

EFFECTIVENESS OF TAX INCENTIVES TO BOOST (RETIREMENT) SAVING: THEORETICAL MOTIVATION AND EMPIRICAL EVIDENCE

Orazio P. Attanasio, James Banks and Matthew Wakefield

TABLE OF CONTENTS

Introduction	146
How are tax incentives supposed to work? a simple theoretical outline.....	147
Evidence on the effectiveness of individual retirement accounts in the United States.....	150
The UK experience	155
Aggregate evidence on the effectiveness of TESSAs	157
Micro-evidence on the effectiveness of ISAs	159
Conclusions and thoughts for further research	166
<i>Appendix.</i> Analysis at the individual level.....	169
Bibliography	172

The authors are from University College London, Institute for Fiscal Studies and National Bureau of Economic Research; University College London and Institute for Fiscal Studies and Institute for Fiscal Studies, respectively. Data from the Family Resources Survey (FRS) were made available by the Department for Work and Pensions; data from the British Household Panel Survey (BHPS) were made available through the UK Data Archive, which is based at the University of Essex and funded by the Economic and Social Research Council. These providers bear no responsibility for the interpretations of the data presented here.

INTRODUCTION

With ageing populations and shorter working lives generating increasing pressure on public finance and pension systems, the adequacy of household saving for retirement has become a policy issue all around the world. As successive reform to public pensions makes such systems less generous, there is an increasing role for individual provision in retirement income, which in turn generates a need for households and individuals to save more.

One way governments have traditionally tried to encourage retirement saving is by providing special tax arrangements, or tax incentives, for funds held in particular types of saving accounts. Whilst these tax advantaged savings products may only provide a tax treatment for saving that is neutral with respect to consumption, it is usually the case that they still provide a tax advantage over other forms of saving such as interest bearing accounts, direct holdings of equity, or intermediated products such as unit trusts or investment trusts and therefore increase the (marginal) rate of return to saving.

In order for such tax incentives to achieve their stated goal of increasing or encouraging saving, it must be the case that the saving level responds positively to an increase in the net rate of return. If that has to be the case, the funds going into such accounts need to have come from individuals reducing their consumption levels as opposed to simply moving money from one form of saving to another. As a result the simple size and take-up of such schemes is not an indicator of their effectiveness as a policy instrument in increasing personal and household saving. Instead the evaluation of their success is somewhat more complicated, as it should consider what the same households that are taking up these schemes would have saved in their absence. Moreover, if one is interested in the effect of these schemes on *national* rather than personal savings, one has also to consider the implications for government saving that is affected through the reduced tax liabilities implied by these schemes. Obviously, these are not simple issues and have been the subject of much empirical research.

In this paper we consider the empirical evidence from two countries, the United States and the United Kingdom, regarding the extent to which funds in some specific tax advantaged accounts represent new saving. Both countries have historically been at the forefront of those adopting tax incentives for saving, with the first products being introduced over fifteen years ago. We begin by setting out

a broad and simplified theoretical framework in order to fix ideas, and then move on to consider the US and the UK experiences in turn.

HOW ARE TAX INCENTIVES SUPPOSED TO WORK? A SIMPLE THEORETICAL OUTLINE

In this section, we provide a simple theoretical framework which we find useful to organise our ideas about the economic rationale for tax incentives to saving. In addition, we briefly discuss other possible channels through which saving schemes that have a tax incentive might be helpful in increasing saving, independently of the strictly economic effects.

Introducing a tax advantage to a saving instrument is, in most cases equivalent to a modest increase in the rate of return to that particular asset. The reason why we say “modest” is that many of these schemes are associated with a tax deferral or at most a partial tax exemption, rather than being entirely exempt from taxes. That is, individuals might be exempt from taxes on the income that is contributed to a given scheme (often within a certain limit), but they are subsequently taxed when resources are drawn from that asset, typically at retirement. Alternatively, if interest payments, capital gains, and withdrawals are untaxed, then contributions to such accounts typically need to be made out of taxed income. The tax favoured accounts we consider in the empirical sections of this paper encompass both types – the former design is that followed by the Individual Retirement Account in the US; the latter is that of the TESSA and ISA in the UK and the “Roth IRA” in the US. Since the two designs are equivalent if marginal tax rates are constant over the life-cycle we do not distinguish between the two in this section.

An increase (modest or otherwise) in the rate of return on a specific asset will have two effects on the decisions of a saver. On the one hand, it might change the total amount that an individual will save. On the other, given the amount saved, it will change the composition of the optimal portfolio and, presumably, will shift resources towards that particular asset.

Whether a tax incentive works in stimulating savings concerns mainly the first effect, that is, if the total amount saved is larger than it would have been in the absence of the tax incentive. In a standard model, the answer to this question is in general, ambiguous. As is well known, there is an income and a substitution effect at play. On the one hand present consumption is more expensive relative to future consumption, which should increase current saving. On the other, the amount that it is necessary to save to achieve a given level of wealth is reduced. When there are more than two periods, there is also a wealth effect that tends to increase savings. To fix ideas, consider the following stylised model. Our consumers live for four periods and face no uncertainty. 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, A_1 . Preferences are given by an intertemporally separable constant relation risk-aversion utility function. The typical household receives income y_1 and y_2 in periods 1 and 2 and nothing in period 3. For simplicity we assume that the discount factor equals the interest rate. This implies that, before tax-free schemes became available, the consumption profile is flat: $c_1 = c_2 = c_3$. To formalise, we have the maximization problem:

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

subject to

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

A first order condition for this problem is given by the following:

$$c_{t+1}^{-\rho} = \frac{(1+\delta)}{(1+r)} c_t^{-\rho}; t = 1,2,3$$

Using these equations and the intertemporal budget constraint one can solve for the level of consumption in each period. Consumption in period 1 is given by the following expression:

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

Equation 1 is useful because it relates consumption in period 1 to the level of resources available to the individual in the future (income in period 2 discounted by the interest rate) and those currently available to her. The latter consist of current income and current assets (which reflect decisions taken in the past). From expression [1] it is apparent that the effect of changes in the interest rate is, in principle, ambiguous. In the case of log utility ($\rho = 1$), income and substitution effects in the denominator exactly cancel out, so that one is only left with the wealth effect (the fact that future income is discounted more heavily at high levels of the interest rate), so that consumption goes down. When the elasticity of intertemporal substitution is higher than 1, the effect of an increase in the interest rate is unambiguously negative.

While this simple model is very stylised it can be used to interpret some of the evidence we present in the next section. There we compare the rate of

consumption growth for consumers who contribute to a tax-favoured scheme in both of the periods that we consider, to that for consumers that start contributing after the first observation of consumption. We will call them “continuing” and “new” contributors, respectively. In terms of our simple model, we compare a consumer who has already contributed to the tax favoured scheme in period 0 to one that starts contributing in period 1. Such a consumer experiences a change of interest rate in period 1 of their life. If there were no contribution limits, such a consumer would shift all of her existing assets to the tax favoured scheme. But she will also change her consumption. The direction and size of this change depends on the size of the intertemporal elasticity of substitution (ρ), but if the tax-favoured scheme is to increase saving then it must be negative. On the other hand, a household which is already contributing to a tax free scheme will not change consumption as it will not experience a change in the interest rate. 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 the tax-favoured scheme is effectively increasing saving. Moreover old contributors would have already shifted existing assets to the tax-favoured scheme. Therefore we are likely to observe a smaller rate of growth (even a negative one) for non-tax favoured financial assets for new contributors than for continuing contributors.

These arguments were used by Attanasio and DeLeire (2002) in their study of the effect of Individual Retirement Accounts (IRAs) in the US, one of the first tax favoured schemes introduced. In the following section we reproduce some of their results. The effect of participation in the IRA program on consumption is ambiguous and depends upon the elasticity of intertemporal substitution. The model, however, can also be used to study the possible effects of different tax favoured schemes. And one can easily introduce contribution limits and the like to the basic structure. The main message that comes out of this framework is that, in terms of economic incentives, the effect that tax favoured schemes might have on consumption is bound to be relatively small and is not even unambiguously positive. Of course this is not to say that tax-favoured schemes of a different nature could not have any other effects on saving. One might want to consider different channels through which these schemes could operate in generating new saving.

First, there is the information and economic literacy channel. One example of a tax-favoured scheme using this channel is those 401(k) accounts in the US that provide financial education and information through employer provided workshops and seminars. These might be effective in increasing saving if some participants in the scheme are made more aware of their financial needs in retirement and of the necessity to direct extra resources to retirement saving. While a discussion of these issues is beyond the scope of this paper, we should note that schemes that work through extensive use of financial literacy seminars and work-

shops over a prolonged time period can be very expensive. Second, if one believes that the model above does not constitute a good approximation to individual behaviour because many individuals behave irrationally, then alternative justifications for the schemes can be important. For instance, if self-control problems contribute to the “inadequacy” of retirement saving patterns, then one might think that consumers with these problems may be induced to save more in such schemes since the penalties typically involved in withdrawing funds from a tax-favoured scheme provide a commitment mechanism. While these ideas move away in a substantive fashion from the model we have considered here, in that they assume that consumers are not fully rational in their intertemporal endeavours, they have received some support in the literature, especially in the case of 401(k)s. The implications of such channels for the design of optimal incentives are not completely obvious. The rate of return advantage, for instance, might not be as crucial as often implied in the policy debate.

EVIDENCE ON THE EFFECTIVENESS OF INDIVIDUAL RETIREMENT ACCOUNTS IN THE UNITED STATES

In this section, we reproduce some of the evidence in Attanasio and DeLeire (2002) on the effectiveness of IRAs in generating new savings. While IRAs had been available for a number of years in the US to self-employed individuals, they were made available to all households in 1982. The scheme involves a considerable tax advantage in that contributions to an IRA are exempt from income taxes. There are however, limits to the contributions that could not exceed \$2 000 per year for a working individual and \$2 250 for a couple with only one employee.

As we mentioned above, Attanasio and DeLeire (2002) compare both the change in consumption and the change in non-IRA financial assets of the “new” and “continuing” contributors. They use data from the Consumer Expenditure Survey, a continuous and reasonably comprehensive survey started in 1980. In this survey each household stays in the sample for four consecutive periods each lasting a quarter of a year, and reports expenditure on all consumption items over the most recent period in a retrospective interview at the end of the quarter. The survey is particularly useful for the purpose at hand, because it is available around the time at which IRAs were made universally available. In the sample, therefore, we are able to see a relatively large number of individuals who join the scheme between the first and fourth interview. One can therefore compare how consumption (and assets) change for “continuing” and “new” contributors. As the legislation changed again in 1986, we focus in particular on the period 1982-1986. Indeed, the first few years were the most interesting both because one can find many “new” contributors in the data as the scheme started to become popular and because differences between “new” and “continuing” contributors in this period were prob-

ably due to information or other random events and are less likely to be driven by systematic differences in consumption growth that might obscure the results.

The empirical specification Attanasio and DeLeire (2002) use to implement the test based on consumption consists of the following simple equation:

$$\Delta \ln C_i = \beta^c X_i^c + \gamma_i^c d_i^{new} \delta_i + \gamma_2^c d_i^{new} (1 - \delta_i) + \varepsilon_i^c \quad [2]$$

where $\Delta \ln(C_i^t)$ is the change in log real consumption for household i observed at time t . 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 and δ_i equals one if household i is observed between 1982 and 1986 and zero if it is observed after 1986. Under the null that the IRA incentives did not generate new saving, the coefficients γ^c should be zero. We allow the coefficients to be different before and after 1986 to reflect changes in the IRA legislation in that year. While equation 2 is in logs, and therefore focuses on proportional changes in consumption, in the empirical application we also estimate it in levels.

In Table 1a, we report the results for the log specification. In addition to the OLS results, we also report the results on three quantile regressions: the 50th, 75th and 90th. For each of these, we report the results of a specification that controls for income and income growth and one that does not. While these variables might be endogenous, it might also be argued that people start contributing to an IRA when they receive a positive income shock.

The coefficient on new contributors is *not* significantly different from zero in any of the specifications that we report. Furthermore, in most specifications the point estimates are positive rather than negative. According to these results,

Table 1a. **Change in log total consumption**

	Not controlling for income and income growth				Controlling for income and income growth			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	50th	75th	90th	OLS	50th	75th	90th
New contributor in 1982-1986	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 1987-1990	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

Note: P-value reports the results of a Wald test that the two coefficients are equal. Standard errors are reported in parentheses. Controls include: number of children; number of seniors; education dummies; race dummy; age dummies; regional, monthly and yearly dummies are not shown.

therefore, we can resoundingly reject the hypothesis that IRAs create new saving: consumers who start contributing to one do not seem to be reducing their consumption. In Table 1b, we report the results for the specification in levels, which are broadly consistent with those in Table 1a. The only exception is the point estimate for the OLS specification, which is negative. However, the coefficient is small from an economic point of view (\$200) and not significantly different from zero.

Table 1b. **Change in levels of total consumption**

	Not controlling for income and income growth				Controlling for income and income growth			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	50th	75th	90th	OLS	50th	75th	90th
New contributor in 1982-1986	-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 1987-1990	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. Controls include: number of children; number of seniors; education dummies; race dummy; age dummies; regional, monthly and yearly dummies are not shown.

In Table 2, we report the results for a regression that puts saving rates (rather than the change in consumption) on the left-hand side. Saving rates are obtained as income minus consumption divided by income. The change in the savings rate combines both the change in consumption and the change in tax liabilities when we control for income growth. In order to isolate the change in tax liabilities and

Table 2. **Change in savings rate**

	Not controlling for income growth				Controlling for income growth			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	50th	75th	90th	OLS	50th	75th	90th
New contributor in 1982-1986	0.165 (0.049)	0.089 (0.034)	0.145 (0.059)	0.226 (0.095)	-0.014 (0.048)	0.003 (0.040)	0.007 (0.043)	-0.126 (0.074)
New contributor in 1987-1990	0.173 (0.053)	0.158 (0.035)	0.173 (0.062)	0.322 (0.107)	-0.031 (0.050)	-0.023 (0.040)	-0.024 (0.044)	-0.022 (0.078)
P-value	0.911	0.150	0.739	0.489	0.796	0.648	0.605	0.322

Note: P-value reports the results of a Wald test that the two coefficients are equal. Standard errors are reported in parentheses. Controls include: number of children; number of seniors; education dummies; race dummy; age dummies; regional, monthly and yearly dummies are not shown.

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 program) 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 the change in saving out of after-tax income is on average 17 per cent 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.

The figures in Tables 1 and 2 are interesting in that they are the only results in the literature on the effectiveness of tax incentives to saving that consider directly the effect of these schemes on consumption. As we mentioned above, however, one can also look at what happens to the change in non-IRA financial assets for new and continuing contributors. In Table 3, we report Attanasio and DeLeire's results for a specification that relates the change in non-IRA financial assets to a number of controls and the "new-contributor" dummy.

The results indicate that new contributors accumulate non-IRAs financial assets significantly less than continuing contributors. The OLS estimates are particularly suggestive, as the difference is about \$2 000, which is the limit to annual IRA contributions.

Table 3. **Change in non-IRA financial assets**

	Not controlling for income and income growth			Controlling for income and income growth		
	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	75th	90th	OLS	75th	90th
New contributor in 1982-1986	-2 038.1 (669.0)	-1 216.6 (513.4)	-3 892.0 (1 106.0)	-1 956.9 (675.8)	-1 194.5 (500.6)	-3 647.1 (1 152.5)
New contributor in 1987-1990	-1 084.3 (736.5)	-827.9 (570.6)	-1 387.5 (1 259.5)	-1 039.7 (740.4)	-584.5 (552.9)	-1 197.2 (1 307.2)
P-value	0.3292	0.6058	0.1274	0.3497	0.4013	0.1471

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

This second test is also particularly interesting as a complement to the consumption based test because it sets the hypothesis of “no new saving” as the alternative, while the consumption based test set it as the null. As it happens, the results of the two tests are both consistent with the hypothesis that the IRAs were ineffective in creating new savings.

Before moving on to evaluating the likely effect that this legislation has had on national savings, it is worth mentioning that the evidence above has sometimes been misinterpreted. In particular, the evidence on the changes in assets has sometimes been interpreted along the following lines. The figures in Table 3 show that continuing contributors accumulate non-IRA assets at a faster rate than new contributors. As new contributors will become continuing contributors (there is evidence that shows that most new contributors will keep contributing) they will start accumulating wealth in the same manner as continuing contributors and therefore save more. Therefore IRAs work as an incentive. Such reasoning is flawed. The only thing that Table 3 shows is that the stock of non-IRAs assets grows less for new contributors than for continuing contributors. The most likely cause is that new contributors are re-shuffling existing portfolios into the IRAs. Continuing contributors may eventually exhaust this source of funds and so begin to substitute within their flow of saving. The crucial evidence is that in Tables 1 and 2: they show that new contributors do not decrease consumption and do not increase saving rates. This directly illustrates an absence of the effects necessary if the tax-favoured scheme is to generate new saving.

The estimates we obtained can be used, together with some strong assumptions, to estimate the effect that the IRA legislation has had on national saving. As we have two different estimates, one obtained from the consumption (and saving) specification and one from the asset specification, we can also check, as an additional test of our approach, whether the figures we obtain are internally consistent. This is what we do in Table 4.

We know that, in our sample, the average contribution to IRAs was equal to \$3 170. These must come from: *i*) a reduction in consumption; *ii*) a reduction in tax liabilities; *iii*) a reduction in pre-existing financial assets. We are able, to impute the marginal tax rate for each household in our sample, and therefore to compute the reduction in tax liabilities induced by their IRA contribution. This turns out to be \$1 110. The most generous estimates for the reduction in the level of consumption (\$275, although not significantly different from zero), implies (together with the other two figures) a reduction in non-IRAs financial assets of \$1 785, which compares to our estimate of \$1 957. Equivalently, if we start from our estimates of the reduction in non-IRA assets, we find an implied reduction in consumption of \$103, which is remarkably similar to our estimates.

Table 4. The composition of IRA contributions

	Method 1 (based on coefficient from the consumption level specification)		Method 2 (based on coefficient from the asset specification)		Method 3 (based on coefficients from the consumption and asset specifications)	
	Real level	% of average contribution	Real level	% of average contribution	Real level	% of average contribution
Average IRA contribution	3 170	100	3 170	100	3 170	100
Reduction in consumption	275	9	103 ¹	3	275	9
Reduction in non-IRA financial assets	1 785 ¹	56	1 957	62	1 957	62
Reduction in tax liabilities	1 110	35	1 110	35	938 ¹	38
<i>Per cent of IRA contributions that are:</i>						
New household saving ² of which is:		44		38		38
New national saving		9		3		9
Reshuffled assets		56		62		62

1. This number is calculated as the remainder.

2. New household saving is the reduction in consumption plus the reduction in tax liabilities.

From these simple computations we can draw two lessons. First, the evidence from the two tests presented is remarkably consistent, making them credible. Second, the contributions of the IRA legislation to new saving is minimal.

THE UK EXPERIENCE

The UK has an extended history of using tax incentives to encourage saving in various forms. Most important for our purpose is the evidence concerning two particular tax-exempt savings vehicles: Tax Exempt Special Savings Accounts (TESSAs) and Individual Savings Accounts (ISAs). Both of these products allow individuals to save over a short time horizon, as well as to fund their retirement. What the schemes have in common with Individual Retirement Accounts is that they allow income and capital gains accruing to funds held in the accounts to be received free of tax, and so aim to promote private saving through increases in the net rate of return.

TESSAs were introduced in 1991 and were subsequently replaced by ISAs in 1999 (*i.e.* no new TESSAs could be opened once ISAs had been introduced). Over the intervening period the remainder of the savings environment was, in the context of recent UK savings policy, relatively stable. In particular, the three-tier system of pension provision was already established. The first tier was comprised of a Basic State Pension, supplemented by means-tested benefits for those with low entitlements and/or little additional income. At the second tier workers could choose between a state earnings related pension or (tax privileged) private provi-

sion either in the form of a personal pension (introduced in 1988) or, where available, an occupational scheme. The third tier consisted of additional voluntary savings to supplement provision from the lower tiers. Though not strictly speaking *retirement* savings vehicles, TESSAs and ISAs are accounts opened voluntarily and over and above any wealth accumulated through the pension system, and so can be considered as part of the third tier. The stability of other savings arrangements, coupled with the relative short time horizons within which we analyse the effects of the TESSAs and ISAs, allows us to consider each policy reform largely in isolation from other institutional factors.

When TESSAs were introduced it was already possible to make some short-term savings in a tax-privileged form using a Personal Equity Plan (PEP). As the name suggests, funds held in PEPs had to be held in equities although these could be held either directly or in trust. Contributions into PEPs were paid from net income, but relative to other means of holding shares the accounts were tax privileged since any interest income or capital gains accruing to the fund was tax exempt.¹ Contributions to TESSAs were also paid from net income but in contrast to PEPs these accounts provided tax relief for interest income accruing to funds held in designated bank or building society deposit accounts. Also unlike PEPs, this tax advantage could be received only if the capital remained untouched for five years; early withdrawals would pay back the tax advantage but attracted no further penalty.


ISAs replaced both TESSAs and PEPs from April 1999. The ISA is a tax privileged savings vehicle for cash deposits, or for holdings of stocks and shares either directly or in trust, or for both cash and equities. Like both TESSAs and PEPs, contributions to ISAs are paid from net income. The absence of a statutorily fixed holding period is the main difference between a cash ISA and a TESSA. The option of holding cash or safe interest bearing accounts and so avoiding stock market risk is what differentiates the product from a PEP. Like TESSAs and PEPs before them, ISAs were restricted in terms of the total amount that could be invested in an account in any one year. Indeed, the introduction of ISAs actually reduced the maximum amount that an individual could save in a tax advantaged non-pension form. An individual holding a TESSA and a PEP could save up to £7 800 a year in these accounts, or £9 000 if the TESSA was in its first year.² ISA saving is limited to £7 000 per year, of which at most £3 000 can be in cash. However, as evidence indicates that the median holding of non-pension financial wealth amongst families in the UK was around £1 000 in 2000 (Banks, Smith and Wakefield, 2002), even the stricter contribution limits of ISAs are unlikely to constrain many savers, particularly amongst those with lower incomes who were a target for the new policy.³ ISAs can be seen as giving expenditure tax treatment on all accessible savings for the majority of households in the UK.

In the context of a discussion of tax-advantaged products and their impact on savings, both TESSAs and ISAs can be seen as an expansion of tax exemptions that increase the interest rate available on some part of a saver's assets. TESSAs were the first product that allowed individuals in the UK to save from net income and pay no tax on returns to a simple cash deposit account. The advent of ISAs meant that savers could for the first time pay no tax on the returns to their savings made from net income without having to either sacrifice liquidity or bear stock market risk. We now look at evidence to try and uncover whether or not these expansions of tax exemptions did increase saving through the mechanism discussed in the theoretical section of the paper. To anticipate, our conclusion is consistent with that from the IRA reforms in the United States: there is little evidence that the policies had strong effects in generating increased saving.

Aggregate evidence on the effectiveness of TESSAs

A lack of detailed individual or household level data on wealth before 1995 makes a detailed analysis of the effects of the introduction of TESSAs infeasible. Instead we simply report summary evidence from the aggregate statistics which suggests a high degree of reshuffling of portfolios. Take-up of TESSAs was initially high, as shown in Figure 1. Immediately on their introduction, 2 million policies were taken out, representing just under one in ten households. Subsequently, the number of accounts rose steadily through the period and by their abolition in 1999 well over 5 million accounts were active. Since each account lasts for a

Figure 1. Take-up of tax exempt special savings accounts




Source: Inland Revenue Statistics.

fixed five-year period the stock of TESSAs held declined steadily in the following years as the accounts expired.

Looking at the time series of TESSA ownership it is clear that there was an early period of rapid adoption, between 1991 and 1993, after which growth in ownership rates was steadier (with the exception of a final short sharp increase in 1999 which is a “pre-abolition” effect). Remembering that there were fixed contribution limits, it is natural to ask what happened to average balances in this period. In the first year of a TESSA individuals were allowed to contribute £3 000, and in each subsequent year contributions were limited to £1 800, although subject to a total five-year limit of £9 000.

Figure 2 uses data on the stock of wealth held in TESSAs, combined with data on the number of active accounts, to show the average size of TESSA balances over the period. A striking pattern emerges. Over the period of early adoption, and indeed in the two subsequent years, there are discrete jumps in the average balances each time a new financial year begins. Although average balances do not increase by the entire amount of the increasing contribution limit it is clear that for the initial years at least, balances are close to the maximum they could be. Without household level data on consumption and/or saving in financial assets of the kind exploited for the US in the previous section, we cannot construct a test of the extent to which these observed changes in average TESSA balances are due to new saving or reshuffling of assets. There is though a plausible mechanism through which reshuffling would lead to the observed patterns, and particularly

Figure 2. Average TESSA balances



Source: Inland Revenue Statistics.

the discrete jumps in balances associated with the relaxation of constraints. Specifically, in the early years of an account the annual contribution limits effectively constrain the overall balances that can be held in the account. Each April an individual becomes eligible for another year of contributions which consequently raises the total amount that can be held in the account and therefore allows more reshuffling to take place, up to the point that balances in other forms of saving are available to be transferred. Such a pattern of large contributions at the beginning of the year, followed by no further contributions for the rest of the fiscal year, would not typically be associated with reduced consumption patterns⁴.

Following this episode, once TESSAs had been established and the first five year accounts were beginning to mature and possibly “roll-over” into follow-on accounts, there was a period of remarkable stability in average balances which continued until the accounts began to expire post 1999. Once again, although robust statements are difficult to make given the lack of micro data and the absence of an appropriate counterfactual, this would be consistent with the build up of TESSA balances having been mainly reshuffling. Given that evidence on the distribution of wealth suggests that stocks of assets are limited relative to five-year contribution limits for the majority of households, if the funds flowing into TESSAs were arising from reshuffling, then TESSA balances would begin to plateau as the funds available begin to run out. This seems to be exactly what happened as the system of TESSAs became more established through the mid 1990s.

The final episode takes place post 1999, and is less easy to interpret directly. Following the abolition of TESSAs and their replacement with ISAs there is a period when average balances fall steadily, in parallel with falls in the number of active accounts. Two possible interpretations present themselves. Firstly, this could be a selection effect – those who took out TESSAs early (and hence whose TESSAs expired early) typically may have had higher balances than those who took them out late. Secondly, however, this decline in average balances could have reflected marginal saving (or even marginal reshuffling) being directed into ISAs even for those who had a TESSA account. Such behaviour would have been rational (given equal pre-tax rates of return) since ISAs provide the same tax advantage as TESSAs but without the liquidity restriction.

The period around the replacement of TESSAs with ISAs is clearly instructive for evaluating tax incentives for saving. This is also a period where some micro data on savings and wealth holding was available in Britain, so it is this evidence which we turn to next.

Micro-evidence on the effectiveness of ISAs

To reiterate what was said above, the introduction of ISAs effectively created expenditure tax treatment on all accessible savings for the majority of households

in the UK. Compared with the tax favoured savings products that preceded them, ISAs offered greater flexibility in terms of accessibility of funds and the ability to hold cash. In order to make an initial analysis of whether or not ISAs encouraged saving, perhaps particularly among those with moderate incomes who were a target for the policy, we present some “group based” descriptive evidence of the limited micro-data on individual and family savings that is available for the UK. We split the sample into education groups defined to roughly measure whether the individual has only compulsory level schooling, advanced level school education, or some college education. Education might be thought to better capture *lifetime* resources than a control for current income, and unlike income, in our sample of non-students education is also not likely to be determined at the same time that current savings choices are made. We also split the sample into ten-year age bands in order to look at whether or not the new savings product attracted younger savers.

Our main data source is the “Family Resources Survey” (FRS). The FRS is an annual survey of households in Great Britain, which collects data on the economic resources of a cross section of households. It takes a representative sample of the population of private households across Great Britain,⁵ and samples around 40 000 individuals from around 28 000 families⁶ or 22 000 households, each year. In addition to its large sample size and the detailed information it provides on personal and family characteristics, the FRS is useful for our purposes because it contains detailed information on which financial assets individuals hold. It also records banded information on whether the individual thinks that the balance of funds in all accounts held between themselves and their spouse/partner is: less than £1 500; between £1 500 and £8 000; between £8 000 and £20 000; or greater than £20 000.

Table 5 shows how family level ownership of tax favoured financial savings products (TESSAs and PEPs, and subsequently ISAs) changed between financial year 1998/9 – immediately before ISAs were introduced – and financial year 2002/03, which is the most recent available year of FRS data. Since the new product could (subject to limits on amounts saved) be used as an almost perfect substitute for TESSAs and PEPs, but also offered greater flexibility than its predecessors, one would (all else equal) expect to see an increase in the proportion of families holding these assets. The data indeed show quite substantial percentage point increases in ownership of these assets, and the data in the appendix show that the pattern is similar if the data is analysed at the individual rather than family level.⁷ These increases do not seem to be confined to lower education groups who were thought to be a key target of the reform, but are seen across all groups. If the increases were considered in percentage (rather than percentage point) terms, then there is evidence of the effect being weakly skewed towards younger and lower education groups: the increases for the two low education groups in the

Table 5. Percentage of families holding tax-free financial assets

Age ceased education	Age						All
	< 30	30s	40s	50s	60s	70+	
1998/99							
School leaving age	3.3	7.3	7.5	18.8	20.5	9.6	10.5
Age 18	7.1	15.6	19.5	33.4	38.8	24.8	23.1
Age 19+	12.0	27.4	37.1	52.3	60.1	56.1	32.9
All	6.5	14.9	20.5	30.9	32.1	17.4	19.2
2002/03							
School leaving age	8.4	14.3	18.6	26.6	32.7	22.0	19.8
Age 18	14.7	29.3	31.3	42.8	54.7	42.3	35.8
Age 19+	28.1	43.3	47.3	63.0	68.9	68.5	46.6
All	16.0	26.9	31.1	41.4	45.4	32.6	31.2
Percentage point increase							
School leaving age	5.1	7.0	11.1	7.9	12.2	12.4	9.3
Age 18	7.5	13.8	11.8	9.5	15.9	17.4	12.7
Age 19+	16.1	16.0	10.2	10.6	8.7	12.4	13.7
All	9.5	12.0	10.6	10.5	13.3	15.2	12.0

Source: Authors' analysis using the Family Resources Survey. As in all subsequent tables using FRS data, totals are grossed up to be more representative of population totals, but this makes little difference to the figures and ungrossed numbers are available on request from the authors.

youngest age band represent a more than doubling in ownership. Evidence of this kind can be, and in some government publications has been (see Chapter 5 of H.M. Treasury, 2000), interpreted as suggesting that ISAs were more attractive to lower income savers than TESSAs and PEPs had been. On the other hand at least one study that has used micro-data to examine the effects of ISAs argued that the products were little better than TESSAs and PEPs at reaching some low-income groups (Paxton, 2003). Both of these studies may have understated the difference in reach of the different products by comparing stock levels of ownership for the more established products (TESSAs and PEPs in the late 1990s) to those for ISAs in only their early years.

If the evidence of an increase in ownership rates for tax favoured savings products were supplemented by evidence of an increase in ownership of assets more generally, then this would support the notion that ISAs might have encouraged some new savers. Table 6 looks precisely at whether or not the ownership of all non-pension financial savings products (all accounts except current accounts) rose between 1998/9 and 2002/3. In this case there is no evidence of an increase, and in most groups ownership rates actually fell over the period (again data presented in the appendix show a similar pattern for individuals as for families). If

Table 6. **Percentage of families holding any financial assets other than a current account**


Age ceased education	Age						All
	< 30	30s	40s	50s	60s	70+	
1998							
School leaving age	39.1	49.5	49.6	62.2	66.9	67.4	55.6
Age 18	56.2	75.7	72.6	82.6	81.9	83.0	75.1
Age 19+	67.7	83.5	88.7	92.2	90.5	91.4	82.5
All	51.0	65.5	70.4	76.1	75.4	73.9	67.6
2002							
School leaving age	33.2	50.1	52.3	59.1	63.7	65.9	53.8
Age 18	48.3	71.4	70.3	78.1	80.6	78.8	71.4
Age 19+	64.5	79.5	85.1	89.9	89.8	91.7	79.8
All	46.7	64.4	68.0	73.9	73.3	72.3	65.6
Percentage point increase							
School leaving age	-6.0	0.7	2.7	-3.1	-3.2	-1.6	-1.7
Age 18	-7.9	-4.4	-2.4	-4.5	-1.3	-4.2	-3.7
Age 19+	-3.3	-4.0	-3.7	-2.2	-0.8	0.3	-2.7
All	-4.3	-1.2	-2.4	-2.2	-2.1	-1.6	-2.1

Source: Authors' analysis using the Family Resources Survey.

anything, the decline was biggest in the youngest, lowest education groups. While it is possible that the decline would have been more dramatic in the absence of ISAs, the combined evidence of Tables 5 and 6 strongly suggests that many people who held ISAs in 2002/03 were simply reallocating or “reshuffling” assets that they would have held in other accounts if ISAs had never been introduced.

Even if it is the case that the advent of ISAs had little impact on asset ownership rates across education and age-groups in the population, the policy could still have been successful in encouraging new saving if those who were saving began to save more after the new instrument was introduced. We therefore need to consider data on the amounts that people save. Figure 3 and Table 7 (with a corresponding analysis at the individual level in Appendix Table A3) exploit the FRS data which records in broad bands the level of family wealth held in non-pension financial assets. Both the aggregate data in the figure and the grouped data in the table show evidence of a slight increase over time in the proportion of families with more than £1 500 in these assets. This increase seems to have been experienced fairly uniformly across groups. There are at least two reasons why it is hard to interpret this data as evidence of increased saving due to ISAs. The first is that over time inflation would have eroded the real value of £1 500 of nominal wealth: to compare families with more than £1 500 of nominal wealth in 2002 to families with the same real value of wealth in 1998, one would need to observe how many

Figure 3. Percentage of the population living in families with more than £1 500 in financial assets



Source: Authors' analysis of Family Resources Survey, various years.

Table 7. Percentage of families holding more than £1 500 in financial assets

Age ceased education	Age						All
	< 30	30s	40s	50s	60s	70+	
1998							
School leaving age	12.8	24.0	27.5	45.2	56.5	58.7	37.4
Age 18	20.8	43.6	48.4	64.8	76.3	75.5	54.0
Age 19+	33.5	60.7	69.6	78.9	84.0	90.9	61.1
All	20.9	40.2	48.7	60.1	67.8	66.6	48.5
2002							
School leaving age	12.8	26.7	33.1	47.1	58.6	61.8	40.3
Age 18	19.3	45.8	52.4	66.6	77.0	78.6	56.8
Age 19+	35.7	63.6	72.8	81.5	90.9	93.2	64.8
All	22.0	43.4	51.7	63.5	70.0	70.4	51.8
Percentage point increase							
School leaving age	0.0	2.6	5.7	1.9	2.1	3.1	2.9
Age 18	-1.6	2.3	4.0	1.9	0.7	3.1	2.7
Age 19+	2.1	2.9	3.3	2.6	6.9	2.3	3.7
All	1.2	3.2	3.1	3.4	2.2	3.8	3.4


Source: Authors' analysis using the Family Resources Survey.

families had more than around £1 390 of wealth in the earlier year,⁸ but this information is not contained in the FRS. Additionally, the FRS data measures a stock of wealth whereas we would like to observe the flow of wealth going into these assets. In order to get a measure of this flow that can be presented in real terms, it is necessary to use a second dataset, the British Household Panel Survey (BHPS).

The British Household Panel Study is a survey of around 10 000 adults in around 5 000 households (or 7 000 families).⁹ The same households are interviewed annually although as in any panel study, some respondents drop out of the panel.¹⁰ Like the FRS the survey asks participants about a wide range of personal and family characteristics including educational attainment and age. Although somewhat less detailed information is elicited on income sources, the BHPS does ask interviewees whether or not they maintain some savings by “putting something away” from their income. Those who answer “yes” to this question are then asked to estimate “About how much on average do you personally manage to save a month?” It is the information contained in answers to this question that we exploit here.

Figure 4 shows that, with the exception of the occasional upwards movement for the highest education group, there is little evidence of growth in the real (2002) value of average monthly saving in financial assets by families in our BHPS sample

Figure 4. **Mean monthly saving among families saving positive amounts**
Age ceased education



1. These averages do disguise some change in the proportion in each group that are saving positive amounts. However, although there is a general upward trend in the proportion of savers during the early years of the data, this levelled off before ISAs were introduced: for all groups 1998 marks the peak in the series for the proportions of savers.


Source: Authors' analysis of the British Household Panel Survey, various years.

between 1991 and 2001. Only for the highest education group do average savings rise above their 1998 level in the years after the introduction of ISAs, and that only in 2001, the final year of our data. Furthermore, for all the education groups, the data underlying this chart show that the proportion of families saving anything at all peaked in 1997 or 1998, just before the introduction of ISAs. On the face of it there seems to be little evidence of an effect of the introduction of ISAs on the amounts that families in the BHPS were saving in financial assets.

Finally, Figure 5 combines FRS and BHPS data in an attempt to provide a visual assessment of how the increase in the number of families holding tax-favoured savings products between 1998 and 2001 was related to the changing average levels of saving by families. There seems to be no clear pattern to the points displayed on the chart, although there are too few data points to scientifically assess whether or not there is a relationship such that average savings went up more (or down by less) in groups with stronger movements into tax favoured assets. Nonetheless, the fact that there are some groups with high take-up of the tax-favoured assets *and* big declines in average savings suggests that the introduction of ISAs certainly was not enough to offset other factors that may have been acting to depress saving levels.

Figure 5. **Percentage increase in amounts saved and numbers with tax free assets, 1998-2001**

Age ceased education, age:



Source: Authors' analysis of Family Resources Survey and British Household Panel Survey, various years.

Overall, this descriptive look at the data suggests that while the take-up of ISAs was quite high, there is no strong evidence that this had much affect on overall ownership of non-pension financial assets or on levels of saving among those with such assets. This evidence, particularly on ownership, is consistent with reshuffling behaviour of the kind observed in the US with IRAs.

CONCLUSIONS AND THOUGHTS FOR FURTHER RESEARCH

The evidence presented in this paper for IRAs in the US, and TESSAs and ISAs in the UK, suggests that, at the most, only relatively small fractions of the funds going into tax-advantaged savings vehicles can be considered to be “new” saving. As such, the best interpretation of the evidence is that such policies are expensive ways of encouraging savings. In addition, to the extent that the reshuffling of assets leads to a reduction in the tax liabilities without any real change in economic behaviour, there is some deadweight loss associated with such policies. Additionally, since those with the greatest reshuffling possibilities are the wealthiest members of society, these policies will typically have some distributional impact. *A priori* we would expect these factors to be biggest for IRAs and TESSAs, where the liquidity restriction would be likely to prevent savers with low wealth from participating. As such the introduction of ISAs in the UK is an important episode, showing, as it does, that even when such accounts do not include minimum holding periods the extent to which they encourage genuinely new saving is limited.

Of course, the arguments we used in this paper are mainly based on a standard economic model that assumes rational and informed behaviour. It is of course possible that such a model is not adequate and that to make the right intertemporal decisions individuals need either information or commitment devices. But if that is the case, and the “if” looms large here, then it is not obvious that rate of return advantages that can be expensive in terms of government revenue, are relevant for these other channels to work. Much work is needed to assess the extent to which individuals make the “wrong” decisions, and the possible importance of employer programmes, self-control, financial education, awareness, advertising, and government endorsements in changing such decisions. Particular types of arrangements may work, and may be responsible for what new saving we do observe in the products we considered above [and particularly in other products such as 401(k) in the US, and Personal or Stakeholder Pensions in the UK which we have not considered here].

The evidence presented above is limited in its scope by the nature of the data and policy variation available to researchers. In recent years, substantial progress has been made in the use of experimental and quasi-experimental methods to evaluate policy reform, particularly with application to labour market, education and training policies. Such approaches have not been applied in the area of

savings policies and yet such methods, coupled with adequate measures of both saving in various forms and of consumption expenditures, would seem a promising direction for evaluating tax incentives more formally. As new data sources, which include important information on household financial behaviour and intertemporal choices (such as the HRS AHEAD in the US and ELSA in the UK) become available, the time is ripe for more work in this area. The role that will be played by these new and innovative data sources will be crucial. They represent a fundamental advance that was only developed in the last few years and that should be developed even further, both in terms of increasing the households covered by these surveys (which have so far been limited to elderly individuals) and the countries in which they are collected.

Notes

1. PEPs also attracted a “dividend tax credit”, as subsequently did ISAs. For an example explaining the relative generosity of the credit in the two policies, see Banks and Tanner (1999), pp. 91-92.
2. More could be saved if a “single company PEP” was held alongside a general PEP.
3. An inland revenue press release (<http://archive.treasury.gov.uk/pub/html/budget97/ir4.html>) produced when the policy was first proposed said that “the new scheme will have a particular emphasis on encouraging those on low incomes to save”.
4. The observed pattern of changes in balances could be associated with some reduction in consumption if, for example, the TESSA contribution limit can easily be exhausted early in the first year, but subsequently extra savings are made from income each month. This could mean that each April a lump of assets cumulated over the previous 11 months would be transferred into the TESSA. However, even this explanation is more plausible if the payments made in the first year come from reshuffling existing assets. Further, the evidence considered in this paper on IRAs in the US and ISAs in the UK, give additional weight to the supposition that changes in TESSA balances might have been due to reshuffling of assets.
5. Excluding the sparsely-populated Scottish islands and area north of the Caledonian Canal. More information on the FRS can be found at www.dwp.gov.uk/asd/frs/.
6. Throughout this section the term “family” refers to the unit of assessment for state benefits, which will be either a single adult or an adult couple, plus any dependent school aged children that live with the adult unit.
7. The similarity to the family analysis is even closer if one works at a household level. This is because the majority of UK households contain only one benefit unit (family). Household level analysis is available from the authors on request.
8. This number is derived using the Retail Prices Index for the middle month (September) of each year of our data.
9. Note that since 1997, there have been a number of non-representative sub-samples added to the BHPS sample. We do not use these households in our analysis. More information on the BHPS can be found at www.iser.essex.ac.uk/bhps/doc/.
10. In total, 62 per cent of respondents who gave a full interview at wave 1 (1991) were present in wave 10 (2000).

Appendix
Analysis at the individual level

Table A1. Percentage of adults holding tax-free financial assets

Age ceased education	Age						All
	< 30	30s	40s	50s	60s	70+	
1998							
School leaving age	3.5	7.4	10.4	19.4	20.8	9.6	11.5
Age 18	7.3	15.2	18.4	33.4	37.8	24.5	22.5
Age 19+	11.7	25.0	34.6	49.7	57.6	50.9	30.2
All	6.4	13.1	18.7	28.3	29.6	15.8	17.9
2002							
School leaving age	8.9	15.4	19.3	28.6	33.4	23.0	21.4
Age 18	15.2	26.9	30.7	42.8	53.7	41.6	35.0
Age 19+	27.5	39.7	43.3	62.6	65.9	66.1	43.8
All	16.0	24.2	28.8	39.6	42.9	31.5	29.7
Percentage point increase							
School leaving age	5.5	8.0	9.0	9.2	12.6	13.4	9.9
Age 18	7.9	11.7	12.3	9.4	15.9	17.1	12.5
Age 19+	15.8	14.7	8.6	12.9	8.3	15.3	13.6
All	9.6	11.1	10.1	11.4	13.2	15.7	11.8

Source: Authors' analysis using the Family Resources Survey.

Table A2. **Percentage of adults holding any financial assets than a current account**

Age ceased education	Age						All
	< 30	30s	40s	50s	60s	70+	
1998							
School leaving age	41.2	53.1	54.6	63.3	66.7	67.6	57.5
Age 18	57.0	74.3	70.5	80.6	80.9	81.4	73.7
Age 19+	67.2	80.7	85.7	89.6	89.0	89.0	80.2
All	51.5	64.2	67.9	72.8	73.4	72.6	66.4
2002							
School leaving age	35.8	52.1	53.7	61.2	64.0	65.8	55.7
Age 18	49.8	69.8	68.9	76.7	80.3	77.6	70.7
Age 19+	63.2	77.2	80.8	88.5	89.8	91.0	77.5
All	47.5	62.6	65.3	71.7	71.6	71.0	64.4
Percentage point increase							
School leaving age	-5.4	-1.0	-0.8	-2.1	-2.8	-1.8	-1.8
Age 18	-7.2	-4.5	-1.6	-3.9	-0.6	-3.8	-3.0
Age 19+	-3.9	-3.6	-5.0	-1.1	0.9	1.9	-2.7
All	-3.9	-1.6	-2.6	-1.1	-1.8	-1.5	-2.5

Source: Authors' analysis using the Family Resources Survey.

Table A3. Percentage of adults living in families holding more than £1 500 in financial assets

Age ceased education	Age						All
	< 30	30s	40s	50s	60s	70+	
1998							
School leaving age	16.6	31.1	36.9	53.4	63.0	62.2	43.1
Age 18	24.6	49.5	54.2	71.0	79.9	77.3	58.3
Age 19+	37.4	64.0	75.5	81.3	86.7	91.0	64.2
All	24.1	43.3	52.7	63.6	71.0	68.3	51.8
2002							
School leaving age	16.5	35.6	41.6	55.9	64.1	65.3	46.8
Age 18	23.9	52.7	59.1	72.3	80.8	81.1	61.7
Age 19+	39.6	67.5	75.1	85.4	91.6	94.2	67.5
All	25.7	47.9	55.8	67.5	72.3	72.3	55.7
Percentage point increase							
School leaving age	-0.1	4.5	4.7	2.5	1.1	3.1	3.7
Age 18	-0.7	3.2	4.9	1.3	0.9	3.8	3.4
Age 19+	2.2	3.5	-0.4	4.1	4.9	3.2	3.3
All	1.6	4.6	3.1	3.9	1.3	4.0	3.9

Source: Authors' analysis using the Family Resources Survey.

BIBLIOGRAPHY

- ATTANASIO, O. and T. DeLEIRE (2002), "The Effect of Individual Retirement Accounts on Household Consumption and National Saving", *The Economic Journal*, 112, pp. 504-538.
- BANKS, J., Z. SMITH and M. WAKEFIELD, (2002), *The Distribution of Financial Wealth in the UK: Evidence from 2000 BHPS Data*, IFS working paper W02/21, www.ifs.org.uk/workingpapers/wp0221.pdf.
- BANKS, J. and S. TANNER (1999), *Household Saving in the UK*, IFS report, Institute for Fiscal Studies, London.
- H.M. TREASURY (2000), *Helping People to Save*, London: HM Treasury (www.hm-treasury.gov.uk/media/8EA/F9/511.pdf).
- PAXTON, W. (2003), *Tax Efficient Saving: The Effectiveness of ISA*, Institute for Public Policy Research, London (www.ippr.org.uk/press/index.php?release=197¤t=2003).