The Economic Journal, **112** (*July*), 504–538. © Royal Economic Society 2002. Published by Blackwell Publishers, 108 Cowley Road, Oxford OX4 1JF, UK and 350 Main Street, Malden, MA 02148, USA.

THE EFFECT OF INDIVIDUAL RETIREMENT ACCOUNTS ON HOUSEHOLD CONSUMPTION AND NATIONAL SAVING*

Orazio P. Attanasio and Thomas DeLeire

A major debate exists on whether expanding tax-favoured savings accounts such as Individual Retirement Accounts (IRAs) will increase national savings. Much of the empirical debate has centred on whether IRA contributions before the Tax Reform Act of 1986 represented new savings or merely reshuffled assets. We find no evidence that households financed their IRA contributions from reductions in consumption, at least initially. We find evidence that households financed their IRA contributions from existing savings or from saving that would have been done anyway. Our results indicate that, at most, 9% of IRA contributions represented net additions to national saving.

The US national savings rate has declined dramatically over the past 10–15 years. This decline might have major implications for the domestic capital accumulation rate, productivity growth and the ability of households to finance their own retirement and has been the subject of much research. The primary reason for falling national savings is the decline in the US personal savings rate. While not as dramatic as in the U.S., personal and national saving has also declined in many other developed countries. In many countries, the policy debate centres on the necessity of increasing the level of personal saving, which is perceived to be 'inadequate'. However, as there is little consensus among economists as to what has caused the decline in personal savings rates, there is also little consensus as to the most effective way of reversing it.

One proposal that has received much attention from both economists and policy makers is the tax-favoured savings account. Tax-favoured savings accounts such as Individual Retirement Accounts (IRAs) or 401(k) plans may provide households an incentive to increase their saving since households can defer paying taxes on both contributions to and interest earned from these accounts. Other countries, such as the U.K., have also been experimenting with various forms of tax incentives to saving.

IRAs were first introduced in 1974 in the U.S., but eligibility for them was initially restricted to individuals without access to a pension. In 1981, eligibility greatly expanded to include nearly all individuals. Subsequently, total IRA contributions surged from about $7 billion in 1981 to over $40 billion in 1982 (in 1993 dollars). Five years later, eligibility for IRAs was again restricted by the Tax Reform Act of 1986. This restriction, combined with the increased popularity of 401(k) plans, led

* We are grateful to Eric Engen, Bill Gale, Erzo Luttmer, Will Manning, Jim Poterba, Jonathan Skinner, Steve Venti, David Wise and seminar participants at the University of Chicago and the Federal Reserve Board of Governors for useful discussions. We also received valuable comments from the Editor and two anonymous referees. Special thanks to Dean Maki who generously provided us with his computer programs.

to a sharp fall in IRA contributions from almost $50 billion in 1986 to about $12 billion in 1987.

The research on the effects of the legislation that gave tax-favoured status to IRA contributions has focused on whether these fiscal incentives have increased the level of personal and national saving. The issue is an important one: policies that create tax-favoured savings accounts are expensive in terms of lost tax revenue and, consequently, in terms of government saving. These policies must, therefore, raise the level of personal saving by more than the lost tax revenue to raise national saving.

Households can finance IRA contributions from three sources: transfers from *previously existing* non-IRA financial asset balances, transfers from non-IRA saving that would have been done even in the absence of IRAs, and from decreases in consumption.[1] The two transfers represent reshuffling, while the stated aim of the IRA legislation was to stimulate new saving. The important policy question addressed by the literature and by this paper is whether tax favoured savings accounts have an effect on personal and, more generally, national savings rates. In other words, one would like to know whether the large increase in IRA contributions observed between 1982 and 1986 represented primarily new saving or reshuffled savings.

There is little consensus on the theoretical model that best describes saving behaviour. Moreover, even simple models of consumption behaviour yield ambiguous predictions about the effect of changes in the return to saving, as income effects may counteract substitution effects. In more complex models that take into account the differentials in risk, liquidity and availability of different assets, it is extremely hard to discern the effects of tax-favoured accounts on consumption and saving behaviour. In general, in most models, the answer is likely to depend on the magnitude of several key preference parameters. Because of this, the issue of the effectiveness of these tax incentive-based plans in raising national saving is mainly an empirical one.

Unfortunately, it is very difficult to assess the effectiveness of these savings plans empirically. From an aggregate point of view, we know that personal saving rates have continued to decline even after Congress implemented fiscal incentives for saving. Of course, it is conceivable that, without these incentives, saving rates would have declined even more. From a microeconomic point of view, we know that households participating in these plans save more than households not participating in them. It is extremely difficult, however, to establish how much the same households would have saved without the tax incentive. A legitimate hypothesis is that participation in the IRA programme is simply an indicator of a high propensity to save and that participants would have saved the same amount without the tax incentives. If so, these tax incentives merely represent a transfer from taxpayers to savers.

[1] As we discuss below, IRA contributions can also be financed by an increased level of saving with constant consumption, if the contributing households save the decreased tax liability. Obviously, the implications for national saving of a reduction in consumption and of an increase in saving originating from a reduction in tax liabilities are very different.

The main problem that has plagued the literature on the effectiveness of the IRA legislation and of other tax incentives for saving is the difficulty in controlling for unobserved heterogeneity in the taste for saving. As the available data are not experimental, one is forced to estimate the effect of the legislation from differences between the behaviour of different individuals. The assumption that is always necessary to make is that observed differences in behaviour are not due to unobserved factors but to the effects of the legislation under study.

In this paper, we use a new empirical test to determine whether IRA contributions represent reshuffling of existing saving[2] or new saving. Our approach is novel in several respects. First, while many of the papers in the literature compare the behaviour of contributors to that of non-contributors, here we compare the behaviour of two groups of IRA contributors: new contributors – households that just opened an IRA account – and continuing contributors – households that previously had made IRA contributions. The difference in the behaviour of these two groups, new contributors and continuing contributors, is what identifies the effect of the tax incentives for retirement saving. The maintained identifying assumption is that the observed differences between the two groups are not due to some unobserved taste characteristics. This assumption will be true, for example, if the only difference between households is their timing for beginning contributing to IRAs. Given that both groups contribute to IRAs, both show a taste for saving. Therefore, the problem of unobserved heterogeneity that has confounded comparisons between the changes in non-IRA financial asset balances of contributors and non-contributors is less likely to arise.[3]

Second, unlike most tests in the literature that use only information on changes in financial assets, our main test uses information on household consumption. The idea is that, if the tax incentives led households to increase their saving, we expect these households to have reduced their consumption to finance contributions when they opened an IRA account.

Comparing the changes in consumption of new and continuing contributors appeals to the fact that, if the savings plan is effective, it should work through a substitution effect. When a household starts participating in such a plan, it should shift to a new (lower) consumption *level* as it faces a different intertemporal price of consumption. Income effects that may be at work even before the household starts participating in the plan are likely to work in the opposite direction. Therefore, the change in consumption that occurs when a household starts participating in a savings plan can provide very valuable information on the effectiveness of the plan in increasing that household's saving. Even though we do not estimate an explicit structural model, our consumption-based test is based on a

[2] That is, existing assets or saving that would have been done regardless of the IRA incentives.

[3] Venti and Wise (1990), for example, use the Consumer Expenditure Survey (CEX) to compare the non-IRA saving of IRA contributors with that of non-contributors. It is likely that IRA contributors are different in their tastes for saving from non-contributors. Since continuing contributors were once new contributors, assuming their tastes are identical is reasonable. Life-cycle effects or differences in needs plausibly could explain why some individuals start contributing before others. Gale and Scholz (1994) compare the non-IRA saving of two types of IRA contributors and provide evidence that IRA contributors do differ in their tastes for saving from non-contributors.

structural model of consumption behaviour and could not be interpreted outside such a structure.

In addition to our main test based on the changes in consumption of new and continuing contributors, we also analyse the changes in financial assets other than IRAs for the same two groups. The idea is that, under the hypothesis that IRA contributions do not represent new saving, new contributors should have a smaller increase in non-IRA financial assets than continuing contributors.

This second test is subject to a greater number of caveats than our first test, which we discuss below. However, it is an interesting test because it sets the 'new saving' hypothesis as the null, while the consumption-based test sets the 'reshuffling' hypothesis as the null. If the data set we use, the Consumer Expenditure Survey (CEX), that has been under-utilised in the study of saving incentives, is of poor quality and dominated by measurement error, we should fail to reject both null hypotheses. It is therefore profitable to check whether the results obtained in the two tests are consistent.

The Tax Reform Act of 1986 changed the IRA legislation removing some of the fiscal incentives for households with annual income above $40,000. Because of this, our tests distinguish the pre- and post-1986 periods. To date this exercise, which is interesting if saving incentives work differently for different income groups, has not been performed.

Our tests are not without problems and limitations. The most important is the possibility that new and continuing contributors differ in some unobserved dimension. As all identifying restrictions, this assumption is not testable. In what follows, however, we check whether the saving behaviour of the two groups is systematically different. We discuss these issues and a number of other important caveats in Section 3.

While these caveats should make one cautious in interpreting our results, our findings constitute *prima facie* evidence that the increase in IRA contributions that occurred from 1982 to 1986 was primarily financed from reshuffled assets and saving that would have been done anyway and was not new saving. In particular, we find no evidence that, during the 1982–6 period, households reduced their consumption to finance IRA contributions in the first nine months following their initial participation in an IRA plan. We also find evidence that these households used existing assets to finance their IRA contributions. This evidence seems to be particularly true before 1986. We fail to reject the hypothesis that changes in assets for new and continuing contributors were the same after 1986, when the Tax Reform Act reduced the eligibility for the IRA tax benefits and, therefore, changed the composition (and the number) of IRA contributors.

1. Literature Review

There is a substantial literature that addresses the issue of whether tax-favoured savings accounts induce households to increase their saving and whether tax-favoured savings accounts can increase national saving. The US literature has examined the impact of two programmes that create tax-favoured savings accounts – the IRA programme and the 401(k) programme. A smaller literature in the U.K. has examined the effects of tax-favoured savings accounts such as Personal Equity Plans (PEPs) and Tax Exempt Special Savings Accounts (TESSAs). All these studies examine data on household assets to determine the effectiveness of tax-favoured savings accounts in increasing household saving. No study to date has examined data on household consumption to test directly whether households financed their contributions to tax-favoured savings accounts from reductions in consumption. Kotlikoff (1990) has noted the lack of consumption studies in the literature, for example. The large literature on the effects of IRAs and 401(k)s has been reviewed in the Fall 1996 issue of the *Journal of Economic Perspectives* in papers by Hubbard and Skinner (1996), Poterba *et al.* (1996) and Engen *et al.* (1996). Additional reviews include Skinner (1991) and Bernheim (1997).

The major papers examining US tax-favoured savings accounts have estimated structural models of household saving behaviour using many different data sets including the Consumer Expenditure Survey (CEX), the Survey of Income and Program Participation (SIPP), and the Survey of Consumer Finances (SCF). These studies include Venti and Wise (1986, 1987, 1990, 1991), Gale and Scholz (1994), and Poterba *et al.* (1997). These papers reach conflicting conclusions with Gale and Scholz finding that the IRA programme did not generate new national savings while the Venti and Wise studies find that most IRA contributions represent new national savings.

Several studies have sought to determine the likelihood that contributions to tax-favoured savings accounts represent new saving by determining whether most contributions were made at the contribution limit or whether contributors had large stocks of financial assets that potentially could be substituted into the tax-favoured account. These studies include Burman *et al.* (1990) and Ozanne (1994) using data from the U.S. and Banks *et al.* (1994) and Banks and Tanner (1996) using UK data. These studies find that contributors typically have stocks of assets that potentially could be substituted into the tax-favoured accounts, raising the possibility that these accounts would generate little new saving. The studies using US data also find that most contributions to IRAs were limit contributions. This finding has been interpreted as showing that the programme could have had little effect on saving behaviour because the interest rate on marginal saving would not have changed. However, this conclusion is unwarranted because unless contributions are made at the limit in every period over the life-cycle, the programme will still affect the life-time budget constraint and potentially have an impact on saving. We show this possibility in our theoretical example in the following section. Engen and Gale (1993) also show this possibility using simulations.

A few papers have used panel data to address the question of the effectiveness of tax-favoured savings accounts. These include Feenberg and Skinner (1989), Engen *et al.* (1994) and Joines and Manegold (1995), all of which use the IRS-Michigan Tax Panel. These papers also find differing effects of tax-favoured savings accounts on national saving.

As mentioned above, none of the studies use consumption data to determine whether IRAs led households to increase their saving. A dimension in which our study adds to the existing literature on tax-favoured savings accounts is by using a panel of consumption data, the CEX. In addition, we compare households that just opened an IRA account with households that previously had made contributions to an IRA to mitigate the effects of unobserved heterogeneity.

2. Did IRA Incentives Increase Saving? A Simple Theoretical Framework

The empirical tests are described in detail in Section 3. However, although we neither test a formal model nor estimate structural parameters, the basic ideas behind our tests are best described with the help of a theoretical model. Our approach is to use available, non-experimental data and to examine the consumption and saving behaviour of new and continuing IRA contributing households. When a household starts participating in a tax favoured savings plan, it experiences a discrete change in the return to its saving. If this change is to have a positive effect on overall saving, it must imply a reduction in consumption. In other words, when a household starts participating in an IRA plan, it has to move to a new consumption trajectory unless it is simply reshuffling existing saving. On the other hand, households that have been contributing to an IRA for some time are *already* on the new consumption trajectory. These considerations imply a simple test: if the IRA programme is to lead households to increase their saving, then consumption changes for new contributors should be lower than consumption changes for continuing contributors. A null hypothesis that the tax incentives of IRAs had no effect on household saving implies that there should be no difference in the changes in consumption between new and continuing contributors.

To formalise this intuition, we consider a life cycle model without uncertainty. While this model is extremely simple, it demonstrates several points, most of which are either well known or have briefly been discussed in the literature. First, we show that, even in the absence of contribution limits, the effect of a change in the after-tax interest rate will have an ambiguous effect on consumption, because of counteracting income, wealth and substitution effects. Therefore, within the theoretical model we consider, the IRA legislation can have a positive effect on saving only if the substitution and wealth effects prevail on the negative income effect. Second, we show how comparing the consumption changes of new and continuing contributors can test the effectiveness of the IRA legislation. Third, we show that in the presence of binding limits to contributions, the chances that substitution and wealth effects prevail are reduced.[4] Finally, we show that if IRAs are to generate new saving, households should reduce consumption when they begin participating in the programme regardless of whether they initially contribute to the contribution limit.

Consider a consumer at the beginning of the second period of her four-period life cycle with wealth A_1. We label the four periods as 0, 1, 2 and 3. Decisions at time 0 have already been taken and are reflected in the initial value of wealth. Preferences are given by an intertemporally separable CRRA utility function.

[4] In passing we mention that, if the IRA contributions of continuing contributors are to represent new saving, then these contributors also should have reduced consumption in the period in which they first began contributing to the IRA programme.

The consumer's problem at time 1 is how to allocate consumption across the three remaining periods of her life. In this example, we assume the household receives income y_1 and y_2 in periods 1 and 2 and nothing in period 3. We will compare the pattern of consumption for an individual who has contributed to IRAs in periods 0, 1 and 2 (a 'continuing' contributor) with one for whom the IRA option became available only period 1 (a 'new' contributor). To simplify the algebra, we assume that the discount factor equals the pre-IRA interest rate. Moreover, we will assume that, in period 0, the new contributor was not aware of the possibility of investing into IRA accounts. This implies that, before IRAs became available, the consumption profile would have been flat, that is, $c_0 = c_1 = c_2 = c_3$. Finally, for simplicity, we start the analysis assuming that there are no limits to IRA contributions. The effects of these limits are discussed below. Given these assumptions, we can start to consider the optimisation problem of an individual who has no IRA available:

$$\max U = (1-\rho)^{-1}[c_1^{1-\rho} + (1-\delta)^{-1}c_2^{1-\rho} + (1+\delta)^{-2}c_3^{1-\rho}] \tag{1}$$

subject to the lifetime budget constraint:

$$c_1 + \frac{c_2}{(1+r)} + \frac{c_3}{(1+r)^2} = A_1 + y_1 + \frac{y_2}{(1+r)}.$$

The Euler equation for this problem is

$$c_{t+1}^{-\rho} = \frac{(1+\delta)}{(1+r)} c_t^{-\rho} \tag{2}$$

and period 1 consumption is given by

$$c_1 = \frac{\left(A_1 + y_1 + \frac{y_2}{1+r}\right)}{\left\{1 + \frac{(1+r)^{\frac{(1-\rho)}{\rho}}}{(1+\delta)^{\frac{1}{\rho}}} + \left[\frac{(1+r)^{\frac{1-\rho}{\rho}}}{(1+\delta)^{\frac{1}{\rho}}}\right]^2\right\}}. \tag{3}$$

The effect of an unexpected change in the interest rate, r, to r^* (which corresponds to the availability of IRA accounts) on consumption in period 1 is ambiguous and, in general, will depend on the parameter of intertemporal elasticity, $1/\rho$. If $1/\rho$ is larger than one, the substitution effect prevails over the income effect and an increase in the interest rate is associated with a decrease in first period consumption and an increase in saving. The only reason we are considering more than two periods is to allow for the wealth effect that arises from the fact that an increase in the interest rate implies that future income is discounted more heavily. Without such an effect, log utility ($\rho = 1$) implies that income and substitution effects cancel out and there is no effect of the interest rate on current consumption.

Period 1 is the period in which the interest rate changes for the new contributors – that is, it is the period in which they start contributing to IRAs. Suppose first that the participation, and therefore the implied change in interest rate was unanticipated. In an IRA programme with no contribution limits, when the household begins participating in the programme, it will move all of its existing assets into the IRA and change its consumption depending on the size of the parameter of intertemporal elasticity. Therefore, if we compare the consumption changes from period 0 to period 1 for the continuing contributor and the new contributor, one would observe a smaller rate of growth for the latter than for the former. If, on the other hand, the change is anticipated, the result is the same. To see why, notice that, if the change is anticipated, there will be an Euler equation holding between all periods. If we consider the Euler equation between periods 0 and 1 for new and continuing contributors, the one for new contributors will imply a lower rate of growth of consumption as the relevant interest rate for this group *between 0 and 1* will be the pre-IRA one. This is the basic idea behind our first test: if participation in an IRA is to increase saving, it has to reduce consumption when an individual starts contributing.

The implicit assumption that makes our test valid is the homogeneity in tastes between new and continuing contributors. If, for instance, new contributors had a different discount factor from continuing contributors, our tests would be invalid.

So far we have assumed, for expositional simplicity, that there are no limits to IRA contributions. In practice, the limits of contributions to the IRA programme are binding. In the presence of limits, the objective function of the household does not change, but the intertemporal budget constraint becomes more complicated. Define I_1 to be the contribution to IRAs in period 1 and I_2 to be the contribution to IRAs in period 2. IRAs earn an interest rate $r^* > r$ – the interest rate earned on non-IRA financial assets. There is a limit to IRA contributions, L, such that $L \geq I_1$ and $L \geq I_2$. There are three cases to consider. The first is the case in which the IRA limit is not binding in any period. This case is identical to the case where there are no IRA limits, discussed above. The second case is one in which the amount of assets the household has to reshuffle into IRAs exceeds twice the IRA contribution limit. The solution in this case is trivial. The household will reshuffle the contribution limit in each period and consumption will not fall in any period. The final case is more interesting. Suppose the contribution limit is such that $2L > A_1 > L$, that is, the household cannot reshuffle all of its assets to IRAs in the first period but also cannot reshuffle to the limit in both period 1 and 2. Recall that, in this simple example, we ignore uncertainty and precautionary saving so that assets serve only to smooth consumption over the lifecycle.

These conditions define the evolution of assets:

$$A_2 = (1 + r^*)I_1 + (1 + r)(A_1 - I_1 + y_1 - c_1) \tag{4}$$

$$A_3 = (1 + r^*)^2 I_1 + (1 + r^*)I_2 + (1 + r)[A_2 - (1 + r^*)I_1 - I_2 + y_2 - c_2] \tag{5}$$

$$A_3 = c_3. \tag{6}$$

We assume the constraints bind in the first period but not the second:

$$\begin{aligned} I_1 &= L \\ I_2 &= A_2 - (1 + r^*)L + y_2 - c_2. \end{aligned} \tag{7}$$

Substituting gives the following formula for the lifetime budget contraint:

$$c_1 + \frac{c_2}{(1+r)} + \frac{c_3}{(1+r^*)(1+r)} = \frac{(r^*-r)}{(1+r)}L + A_1 + y_1 + \frac{y_2}{(1+r)}. \tag{8}$$

Thus, the Euler equations for this problem are

$$\begin{aligned} c_2^{-\rho} &= \frac{(1+\delta)}{(1+r)} c_1^{-\rho} \\ c_3^{-\rho} &= \frac{(1+\delta)}{(1+r^*)} c_2^{-\rho} \end{aligned} \tag{9}$$

and consumption in the first period is given by

$$c_1 = \frac{\left[\frac{(r^*-r)}{(1+r)}L + A_1 + y_1 + \frac{y_2}{1+r}\right]}{\left\{1 + \frac{(1+r)\frac{1-\rho}{\rho}}{(1+\delta)\frac{1}{\rho}} + \left[\frac{(1+r)\frac{1-\rho}{\rho}}{(1+\delta)\frac{1}{\rho}}\right]\left[\frac{(1+r^*)\frac{1-\rho}{\rho}}{(1+\delta)\frac{1}{\rho}}\right]\right\}}. \tag{10}$$

As in the case with no IRA limit, the effect of participation in the IRA programme on consumption in period 1 is ambiguous and depends on the parameter of intertemporal elasticity. But now, even if the effect is positive, it will be considerably smaller than in the case without limits, as only one element of denominator will be affected by the change. Moreover, even though the marginal tax rate does not change until period 2, the household will reduce its consumption in period 1 if it is going to reduce its consumption at all. Also note that, if the constraint is binding in both periods, the effect of the change in the interest rate will be zero as r^* is an inframarginal tax rate.

From this theoretical example, we can see three points. First, the effect of participation in the IRA programme on consumption is ambiguous and depends on the elasticity of intertemporal substitution. Second, if IRAs generate new saving, new contributors will reduce their consumption even in the presence of sometimes binding contribution limits. Third, the presence of contribution limits reduces the amount of new saving the IRA programme could generate.

3. The Empirical Model

The question we address in this paper is whether the tax incentives of the IRA programme led to an increase in national savings through a reduction in household consumption. To tease the answer out of the available, non-experimental data, we use the CEX, a representative sample of US households available on a continual basis since the early 1980s. As discussed below, in the CEX, households are interviewed four times over a period of a year and are asked information on consumption expenditure, income, demographic variables, and whether and how much they contribute to an IRA. We try to minimise the effect of unobserved heterogeneity in the taste for saving by focusing only on households that were contributing to IRAs when last observed in our sample. We then examine the consumption and saving behaviour of new and continuing IRA contributing households. That is, we test the implications about different hypotheses for those households that started contributing to an IRA during the nine months preceding the last interview and those who were contributing both at the first and last interview.[5]

We compare these two groups to perform two tests of the hypothesis that IRA incentives stimulated saving. Our first test, whose logic we discussed in the previous section, is based on a comparison of their consumption changes. The second is based on a comparison of their changes in non-IRA financial assets. Once again, the key identifying assumption we make is that there is no unobserved heterogeneity between new and continuing contributors.

3.1. *A Test Based on Consumption Changes*

In our first test, the null hypothesis is that the tax incentives of IRAs had no effect on household saving. This implies that there should be no difference in the changes in consumption between new and continuing contributors. The empirical specification we use to implement such test consists of the simple equation:

$$\Delta \ln C_i = \boldsymbol{\beta}^c \mathbf{X}_i^c + \gamma^c d_i^{new} + \varepsilon_i^c \qquad (11)$$

where $\mathbf{X}_i^c$ is a vector of controls discussed in the results section and d_i^{new} is a dummy variable that equals one for the households that started contributing during the interview period. Under the null that the IRA incentives did not generate new saving, the coefficient γ^c should be zero. While (11) is in logs, and therefore focuses on proportional changes in consumption, in the empirical application we also estimate it in levels.

As a consequence of the Tax Reform Act of 1986, the number of households contributing to IRAs dropped dramatically and the composition of contributors changed. After the Tax Reform Act of 1986, only households with relatively low incomes or households without access to pension plans were eligible to make tax-deductible contributions to an IRA. These lower income households tend to have very low levels of non-IRA financial asset balances. Therefore, it is likely that these households would be less able to take advantage of the tax incentives of the IRA programme by financing their IRA contributions from existing savings or saving that would have been done anyway than high income households would.

Because of the difference in the tax status of IRAs before and after 1986, we allow the coefficient on the new contributor dummy to be different in the two periods in all our equations. If the reshuffling hypothesis is true, we expect that new contributors in the 1982–6 period would finance their IRA contributions from existing savings or from saving that would have been done anyway to a greater extent than would new contributors in the post-1986 period. We therefore modify (11) as follows:

[5] While we stress that our results only apply to consumption and saving in the first nine months after a household starts making contributions, we should also note that much of the evidence in the literature on the effectiveness of IRAs (Venti and Wise 1987, 1990, 1991) is based on one year's worth of IRA contributions. While it is a limitation, focusing on the first nine months allows for a comparison with the results of the previous literature.

$$\Delta \ln C_i = \boldsymbol{\beta}^c \mathbf{X}_i^c + \gamma_i^c d_i^{new} \delta_i + \gamma_2^c d_i^{new}(1 - \delta_i) + \varepsilon_i^c \quad (11')$$

where δ_i equals one if household i is observed between 1982 and 1986 and zero if it is observed after 1986.

This paper is the first in the literature to exploit the information contained in the consumption behaviour of IRA participants to determine whether their contributions represent new saving. However, our consumption-based test is not exempt from problems. Here we discuss some important caveats that have to be kept in mind when interpreting our results.

First, it is possible that we fail to reject the hypothesis that the coefficient γ^c is equal to zero simply because the consumption information we use is affected by measurement error. Obviously, the presence of noise in the data reduces the power of our test. However, we will also put confidence intervals around γ^c and thus on the potential amount of new national saving. In addition, we should stress that our second test, conducted on the same sample, sets as a null the hypothesis that the tax incentives are effective. Therefore, if we fail to reject the null in the first test because of lack of power caused by noisy data, we should also fail to reject the null in the second.

Second, it is possible that consumers who start participating in an IRA plan initially reshuffle their existing assets and, after a period of adjustment, reduce their consumption. If this reduction occurs after nine months, our test would miss it. Given the available data, we have no way to know what households do after they leave the sample. Should our test fail to reject the null, therefore, it would only indicate that households participating in IRAs do not increase their overall saving in the first nine months. It is not obvious, however, how to write an optimisation model in which households reduce their consumption (and increase their saving) some time after entering the programme, that is, after experiencing an increase in the rate of return to their saving. In fact, as our theoretical example demonstrates, it is not possible (or at least very difficult) to do so in the context of the life-cycle model.

Third, it is possible that IRA contributions are financed neither from existing assets (or saving that would have been done regardless of the presence of fiscal incentives) nor from reduced consumption. An alternative source is the reduction in tax liabilities arising from the participation in an IRA plan. Because of this possibility, in addition to the test based on consumption, we run an additional regression where the dependent variable is the change in saving rates out of after-tax total income.

$$\Delta(S/C)_i = \boldsymbol{\beta}^s \mathbf{X}_i^s + \gamma_i^s d_i^{new} \delta_i + \gamma_2^s d_i^{new}(1 - \delta_i) + \varepsilon_i^s \quad (12)$$

where S/C is the ratio of saving to consumption. We prefer to divide saving by consumption to control for differences in the total amount of resources available to individual households. If one controls for *before tax* income changes, the change in this savings rate measure includes the change in consumption and the change in tax liabilities linked to the participation in an IRA plan. The implications for national saving of an increase in saving arising from a reduction in consumption or from a reduction in tax liabilities are obviously very different. In the latter case, the increase in private saving is counterbalanced by a reduction in government saving.

3.2. *A Test Based on Changes in Financial Assets*

While our first test is based on changes in consumption, most of the literature, instead, has examined changes in asset balances to determine the effects of IRA incentives. Income, consumption and changes in assets are obviously linked by simple accounting identities. Thus, to complement our first test, we also provide an additional test of the hypothesis that the IRA legislation was effective, based on the comparison of changes in non-IRA financial assets for new and continuing contributors. If there was no reshuffling, there is no reason to believe that the change in non-IRA financial assets should be smaller for new contributors than for continuing ones. On the contrary, if IRA contributions do not represent new saving, one would expect the change in non-IRA financial assets to be smaller for the new contributors than for continuing contributors. This test also complements our consumption test because its null hypothesis is no reshuffling as opposed to no new saving.

Another way to see this last point is to assume that there is an equilibrium portfolio of IRA and non-IRA financial assets (the former having a higher return but less liquidity). If IRA contributions do not represent new saving, one would observe that new contributors immediately reduce the stock of non-IRA financial assets to build IRA assets up to the equilibrium position. On the other hand, continuing contributors are likely to have already reached the equilibrium and would be allocating their flow of saving between IRA and non-IRA saving.

A simple example can help us to make the logic behind our second test clear. Suppose a household saves \$300 a year and, when the IRA programme starts (or the household decides to join it), it has a stock of financial saving of \$3,000 and decides to keep \$800 in relatively liquid financial assets and move the rest to an IRA account. Furthermore, we assume that the household has a limit contribution of \$2,500. If the IRA programme does not generate new saving, the household still saves \$300 a year. On the first year, however, it puts \$2,500 in the IRA account, therefore, decreasing the stock of financial assets by \$2,200. Having reached the desired portfolio composition, in the second year, the household in question will put the \$300 it saves in the IRA account and will not change the stock of non-IRA financial assets. Therefore, under the hypothesis that IRA contributions are not new saving, the changes of non-IRA financial assets would be smaller (−2,200 vs. 0 in the example above) for a new contributor than for a continuing one. For two households like the one in the example but observed as they join and after a year, the difference in the change in non-IRA financial assets will be \$2,200.

The empirical equation that allows us to perform this test is therefore

$$\Delta A_i = \boldsymbol{\beta}^a \mathbf{X}_i^c + \gamma_1^a d_i^{new} \delta_i + \gamma_2^a d_i^{new} (1 - \delta_i) + \varepsilon_i^a \tag{13}$$

where A_i is the stock of non-IRA financial assets. The null hypothesis that the parameters γ^a are equal to zero corresponds to the hypothesis that the saving incentives do stimulate new saving. As mentioned above, if households finance IRA contributions from existing assets, we would expect the parameters γ^a to be less than zero. Once again, we allow for the possibility of different effects before and after 1986.

The results of the test based on (13) should be interpreted with caution for a number of reasons. First, even if the IRA legislation generates new saving, new contributors would readjust their existing portfolios. That is, even if there were new saving, some reshuffling would happen if the change in the structure of returns causes a different portfolio to be optimal. Using the same example as before, suppose that the household considered increasing its annual saving from \$300 to \$800 but still wants to keep a balance of liquid assets of \$800 and, as before, has an IRA contribution limit of \$2,500. In the first year, it will decrease the stock of financial assets by \$1,700 and, in the second year, by another \$500. In the following year, it will contribute the \$800 to the IRA and the change in the stock of non-IRA financial assets will be zero. The difference between new and continuing contributors will now be substantially smaller than the contribution to IRAs (1,700 − 500 = 1,200 if we consider as the typical continuing contributor one that has contributed for one year and 1,700 if we consider it as one who has contributed for two years). The higher the proportion of IRA contributions that is new saving, the smaller will be the difference between the changes in non-IRA financial assets between new and continuing contributors. These considerations are important for the exercise we perform in Section 5.

One might argue that because of the limits to contributions, it might take several quarters or years for IRA contributors to reshuffle their non-IRA financial assets. Thus, a comparison of the non-IRA financial asset growth of new and continuing contributors may understate the true amount of reshuffling going on. Furthermore, one might argue that the timing of the observations is quite important. If the households which we define as 'new' contributors have already reshuffled their assets, they would look like continuing contributors. These arguments, however, are about the power of the test that compares the non-IRA financial asset growth of new and continuing contributors and is not a major worry if we do reject the null hypothesis.[6]

Overall, the first test we propose, based on changes in consumption, is a stronger and more direct test and is a direct implication of a standard optimisation model. Our second test, however, constitutes a useful complement to the first one. As mentioned above, while the first test we propose has the reshuffling hypothesis as a null, the opposite is true for the second test. When comparing the change in non-IRA financial assets between new and continuing contributors, the null is the hypothesis of new saving. Importantly, the consumption-based test is not affected by problems associated with changes in the optimal portfolio; it is based on the simple idea that IRA incentives can only increase national saving if they reduce household consumption.

[6] We should stress that, given the CEX sampling frame, in the last interview, we observe the level of non-IRA financial assets and its change *in the last 12 months.* We also know whether the household was contributing or not 12 months before the interview. Therefore, for new contributors, we observe the change in non-IRA financial assets during the period over which the household opened an IRA account.

3.3. *Testing for Differences in Saving Behaviour between New and Continuing Contributors*

The tests discussed in the preceding sub-sections are based on the comparison of new and continuing contributors. The implicit identification assumption, therefore, is that these two groups do not differ in their 'unobserved taste for saving'. The issue that unobservable tastes for saving of IRA contributors may differ from those of non-contributors has plagued earlier studies comparing the savings behaviour of these groups. Because both groups choose to contribute to IRAs and because continuing contributors were once new contributors, it is likely that this problem is less serious in our study; these groups are likely to have similar unobservable tastes for saving. Of course, the issue is why some households choose to start contributing to IRA plans before others. It is possible to think of plausible explanations of these observed differences that would not jeopardise the logic of our tests. One possibility, for instance, is that there is a life cycle pattern for portfolio composition that, regardless of unobserved tastes for saving, would generate differences across households. Indeed, we do observe that continuing contributors are, on average, older than new contributors. In the end, however, as with all identification restrictions, ours is not testable. However, it is worth investigating whether the two groups under consideration present significantly different saving behaviour.

If our two groups do have different unobservable tastes for saving, we also would expect to observe a systematic difference in saving rates between the two groups. Therefore, as an indication as to the extent to which unobservable heterogeneity may be an issue, we test the hypothesis that the two groups have the same saving rates by estimating

$$SR_i = \boldsymbol{\beta}^s \mathbf{X}_i^s + \gamma_1^s d_i^{new} \delta_i + \gamma_2^s d_i^{new} (1 - \delta_i) + \varepsilon_i^s \tag{14}$$

where SR_i is the ratio of total saving (including IRA contributions) to non-durable consumption for household i.

Another indication of the extent to which unobservable tastes for saving may differ between new and continuing contributors is the stock of non-IRA financial asset balances. Once again, if these two groups have similar unobservable tastes for saving, we would expect them to have similar stocks of non-IRA financial assets. Testing this involves estimating

$$K_i = \boldsymbol{\beta}^k \mathbf{X}_i^k + \gamma_1^k d_i^{new} \delta_i + \gamma_2^k d_i^{new} (1 - \delta_i) + \varepsilon_i^k \tag{15}$$

where K_i is non-IRA financial asset balances for household i divided by non-durable consumption.[7] Ideally, we would like to compare the total asset balances of new and continuing contributors, that is, IRA balances plus non-IRA financial asset balances. Unfortunately, as discussed in Section 4, we do not have information on IRA balances, only IRA contributions. Comparing new and continuing contributors using the savings variable K_i, as in (15), presents an advantage and a disadvantage relative to using the saving rate, as in (14). Both the advantage and the

[7] We prefer to divide the stock of saving on the left-hand side of (15) by consumption to control for differences in permanent income. For example, the assets of a college graduate are bound to be larger than those of a high school dropout, regardless of their 'taste for saving.'

disadvantage arise from the fact that K_i represents a stock of savings while the savings rate is a flow. Using K_i is an advantage because the total stock of savings reflects a household's past history of saving; a regression based on the stock of savings may be more powerful in picking up differences in the taste for saving than a regression based on the flow of saving. On the other hand, using K_i is a disadvantage because the presence of household specific fixed effects correlated with the new contributor dummy variable would bias our results. If not for the exclusion of IRA balances from K_i, first differences of (15) would yield (14) and eliminate the fixed effects.

4. Data and Sample Characteristics

The CEX is the only data set that contains information on the consumption, IRA contributions and non-IRA financial asset growth of a large sample of US households. Thus, it is an ideal data set on which to perform our tests of the new saving and reshuffling hypotheses.

The CEX is a revolving panel data set.[8] Each household in the sample for a given year is interviewed four times and, each quarter, a new sample of households is introduced. Each interview contains detailed data on household consumption expenditure in the three months preceding the interview as well as other demographic information. The first and fourth interviews contain information on income and IRA contributions over the twelve months preceding the interview. The fourth interview also contains information on non-IRA financial asset balances and changes in these balances over the past year. Some households fail to complete all four interviews. Because of the nature of our test, we restrict our sample to include only those households that complete all of the interviews.[9]

Consider those households that reported making IRA contributions in the fourth interview. We can divide this group into new contributors – households that reported making IRA contributions in the fourth but not the first interview – and continuing contributors – households that reported making IRA contributions in both the fourth and first interviews.

There is an issue as to whether we are correctly classifying IRA contributing households as new or continuing contributors. It is possible that some households that already have IRA assets may have failed to contribute to their IRA account in the twelve months preceding the first interview but make a contribution in the twelve months preceding the fourth interview. These intermittent contributors, if they exist, would be falsely classified as new contributors. It is not clear to us what kind of bias this problem introduces. Fortunately, the evidence from other studies suggests that only a very few IRA contributors are intermittent contributors.[10]

[8] For an extensive discussion of the CEX, see Attanasio (1994).

[9] While we are aware of the possibility that this procedure could introduce sample selection bias, we ignore this problem. The evidence in Nelson (1994) shows that the households that complete more interviews are likely to contain more educated and wealthier individuals.

[10] Skinner (1991) and Joines and Manegold (1995) provide evidence from the Michigan-IRS Tax Panel on the persistence of IRA contributions. Joines and Manegold find that, from 1982 and 1986 when the tax benefits of IRAs were universally available, 24–40% of contributors made a contribution in all five years. Skinner also finds a 'strong persistence' in IRA contributions.

There is a third group of IRA contributors that we do not consider in our analysis. This group consists of households that reported making an IRA contribution in the first interview but not the fourth. These households exited from the IRA programme and thus we label them 'drop-outs'.[11]

In what follows, we deflate nominal measures in two different ways. In the first, we report the measure as a ratio to non-durable consumption, which is defined as total consumption expenditure minus expenditure on durables, health, education and housing. The division by non-durable consumption is justified by the necessity to control for different levels of permanent income. In the second, we convert nominal dollars into 1982–4 dollars using a Stone price index for consumption constructed from the detailed BLS price indexes using individual household expenditure shares as weights.

We calculate consumption growth as the change in log real total consumption expenditure (deflated by the Stone price index described above) between the quarters preceding the first and fourth interviews. We compute saving rates as the difference between annual income over the 12 months preceding the last interview and total consumption expenditure, divided by consumption. This is a monotonic transformation of the traditional saving rate defined as the ratio of saving and income.

We calculate non-IRA financial assets as the sum of checking, savings, stock portfolio and bond portfolio balances at the time of the last interview. In addition to these balances, the CEX also contains information on the changes in these balances over the year preceding the last interview. We use these data to construct the dependent variable of the test based on asset changes.

Our calculation of the change in the savings rate is more problematic. The CEX has data on consumption expenditure for the 3 months preceding each interview and income data for the 12 months preceding the first and last interviews. Thus, the data on income span 21 months while the data on consumption expenditure span only 12 months. Furthermore, the annual income data from the last interview overlap with all 12 months of consumption data while the annual income data from the first interview overlap with only the first 3 months of consumption data. We calculate the 3-month savings rate for the last interview as the difference between one-fourth of annual income in the last interview and consumption expenditure in the 3 months preceding the last interview. We define the 3-month savings rate for the first interview analogously. The change in the savings rate is calculated as the difference between these two 3-month savings rates.

In Table 1, we report the number and percentage of each type of IRA contributor in each year of our sample. The percentage of all IRA contributors combined increased dramatically in 1982 following the universal expansion of eligibility, and declined after eligibility was curtailed for many individuals in 1986. The percentage of new contributors exceeded that of continuing contributors in 1982 and 1983 and fell thereafter to about 5%. The percentage of continuing contributors closely follows the trend for all contributors. The percentage of IRA

[11] We have considered 'drop-outs' in additional analyses not presented in the paper. Drop-outs tend to have lower incomes and negative income growth relative to continuing contributors.

Table 1
Types of IRA Contributors

Survey year	Sample size	New contributors N	New contributors %	Continuing contributors N	Continuing contributors %	Drop-outs N	Drop-outs %
1982	2,258	150	6.6	49	2.2	28	1.2
1983	3,136	238	7.6	198	6.3	62	2.0
1984	3,229	173	5.4	333	10.3	109	3.4
1985	3,555	188	5.3	423	11.9	140	3.9
1986	1,521	75	4.9	198	13.0	78	5.1
1987	4,065	199	4.9	455	11.2	231	5.7
1988	3,540	133	3.8	243	6.9	257	7.3
1989	3,603	183	5.1	213	5.9	174	4.8
1990	3,578	193	5.4	193	5.4	174	4.9
All years	28,485	1,532	5.4	2,305	8.1	1,253	4.4

Source: Consumer Expenditure Survey.
Note: Only households that complete all four interviews are included in the sample.

drop-outs rose steadily until 1988. In fact, from 1986 to 1988, there were more drop-outs than there were new contributors. The data reported in Table 1 is consistent with that reported in Skinner (1991) which finds that, from 1982 to 1986, the probability that an individual contributes to an IRA, given that the individual contributed in the previous year, is about 70%.

In Table 2, we report summary statistics on the observable characteristics of new and continuing contributors and test whether these statistics are equal across the two groups. In the third column, we report the *t*-statistic for the difference between the two means (or medians). New and continuing contributors differ on observable characteristics in several ways. On average, new contributors are less educated, have less income, have less wealth, are more likely to be minority, and have more children than continuing contributors. These findings suggest, but do not ensure, that there may be differences in unobservable characteristics between the two groups. Furthermore, given the evidence in Nelson (1994) which suggests that households which complete all four interviews tend to have higher incomes than those which do not, these findings suggest that our sample may under-represent new contributors.

5. Results

We estimate (11′)–(15) using consumption, asset, savings and demographic data from our sample of IRA contributors from the CEX. Because consumption and savings behaviour can change dramatically over the life cycle and the business cycle, we include in the estimating equations several controls meant to capture both life cycle and business cycle effects. These include five age dummies, indicators for both the number of children and the number of senior citizens in the household, a dummy for the presence of a spouse, and annual time dummies. Because the change in consumption between the first and the fourth interviews spans less than a year, the equations for consumption growth and saving

Table 2
New versus Continuing Contributors

Means (standard deviation)	New contributors $N = 1{,}532$	Continuing contributors $N = 2{,}766$	t-test for equality of means
Real annual income	35,141	38,414	−4.300
	(21,675)	(23,194)	
Median real annual income	32,489	35,280	−2.828
Real annual consumption	28,359	29,922	−0.823
	(15,490)	(17,825)	
Median real annual consumption	25,316	26,503	−1.826
Real wealth	19,128	28,371	−9.315
	(31,402)	(39,644)	
Median real wealth	6,632	11,437	−5.962
Real IRA contribution	2,762	2,381	2.041
	(8,502)	(2,332)	
Median real IRA contribution	1,807	1,917	−4.054
Minority	0.0440	0.0227	3.684
	(0.2051)	(0.1490)	
Married	0.7507	0.7358	−0.127
	(0.4327)	(0.4410)	
High school drop-out	0.0892	0.0892	0.799
	(0.3122)	(0.3033)	
High school graduate	0.2963	0.2385	3.797
	(0.4568)	(0.4263)	
Some college	0.2096	0.2120	−0.531
	(0.4071)	(0.4088)	
College graduate	0.4049	0.4603	−4.322
	(0.4910)	(0.4985)	
Age	46.98	48.30	−5.496
	(12.27)	(11.42)	
Number of children	0.7173	0.5671	5.207
	(1.0242)	(0.9214)	
Number of senior citizens	0.1347	0.1227	1.220
	(0.4113)	(0.3918)	

Source: CEX.
Note: Statistics are computed using sampling weights from the CEX.

also include monthly seasonal dummies. Finally, we control for education level, minority status and region of residence. Due to the presence of top coding and a potentially skewed distribution of wealth, income and consumption in our sample, we estimate these equations using both least squares and quantile regressions.

5.1. *The Consumption Test*

In Table 3, we report the estimates of the key parameters of (11′) – the coefficients on the 'new contributor' dummies in the two periods.[12] The dependent variable in panel *a* is the change in log real total consumption between the quarters preceding the first and fourth interviews. The dependent variable in

[12] We report the complete results of both the OLS and quantile regressions (including the 50th, 75th and 90th percentiles) of the consumption growth equation, (11′), in Tables A1 and A2 of the Appendix.

Table 3

Change in Total Consumption

	Not controlling for income and income growth				Controlling for income and income growth			
	(1) OLS	(2) 50th	(3) 75th	(4) 90th	(5) OLS	(6) 50th	(7) 75th	(8) 90th
(a) Log								
New contributor in 1982–86	0.011 (0.024)	0.002 (0.021)	0.031 (0.028)	−0.040 (0.050)	−0.008 (0.026)	0.008 (0.021)	0.017 (0.028)	−0.069 (0.061)
New contributor in 1987–90	0.019 (0.027)	0.030 (0.024)	0.001 (0.032)	−0.040 (0.055)	0.013 (0.028)	0.027 (0.022)	−0.002 (0.031)	−0.012 (0.065)
p-value	0.831	0.366	0.474	0.993	0.558	0.524	0.636	0.517
(b) Levels								
New contributor in 1982–86	−175.8 (257.1)	40.4 (111.4)	224.7 (199.5)	−203.0 (538.9)	−275.0 (272.1)	64.6 (108.6)	176.4 (147.0)	−761.7 (633.5)
New contributor in 1987–90	131.4 (281.4)	154.4 (121.6)	127.8 (220.0)	−504.0 (592.2)	97.1 (291.8)	87.0 (116.0)	82.0 (160.0)	−454.3 (669.4)
p-value	0.412	0.481	0.739	0.700	0.341	0.885	0.657	0.733

Note: p-value reports the results of a Wald test that the two coefficients are equal. Standard errors are reported in parentheses.

panel *b* is difference in real levels of total consumption between these two quarters. As noted above, because eligibility for the IRA programme was restricted for high-income households as a result of the Tax Reform Act of 1986, we allow the coefficient on the 'new contributor' dummy to differ between the 1982–6 period and the 1987–90 period. In the first four columns of Table 3, we do not control for income and income growth, while we do in the last four.

The results in panel *a* of Table 3 indicate that new contributors have slightly higher consumption growth than continuing contributors on average, at the median, and at the 75th percentile when we do not control for income or income growth. While these variables might be endogenous, not controlling for them might fail to capture the fact that some shocks might lead households to increase their consumption and to begin participating in the IRA programme. When we introduce income and income growth among the regressors, the results indicate that new contributors have slightly lower consumption growth than continuing contributors on average and at the 90th percentile. However, in all of these results, the difference in consumption growth between new and continuing contributors is never significantly different from zero. Moreover, the point estimates are economically small. Thus, these results do not present any evidence that consumption growth was lower for new contributors than for continuing contributors in either period. The results in panel *b* are similar.[13]

[13] Table 3 also reports the results of a test that the coefficient on the new contributor dummy was different in the 1982–6 period from that in the 1987–90 period. In no specification do these coefficients differ significantly.

These results represent a piece of evidence against the hypothesis that the IRA legislation was effective in generating new saving. The only way in which such legislation could generate a positive effect on national saving is by reducing household consumption. As argued in Section 3, it is conceivable that consumers who start participating in an IRA plan initially reshuffle their assets and only after some time reduce their consumption. However, because we observe the households in our sample for only a year, we have no evidence on what they do after that period of time. We should stress, however, that in most optimisation models including the life-cycle model, if the IRA programme is to generate new national saving, consumption should decrease when households began participating in the IRA programme.[14]

Of course, as we fail to reject the null hypothesis that there was no new saving, our results could be due to the low power of our test. Such a low power could be caused by measurement error in our consumption data. While this is obviously a possibility, we should stress that even the point estimates of the parameters of interest tend to point in the wrong direction. Moreover, the estimates are precise enough to rule out a large consumption effect.

As noted in Section 2, it is possible that IRA contributions are financed in part from reduced tax liabilities that result from participating in the plan. If so, household saving could increase even if household consumption did not fall. Of course, in this case, this increase in household saving would be financed completely by a reduction in government revenues and, thus, national saving would not increase. To test for this possibility, we estimate (12); we regress the change in the saving rate between the first and fourth interviews on a set of control variables and our new contributor dummies. We report the key parameter estimates in Table 4. As in the consumption equation, we present OLS, median, 75th, and 90th percentile specifications both with and without controls for income and income growth.[15]

The change in the savings rate combines both the change in consumption and the change in tax liabilities when we control for income growth. To isolate the change in tax liabilities and the change in consumption from any change in income, it is essential to control for income growth. Thus, we favour the specifications that do so. However, in an attempt to isolate the effect of a reduction in tax liabilities (that are directly related to the participation in the IRA programme) from other changes in income, we control for growth in *before tax income.*

The results are somewhat mixed. The first four columns, which do not control for income growth, show that change in saving out of after-tax income is on average

[14] Our results are robust to changes in the definition of the growth rate in consumption. In our consumption test, we define the growth rate of consumption as the change in log total consumption between the first and fourth interviews. One could also define the growth rate to be the change in total consumption between the first and second, first and third, second and third, second and fourth, or third and fourth interviews. Doing so does not change our results in any way. We also have tried alternative definitions of consumption. For instance, we have excluded expenditure on durable goods (which has a saving component) from the definition of consumption we use for our exercise. The results were, once again, unaffected.

[15] The full set of results is reported in Table A3 of the Appendix.

Table 4

Change in Savings Rate

	Not controlling for income growth				Controlling for income growth			
	(1) OLS	(2) 50th	(3) 75th	(4) 90th	(5) OLS	(6) 50th	(7) 75th	(8) 90th
New contributor in 1982–6	0.165	0.089	0.145	0.226	−0.014	0.003	0.007	−0.126
	(0.049)	(0.034)	(0.059)	(0.095)	(0.048)	(0.040)	(0.043)	(0.074)
New contributor in 1987–90	0.173	0.158	0.173	0.322	−0.031	−0.023	−0.024	−0.022
	(0.053)	(0.035)	(0.062)	(0.107)	(0.050)	(0.040)	(0.044)	(0.078)
p-value	0.911	0.150	0.739	0.489	0.796	0.648	0.605	0.322

Notes: p-value reports the results of a Wald test that the two coefficients are equal. Standard errors are reported in parentheses.

17% greater for new contributors than for continuing contributors. The median, 75th percentile and 90th percentile changes are also large and are all statistically significant. The last four columns, however, control for before tax-income growth. In all of these specifications, the change in saving out of after-tax income is not statistically different for new and continuing contributors. While it is logically possible that households financed a portion of their IRA contribution out of reduced tax liabilities, we can only find weakly suggestive evidence that this is the case.

5.2. *Asset Test*

Having not been able to reject the null hypothesis that households did not finance their IRA contributions from reductions in consumption, we now turn to the null hypothesis that households did not finance their IRA contributions from reshuffling existing saving. The null is now that no reshuffling occurs while in the consumption test the null is that no new saving occurs. In Table 5, we report the estimates of the key parameters of (13), the change in non-IRA financial assets equation.[16]

As mentioned in Section 4, non-IRA financial asset balances are defined as the sum of the household's stocks, bonds, savings accounts and checking accounts. In panel *a* of Table 5, we divide non-IRA financial asset balance growth by non-durable consumption while, in panel *b*, we deflate non-IRA financial asset balances by the Stone price index.[17] Again, we report OLS estimates as well as estimates from quantile regressions. As in the consumption growth equation, we allow the coefficient on the new contributor dummy variable to change between the 1982–86 period and the 1987–90 period.

[16] The full set of results is reported in Tables A4 and A5 of the Appendix.

[17] In particular, we define the real change in non-IRA financial assets as $[A_{t+1} - A_t(1 + \pi_{t+1})]/PI_{t+1}$ where PI_{t+1} is the price index in period $t+1$ and B_{t+1} is the rate of change in the price index from period t to period $t+1$.

Table 5

	Not controlling for income and income growth			Controlling for income and income growth		
	(1) OLS	(2) 75th	(3) 90th	(4) OLS	(5) 75th	(6) 90th
(a) Change in Non-IRA Financial Assets/Consumption						
New contributor	−0.0737	−0.0847	−0.1212	−0.0681	−0.0775	−0.1030
in 1982–6	(0.0494)	(0.0211)	(0.0677)	(0.0512)	(0.0199)	(0.0694)
New contributor	0.0421	−0.0471	−0.0199	0.0413	−0.0293	−0.0087
in 1987–90	(0.0545)	(0.0232)	(0.0751)	(0.0561)	(0.0220)	(0.0775)
p-value	0.1098	0.2220	0.3051	0.1407	0.0960	0.3525
(b) Change in Non-IRA Financial Assets						
New contributor	−2038.1	−1216.6	−3892.0	−1956.9	−1194.5	−3647.1
in 1982–6	(669.0)	(513.4)	(1106.0)	(675.8)	(500.6)	(1152.5)
New contributor	−1084.3	−827.9	−1387.5	−039.7	−584.5	−1197.2
in 1987–90	(736.5)	(570.6)	(1259.5)	(740.4)	(552.9)	(1307.2)
p-value	0.3292	0.6058	0.1273	0.3497	0.4013	0.1471

Notes: p-value reports the results of a Wald test that the two coefficients are equal. Standard errors are reported in parentheses.

The results in panel *a* of Table 5 indicate that new contributors in the 1982–86 period had much slower growth in their non-IRA financial asset balances than did continuing contributors and that this difference is evident on average, at the 75th, and at the 90th percentiles. The difference in the rate of growth of non-IRA financial assets of new and continuing contributors is large – more than 7% of non-durable consumption when we do not control for income and income growth, and almost 7% of non-durable consumption when we do. Furthermore, all of the differences are statistically significant except for those based on the OLS results that are only marginally significant or are insignificant.

The change in the level of non-IRA financial assets in constant dollars gives a real dollar estimate of the difference in non-IRA financial asset growth between new and continuing contributors. As reported in panel *b* of Table 5, this difference is, on average, over $2,000 dollars when we do not control for income and income growth and over $1,900 when we do. Moreover, all of these differences are statistically significant. In the 1987–90 period, on the other hand, the difference between the non-IRA financial asset growth of new and continuing contributors is much smaller, especially at the 90th percentile.

Because of the eligibility restrictions imposed on high income households by the Tax Reform Act of 1986, we should expect greater amounts of reshuffling in the 1982–6 period than in the 1987–90 period. This difference is because, in the 1987–90 period, we would expect the composition of households contributing to IRAs to have changed as the tax incentives of IRAs were removed for most high-income households. Our estimates show exactly this pattern. The difference in non-IRA financial asset growth between new and continuing contributors was much higher in the 1982–6 period. In Table 5, we report results of a test of whether the coefficients on the new contributor dummy variable are equal in the two time periods. However, we can reject the null at a 95% confidence level in any specification.

The differences in the changes in non-IRA financial assets for new and continuing contributors, as reported in panel *b* of Table 5, are quite large relative to their average and median real IRA contribution, which equal $2,762, and $1,807 respectively, indicating reshuffling behaviour by new contributors. However, as discussed in Section 4, there can be a substantial amount of reshuffling even when the household increases its saving. However, the larger is the difference between the changes in the non-IRA financial assets between new and continuing contributors relative to IRA contributions, the smaller the proportion of new IRA contributions that can be new saving. While the results reported in Table 5 indicate a substantial amount of reshuffling on the part of new contributors, these results could be consistent with some new saving. Note that, to the extent that continuing contributors are still reshuffling assets into IRA accounts, the point estimates in Table 5 reflect lower bounds on the amount of reshuffling.[18]

In summary, while the results of both the consumption and asset tests for the 1982–6 period are not inconsistent with the reshuffling hypotheses, for the 1987–90 period, the asset test cannot always reject the hypothesis of no reshuffling. A possible interpretation of this result is that saving incentives were more effective in stimulating the saving of contributors who participated in IRA plans after 1986. It should be remembered that the 1986 Tax Reform Act, by changing eligibility rules, changed substantially the number and composition of IRA contributors.

5.3. *Heterogeneity in Saving Rates and in Stocks of Saving*

A legitimate concern is whether the results we obtain are due to unobserved heterogeneity. To help to explore this concern, we estimate (14) and (15): (14) is a saving behaviour equation; (15) is a non-IRA financial asset balance equation. Estimating these equations indicates whether new and continuing contributors have observable differences in their saving behaviour and may indicate the extent to which differences in the taste for saving between the two groups are important.

We report OLS and median regression estimates of the key parameters of (14) in the first two columns of Table 6. We define saving as the difference between after-tax income and total consumption. To control for possible differences in current and permanent income between the two groups, we consider the ratio of saving to consumption and report these results in Table 6.[19]

[18] In consumption and asset tests that also allow for 'drop-outs' from the IRA programme, none of the results for new contributors change. Moreover, none of the results for drop-outs in the period in which IRAs were universally available are consistent with IRA contributions representing new saving. In particular, consumption growth for drop-outs tends to be negative rather than positive and non-IRA financial asset growth tends to be negative rather than zero. Since dropping out should have no effect on non-IRA financial asset growth under either the new saving or the reshuffling hypothesis, it is unclear what to make of this finding other than drop-outs are not randomly selected from the group of IRA contributors.

[19] This variable is a monotonic transformation of the saving rate and has the advantage of being defined even when income is equal to zero.

Table 6

Saving and Non-IRA Financial Assets as a Ratio of Consumption

	Saving/Consumption		Non-IRA financial asset balances/Consumption	
	(1) OLS	(2) 50th	(3) OLS	(4) 50th
New contributor in 1982–6	−0.064	−0.064	−0.276	−0.084
	(0.034)	(0.036)	(0.064)	(0.025)
New contributor in 1987–90	−0.045	−0.034	−0.177	−0.109
	(0.038)	(0.039)	(0.070)	(0.028)
p-value	0.700	0.556	0.290	0.503

Notes: p-value reports the results of a Wald test that the two coefficients are equal. Standard errors are reported in parentheses.

The first two columns of Table 6 indicate that, in both periods, new contributors have lower savings rates than continuing contributors, although the point estimates are not statistically different from zero. This suggests that unobservable differences between the two groups may not be a problem. We report estimates of the key parameters of (15) in the last two columns of Table 6 and consider the ratio of non-IRA financial asset balances to consumption. As discussed in the previous section, this equation could be more powerful in detecting systematic differences in saving behaviour between the two groups considered but, at the same time, could present problems if individual fixed effects are important. The results in the last two columns of Table 6 indicate that the initial level of non-IRA financial asset balances (relative to consumption), keeping other variables constant, was systematically lower for new relative to continuing contributors.

For our consumption-based and asset-based tests to be valid, we must assume that new and continuing contributors do not differ in terms of their 'unobserved tastes' for saving. While we cannot test this identifying assumption, we have tested whether two measures of *observable* saving differ for the two groups. These tests yield conflicting results; new and continuing contributors do not differ statistically in terms of saving rates but do differ in terms of non-IRA financial asset balances. Moreover, as in Section 3.3, it is unclear which test is the better one. In the end, we must rely on the validity of our identifying restriction, despite the concern that the two groups may differ in unobservable ways that are correlated with saving and consumption behaviour.

5.4. *Changes in the Composition of New Contributors*

One might argue that the reason there was much less reshuffling of existing saving into IRAs after 1986 than before is because, at this time, the composition of IRA contributors changed substantially. If the change in the eligibility rules shifted the composition of IRA contributors towards lower income households which may be less likely to reshuffle, one might suspect that IRA contributions made after 1987 would be more likely to represent new saving. It does not necessarily follow, however, that the IRA contributions made after 1987 represented new saving. If they did, we should expect to see a rejection of the consumption growth based test for the 1987–90 period, which we do not. Another possible explanation for the results in the second period is that the number of contributors (and especially the number of new contributors) dropped dramatically after 1986, making the test in the second period a relatively weak one.

To check further whether the composition of new contributors differed in the 1987–90 period from that in the 1982–86 period, Table 7 reports some observable characteristics of new contributors in both periods. In the 1982–6 period, new contributors had lower average and median incomes, had lower average consumption, had less average wealth, were less likely to be married and were less likely to have senior citizens in the household than were new contributors in the 1987–90 period. Overall, the evidence is inconclusive on whether the composition of new contributors might have changed in the 1987–90 period towards households with contributions that were more likely to generate new saving. The fact

Table 7
New Contributors over Time

Mean (standard deviation)	1982–6 $N = 871$	1987–90 $N = 661$	t-test for equal of means
Real annual income	33,076.7	37,881.4	−4.320
	(19,689.3)	(23,801.3)	
Median real annual income	31,166.4	34,493.9	−1.932
Real annual consumption	27,457.3	29,557.0	−2.632
	(13,732.6)	(17,492.1)	
Median real annual consumption	25,390.9	24,990.2	0.051
Real wealth	16,979.2	21,935.7	−2.973
	(27,315.1)	(35,881.2)	
Median real wealth	6,378.5	6,679.3	−0.602
Real IRA contribution	3,169.7	2,220.5	2.166
	(10,462.7)	(4,741.6)	
Median real IRA contribution	1,931.0	1,593.6	7.597
Minority	0.043	0.046	−0.277
	(0.202)	(0.209)	
Not married	0.769	0.726	1.928
	(0.422)	(0.446)	
High school drop-out	0.099	0.077	1.483
	(0.298)	(0.266)	
High school graduate	0.296	0.296	−0.013
	(0.457)	(0.457)	
Some college	0.194	0.231	−1.762
	(0.395)	(0.422)	
College graduate	0.412	0.396	0.611
	(0.492)	(0.489)	
Age	47.29	46.57	1.135
	(12.15)	(12.43)	
Number of children	0.711	0.726	−0.291
	(1.018)	(1.033)	
Number of senior citizens	0.116	0.160	−2.056
	(0.374)	(0.456)	

Source: CEX; Notes: Statistics are computer using sampling weights from the CEX.

that many of the observable characteristics of new contributors change across these time periods, however, adds justification to our separating these effects in our empirical analysis.

A similar concern is what Bernheim (1994) – in his discussion of Engen *et al.* (1994) – called the 'dilution effect'. The argument is that households with strong unobserved tastes for saving may begin participating in a tax-favoured saving programme before other households. Thus, as the programme matures, the average assets of contributors would tend to fall over time as households with less taste for saving begin to participate. If this were the case in the IRA programme, we would expect new contributors to have slower non-IRA financial asset growth relative to continuing contributors, which is exactly what we find. However, if this dilution effect exists, we would expect to observe the non-IRA financial asset balances of IRA contributors to fall over time. Table 8 reports the ratio of IRA contributors' non-IRA financial asset balances to consumption and the level of non-IRA financial asset balances over time in. Neither series has a downward trend. Thus, it is unlikely that the dilution effect is strong enough to explain our results.

6. Assessing the Effect of IRAs on Personal and National Saving

With the help of some additional information and assumptions, we can use the point estimates from our empirical analyses to calculate what fraction of IRA contributions made during the first nine months of participation represent new national saving, new household saving and reshuffled pre-existing assets during the 1982–6 period. Obviously, if we interpret the results of our consumption test at face value, our estimate of the effect of the IRA programme on national saving is zero or negative. If, however, measurement error in consumption inflates our standard errors, one might want to use the point estimates from Table 3 even if they are statistically not different from zero. Moreover, one can use the coefficients in Table 5 to compute an alternative measure of the effect

Table 8

Non-IRA Financial Asset Balances of IRA Contributors

Year	Ratio of non-IRA financial asset balances to consumption	Levels of real non-IRA financial asset balances
1982	1.25	17,782.7
1983	1.25	16,901.4
1984	1.00	16,014.2
1985	1.48	18,071.6
1986	1.22	20,692.6
1987	1.32	20,190.6
1988	1.32	19,535.4
1989	1.23	19,432.4
1990	1.03	17,350.4

Source: CEX.
Note: Statistics are computed using sampling weights from the CEX.

of the IRA programme. The fact that we find substantial amounts of reshuffling when households begin participating in the IRA programme does not preclude their being some new saving and 'optimal' portfolio adjustment. In light of this, we attempt to divide the contributions made by new households into reshuffled saving, new saving and reductions in tax liabilities (ie, reductions in government saving).

We can calculate an estimate of the amount of new national saving, new household saving and reshuffled assets using three methods. In the first, we use the point estimate on the change in consumption (from Table 3 panel *b*, Column (5)), despite its insignificance. In the second, we use our point estimate of the difference in the change in non-IRA financial assets between new and continuing contributors (from Table 5, panel *b*, Column (4)) along with a calculation of the average marginal tax rate for new contributors. In the third, we use the estimates from panel *b* in both Table 3 and Table 5. Comparisons between these three methods are also indicative of the degree of consistency between the results of the consumption and asset-based tests.

Our first method uses the point estimate on the change in consumption from Table 3, panel *b*, Column (5). While this estimate is not significantly different from zero, this calculation should conservatively measure the amount of reshuffling and provide a 'generous' estimate on the amount of new national saving. This point estimate is –275 which implies that only 9% (= 275/3,170) of IRA contributions represent new national saving. The average contribution for new contributors in the 1982–6 period was \$3,170 while the average marginal (federal) tax rate for new contributors was 31%.[20] We conservatively add an additional 4% to account for state taxes. Thus, our estimate of the average reduction in tax liabilities is \$1,110 (= 35% of \$3,170). In this method, we do not use our estimate of the amount of reshuffled non-IRA financial assets but assume they make up the remainder of the IRA contribution. These calculations imply that 44% (= (1,110 + 275)/3,170) of IRA contributions represent new household saving of which only 9% (= 275/3,170) represent new national saving. The remainder (56%) represents reshuffled assets. These calculations are outlined in first two columns of Table 9.

Our second method is to use the point estimate on the difference in changes in non-IRA financial assets between new and contributing contributors from Table 5, panel *b*, Column(4) – \$1,957. If continuing contributors are not reshuffling existing assets but are reallocating the flow of saving, this amount can be interpreted as the amount of reshuffling linked to new IRA contributions. Thus, the fraction which are reshuffled assets is 62% (= 1,957/3,170). Adding the estimate of the effect of participation in the IRA programme on tax liabilities to that on assets yields \$3,066.9, a figure that accounts for 97% of the average IRA contribution. These calculations are summarised in the middle two columns of Table 9. Thus, the fraction of IRA contributions made by new contributors which are new national saving is estimated to be 3%. Since the

[20] We follow the procedure described by Maki (1997) to calculate marginal tax rates for the CEX households. We are grateful to Dean Maki for providing us with his computer code.

Table 9

The Composition of IRA Contributions

	Method 1 (based on coefficient from Table 3, panel *b*)		Method 2 (based on coefficient from Table 5, panel *b*)		Method 3 (based on coefficients from panel *b* in both Table 3 and 5)	
	Real level	Percentage of average contribution	Real level	Percentage of average contribution	Real level	Percentage of average contribution
Average IRA contribution	3,170	100	3,170	100	3,170	100
Reduction in consumption	275	9	103*	3	275	9
Reduction in non-IRA Financial assets	1,785*	56	1,957	62	1,957	62
Reduction in tax liabilities	1,110	35	1,110	35	938*	38
Percent of IRA Contributions that are:						
New household saving†	44		38		38	
Of which is new national saving	9		3		9	
Reshuffled assets	56		62		62	

*This number is calculated as the remainder.
†New household saving is the reduction in consumption plus the reduction in tax liabilities.

change in tax liabilities represent new household saving even if they do not represent new national saving, the faction of IRA contributions that represent new household saving is 38%.

Our third method uses both the point estimate on the change in consumption from Table 3, panel *b*, Column (5) and the point estimate on the change in non-IRA financial assets from Table 5, panel *b*, Column (4). In this method, we do not use our estimate of the average reduction in tax liabilities but assume it to be the remainder. These calculations again imply that 62% (= 1,957/3,170) of IRA contributions represent reshuffled assets. The rest, 38%, represent new household saving of which 9% (= 275/3,170) represent new national saving. These calculations are outlined in last two columns of Table 9.

One can learn two things from Table 9. First, the effect of the IRA programme on personal and national saving has been modest. Second, the three methods yield roughly similar estimates, showing that the results from our consumption test and our asset test are consistent with each other.

7. Conclusion

We approach the question of whether the IRA programme increased national saving using a new yet simple approach. Using consumption data from a sample of households that either just opened an IRA account or already had an IRA account, we can directly test whether the IRA programme led to increased national saving by determining whether newly contributing households decreased their consumption when they began participating in the programme.

We find no evidence that these households decreased their consumption in the first nine months after they started contributing to an IRA plan. While it is possible to argue that they subsequently decreased their consumption (and increase their saving), it is not possible for them to do so in the context of a life-cycle model. In any case, our sample precludes us from investigating this possibility.

The data we use also enable us to examine the changes in the non-IRA financial assets of IRA contributors, as much of the literature on the effects of the IRA programme has done. While this test is subject to a number of caveats, we find that new contributors, on average, reduced their non-IRA financial assets by over $1,900 relative to continuing contributors who are more likely to have reached an equilibrium portfolio allocation.

While participation in the IRA programme may have increased household saving (through reduced tax liabilities), our estimates based on our consumption test indicate that there was little or no increase in national saving in the 1982–6 period. Using our largest (but still insignificant) point estimate of new contributors' reduction in consumption, we estimate that at most 9% of the IRA contributions of new contributors represent new national saving. This estimate is on the low end of the estimates found in the literature. Using our point estimate of new contributors' reductions in non-IRA financial assets, we estimate that only 3% of the IRA contributions of new contributors represent new national saving.

Our estimates from our asset test do seem to indicate, however, that there were some differences before and after the 1986 tax reform and to suggest that there was less reshuffling in the post-1986 period. The results are mixed, however, as the estimates from our consumption test indicate that new contributions were not new saving in this latter period as well as in the pre-1986 period. This interesting result does suggest that it is worth asking the question of whether tax incentives for saving are, in general, more effective for lower income households (Engen and Gale, 2000).

A legitimate question is whether 9% of IRA contributions being new national saving is, in fact, a large enough percentage to justify the IRA programme. Hubbard and Skinner (1996) suggest that it might be. However, the main contribution of this study is that, first, using data on consumption and, second, using continuing contributors as a control group for households that just opened an IRA account, one finds that only a small fraction of IRA contributions made in the first nine months of participation represent new national saving. This fraction is smaller than that found by most of the literature using only data on assets and which compares IRA contributors with non-contributors.

University College London, IFS, and NBER
University of Chicago

Date of receipt of first submission: December 1999
Date of receipt of final typescript: May 2001

Appendix

Table A1

Change in Log Total Consumption

	(1) OLS	(2) 50th	(3) 75th	(4) 90th	(5) OLS	(6) 50th	(7) 75th	(8) 90th
Constant	0.035	0.325	0.561	0.595	0.325	0.430	0.422	0.606
	(0.140)	(0.121)	(0.159)	(0.238)	(0.196)	(0.150)	(0.207)	(0.434)
New cont. 1982–6	0.011	0.002	0.031	−0.040	−0.008	0.008	0.018	−0.069
	(0.024)	(0.021)	(0.028)	(0.050)	(0.026)	(0.021)	(0.028)	(0.061)
New cont. 1987–90	0.019	0.030	0.001	−0.040	0.013	0.027	−0.002	−0.013
	(0.027)	(0.024)	(0.032)	(0.055)	(0.028)	(0.023)	(0.031)	(0.065)
Income	–	–	–	–	−0.029	−0.013	0.014	−0.008
					(0.014)	(0.011)	(0.015)	(0.034)
Change in income	–	–	–	–	0.047	0.040	0.018	0.047
					(0.012)	(0.010)	(0.013)	(0.027)
Number of children	0.010	0.012	0.011	−0.022	0.008	0.015	0.013	−0.016
	(0.010)	(0.009)	(0.013)	(0.025)	(0.011)	(0.009)	(0.012)	(0.028)
Number of seniors	−0.047	−0.038	−0.022	−0.087	−0.063	−0.039	−0.041	−0.091
	(0.035)	(0.031)	(0.042)	(0.073)	(0.035)	(0.029)	(0.040)	(0.080)
Minority	−0.060	−0.011	−0.041	0.051	−0.049	0.004	−0.046	0.011
	(0.055)	(0.048)	(0.065)	(0.108)	(0.057)	(0.046)	(0.062)	(0.136)
Married	−0.035	−0.027	−0.004	−0.009	−0.016	−0.021	−0.014	0.007
	(0.021)	(0.019)	(0.025)	(0.045)	(0.023)	(0.018)	(0.025)	(0.055)
High school graduate	0.051	0.032	0.050	0.026	0.066	0.041	0.066	0.054
	(0.034)	(0.029)	(0.040)	(0.070)	(0.035)	(0.027)	(0.038)	(0.081)
Some college	0.029	0.015	0.032	−0.004	0.048	0.021	0.048	−0.007
	(0.035)	(0.031)	(0.041)	(0.073)	(0.036)	(0.029)	(0.040)	(0.085)
College graduate	0.064	0.033	0.060	0.020	0.082	0.044	0.068	0.051
	(0.032)	(0.029)	(0.038)	(0.068)	(0.034)	(0.027)	(0.038)	(0.079)
Age: 25–34	−0.214	−0.494	−0.456	−0.136	−0.226	−0.490	−0.461	−0.086
	(0.131)	(0.113)	(0.146)	(0.215)	(0.138)	(0.103)	(0.135)	(0.264)
Age: 35–44	−0.165	−0.456	−0.445	−0.158	−0.172	−0.441	−0.451	−0.129
	(0.131)	(0.113)	(0.146)	(0.213)	(0.137)	(0.103)	(0.134)	(0.262)
Age: 45–58	−0.184	−0.441	−0.423	−0.235	−0.193	−0.429	−0.416	−0.197
	(0.130)	(0.112)	(0.145)	(0.212)	(0.137)	(0.102)	(0.134)	(0.261)
Age: 59–64	−0.136	−0.402	−0.385	−0.184	−0.143	−0.380	−0.371	−0.109
	(0.132)	(0.113)	(0.147)	(0.216)	(0.138)	(0.103)	(0.135)	(0.265)
Age: $\geqslant$ 65	−0.038	−0.351	−0.291	0.077	−0.032	−0.337	−0.286	−0.153
	(0.141)	(0.122)	(0.158)	(0.240)	(0.147)	(0.111)	(0.147)	(0.290)

Notes: Coefficients on regional, monthly and yearly dummies are not shown. Standard errors are in parentheses.

Table A2

Change in Levels of Total Consumption

	(1) OLS	(2) 50th	(3) 75th	(4) 90th	(5) OLS	(6) 50th	(7) 75th	(8) 90th
Constant	−585.6	356.5	2492.8	5346.9	1509.8	909.1	−1055	−2255.1
	(1478.1)	(627.5)	(1108.9)	(2550.1)	(2053)	(803.8)	(1122)	(4985)
New cont. 1982–6	−175.8	40.4	224.7	−203.0	−275.0	64.6	176.4	−761.7
	(257.1)	(111.4)	(199.5)	(538.9)	(272.1)	(108.6)	(147.0)	(633.5)
New cont. 1987–90	131.4	154.4	127.8	−504.0	97.1	87.0	−82.0	−454.3
	(281.4)	(121.6)	(220.0)	(592.2)	(291.8)	(116.0)	(159.5)	(669.4)
Income	–	–	–	–	−252.2	−51.4	350.5	738.0
					(143.6)	(56.5)	(88.6)	(422.3)
Change in income	–	–	–	–	322.1	186.3	8.3	−93.0
					(123.5)	(49.1)	(72.7)	(320.6)
Number of children	78.4	94.0	196.1	164.5	35.0	88.6	146.7	133.5
	(109.4)	(46.7)	(89.4)	(254.6)	(111.8)	(44.5)	(63.7)	(284.4)
Number of seniors	−643.4	−93.5	−52.0	−1147.2	−793.0	−90.3	−107.1	−1023.1
	(368.3)	(159.5)	(291.4)	(835.0)	(374.0)	(148.9)	(202.9)	(917.7)
Minority	−586.9	−51.4	−189.3	−858.0	−296.4	24.1	15.8	−352.2
	(577.0)	(248.6)	(444.0)	(1129.0)	(591.6)	(235.5)	(321.9)	(1212)
Married	−238.6	−66.7	505.7	1575.8	−40.3	−46.1	329.5	1403.1
	(224.5)	(97.0)	(176.1)	(468.0)	(238.2)	(94.6)	(132.9)	(571.3)
High school graduate	489.4	112.2	285.8	564.1	586.0	230.0	281.1	716.5
	(353.4)	(153.0)	(275.3)	(749.5)	(362.4)	(144.0)	(197.2)	(840.4)
Some college	294.1	59.4	521.5	502.4	478.9	122.3	454.3	981.4
	(368.7)	(159.6)	(287.0)	(775.4)	(379.5)	(150.8)	(206.3)	(881.7)
College graduate	528.3	194.2	924.5	1782.8	658.5	247.8	454.3	1286.3
	(343.5)	(148.8)	(266.6)	(715.2)	(357.6)	(142.3)	(206.3)	(828.8)
Age: 25–34	−1224.8	−1252.0	−3111.6	−5085.8	−1041.0	−1301.4	−2615	−4808.4
	(1383.6)	(586.2)	(1021.9)	(2281.9)	(1442)	(560.3)	(695.8)	(2746)
Age: 35–44	−615.6	−1044.5	−2839.4	−4405.7	−377.4	−1100.5	−2466	−4541.8
	(1378.9)	(584.2)	(1017.2)	(2266.2)	(1438)	(558.8)	(693.6)	(2729)
Age: 45–58	−872.6	−933.8	−2638.2	−4550.0	−673.8	−1004.1	−2214	−4497.8
	(1372.0)	(581.1)	(1012.0)	(2243.5)	(1432)	(556.0)	(689.9)	(2717)
Age: 59–64	−574.5	−781.7	−2642.0	−4545.7	−327.6	−829.0	−2127	−4393.0
	(1385.1)	(586.8)	(1023.2)	(2275.4)	(1445)	(561.4)	(698.0)	(2755)
Age: ⩾ 65	717.2	−712.3	−2274.9	−2121.3	1059.6	−801.9	−1966	−2702.7
Age: ⩾ 65	(1483.7)	(630.3)	(1108.7)	(2588.6)	(1542)	(600.7)	(759.1)	(3090)

Notes: Coefficients on regional, monthly and yearly dummies are not shown. Standard errors are in parentheses.

Table A3

Change in Savings Rate

	(1) OLS	(2) 50th	(3) 75th	(4) 90th	(5) OLS	(6) 50th	(7) 75th	(8) 90th
Constant	0.2706	0.1604	0.5095	0.7365	0.2741	0.1565	0.4716	0.7707
	(0.2765)	(0.1683)	(0.3251)	(0.3789)	(0.2646)	(0.1954)	(0.1990)	(0.4046)
New cont. 1982–6	0.1647	0.0895	0.1446	0.2257	−0.0138	0.0026	0.0073	−0.1259
	(0.0490)	(0.0338)	(0.0589)	(0.0950)	(0.0476)	(0.0404)	(0.0435)	(0.0742)
New cont. 1987–90	0.1726	0.1581	0.1725	0.3223	−0.0313	−0.0229	−0.0244	−0.0216
	(0.0533)	(0.0347)	(0.0619)	(0.1067)	(0.0500)	(0.0399)	(0.0441)	(0.0779)
Change in income	–	–	–	–	0.5246	0.5560	0.5920	0.5341
					(0.0194)	(0.0156)	(0.0192)	(0.0440)
No of children	−0.0016	−0.0039	−0.0406	−0.1210	−0.0028	−0.0070	−0.0178	−0.0998
	(0.0211)	(0.0144)	(0.0252)	(0.0403)	(0.0198)	(0.0163)	(0.0179)	(0.0311)
Number of seniors	0.0663	0.0180	0.1484	0.1002	0.0767	0.0170	0.0952	0.1061
	(0.0725)	(0.0460)	(0.0873)	(0.1520)	(0.0685)	(0.0534)	(0.0578)	(0.0914)
Minority	0.1635	−0.0885	0.1539	0.4294	0.1541	0.0259	−0.0406	0.2798
	(0.1042)	(0.0705)	(0.1265)	(0.2115)	(0.1012)	(0.0828)	(0.0871)	(0.1492)
Married	0.0608	0.0288	0.0995	0.1050	0.0402	0.0232	0.0568	0.0612
	(0.0423)	(0.0284)	(0.1265)	(0.0851)	(0.0399)	(0.0326)	(0.0357)	(0.0648)
High sch. graduate	−0.1435	−0.0623	0.0091	0.0025	−0.1323	−0.1019	0.0009	0.0172
	(0.0659)	(0.0466)	(0.0822)	(0.1317)	(0.0621)	(0.0538)	(0.0568)	(0.0991)
Some college	−0.0565	−0.0233	0.0245	−0.0051	−0.0651	−0.0097	0.0643	0.0061
	(0.0690)	(0.0483)	(0.0842)	(0.1391)	(0.0651)	(0.0560)	(0.0592)	(0.1050)
College graduate	−0.0738	−0.0164	0.0091	0.0594	−0.0837	−0.0235	0.0500	0.0976
	(0.0641)	(0.0455)	(0.0794)	(0.1316)	(0.0605)	(0.0526)	(0.0553)	(0.0986)
Age: 25–34	−0.0977	−0.0563	0.0283	0.3084	0.0509	0.0860	0.1038	0.3631
	(0.2596)	(0.1552)	(0.3061)	(0.3142)	(0.2471)	(0.1791)	(0.1798)	(0.3783)
Age: 35–44	−0.1648	−0.1251	−0.0320	0.2409	−0.0239	0.0281	0.0124	0.3030
	(0.2585)	(0.1547)	(0.3050)	(0.3107)	(0.2462)	(0.1785)	(0.1792)	(0.3785)
Age: 45–58	−0.0716	−0.0757	0.0803	0.3135	0.0250	0.0578	0.0898	0.3869
	(0.2573)	(0.1537)	(0.3048)	(0.3069)	(0.2451)	(0.1775)	(0.1783)	(0.3745)
Age: 59–64	−0.0239	−0.0458	0.0626	−0.4296	−0.0190	0.0287	0.0547	0.3780
	(0.2599)	(0.1555)	(0.3076)	(0.3149)	(0.2476)	(0.1795)	(0.1803)	(0.3782)
Age: ≥ 65	−0.3410	−0.1765	−0.2417	0.1300	−0.1990	−0.0584	−0.0635	0.0256
	(0.2792)	(0.1687)	(0.3293)	(0.3796)	(0.2655)	(0.1951)	(0.1971)	(0.3950)

Notes: Coefficients on regional, monthly and yearly dummies are not shown. Standard errors are in parentheses.

Table A4

Change in Non-IRA Financial Assets/Non-Durable Consumption

	(1) OLS	(2) 75th	(3) 90th	(4) OLS	(5) 75th	(6) 90th
Constant	−0.2098	0.0980	0.2467	−0.8477	−0.2778	−0.5830
	(0.2722)	(0.1115)	(0.3080)	(0.3658)	(0.1517)	(0.4847)
New contributor in 1982–6	−0.0737	−0.0847	−0.1212	−0.0681	−0.0775	−0.1030
	(0.0495)	(0.0211)	(0.0677)	(0.0512)	(0.0199)	(0.0694)
New contributor in 1987–90	0.0421	−0.0471	−0.0199	0.0413	−0.0293	−0.0087
	(0.0545)	(0.0232)	(0.0751)	(0.0561)	(0.0220)	(0.0775)
Change in annual income	–	–	–	0.000	0.000	0.000
				(0.000)	(0.000)	(0.000)
Annual income	–	–	–	0.066	0.036	0.087
				(0.025)	(0.012)	(0.047)
Number of children	−0.0165	−0.0235	−0.0533	−0.0166	−0.0225	−0.0536
	(0.0210)	(0.0097)	(0.0323)	(0.0215)	(0.0091)	(0.0337)
Number of senior citizens	0.0918	0.0663	0.0192	0.0996	0.0808	−0.0226
	(0.0713)	(0.0300)	(0.0974)	(0.0724)	(0.0276)	(0.0953)
Minority	−0.0066	−0.0494	−0.2385	−0.0053	−0.0677	−0.2324
	(0.1071)	(0.0448)	(0.1447)	(0.1094)	(0.0426)	(0.1423)
Married	0.0062	−0.0173	−0.0792	−0.0210	−0.0261	−0.1212
	(0.0427)	(0.0182)	(0.0592)	(0.0448)	(0.0178)	(0.0640)
High school graduate	0.0902	0.0262	0.1104	0.0864	0.0358	0.1382
	(0.0686)	(0.0291)	(0.0955)	(0.0699)	(0.0273)	(0.0976)
Some college	0.0239	0.0257	0.1245	0.0092	0.0274	0.1420
	(0.0717)	(0.0304)	(0.1001)	(0.0733)	(0.0285)	(0.1028)
College graduate	0.0963	0.0789	0.3117	0.0698	0.0784	0.3248
	(0.0669)	(0.0284)	(0.0944)	(0.0687)	(0.0269)	(0.0982)
Age: 25–34	0.0410	0.0459	0.0922	0.0329	0.0585	0.0646
	(0.2602)	(0.1065)	(0.2847)	(0.2633)	(0.0990)	(0.1332)
Age: 35–44	0.0572	0.0656	0.0966	0.0372	0.0596	0.0413
	(0.2594)	(0.1062)	(0.2844)	(0.2626)	(0.0988)	(0.1345)
Age: 45–58	0.0596	0.0161	0.1070	0.0432	0.0164	0.0741
	(0.2579)	(0.1054)	(0.2806)	(0.2611)	(0.0980)	(0.1239)
Age: 59–64	0.0708	0.0185	0.3012	0.0650	0.0247	0.2870
	(0.2605)	(0.1066)	(0.2855)	(0.2637)	(0.0991)	(0.1345)
Age: $\geqslant$ 65	−0.1382	−0.0315	0.2429	−0.1534	−0.0332	0.3332
	(0.2808)	(0.1146)	(0.3151)	(0.2842)	(0.1062)	(0.1864)

Notes: Coefficients on regional and yearly dummies are not shown. Standard errors are in parentheses.

Table A5

Change in Real Non-IRA Financial Assets

	(1) OLS	(2) 75th	(3) 90th	(4) OLS	(5) 75th	(6) 90th
Constant	−509.2	313.8	301.8	−24005.4	−11209.7	−21103.7
	(3663.6)	(2792.2)	(5314.4)	(4937.0)	(4267.7)	(10506.5)
New contributor in 1982–6	−2038.1	−1216.6	−3892.0	−1956.9	−1194.5	−3647.1
	(669.0)	(513.4)	(1106.0)	(675.8)	(500.6)	(1152.5)
New contributor in 1987–90	−1084.3	−827.9	−1387.5	1039.7	−584.5	−1197.2
	(736.5)	(570.6)	(1259.5)	(740.4)	(552.9)	(1307.2)
Change in annual income	–	–	–	−0.010	−0.012	−0.005
				(0.012)	(0.010)	(0.030)
Annual income	–	–	–	2435.0	1318.0	2137.6
				(347.5)	(346.1)	(940.4)
Number of children	−304.9	−512.7	−1179.3	−322.8	−517.2	−1184.3
	(283.2)	(224.9)	(495.9)	(282.7)	(211.8)	(507.4)
Number of senior citizens	194.2	−192.6	337.0	469.3	147.0	927.6
	(975.4)	(761.3)	(1695.6)	(970.7)	(732.0)	(1607.8)
Minority	−2340.8	−1480.9	−4050.2	−2225.3	−2081.9	−3882.1
	(1457.0)	(1094.8)	(2356.1)	(1454.4)	(1073.8)	(2529.2)
Married	459.6	1071.8	1818.6	−542.6	309.5	702.3
	(574.6)	(445.1)	(1012.6)	(594.2)	(459.0)	(1113.0)
High school graduate	1047.0	1447.9	1997.3	929.1	651.0	2918.8
	(924.9)	(721.9)	(1666.1)	(924.1)	(697.3)	(1692.4)
Some college	2481.9	1775.3	2995.1	2180.4	1350.3	3214.8
	(966.9)	(753.5)	(1752.2)	(968.6)	(729.5)	(1799.1)
College graduate	4418.4	3567.6	9565.1	3680.2	3072.0	9810.1
	(899.8)	(703.9)	(1658.0)	(908.1)	(687.9)	(1724.3)
Age: 25–34	1428.4	484.4	3670.9	1169.8	−216.7	3304.6
	(3499.2)	(2654.6)	(4935.4)	(3475.3)	(2549.7)	(5038.1)
Age: 35–44	2278.8	1870.6	5892.2	1588.2	667.7	4901.1
	(3499.2)	(2647.0)	(4921.6)	(3466.9)	(2543.5)	(5030.6)
Age: 45–58	2645.7	905.2	5713.3	1974.9	293.6	5453.6
	(3469.6)	(2628.7)	(4872.7)	(3446.8)	(2527.4)	(4981.2)
Age: 59–64	3472.1	1583.5	7158.3	3197.7	1013.5	6679.6
	(3506.0)	(2660.3)	(4972.8)	(3482.7)	(2556.4)	(5076.3)
Age: ⩾ 65	5745.6	2102.5	12025.4	5180.9	670.7	11906.9
	(3790.2)	(2879.9)	(5536.5)	(3764.2)	(2768.7)	(5572.1)

Notes: Coefficients on regional and yearly dummies are not shown. Standard errors are in parentheses.

References

Attanasio, Orazio P. (1994). 'Personal saving in the U.S', in (James M. Poterba, ed.), *International Saving Comparisons*, Chicago: University of Chicago Press, pp. 57–123.

Banks, James, Dilnot, Andrew and Low, H. (1994). 'The distribution of wealth in the U.K.', mimeo, Institute for Fiscal Studies.

Banks, James and Tanner, Sarah (1996). 'Savings and wealth in the U.K.: evidence from micro-data', *Fiscal Studies*, vol. 17(2), pp. 37–64.

Bernheim, B. Douglas (1994). 'Comment on: do saving incentives work?', *Brookings Papers on Economic Activity*, pp. 152–66.

Bernheim, B. Douglas (1997). 'Rethinking saving incentives', in (Alan J. Auerbach, ed.), *Fiscal Policy: Lessons from Economic Research,* Cambridge, MA: MIT Press, pp. 259–313.

Burman, Leonard, Cordes, Joseph and Ozanne, Larry (1990). 'IRAs and national savings', *National Tax Journal*, vol 43, pp. 259–83.

Engen, Eric M. and Gale, William G. (1993). 'IRAs and saving in a stochastic life-cycle model', mimeo, UCLA.

Engen, Eric M. and Gale, William G. (2000). 'The effects of 401(k) plans on household wealth: differences across earnings groups', Working Paper no. 8032, National Bureau of Economic Research.

Engen, Eric M., Gale, William G. and Scholz, John Karl (1994). 'Do savings incentives work?', *Brookings Papers on Economic Activity*, pp. 85–151.

Engen, Eric M., Gale, William G. and Scholz, John Karl (1996). 'The illusory effects of saving incentives on saving', *Journal of Economic Perspectives*, vol. 10(4), pp. 113–38.

Feenberg, Daniel and Skinner, Jonathan (1989). 'Sources of IRA savings', in (Lawrence H. Summers, ed.), *Tax Policy and the Economy*, Cambridge, MA: MIT Press, pp. 25–46.

Gale, William G. and Scholz, John Karl (1994). 'IRAs and household savings', *American Economic Review*, vol 84, pp 1233–60.

Hubbard, R. Glenn and Skinner, Jonathan S. (1996). 'Assessing the effectiveness of saving incentives', *Journal of Economic Perspectives*, vol. 10(4), pp. 73–90.

Joines, Douglas H. and Manegold, James G. (1995). 'IRAs and saving: evidence from a panel of taxpayers', mimeo, University of Southern California.

Kotlikoff, Laurence J. (1990). 'The crisis in U.S. savings and proposals to address the crisis', *National Tax Journal*, vol. 43, pp. 233–46.

Maki, Dean (1997). 'Household debt and the Tax Reform Act of 1986', mimeo, Board of Governors of the Federal Reserve System.

Nelson, Julie (1994). 'On testing for full insurance using consumer expenditure survey data: comment', *Journal of Political Economy*, vol. 102, pp. 384–94.

Ozanne, Larry J. (1994). 'Past experience and current proposals for IRAs', in (M. Kosters, ed.), *Personal Saving, Consumption, and Tax Policy*, Washington, DC: AEI Press.

Poterba, James M., Venti, Steven F. and Wise, David A. (1996). 'How retirement saving programs increase saving', *Journal of Economic Perspectives*, vol. 10(4), pp. 91–112.

Poterba, James M., Venti, Steven F. and Wise, David A. (1997). 'A personal retirement saving programs and asset accumulation: reconciling the evidence', in (David A. Wise, ed.), *Frontiers in the Economics of Aging*, Chicago: University of Chicago Press, pp. 23–106.

Skinner, Jonathan (1991). 'Individual retirement accounts: a review of the evidence', Working Paper no. 3938, National Bureau of Economic Research.

Venti, Steven F. and Wise, David A. (1986). 'Tax-deferred accounts, constrained choice, and estimation of individual saving', *Review of Economic Studies*, vol. 53, pp. 579–601.

Venti, Steven F. and Wise, David A. (1987). 'IRAs and saving', in (Martin A. Feldstein, ed.), *The Effects of Taxation on Capital Accumulation*, Chicago: University of Chicago Press.

Venti, Steven F. and Wise, David A. (1990). 'Have IRAs increased U.S. savings? Evidence from consumer expenditure surveys', *Quarterly Journal of Economics*, vol. 105, pp. 661–98.

Venti, Steven F. and Wise, David A. (1991). 'The saving effect of tax-deferred retirement accounts: evidence from SIPP', in (B. Douglas Bernheim and John B. Shoven, eds.), *National Savings and Economic Performance*, Chicago: University of Chicago Press.