saving and debt; the data provide no evidence that IRA saving is accompanied by increased debt. Apparently, IRAs are not typically funded by borrowing. And there is no indication of substitution away from other financial asset saving. As emphasized above, the positive relationship between the two forms of saving is likely to reflect individual-specific savings effects. There is, however, no guarantee that inducement to fund an IRA account does not at the same time lead to increased consideration of future needs and thus to increased saving in other forms as well. In general, the virtual absence of saving among a large proportion of the population seems inconsistent with careful life-cycle planning.

Descriptive Regressions

The relationship between annual IRA saving and saving in other forms can be summarized by simple regressions of the change in other asset balances on IRA saving, controlling for other individual attributes. The results are shown in table 4.7, panels A and B, for households and families, respectively. Again these relationship show little substitution of IRA for other forms of saving. For example, the coefficient on total wealth (excluding IRAs) is 0.65, the coefficient on nonhousing wealth is 0.42, and the coefficient on debt is 0.07. The results for households and families are very similar. Because the regressions control for several individual attributes, the effect of individual-specific saving effects is less likely to have an important effect on these results than on the tabulations above.

4.2 The Incentive Effects of IRAs[10]

4.2.1 Promotion of IRAs

The widespread promotion of IRAs may have been the most important reason for their use, as emphasized in our previous work. Advertising of IRAs has typically emphasized the avoidance of current taxes, as well as the importance of prudent planning for retirement. They are available through almost any bank and through many other financial institutions. Recent evidence lends support to the speculation that promotion has been an important determinant of IRA purchasing behavior. First, according to preliminary IRS data, only 7.2 percent of those filing tax returns contributed to IRAs in 1987; in the previous year over 15 percent contributed. The reduction evidently reflects contributor misperceptions about the eligibility changes in the Tax Reform Act of 1986. Although the law affected IRA tax deductibility only for families who have both qualified pensions and incomes over $40,000, reporting of the 1986 Tax Reform Act and the less intense promotion by financial institutions has apparently left the widespread impression that the IRA has been eliminated. Indeed, a recent survey shows that about half of all persons who are in fact still eligible to contribute to an IRA think they are not.[11]

Another indication that promotion plays an important role is provided by Feenberg and Skinner (1989). Their data on tax returns suggest that families are often unaware of the actual contribution limits. A large fraction of families with legal limits of either $2,250 or $4,000 contribute exactly $2,000. In their view "the most compelling explanation for the false $2,000 limits is that the advertisements and brochures for IRAs common during the early 1980s made both a positive impression on consumers (encouraging them to buy IRAs) and a negative impression (that $2,000 was the legal limit)."[12]

Evidence on the role of promotion is also provided by the timing of IRA contributions. Contributors transferring assets from one account to another and seeking only to maximize the tax advantage of an IRA should contribute in January. Yet typically 40 to 50 percent of all contributions are made in March or April of the following year (Summers 1986). Such a response is undoubtably influenced by the intense advertising that coincided with the tax filing deadline.

4.2.2 Simple Economic Incentives

Two aspects of IRAs provide more traditional economic incentives to save through their use: one is that the contribution itself is tax deductible, the other is that the interest on the contribution accumulates tax free, with taxes paid only when funds are withdrawn from the account. On the other hand, once money is placed in an IRA account there is a 10 percent penalty for withdrawal before the age of 59.5. (The penalty is now 15 percent.) In this sense, the IRA is less liquid than a conventional account.

Some persons of course may consider the illiquidity of IRAs an advantage; it may help to ensure behavior that would not otherwise be followed. It may be a means of self-control. The fact that the opportunity is lost if a contribution is not made in the current year may serve the same purpose. One cannot, as with conventional saving, put it off—possibly a self-delusion—until the next year.[13]

On the other hand, because of the higher return on IRAs, to achieve any given level of retirement income requires less saving if funds are placed in an IRA account than if they are placed in a conventional account. This "income" effect raises the possibility that there could in fact be less saving with than without IRAs. The effect of IRAs on saving is the net result of all of these factors, including their promotion, and will depend on the distinction that investors make between IRA saving for retirement and other saving, as explained below.

4.2.3 One Form of Saving or Two

It may be tempting to think of IRAs and conventional saving accounts as equivalent assets, or goods, simply with different prices, in which case one might think of IRAs as only a price subsidy of conventional saving with a limit on the quantity that can be had at the subsidized price. But to the extent that consumers treat them as different assets or goods—possibly because one is intended for retirement and the other for short-term saving or because one is less liquid than the other—and to the extent that the promotion has influenced their use, this view will not yield an adequate representation or forecast of the saving effect of IRAs. Indeed, our previous work indicates quite strongly that the two are not treated as equivalent by consumers.[14]

The idea may be made clear by the use of two graphs. Figure 4.1 is intended to represent a simple view of the effect of IRAs on saving. It shows the tradeoff between the allocation of current income to current consumption versus saving for future consumption, for three current income levels. The dashed lines represent budget constraints without the IRA program and the solid lines the budget constraints with the program. In the latter case, saving is subsidized up to the IRA contribution limit, say $2,000. The more steeply sloped segment represents the availability of tax-advantaged saving up to the limit:

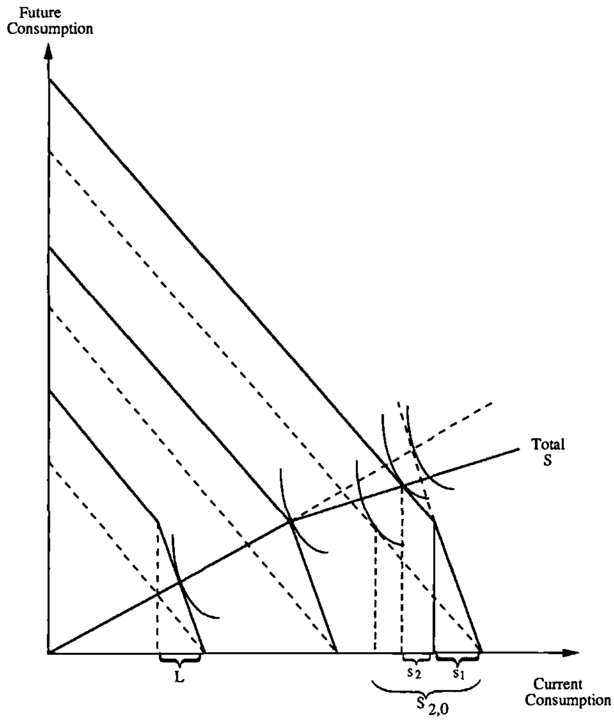

Fig. 4.1 IRAs and saving: A simple view

each dollar of consumption forgone yields more than $1 of IRA saving. The line labeled "Total S" shows the relationship between income and saving. A family at the highest income level would, in the absence of the IRA program, save more than the IRA limit ($S_{2,0}$ measured from the intersection of the budget constraint with the horizontal axis). As the graph is drawn, the IRA program reduces saving out of current income, although retirement consumption is also increased. This is the income effect of the program. Without the program, non-IRA saving would have been $S_{2,0}$. With the program, IRA saving is S_1 and non-IRA saving $(S_1 + S_2) - S_1$. The addition of the IRA saving is more than offset by the reduction in non-IRA saving.

There are two potential flaws in this stylized reasoning. The first is the assumption that saving for retirement is equivalent to any other form of saving; that they are equivalent goods and treated as such by consumers. As emphasized above, they may not be. Indeed, the fact that IRAs are much less liquid than other forms of saving suggests in itself that they will not be treated as equivalent. Second, this simple view ignores the potential effect of the enormous promotion of IRAs discussed above.

In addition, other evidence suggests that personal saving behavior cannot

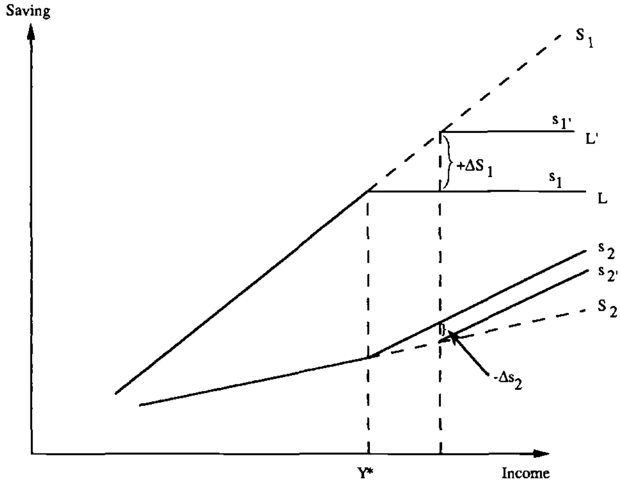

Fig. 4.2 IRAs and saving: A more general view

be explained by price effects, through the interest rate or tax laws. In general, the empirical evidence that saving behavior is noticeably affected by changes in the interest rate, at least over the range observed in the United States, is weak. In principal, whatever the effect of changes in the interest rate, the effect should also be reflected in the relationship between saving and the marginal tax rate, where interest payments are tax deductible. This reasoning would apply in particular to IRAs. The U.S. data, however, reveal mixed evidence on the effect of *existing differences* in marginal tax rates, after controlling for income.[15] Although direct evidence for IRAs is weak, the Canadian experience provides much stronger evidence. Analysis by Wise (1984) shows a very strong effect of income but the most appealing functional form specification shows no marginal tax rate effect, although functional forms that do not fit the data give the impression of a substantial effect.[16] Thus exclusive emphasis on price effects, through the marginal tax rate, may in general be misplaced.

Our analysis relaxes the assumptions reflected in figure 4.1. The two forms of saving are allowed to be treated as two goods. The IRA program may present a bargain on a distinct good, saving for retirement, not just a subsidized price on the one and only form of saving. But the general specification used in the analysis allows the data to reveal that they are treated as a single good, if that possibility is more consistent with observed behavior. This approach is summarized in figure 4.2. Here, IRA and non-IRA saving are treated as separate goods, S_1 and S_2, respectively. The heavy solid lines represent the saving that in figure 4.1 is represented by the single line "S." If the IRA limit were increased from L to L', persons with incomes below Y^* would be unaffected, since they are not constrained by the lower limit. If the increase were small, those with incomes above Y^* would increase IRA saving by ΔS_1 and would reduce non-IRA saving by ΔS_2. Our analysis is structured to determine to what extent the latter reduction offsets the increase in the former. The analysis takes account of the IRA limit and makes important use of the non-IRA saving of persons who are, as compared to those who are not, constrained by the IRA limits (either 0 or L). Our prior estimates strongly reject the figure 4.1 view.

4.3 Formal Estimates Based on the SIPP Data

Using the SIPP data we have obtained estimates based on the same model specification that we used in our prior analysis of SCF and CES data. The specification is summarized here, with further details in the appendix.

4.3.1 The Model

We concentrate on the potential substitution between IRAs and other liquid financial asset saving, assuming that, in the short run at least, IRAs are unlikely to be substituted for nonliquid wealth like housing. There are three key features of the model. First, the analysis uses individual attributes like age, income, and past saving behavior—as measured by accumulated assets—to control for individual-specific saving effects. Second, controlling for these attributes, the function S_1 and S_2 are estimated. Third, having determined S_1 and S_2, the results are summarized by the estimated change in the two forms of saving—ΔS_1 and ΔS_2—when the limit is increased. More formally, the budget constraint is given by

$$(1) \qquad C = Y - T - P_1 S_1 - P_2 S_2 = Y - T - (1-t)S_1 - S_2,$$ [17]

where T represents taxes before saving, $P_1 = 1 - t$ is the price of IRA saving in terms of current consumption, and $P_2 = 1$ is the price of other saving in terms of current consumption, where t is the marginal tax rate. At times $Y - T$ is denoted by Y_T. Desired but not observed S_1 and desired as well as observed S_2 are allowed to be negative. In addition, the potential substitution between S_1 and S_2 is allowed to be quite flexible and distinct from the substitution between either form of saving and current consumption. Given current income, a decision function with these characteristics is

$$(2) \qquad V = [C]^{1-\beta}\{[\alpha(S_1 - a_1)^k + (1-\alpha)(S_2 - a_2)^k]^{1/k}\}^{\beta}.$$

This function has a tree structure with one branch current expenditure and the other saving. These two components are evaluated in a Cobb-Douglas manner with the preference parameter β. The two forms of saving are evaluated according to a constant elasticity of substitution subfunction.[18] The parameter α indicates the relative preference for S_1 versus S_2; if $\alpha = .5$, total saving is split equally between the two forms. The important feature of this functional form is that it allows greater substitution between the two forms of saving than between either of these and current consumption. The elasticity of substitution between S_1 and S_2 is $1/(1 - k)$.

It also allows the IRA advantage to be reflected first in a lower cost of saving in terms of current income, through the current budget constraint, and, in addition, through different preferences for the two assets, possibly reflecting the different rates of return. Although the distinction between current cost and return may be an artificial one in strict economic terms—that the ultimate difference is one of yield only—consumers may understand better, and be influenced to a greater extent, by the current tax saving than by the tax-free compounding of interest. Certainly the promotion of IRAs has tended to highlight the former. In practice, it is not possible to distinguish the quantitative effect of one from that of the other. Indeed, in practice it is not possible to distinguish with any precision the effect of the tax rate from the effect of other variables, income in particular. Nonetheless, both features of IRAs, as well as any effects of advertising or the contract-like nature of IRA saving provisions, are allowed to determine individual choices.

Maximization of (2) subject to the budget constraint yields unconstrained desired levels of S_1 and S_2

$$
\begin{aligned}
S_1 &= a_1 + d_1(Y_T - P_1 a_1 - P_2 a_2) \\
S_2 &= a_2 + d_2(Y_T - P_1 a_1 - P_2 a_2) \\
d_1 &= \frac{(P_1/\alpha)^{1/(k-1)}}{P_1(P_1/\alpha)^{1/(k-1)} + P_2[P_2/(1-\alpha)]^{1/(k-1)}}\beta \\
d_2 &= (\beta - d_1 P_1)/P_2
\end{aligned}
\tag{3}
$$

Two limiting versions of the specification are of special interest.

If k = *0*

The limiting case of (2) as k goes to 0—yielding a Cobb-Douglas, or more precisely, a Stone-Geary specification—is a simpler model than the general one and is much easier to estimate. In fact, the estimated value of k is less than zero—indicating less substitution than a Cobb-Douglas specification would imply—and for simplicity many of the results are described assuming that it is zero. This specification is easily compared with the illustration in the previous section, graphed in figure 4.2. This case yields desired levels of S_1 and S_2, given by:

$$
\begin{aligned}
S_1 &= a_1 + \frac{\alpha\beta}{P_1} \cdot [Y_T - P_1 a_1 - P_2 a_2] \\
S_2 &= a_2 + \frac{(1-\alpha)\beta}{P_2} \cdot [Y_T - P_1 a_1 - P_2 a_2],
\end{aligned}
\tag{4}
$$

and observed levels by:[19]

(5)
$$s_1 = \begin{cases} 0 & \text{if } S_1 < 0, \\ a_1 + \dfrac{\alpha\beta}{P_1} \cdot [Y_T - P_1a_1 - P_2a_2] & \text{if } 0 < S_1 < L, \\ \text{L} & \text{if } L < S_1; \end{cases}$$

$$s_2 = \begin{cases} a_2 + \dfrac{(1-\alpha)\beta}{(1-\alpha\beta)P_2} \cdot [Y_T - P_2a_2] & \text{if } S_1 < 0, \\ a_2 + \dfrac{(1-\alpha)\beta}{P_2} \cdot [Y_T - P_1a_1 - P_2a_2] & \text{if } 0 < S_1 < L, \\ a_2 + \dfrac{(1-\alpha)\beta}{(1-\alpha\beta)P_2} \cdot [Y_T - P_1L - P_2a_2] & \text{if } L < S_1. \end{cases}$$

Here, abstracting from the prices, β is the total marginal saving rate and a is the proportion allocated to IRA saving. The lower-case s's represent actual saving and the upper-case S's, desired saving. The parameter β is the proportion of marginal income that is saved; α is the proportion of saving allocated to IRAs. The term $[(1 - \alpha)\beta]/[(1 - \alpha\beta)P_2]$ represents the marginal saving rate in the non-IRA form once the IRA limit L has been reached.

If the limit L is increased by one unit, the IRA saving of persons at the limit will be increased by $\Delta S_1 = 1$. Other saving will be reduced by $\Delta S_2 = -[P_1/P_2][(1 - \alpha)\beta]/(1 - \alpha\beta)$. If $\alpha = 1$, $\Delta S_2 = 0$. If $\alpha = 0$, $\Delta S_2 = -[P_1/P_2]\beta$.

If k $=$ *1 and* α $=$ *.5*

Under this assumption, the elasticity of substitution between S_1 and S_2 is infinite, and they are given equal weight in the preference function; they are perfect substitutes and are treated as a single asset. Because the price of IRA saving is lower, saving is only through S_1 if $S_1 < L$ and thereafter is through S_2. In this case, the IRA tax advantage simply creates a kink in the intertemporal budget constraint describing the relationship between forgone current consumption and future consumption, and inframarginal arguments could be used to represent the incentive effects of IRAs on persons who would in their absence save more than the IRA limit. This possibility is clearly rejected by the data, however. It is clear from the summary data that this extreme case is inconsistent with actual behavior; a large fraction of persons who have no IRA saving do have some non-IRA (S_2) saving. Saving behavior under this assumption is described in detail in Venti and Wise (1987b) and the relevant sections from that paper are reproduced as an appendix to this paper.

Other values of k

Unlike the $k = 0$ or $k = 1$ cases, there is no closed-form solution to the constrained S_2 function for other values of k. In this case, the constrained functions, $S_2^*(0)$ when $S_1 < 0$ and $S_2^*(L)$ when $S_2 > L$, are defined only implicitly, as described in the appendix

4.3.2 Parameterization of α and β and the Stochastic Specification

To capture the wide variation in saving behavior among individuals, α and β are allowed to depend on individual attributes X. In particular, we attempt to control for individual-specific saving behavior by using past saving behavior, as well as other attributes, to predict β. Both parameters are also restricted to be between 0 and 1 by using the form

$$\beta = F[X\underline{\beta}] \qquad (6)$$
$$\alpha = F[X\underline{\alpha}],$$

where $F(\cdot)$ is the standard normal distribution function and $X\underline{\alpha}$ and $X\underline{\beta}$ are vectors of parameters.

Finally, we allow the S_1 and S_2 functions to be shifted by additive disturbances, ε_1 and ε_2, respectively.[20] The disturbances are assumed to be distributed bivariate normal with standard deviations σ_1 and σ_2, respectively, and correlation r.

There are three possibilities for the observed values of S_1: 0, between 0 and L, and L. A continuously measured value of S_2 is available for each person, yielding three possible joint outcomes for each observation. Estimation, based on these probabilities, is by maximum likelihood.

4.3.3 Results

Parameter Estimates

Estimation with k free to vary yields an estimated k of -1.67 with a standard error of 0.40, as shown in table 4.8b. Thus, although the data do not allow precise estimation of k, large values are clearly rejected.[21] In particular, the data are inconsistent with the limiting case of $k = 1$, which would indicate that the two forms of saving are perfect substitutes.[22] Thus to facilitate calculation, we concentrate on the simpler model, with k set to zero.

Parameter estimates with k set to zero are shown in table 4.8a. Several features of the results stand out: first, there is no relationship between the two forms of saving once family attributes, including past saving behavior, are controlled for. The correlation r between the disturbance terms in the two equations is essentially zero (.02). In particular, the data do not show that families who save more than the typical family in one form save less in the other.

Second, there is a wide range in saving behavior among families. This is summarized by the estimated values of β that range from .022 to .677. Recall that β is the total *desired* marginal saving rate. (Because the constant term a_1 is negative, however, estimated desired saving is negative for a large fraction of families.) Recall that to control for individual-specific saving behavior we have *predicted* β on the basis of individual attributes, including past saving behavior as measured by liquid and nonliquid assets. It is clear that these data

Table 4.8a Parameter Estimates with $k = 0$

Variable	Estimate (Asymptotic Standard Error)				
Covariance terms:					
σ_1			4.68	(.13)	
σ_2			7.11	(.03)	
r			.02	(.01)	
Origin Parameters:					
a_1			−13.00	(.79)	
a_2			.02	(.07)	
Determinants of β and α:	β		α		
Income	−.0125	(.0006)	−.0136	(.0019)	
Age	.0098	(.0007)	−.0456	(.0043)	
Liquid Assets	.0041	(.0002)	−.0048	(.0004)	
Nonliquid Assets	.0013	(.0001)	−.0033	(.0003)	
Pension	.0344	(.0152)	.2156	(.0670)	
Education	.0087	(.0011)	−.0202	(.0056)	
Children	−.0442	(.0075)	−.0219	(.0440)	
Unmarried	−.0161	(.0171)	−.0322	(.0720)	
Constant	−1.1441	(.0088)	4.7302	(.3128)	
Predicted oversample:	Mean	Median	SD	Minimum	Maximum
Parameter:					
β	.210	.198	.060	.022	.677
α	.965	.991	.068	.067	1.000
d_1	.256	.246	.055	.038	.600
d_2	.010	.002	.026	.000	.612
For families predicted to be at the IRA limit:					
Parameter:	Mean	Median			
d_1	.296	.293			
d_2	.028	.013			
d_2^*	.037	.017			

Log likelihood = −40,685.8
Number of observations = 9,524

do that to a substantial extent. Thus while the simple regressions above show a strong relationship between saving in one form and saving in the other, once the individual attributes that explain this relationship are controlled for there is no relationship between the amount of *new* IRA saving and *new* financial saving in other assets.

Third, like the total saving rate, the estimated desired IRA marginal saving rates, $d_1 = \alpha\beta$, also vary widely. The median is .246, the minimum is .038, and the maximum .600.

Table 4.8b Parameters with k Estimated

Variable	Estimate (Asymptotic Standard Error)	
Covariance terms:		
σ_1	4.69	(.10)
σ_2	7.14	(.03)
r	.04	(.01)
Elasticity Parameter, k	-1.67	(.40)
Origin Parameters:		
a_1	−13.20	(.12)
a_2	−.01	(.08)

Determinants of β and α:	β		α	
Income	−.0125	(.0002)	−.0239	(.0036)
Age	.0095	(.0006)	−.0748	(.0102)
Liquid Assets	.0045	(.0002)	−.0101	(.0014)
Nonliquid Assets	.0014	(.0001)	−.0069	(.0009)
Pension	.0310	(.0147)	.4522	(.1195)
Education	.0093	(.0010)	−.0415	(.0099)
Children	−.0468	(.0072)	.0182	(.0734)
Unmarried	−.0128	(.0158)	−.0647	(.1002)
Constant	−1.1364	(.0344)	8.9030	(1.0039)

Predicted oversample:	Mean	Median	SD	Minimum	Maximum
Parameter:					
β	.214	.202	.061	.023	.705
α	.996	.999	.045	.005	1.000
d_1	.259	.249	.057	.038	.583
d_2	.011	.003	.027	.001	.635

For families predicted to be at the IRA limit:

Parameter:	Mean	Median
d_1	.301	.299
d_2	.033	.012
d_2^*	.022	.007

Log likelihood = −40,672.1

Number of observations = 9,524

Fourth, the desired marginal saving rate in other financial assets is typically very small, consistent with the low saving rates revealed by the summary data. Thus, according to these results, if it were not for IRAs, financial asset saving would be much smaller than it is.[23] Even among families predicted to be at the IRA limit, predicted marginal saving in other financial assets is very small on average, .028, and it does not change much when the possibility for IRA saving is exhausted. The estimated rate after the IRA limit is reached, d_2^*, is

.037. The small difference between these latter two estimates is important because it reveals the extent to which increased IRA saving—due to an increase in the IRA contribution limit, for example—would be offset by a reduction in other saving. (Fig. 4.2 makes this clear.)

Independent Estimates of Non-IRA Saving

Because the relationship between d_2 and d_2^* is fundamental to the results, it is informative to demonstrate that the result is not simply due to the functional form used in the analysis. An unconstrained version of the S_2 function, motivated by the piecewise linear illustration in figure 4.2, can be estimated by ordinary least squares. Let Y_2^* be the income at which a family with attributes X_i reaches the IRA saving limit. It is determined from the S_1 function estimates presented in table 4.8a and includes a randomly selected disturbance term for each family.[24] Define $Y_1 = Y$ if $Y < Y^*$ and Y^* otherwise, and $Y_2 = 0$ if $Y < Y^*$ and $Y - Y^*$ if $Y > Y^*$. Then $S_2 = c + \delta_2 Y_1 + \delta_2^* Y_2 + \nu$, where ν is a disturbance term and the δ's are both linear functions of the same variables listed in table 4.8a. Thus δ_2 and δ_2^* correspond roughly to d_2 and d_2^*, respectively. The mean of the predicted values of δ_2 for families with $Y < Y^*$ is .060. For families with $Y > Y^*$, the mean of δ_2 is .066 and the mean of δ_2^* is .090. (The estimated value of c is $-.34$.) It is clear from this unconstrained approximation to the model specification that there is only a limited increase in S_2 saving after the S_1 limit has been reached, as the model estimates show.[25] Again, as a rough approximation, using the mean d_1 for families at the IRA limit (.296 from table 4.8a) and the δ_2 and δ_2^* estimates for families at the limit, an increase of .296 in IRA saving is associated with a .024 reduction (.090 − .066) in S_2 saving, about 8 percent of the IRA increase. If the elasticity parameter k were larger, the difference between d_2 and d_2^* (or between δ_2 and δ_2^*) would be larger and the substitution of IRA for non-IRA saving would be greater.

Simulations

To summarize the implications of the model estimates, we have simulated the effect of raising each families's IRA limit by $1,000. Only families predicted to be at the IRA limit are affected by the increase. For those at the limit, the simulated mean changes in consumption, taxes, and other saving associated with the IRA increase are as follows:

Change in IRA saving	+$856	100.0%
Change in other saving	−$22	−2.6%
Change in consumption	−$565	−66.0%
Change in taxes	−$269	−31.4%

Most of the new IRA saving resulting from an increase in the limit would represent a net increase in total saving; there would be little substitution away from other saving.[26] The average IRA saving of families at the limit before the increase is \$3,174; saving in other financial assets is only \$1,497. After the increase, IRA saving is \$4,030 and non-IRA saving \$1,475.

4.4 Conclusions

The SIPP data confirm that, with the exception of housing, the typical American family saves very little. In particular, financial asset saving of a very large proportion of families is close to zero. These data also indicate that most IRA saving is net new saving; it is not funded by substitution away from other saving or by increased debt. Thus if it were not for IRAs, personal saving would be even lower than it is. If the IRA limit were increased, most of the increase in contributions would be new saving. The model prediction of little substitution is consistent with the descriptive data that show very little non-IRA financial asset saving; there is little to substitute away from. These results are very similar to our findings based on the SCF and the CES. They are also consistent with the recent conclusions of Feenberg and Skinner (1989), based on panel tax data.

If the relevant data are released, the panel nature of the SIPP will allow control for individual-specific saving effects that is potentially better than the correction based on past saving behavior, the procedure followed in this paper. Judging from the work of Feenberg and Skinner (1989), however, it seems unlikely that the conclusions of this paper will be altered substantially. We can think of no reason why extensive substitution would not be revealed by the data.

Appendix
Special Cases of the Estimated Model

In addition to the limiting version of the model detailed in the text, two others are of interest. They are described under the second and third headings below.

1. If $k = 0$

This is the limiting case detailed in the text.

2. If $k = 1$ and $\alpha = .5$

Under this assumption, the elasticity of substitution between S_1 and S_2 is infinite, and they are given equal weight in the preference function; they are perfect substitutes and are treated as a single asset. The decision function (2) becomes

$$V = [C]^{1-\beta}[S_1 + S_2 - (a_1 + a_2)]^{\beta}. \quad \text{(A1)}$$

Because the price of IRA saving is lower, saving is only through S_1 if $S_1 < L$ and thereafter is through S_2, with

$$\text{(A2)}\quad \begin{aligned} s_1 &= \begin{cases} 0 & \text{if } S_1 < 0, \\ (a_1 + a_2) + \frac{\beta}{P}[Y_T - P_1(a_1 + a_2)] & \text{if } 0 < S_1 < L, \\ L & \text{if } L < S_1; \end{cases} \\ s_2 &= \begin{cases} 0 & \text{if } S_1 < L, \\ (a_1 + a_2 - L) + \frac{\beta}{P_2}[Y_T - P_1 L - P_2(a_1 + a_2 - L)] & \text{if } L < S_1. \end{cases} \end{aligned}$$

In this case, the IRA tax advantage simply creates a kink in the intertemporal budget constraint describing the relationship between forgone current consumption and future consumption, and inframarginal arguments could be used to represent the incentive effects of IRAs on persons who would in their absence save more than the IRA limit.

3. Other values of *k*

Unlike the $k = 0$ or $k = 1$ cases, there is no closed form solution to the constrained S_2 function for other values of k. In this case, the constrained functions, $S_2^*(0)$ when $S_1 < 0$ and $S_2^*(L)$ when $S_2 > L$, are defined only implicitly by the relationship

$$\text{(A3)}\quad \frac{P_2(1-\beta)[\alpha(m-a_1)^k + (1-\alpha)(S_2^* - a_2)^k]}{(1-\alpha)(S_2^* - a_2)^{k-1}} = Y_T - P_1 m - P_2 S_2^*,$$

where m *is* either 0 or L. It is derived by maximizing (2) subject to the budget constraint and with the additional constraint that $S_1 = m$. The observed levels of saving are

$$\text{(A4)}\quad \begin{aligned} s_1 &= \begin{cases} 0 & \text{if } S_1 < 0, \\ a_1 + d_1 (Y_T - P_1 a_1 - P_2 a_2) & \text{if } 0 < S_1 < L, \\ L & \text{if } L < S_1; \end{cases} \\ s_2 &= \begin{cases} S_2^*(0) & \text{if } S_1 < 0, \\ a_2 + d_2(Y_T - P_1 a_1 - P_2 a_2) & \text{if } 0 < S_1 < 0, \\ S_2^*(L) & \text{if } L < S_1. \end{cases} \end{aligned}$$

Notes

1. Analysis based on two consecutive years will be undertaken when the data are released.
2. More precisely, families with heads are who self-employed or over 65 or under 21 have been excluded. Household data are also considered. In that case the household head is used to determine whether the household is included.
3. The asset categories are defined as follows: *Housing equity:* Current market

value of home (including mobile homes) less the principal owed on remaining mortgage. *Financial assets excluding stocks and bonds:* Regular (passbook) saving accounts, money market deposit accounts, certificates of deposit or other saving certificates, NOW or other interest bearing saving accounts, money market funds, U.S. government securities, municipal or corporate bonds, other interest earning assets, noninterest bearing checking accounts. *Financial assets including stocks and bonds:* The above category plus the market value of stocks and mutual funds (less debt or margin account) and the face value of U.S. savings bonds. *Debt:* Store bills; credit card bills; bills from doctors, dentists, hospitals, or nursing homes that are not covered by insurance; money owed to individuals outside the family; loans owed to banks, credit unions, or other financial establishments (excluding loans to secure homes, vehicles, or stock and mutual fund shares); other money owed. *Nonhousing assets:* Financial assets including stocks and bonds (see above) plus motor vehicle equity, business equity, net equity in other property (vacation, commercial, or rental), money owed (including mortgages held), and equity in other financial investments, and less debt as described above. IRAs and Keoghs are not included unless otherwise noted. *Total wealth:* Housing equity plus Nonhousing assets as described above. IRAs and Keoghs are not included unless otherwise noted.

4. Based on the 1983 SCF, the median was $22,900.

5. The skewness of the distribution of wealth is reflected in the difference between the medians and the means. The total wealth mean is $48,241, housing equity is $29,398, financial assets including stocks and bonds $13,178, financial assets excluding stocks and bonds $8,395, and debt $3,035.

6. The family unit used is the IRS definition of the tax unit. Thus adult members of the household who are neither the household head nor spouse are classified as separate families. By this definition, there are approximately 40 percent more families than households in the SIPP.

7. Most of the medians are zero and are thus not shown.

8. See Venti and Wise (1986, 1987b).

9. The noncontributor category includes some cases where the difference in reported IRA balances between the two years is negative.

10. Much of the following discussion is drawn from our previous papers. See Venti and Wise (1987b) and Wise (1988).

11. *IRA Reporter,* vol. 6, no. 9 (September 30) 1988.

12. Feenberg and Skinner (1989, 12).

13. One might, for example, have a scheme in which the limit for the current year is added to next year's limit if a contribution is not made in the current year. Or, the contribution limit could cumulate more generally over time if contributions are not made during some period.

14. Especially Venti and Wise (1987b).

15. This may reflect in part an empirical identification problem. Income and marginal tax rates are closely related—although the correlation is by no means perfect—and most data do not provide accurate tax rates. Estimates are very sensitive to functional form. Venti and Wise (1988a) find little effect of the marginal tax rate. However, Feenberg and Skinner (1989) find a significant positive effect.

16. The analysis in Wise (1985) is based on tax records and thus very accurate marginal tax rates, which vary substantially given income. While there is some evidence that the marginal tax rate may affect whether a person contributes to an RRSP, there seems to be no effect on the amount of the contribution.

17. In principle, the marginal tax rate is determined in part by IRA contributions. But since the IRA limits narrowly restrict this influence, we treat t as exogenous.

18. This specification turns out to be a variant of the "S-branch" utility tree described by Brown and Heien (1972). See also Sato (1967) and Blackorby, Boyce, and Russell (1978).

19. Although it is illegal to borrow against an IRA, funds can be withdrawn subject to the 10 percent penalty. But since negative contributions are not observed in the data set, we adopt the assumption of a zero lower limit.

20. A random preference stochastic specification that makes each individual's choices formally consistent with the decision function (2) is obtained if a_1 and a_2 are assumed to be random, with additive disturbances. This specification is not tractable, however, when S_2^* must be solved for implicitly. Experience with both forms in Venti and Wise (1986, 1987a) shows that the results are not appreciably affected by this choice.

21. Because the likelihood function is rather "flat" with respect to k at its estimated value, it is informative to consider likelihood values at other selected values of k. For example, the value is $-40{,}707.1$ at $k \equiv .5$, $-40{,}685.8$ at $k \equiv 0$, $-40{,}677.5$ at $k \equiv -.5$, $-40{,}672.7$ at $k \equiv -1$, and $-40{,}672.1$ at $k \equiv -1.67$ (the maximum-likelihood estimate). Thus a likelihood ratio test rejects the hypothesis of $k = 0$ with a χ^2 statistic of 26. And of course larger values would also be rejected.

22. Two kinds of information in the data provide information on the degree of substitution between S_1 and S_2: one is the extent to which families who have no IRA saving, or who have IRA saving below the limit, save in other financial asset forms. Desired levels of saving S_1 and S_2 are observed as long as S_1 is less than the IRA limit. In addition, the degree of substitution between these two forms of saving is revealed in the data by comparing the share of the marginal dollar of income allocated to S_2 when the family is free to vary S_1 (that is, when S_1 is below the IRA limit) with the income share allocated to S_2 when desired IRA saving is constrained by the upper limit. That is, the extent of a "spillover" of desired IRA saving into non-IRA saving when the limit is reached also provides information on the degree of substitution between these two forms of saving.

23. In a previous paper based on CES data (Venti and Wise 1987b), we obtained very similar results on non-IRA saving. With those data we were able to test the model by using estimates based on the post-IRA period data (1982 and later) to predict saving before IRAs were introduced on a broad scale (1980–81). These estimates matched very closely the actual pre-IRA saving behavior.

24. It is the Y that solves the equation

$$s_1 = a_1 + [\alpha\beta/P_1]\,[Y^* - P_1 a_1 - P_2 a_2] + \varepsilon = L,$$

where ε is randomly drawn from the estimated distribution of ε_1, and α and β depend on family attributes X.

25. The results can not be expected to be the same as those from the model because the simple regression version does not account for the price of S_1 saving nor for the linear expenditure system parameters a_1 and a_2, as shown in eq. (5).

26. Even the very limited substitution suggested by these estimates is more than the data actually reveal. The data suggest a substitution parameter k that is in fact lower than the zero value used in making these calculations.

References

Blackorby, Charles, Richard Boyce, and R. Robert Russell. 1978. Estimation of Demand Systems Generated by the Gorman Polar Form: A Generalization of the S-Branch Utility Tree. *Econometrica* 46, no. 2 (March): 345–63.

Brown, Murray, and Dale Heien. 1972. The S-Branch Utility Tree: A Generalization of the Linear Expenditure System. *Econometrica* 40, no. 4 (July): 737–47.

Carroll, Chris, and Lawrence Summers. 1987. Why Have Private Savings Rates in the United States and Canada Diverged? *Journal of Monetary Economics* 20, no. 2 (September): 249–79.

Feenberg, Daniel, and Jonathan Skinner. 1989. Sources of IRA Saving. *Tax Policy and the Economy* 3:25–46.

Sato, Kazuo. 1967. A Two-Level-Constant-Elasticity-of-Substitution Production Function. *Review of Economic Studies* 34, no. 98 (April): 201–18.

Summers, Lawrence. 1986. Individual Retirement Accounts: Facts and Issues. *Tax Notes* 31, no. 10 (June 9): 1014–16.

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

———. 1987a. IRAs and Saving. In *The Effects of Taxation on Capital Accumulation*, ed. M. Feldstein, 7–48. Chicago: University of Chicago Press.

———. 1987b. Have IRAs Increased U.S. Saving? Evidence from Consumer Expenditure Surveys. *Quarterly Journal of Economics* 105 (August): 661–98.

———. 1988a. The Determinants of IRA Contributions and the Effect of Limit Changes. In *Pensions in the U.S. Economy*, ed. Z. Bodie, J. Shoven, and D. Wise, 9–47. Chicago: University of Chicago Press.

———. 1988b. The Evidence on IRAs. *Tax Notes* 38, no. 4 (January 25): 411–16.

Wise, David A. 1984. The Effects of Policy Change on RRSP Contributions. Prepared for the Tax Policy and Legislation Branch of the Canadian Department of Finance. Cambridge, Mass. Mimeograph.

———. 1985. Contributors and Contributions to Registered Retirement Saving Plans. Prepared for the Tax Policy and Legislation Branch of the Canadian Department of Finance. Cambridge, Mass. Mimeograph.

———. 1987. Individual Retirement Accounts and Saving. In *Taxes and Capital Formation*, ed. M. Feldstein, 3–15. Chicago: University of Chicago Press.

———. 1988. Saving for Retirement: The U.S. Case. *Journal of Japanese and International Economics* 2:385–416.

Comment Michael Rothschild

This is one of a series of papers in which Venti and Wise have investigated effects of tax-deferred retirement accounts on aggregate savings. In the United States, almost all workers could make tax deductible contributions to individual retirement accounts (IRAs) from 1982 to 1986. Venti and Wise argue that the great bulk of contributions to IRA accounts during this period were net additions to aggregate saving. This seems odd, as the first theory that any economist would propose is that both IRAs and ordinary savings are ways of purchasing future consumption; the most important difference between the two is that IRAs permit this purchase on more favorable terms. Since the two goods are almost perfect substitutes, IRA saving should come in the first instance at the expense of other saving.

Michael Rothschild is professor of economics and dean of Social Sciences at the University of California at San Diego; he is also a research associate of the National Bureau of Economic Research.

The simple theory has, at least, three difficulties. First it is not strictly true. IRAs are different goods from savings accounts. IRA contributions are not as liquid as ordinary investment because the IRS discourages (with penalties and, more recently, with absolute prohibitions) their premature conversion into consumption. While IRAs are a good way to finance retirement, they are less attractive as a method of saving for one's children's education. Perhaps as important, but certainly less comfortable for the economist, banks and other financial institutions heavily promoted IRAs. Barrages of advertising may have made IRAs different goods in the minds of consumers from ordinary savings—advertising may indeed have made them appear to be a different good from what they actually were.

A second difficulty of the simple theory is that if it is true, then individual saving should be responsive to changes in interest rates. While some believe this to be true it is not one of the better-established truths of economic science.

A final difficulty with the simple theory is that if it is true it must be phrased in such a way as to take account of the great differences in personal saving behavior that survey and other data reveal. Venti and Wise convincingly demonstrate that different people save vastly different amounts. While controlling for demographic characteristics reduces some of the variability, it still remains true that some people save a lot, some save a little, and most save barely at all. Any theory of saving that is to survive a confrontation with cross-section data will have to allow for (and hopefully explain) individual differences.

The authors present and estimate an ingenious and illuminating theory. Unfortunately, their theory does not permit a clean test of the simple proposition that the two kinds of savings are perfect substitutes. A reformulation of their theory would allow such a test.

Venti and Wise's theory is a variant of the linear expenditure system.[1] In the standard linear expenditure system each good has associated with it a required level of consumption α_i and a share τ_i where $\Sigma\tau_i = 1$. For each good i the consumer must buy an amount α_i of good i. If income is Y, then disposable income is $Y - \Sigma\alpha_i$; the share of disposable income that the consumer spends on good i is τ_i. Venti and Wise use a variant of this system to estimate saving behavior. In their setup there are three goods: C, current consumption, S_1, IRA saving, and S_2, other saving. It is assumed that the intercept for C is zero, while those for the two kinds of saving are to be estimated. In their empirical work the intercept for IRA saving is negative. This implies that at low levels of income people take money out of their IRA accounts. However, Venti and Wise do not allow this; instead they require IRA saving to fall in the interval $[0, L]$ where L is the maximum contribution allowed by law. Heterogeneity is accounted for by allowing the share parameters, the τ's, to be determined by demographic characteristics, wealth, and income.

This functional form does not allow a test of the hypothesis that IRA saving

1. If the parameter k in eq. (2) is not equal to zero, then the form they estimate is more complicated; in their paper, most analysis focuses on the simpler case, when k equals zero.

and other saving are perfect substitutes. Ignore the nonlinearity introduced by letting the share coefficients, the τ's, depend on income. The Venti-Wise theory states that saving is a piecewise linear function of income. IRA saving is positive with a slope of τ_1 over an interval $I = [-\alpha_1, -\alpha_1 + (L/\tau_1)]$.[2] When income lies outside I there is no IRA saving. The slope of non-IRA saving depends on whether income is in I or outside of it. When income is in I, the share of the marginal dollar that goes to other savings is $\delta = \tau_2$, when income is outside I the share is $\pi = \tau_2/(1-\tau_1)$. The theory that the two kinds of saving are perfect substitutes would have $\tau = 0$. However, it is clear that if $\pi = 0$, then $\delta = 0$. Thus, Venti and Wise do not really test the simple-minded perfect-substitutes theory.

I believe that this could be remedied by estimating the marginal share coefficients in the regression more freely and not constraining π to be a scaler multiple of δ. In a similar vein, the other assumptions of the theory could have been tested somewhat more. As we saw above, IRA saving is constrained to be zero when income is below $-\alpha_1$, the lower limit of the interval I, and when income is above $-\alpha_1 + (L/\tau_1)$, the upper limit of I. The Venti-Wise theory implies that the share of the marginal dollar spent on non-IRA saving is the same in both these regions. This is an easily tested hypothesis.

Finally, I wonder whether it would have been possible to allow for individual effects in the intercept terms (the α_1s) as well as in the slope terms.

Despite these quibbles, I want to stress that Venti and Wise have written a most interesting and ingenious paper. Simple aggregate models—which are the ones easiest to use to address such topics as savings—blind us to the great variety of economic situations in which people find themselves and the great variety of behavior that they exhibit. Venti and Wise have shown how this heterogeneity can be comprehended in a model that also has implications for aggregate policy.

2. Recall that α_1, required IRA savings, is negative, while IRA savings is restricted to be nonnegative. The price of IRA savings is unity.