BURCH CENTER WORKING PAPER SERIES

University of California, Berkeley
Department of Economics - #3880
Berkeley, CA 94720-3880

Rethinking Saving Incentives

B. Douglas Bernheim
Stanford University and NBER

March 1996

Burch Working Paper No. **B96-17**

JEL Classification: H31

This paper was prepared for a conference titled "Fiscal Policy: Lessons from Economic Research," sponsored by the Burch Center for Tax Policy and Public Finance at the University of California at Berkeley, and held on February 2-3, 1996. I am grateful to the National Science Foundation (Grant Number SBR95-11321 and Grant Number SBR94-09043) for financial support. William Gale and Jonathan Skinner provided helpful comments.

No. B96-17

RETHINKING SAVING INCENTIVES

Abstract

This paper undertakes a review of the existing academic work on tax incentives and personal saving. Its central conclusions are as follows. First, the traditional life cycle hypothesis has had an excessive influence on the design and conceptualization of empirical investigations concerning taxation and saving. Second, there is little reason to believe that households increase their saving significantly in response to a *generic* increase in the after-tax rate of return. Third, the literature on the relation between Individual Retirement Accounts (IRAs) and personal saving is inconclusive. Fourth, one can be moderately confident that, all else equal, eligibility for a 401(k) plan significantly stimulates personal saving. Fifth, tax incentives probably have important effects on personal saving through "third party" or institutional channels. Sixth, there is, overall, considerable uncertainty about the effects of policies designed to promote saving, particularly in cases where these policies have the potential to induce significant institutional change (such as a consumption tax).

B. Douglas Bernheim
Stanford University

I. Introduction

Low rates of saving in the United States have created widespread concern over investment, growth, the balance of payments, and the financial security of individual households. These concerns have prompted a variety of proposal designed to stimulate saving through tax incentives, ranging from narrowly focused tax-favored savings accounts, to broad-based consumption taxation. Before attempting to weigh the advantages and disadvantages of these proposals, it is important to marshall the available evidence, and review the lessons that economists have learned about the relation (or relations) between tax incentives and personal saving.

This paper undertakes a review of the existing academic work in this area. It attempts both to evaluate the evidence, and to identify the important open questions. Although it does not include a full evaluation of any particular saving incentive proposal, it does identify relevant implications.

The central conclusions of this review are as follows. First, the traditional life cycle hypothesis has had an excessive influence on the design and conceptualization of empirical investigations concerning the relation between taxation and saving. While other behavioral hypotheses are mentioned in the literature with increasing frequency, this usually occurs in the course of explaining an anomalous result, rather than at the stage of designing an empirical strategy, or even at the stage of evaluating results that do not appear to be anomalous from a life-cycle perspective. In part, this is no doubt attributable to the absence of sufficiently well-developed organizing principles for a compelling behavioral alternative. But even in the absence of an intellectually satisfying alternative, it is important to be aware of the potential for reaching misleading conclusions by imposing potentially false structure on the data.

Second, there is little reason to believe that households increase their saving significantly in response to a *generic* increase in the after-tax rate of return. Since the evidence is quite poor, there is still considerable uncertainty on this point. However, it is difficult to identify any robust empirical pattern that is suggestive of a high elasticity.

Third, the literature on the relation between Individual Retirement Accounts (IRAs) and personal saving is inconclusive. Studies that point to a large effect on personal saving contain identifiable biases that overstate this effect. Likewise, studies that find little or no effect on personal saving contain identifiable biases that understate the effect. Due to the nature of the IRA program and the characteristics of the available data, reliable estimation of the effect on personal saving may well be impossible. However, the IRA experience does provide reasonably strong evidence that households' responses to narrowly-focused tax incentives are governed, at least in part, by forces that are not considered in the traditional life cycle framework.

Fourth, the available evidence on 401(k)s allows one to conclude with moderate confidence that, all else equal, eligibility for a 401(k) plan significantly stimulates personal saving. Econometric identification of this effect is facilitated by the fact that, in contrast to IRAs, there is considerable variation in 401(k) eligibility across the population, and some of this variation almost certainly arises from exogenous sources. While no existing study corresponds to the ideal statistical experiment, at least one sensible approach concludes that 401(k)s do not displace other saving despite the fact that its primary shortcoming probably creates a bias in the opposite direction. In contrast, important biases obscure the interpretation of studies that reach the opposite conclusion.

Fifth, tax incentives probably have important effects on personal saving through "third party" or institutional channels. "Third parties" refer to entities other than the households that benefit directly from the tax provisions, such as employers or vendors of tax-favored investments products. Examples of these effects include stimulation of pension provision, indirect encouragement of employer-based financial education, and the advertisement and promotion of opportunities for tax-favored saving. Since little is known about the relation between tax incentives and either the scope or characteristics of institutions that promote saving, this should be a priority for future research.

Sixth, there is, overall, considerable uncertainty about the effects of policies designed to

promote saving. These uncertainties are greatest for policies that have the potential to induce significant institutional change, such as the implementation of a broad-based consumption tax. If, for example, the generic interest elasticity of saving is small (as the evidence seems to suggest), and if targeted tax incentives sustain vital pro-saving institutions (such as the pensions system), then consumption taxation could have disastrous implications for saving. In contrast, narrowly targeted incentives could bolster the institutional structures that encourage saving.

The remainder of this paper is organized as follows. Section II analyzes the role of the life cycle hypothesis in empirical analyses of saving incentives. Section III evaluates the generic response of saving to changes in the after-tax rate of return. Section IV reviews the evidence on tax-favored savings accounts, while section V explores "third-party" effects. Section VI briefly concludes.

II. What are the links between tax policy and personal saving?

A. Is the life cycle hypothesis an adequate organizing framework?

For more than fifty years, the framework of intertemporal utility maximization has dominated economists' thinking about saving incentives. This framework comes with an exceptional pedigree. It traces its roots to Irving Fisher (1930), and lies at the heart of the Life Cycle Hypothesis (LCH) articulated by Modigliani and Brumberg (1954). In recent years, the validity of this framework has become controversial, and an increasing number of economists have expressed doubts concerning its general validity (see e.g. Bernheim, 1994a). Although various authors have suggested provocative alternative paradigms (e.g. Shefrin and Thaler, 1988), no serious competitor for the attention and imagination of most professional economists has yet emerged. As a result, the LCH has remained the preeminent theoretical tool for analyzing saving incentives -- in many cases, with serious reservations on the part of those who use it.

A comprehensive review of the evidence for and against the LCH is beyond the scope of this paper. Suffice it to say that the evidence is mixed. While it would be rash to dismiss the

many empirical successes of the LCH and discard it unconditionally, it is equally rash (in light of its empirical failures and well-founded skepticism about its underlying premises) to employ this theory as the sole organizing principle for understanding saving incentives. Any conceptual framework acts as a filter on reality. It shapes the questions that we pose, and influences the manner in which we interpret evidence that bears on these questions. By treating the LCH as the universal maintained hypothesis, we run the risk of overlooking important evidence concerning the efficacy of saving incentives.

It is, perhaps, an exaggeration to say that economists studying saving incentives have been straight-jacketed by the LCH. Yet references to alternative behavioral hypotheses are somewhat haphazard; they are usually mentioned in a rather *ad hoc* way as possible explanations for otherwise puzzling phenomena (see e.g. Venti and Wise, 1992, or Feenberg and Skinner, 1989). With rare exceptions, alternative behavioral hypothesis have not been used as frameworks for organizing lines of empirical inquiry concerning the effects of tax incentives, or even for interpreting evidence that does not appear to be anomalous from the perspective of the LCH. The scarcity of well-developed organizing principles for these alternative hypotheses (aside from a few relatively imprecise notions such as "mental accounting" and "self-discipline") is a serious shortcoming of the existing literature.

B. Possible links between tax policy and personal saving

The traditional focus on the LCH has generated the perception that the task of evaluating the effects of saving incentives is fundamentally a matter of determining the extent to which these incentives affect the marginal dollar of saving, and of measuring the interest elasticity of saving. Consideration of alternative hypotheses broadens this line of inquiry considerably, to include all of the following concerns.

1. Perceptions of the costs and benefits from saving. The field of psychology appears to offer relatively few organizing principles that are directly useful for evaluating the effects of

saving incentives (see e.g. the discussion in Chapter 9 of Furnham and Lewis, 1986). One notable and pervasive exception is the principle of self-control (see e.g. Ainslie, 1975, 1982, 1984, Maital, 1986, Schelling, 1984, Thaler and Shefrin, 1981, Shefrin and Thaler, 1988, and Hoch and Lowenstein, 1991). Thaler and Shefrin equate the notion of self-control with the view that each individual has more than one set of preferences (or "selves") that conflict at a point in time, and that individuals must expend energy and resources to restrict the kinds of behavior that are motivated by their short-sighted selves. In this view, aspects of the economic environment that tip the internal balance of power toward patient selves have the potential to stimulate saving.

One way to shift this balance of power is to alter the perceived costs and benefits of saving. Certain kinds of saving incentives have the theoretical capacity to precipitate such a shift through a variety of channels. When saving incentives are in place, individuals may be more likely to learn that others (including both "authorities" and peers) regard the benefits of saving as important.[1] By segmenting retirement saving from other forms of saving, certain kinds of tax-favored accounts also make it easier to monitor progress towards long-term objectives.[2] The existence of saving incentives may reinforce the notion that, as something worthy of encouragement, saving is an intrinsically desirable or virtuous activity, entirely apart from explicit consideration of future consequences (see e.g. Scitovsky, 1976, or Katona, 1975). Under certain circumstances, contributions to tax-favored accounts may also instill the perception that saving yields concrete short-run benefits, such as reductions in current-year tax liabilities (Feenberg and Skinner, 1989).

2. Private rules. The psychological literature on self-control also emphasizes the use of "private rules" (see e.g. Hoch and Lowenstein, 1991). Saving incentives may facilitate the formation of effective private rules in two ways. First, they may provide a natural context for developing rules concerning the level of saving, such as "maxing out" on contributions. Certain plans, such as 401(k)s, actually provide participants with limited ability to commit themselves to these rules for short periods of time. Second, individuals may also develop private rules

regarding the allowable uses of funds that they have previously placed in tax-favored accounts. This phenomenon relates to the notion of "mental accounting" discussed by Shefrin and Thaler (1988). The existence of penalties for early withdrawal may help the individual establish and enforce barriers around tax-favored accounts.

3. Third-party activities. A uniform tax on all forms of capital income would create a neutral environment with respect to the form of saving. This neutrality would be preserved if all capital income taxes were eliminated (e.g. through the implementation of a consumption tax). In contrast, limited saving incentives are usually non-neutral, in the sense that they favor saving in certain forms. These non-neutralities may stimulate activities by "third parties" -- that is, parties other than the individuals who benefit directly from the tax provisions, such as employers or vendors of tax-favored investments products -- and these activities may in turn affect the level of personal saving through either life cycle or psychological channels. Perhaps the most obvious example this phenomenon is the development of the pension system. Selective saving incentives may also have subtle effects on the choices that employers make in the context of pension plans. For example, Garrett (1995) argues that non-discrimination requirements induce employers to take steps that stimulate 401(k) contributions among non-highly compensated employees (e.g. through matching contributions or retirement education). Selective incentives may also encourage the vendors of tax-favored savings vehicles to advertise and promote their products actively.

C. Why is the nature of the link important?

A pure (and somewhat narrow-minded) empiricist might dismiss the importance of the nature of the link between tax incentives and saving on the grounds that one can measure the quantitative effect of these incentives without identifying the underlying mechanism. In this section, I argue that the nature of the mechanism is important for three reasons.

First, the possible justifications for attempting to stimulate saving, and the manner in

which we think about the social costs and benefits of saving incentives, differ across behavioral hypotheses. If, for example, profligacy results from a failure to understand financial vulnerabilities, or from an unintended break-down of self-control, then the potential welfare gains associated with pro-saving policies are likely to be much larger than those implied by the LCH.

Second, the nature of the link between tax incentives and saving significantly affects our expectations concerning the likely effects of different kinds of pro-saving policies. As an example, consider the choice between broad-based and narrowly focused strategies for promoting saving. The widespread (though certainly not universal) support for consumption taxation among economists is almost certainly attributable to the influence of the LCH as the preeminent (if not exclusive) conceptual framework. From the perspective of a neoclassical maximizer of intertemporal utility, consumption taxation accomplishes the prime objective: it raises the marginal after-tax rate of return to saving, and reduces the "tax wedge" between the pre-tax return of investment and the after-tax return on saving. Yet the psychological considerations discussed in section B generally suggest that narrowly focused tax incentives have the potential to be much more effective than broad-based measures. Narrow measures focus attention on the key issue, expose individuals to information concerning the importance of saving, provide a natural context for the development and enforcement of private rules, and promote the growth of pro-saving institutions. In contrast, a broad-based consumption tax could undermine the narrow focus on specific objectives that may be essential for the exercise of self-control. As discussed above, a true consumption tax would be institutionally neutral. It would remove one of the primary reasons for compensating workers through pension plans, and it would eliminate the special feature of particular financial instruments (such as IRAs and life insurance policies) that make them especially marketable. It would also eliminate the quirky aspects of the tax system that subtly promote activities such as employee retirement education.

Third, the nature of the links between saving incentives and behavior profoundly affects

the lessons that we might be inclined to draw from a variety of possible empirical patterns. As an example, consider the generalizability of evidence on the interest elasticity of saving. Within the context of the LCH, all saving incentives motivate changes in behavior through the same fundamental mechanism: an increase in the after-tax rate of return. Measurement of the interest elasticity of saving therefore emerges as a central research priority. Alternative behavioral hypotheses allow for the possibility that the interest elasticity of saving may vary according to context, and that households may respond (both positively and negatively) to aspects of tax incentive programs that are not directly related to the after-tax rate of return.

III. Direct responses to generic changes in the after-tax rate of return

Much of the literature on the relation between tax incentives and personal saving concerns the extent to which saving responds, in general, to a change in the after-tax rate of return. As is well-known, the sign of the interest elasticity of saving is theoretically ambiguous in simple life cycle models, due to the existence of opposing income and substitution effects. While certain kinds of life cycle simulations tend to generate substantial positive elasticities (Summers, 1981), this result is highly sensitive to assumptions about utility (Starrett, 1988) and uncertainty (Engen, 1994), and is easily reversed under alternative behavioral hypotheses (e.g. "target" saving behavior). Thus, the size of the interest elasticity of saving is fundamentally an empirical issue. Unfortunately, as explained in this section, the existing methods of measuring the interest elasticity of saving are critically flawed. Despite these flaws, studies based on different approaches typically reach strikingly similar conclusions: evidence of a substantial elasticity is the exception, rather than the rule. Based on this evidence, one is hard-pressed to make the case that saving is *in general* highly sensitive to the after-tax rate of return.

A. The consumption/saving function approach

The earliest approach to measuring the interest elasticity of saving involved the estimation of a consumption function or saving function featuring an interest rate among the list

of explanatory variables. Since the initial work of Wright (1969), this approach has yielded a variety of elasticity estimates, ranging from essentially zero (Blinder, 1975, Howrey and Hymans, 1978, Skinner and Feenberg, 1989) to 0.4 (Boskin, 1978, Boskin and Lau, 1978). This range is somewhat misleading, since the estimates tend to cluster near zero. There has been considerable discussion in the literature concerning the sources of the discrepancies between these various estimates (see e.g. Sandmo, 1985), with particular attention being given to the proper measurement of the real after-tax rate of return. However, the problems with this approach are more fundamental.

As noted by Lucas (1976), the relation between consumption (saving) and interest rates may depend on the wider macroeconomic context; consequently, there may not exist anything that one could properly regard as a stable saving or consumption function. Low elasticity estimates are largely attributable to developments during the 1970s (saving was relatively high and *ex post* real rates of return were very low). Since it is not at all obvious that the 1970s constitute a "clean" macroeconomic experiment for inferring the interest elasticity of saving, it is doubtful that these kinds of estimates have the potential to tell us very much about the likely response of saving to a change in tax policy.

B. The Euler equation approach

A second approach to estimating the interest elasticity of saving finesses the Lucas critique. This approach can be motivated as follows. To the extent a change in the interest rate alters saving, this is accomplished by inducing the individual to revise consumption. We can decompose the change in consumption into two pieces: a change in shape of the age-consumption profile, and a change in the overall level of this profile. The consumption/saving function approach makes no attempt to distinguish between these two effects. Yet it is conceivable that the level of consumption might be quite sensitive to aspects of the economic environment (through effects on expectations about future income, among other channels),

whereas the relation between consumption growth and the interest rate might be stable. In that case, one could estimate this second relation reliably. Using the estimated relation along with the budget constraint, one could then infer the structural effect of the interest rate on the level of consumption and/or saving, thereby recovering the true interest elasticity of saving.

It is possible to justify this approach formally in the context of a standard life cycle planning model. Specifically, Hall (1988) notes that, under appropriate parametric assumptions, the following Euler equation (governing the allocation between consumption in two consecutive periods, t and t+1) is a necessary condition for optimization:

$$\ln\left(\frac{c_t}{c_{t-1}}\right) = k + \sigma r + \epsilon_t \,, \tag{1}$$

where k is a constant, σ is the elasticity of substitution in consumption, and ϵ_t is a random variable that is uncorrelated with information available prior to t. Estimation of this equation permits one to recover the key parameter of the utility function that governs intertemporal tradeoffs, rather than the parameters of a potentially unstable consumption function. The value of the estimated parameter is directly relevant to an analysis of the welfare effects of capital income taxation. Moreover, armed with a parameterized utility function and a budget constraint, one can infer the interest elasticity of saving. For example, $\sigma = 1$ corresponds to Cobb-Douglas utility, which in turn implies fixed expenditure shares. If one estimates $\sigma = 1$, then, abstracting from the effect of changes in the interest rate on the present discounted value of wage income, the uncompensated interest elasticity of saving would be zero.

Hall demonstrates that the rate of change of consumption has been remarkably insensitive to changes in the rate of return. His estimates of σ are correspondingly small; he concludes that σ is almost certainly below 0.1, and possibly equal to zero (Leontief preferences). All of his estimates are well below unity, which corresponds to the Cobb-Douglas case. Thus, abstracting from the impact of the rate of return on the present discounted value of future earnings, one

would infer from these results that the interest elasticity of saving (ϵ_{sr}) is negative. While more recent research suggests that Hall may have underestimated σ due to aggregation problems and other concerns, most authors still place σ (as inferred from Euler equation estimates) significantly below unity, implying (for the case where the individual has no future earnings) $\epsilon_{sr} < 0$ (see e.g. Attanasio and Weber, 1993).

Unfortunately, the interpretation of σ as the intertemporal elasticity of substitution is highly model-specific. As Hall himself notes, the standard life cycle model makes an automatic connection between this intertemporal elasticity and the coefficient of risk aversion, whereas no connection appears to exist in practice. Although Hall exhibits one specification of utility that breaks this connection while still generating an Euler equation with an identical interpretation, there is no guarantee that Hall's result would remain valid under other specifications of utility. For example, a model with "buffer stock" saving might yield an Euler equation with a very different structural interpretation. And if one moves to other behavioral hypotheses, it is not obvious that an estimate of σ has any structural significance at all.

It is nevertheless possible to infer an interest elasticity of saving directly from estimates of (1) (the Euler equation) without specifying a behavioral hypothesis, provided that one is willing to assume the stability of the estimated relation. This is accomplished through a mechanical application of the budget constraint. Suppose in particular that

$$c_t = e^{g(r)t} c_0 \quad . \tag{2}$$

Imagine that a consumer lives for a total of T years, earning a wage flow of w_t through time R, and nothing from time R through time T. Let $W \equiv W_0 + \int_0^R w_\tau e^{-r\tau} d\tau$, where W_0 denotes initial wealth. Then the budget constraint requires

$$c_0 \int_0^T e^{[g(r) - r]t} dt = W \quad . \tag{3}$$

Using equation (3), it is easy to verify that the interest elasticity of saving at time zero (defined as $s_0 = w_0 - c_0$) is given by

$$\epsilon_{sr} = \left(\frac{r c_0}{s_0} \right) \left(\left[\frac{W - W_0}{W} \right] D(w) + [g'(r) - 1]\, D(e^T, r - g(r)) \right) \tag{4}$$

where, for any income stream y_t ($t \in [0, \infty)$), D(y,r) denotes the *duration* of y,

$$D(y,r) \equiv \int_0^\infty \tau \left[\frac{y_t e^{-r\tau} d\tau}{\int_0^\infty y_t e^{-r\tau} d\tau} \right], \tag{5}$$

and where e^T denotes a stream of income that is constant through time T. If one abstracts from the effect of the interest rate on the present discounted value of future earnings (i.e. assume $W_0 = W$), then for the case of $g'(r) = 1$, one obtains $\epsilon_{sr} = 0$. This is intuitive since, in the context of the standard life cycle model, $g'(r) = \sigma = 1$ corresponds to Cobb-Douglas preferences.

Unfortunately, when one departs from the standard life cycle framework, the stability of the function g(r) is no longer assured. With "buffer stock" saving, it is no longer obvious that the Euler equation approach successfully finesses the Lucas critique, since the expected desirability next period's consumption depends upon next period's resources (through the probability that the individual will run out of liquid wealth). The procedure described above may also yield misleading estimates of the interest elasticity of saving under other alternative hypotheses.

To illustrate, imagine that an individual follows a simple rule of thumb that specifies a constant level of consumption for each interest rate r, $c_W(r)$, throughout the pre-retirement period, and a potentially different constant level of consumption, $c_R(r)$, throughout the post-retirement period. For simplicity, also assume that $W_0 = W$. Observing this individual either before or after retirement, one would conclude (based on the Euler equation approach) that $\epsilon_{sr} = 0$. However, it

is entirely possible that $c_R(r)$ and $c_W(r)$ are, respectively, increasing and decreasing in the interest rate. In that case, the actual interest elasticity of saving would be positive. The validity of the procedure for recovering ϵ_{sr} breaks down because the function $g(\cdot)$ depends on t as well as r.

C. Other evidence

In light of the difficulties associated with estimating interest elasticities through either of the two approaches discussed above, it is appropriate to seek out other evidence that could shed light on this issue. Ideally, one would like to find exogenous variation in the rates of return earned by different individuals, and to relate this variation to differences in rates of saving. One obvious candidate for the source of this variation is eligibility to make contributions to a tax-favored savings account, such as an IRA or 401(k). I explore this possibility in the next section.

Another possibility is to identify the interest elasticity of saving by examining variation in the after-tax rate of return within a given class of tax-favored savings plans, such as 401(k)s. One important source of this variation arises from employer matching provisions. Although an increase in the employer match is not mathematically equivalent to an increase in the after-tax rate of return, both have the effect of reducing the price of future consumption, relative to current consumption. Indeed, since the mathematics of an employer match are easier to understand than the mathematics of compound interest, it seems doubtful that the interest elasticity of saving could be positive unless saving responds positively to employer matching provisions within the context of 401(k)s.

Relatively few studies have attempted to relate 401(k) plan provisions, such as employer matches, to the choices of employees. Moreover, the existing studies focus exclusively on 401(k) contributions. Even if 401(k) contributions respond strongly to employer matching provisions, it is conceivable that this response could reflect asset shifting, rather than new accumulation. Thus, a high elasticity of contributions with respect to the match rate would not necessarily establish that the interest elasticity of saving is also substantial. However, it would

be very difficult to reconcile the existence of a substantial positive interest elasticity of saving with the finding that 401(k) contributions do not respond to employer match rates.

The evidence on the effect of 401(k) match rates is mixed. Using survey data gathered by the General Accounting Office, Poterba, Venti, and Wise (1992) conclude that the existence of a match rate is correlated with higher participation, but that the level of the match has little effect. Papke, Petersen, and Poterba (1993) survey a small sample of firms and corroborate this finding. Papke (1992) analyzes data drawn from IRS Form 5500 filings, and finds that the effect of higher match rates is positive at low match rates, but negative at high match rates. Her results are also somewhat sensitive to the introduction of fixed effects. Andrews (1992) studies household level data from the May 1988 Current Population Survey, and concludes that, while the existence of a match increases participation, there is actually a negative relation between the match rate and contributions. Kusko, Poterba, and Wilcox (1993) analyze employee-level data for a single company, and find that contributions and participation are relatively insensitive to changes in the matching rate through time. Scott (1994) argues that most of the negative results on the effects of matching provisions are attributable to the use of *ex post* rather than *ex ante* match rates. Using the 1985-89 Employee Benefit Surveys (for which *ex ante* match rates are available), he finds some evidence that the size of the match matters; however, even Scott's results indicate that most of the effect is attributable to the existence of the match, rather than to its magnitude.

The evidence on match rates is therefore somewhat puzzling. Within the context of the traditional life cycle hypothesis, it is difficult to imagine why employees would respond differently to match rates of 0% and 5%, but behave almost identically with match rates of 5% and 100%. One possibility is that the existence of a match is correlated with the underlying preferences of the employees. This would occur if, for example, high-saving workers sorted themselves into plans with match rates, or demanded that their employers provide matches. There are two problems with this explanation. First, it is hard to understand why the same considerations would not induce a correlation between contributions and the size of the match.

Second, as discussed in section V, there is some reason to believe that matching provisions are adopted as remedial measures to stimulate contributions in instances where employees are predisposed against saving.

The evidence on matching provisions is more easily reconciled with alternative behavioral hypotheses. One can easily imagine that the availability of a match focuses employee attention on the 401(k) plan, authoritatively validates the importance of long-term saving objectives, undermines the resistance of impatient selves (due to its immediacy), and provides additional impetus for establishing a private rule. While higher match rates might add to this impetus, one might well expect the observed discontinuity when match rates are introduced.

IV. Direct responses to tax-favored savings accounts

Even if saving does not in general respond significantly to changes in the after-tax rate of return, this does not necessarily imply that tax-favored savings accounts are ineffective. It is certainly conceivable that individuals respond strongly to tax incentives in the context of certain institutional settings, even if they do not respond more generally. However, to defend this position, one must be prepared to embrace the view that saving is driven in large part by the psychological and institutional factors discussed in section II, and not exclusively by pristine life cycle forces.

The existing literature on tax-favored savings accounts focuses primarily on Individual Retirement Accounts (IRAs) and 401(k)s. I will discuss the available evidence on the efficacy of both. Before doing so, however, it is important to be clear about the nature of the question posed in this section. Specifically, I am attempting to determine whether the saving of those individuals who contributed to IRAs and/or 401(k)s would have been lower had these programs not existed, *all else equal*. I am not, at this point, attempting to learn whether these programs contributed to total personal saving. These are different issues, because the creation or elimination of tax-favored accounts need not leave all else equal. For example, it is conceivable

that those with 401(k)s would, in the absence of statutes authorizing 401(k)s, have had other pensions. I take these kinds of issues up in section V. For the moment, I am using IRAs and 401(k)s merely as opportunities to learn about the "partial equilibrium" response of individuals to focused tax incentives.

A. Individual Retirement Accounts

Individuals without pensions were first permitted to open Individual Retirement Accounts (IRAs) in 1974. These accounts featured tax deductible contributions up to a fixed limit, tax-free accumulation, taxation of principle and interest on withdrawal, and penalties for early withdrawal. Congress extended eligibility to all workers in 1981, and raised annual contribution limits to $2,000 for a single worker, or $2,250 for a married couple with one earner. The Tax Reform Act of 1986 restricted eligibility for *deductible* contributions, based on adjusted gross income (AGI). Deductibility was phased out between AGI of $40,000 and $50,000 for joint filers and between $25,000 and $35,000 for single filers. Individuals with higher levels of AGI remained eligible to make non-deductible contributions up to the same annual limits, and continued to benefit from tax-free accumulation.

Prior to the Tax Reform Act of 1986, IRAs had become quite popular. Annual contributions grew from roughly $5 billion in 1981 to roughly $38 billion in 1986, which represented approximately 20 percent of total personal saving. Contributions plummeted after 1986, falling to less than $10 billion in 1990. While it is indisputable that flows of saving through IRAs were substantial, there is considerable controversy concerning to the extent to which these flows represented new saving. The existing evidence on the efficacy of IRAs falls into five general categories.

1. Direct survey evidence. One approach to measuring the effect of IRAs on saving is simply to ask people how they funded their contributions. In one such survey (Johnson, 1985), about half of respondents said that they would have saved the money anyway, about 10 percent

said that they would have spent all of it, and about 40 percent said that they would have saved some and spent some. Johnson concludes that, on average, individuals reduced consumption by roughly 32 cents to fund each dollar of IRA saving.

These findings should be regarded as, at most, vaguely suggestive. The relevant survey question asks individuals to imagine what they would have done in a counterfactual and purely hypothetical situation. This raises a variety of problems. Respondents may not think very hard about the hypothetical. If they think about it, they may assess the costs and benefits of various decisions differently than they would have in practice. They may accurately report what their *intention* would have been in the hypothetical situation, but actions and intentions do not always coincide. They may also misrepresent their probable intentions in the hypothetical situation, particularly if they believe that some answer is more "virtuous," or if they think that the interviewer is looking for a particular response.

2. Evidence on the frequency of limit contributions. If one adopts a traditional life cycle perspective, then there is ample reason to be concerned about the efficacy of IRAs. Under this view, IRAs have the potential to stimulate saving only if they raise the after-tax rate of return earned on the marginal dollar of saving. But for individuals who save more than the IRA contribution limit, IRAs only affect the rate of return earned on *inframarginal* saving. This creates an "income effect" that increases consumption and decreases saving (both under the LCH and various alternatives -- see section II.B), without an offsetting substitution effect. If, in addition, the interest elasticity of saving is small, then the increase in saving for those saving less than the IRA contribution limit will be modest. It is therefore conceivable that the net effect of the program could be to reduce personal saving, particularly if a large fraction of the population saves more than the contribution limit. To make matters still worse, individuals would have an incentive to reach the limit by shifting existing savings, or by borrowing to fund their contributions.

In light of this argument, many observers have been troubled by the fact that roughly

three-quarters of all contributors saved at exactly the contribution limit (Burman, Cordes, and Ozanne, 1990, and Gravelle, 1991). Others find comfort in the fact that only 30% of IRA contributors reached the contribution limit in each of three consecutive years (Hubbard and Skinner, 1995).

In my view, this focus on the frequency with which IRA participants reached the contribution limit is misplaced, for two reasons. First, the existence of the contribution limit induces a kink in the individual's budget constraint. Many limit contributors may have been at the kink, in the sense that they did not undertake any other long-term saving for which IRA contributions would have been a good substitute. For individuals at the kink, the marginal after-tax rate of return is not well-defined. It is entirely possible, even in the context of the LCH, that the long-term saving of these individuals would have been significantly less than the contribution limit in the absence of the IRA program.

Second, the pessimistic assessment of the data on limit contributors is predicated on the view that the LCH accurately depicts behavior. In section II, I have suggested a variety of channels through which IRAs may have affected saving, besides changing the marginal after-tax rate of return. For example, IRAs may have increased awareness of the need for retirement saving, or they may have validated specific saving targets. Another possibility is that IRAs may have enhanced self-discipline. Even if they did not stimulate *inflows* into households' long-term savings, they may have deterred *outflows* (Thaler, 1994).

3. Correlations between IRA and non-IRA saving. A number of authors have attempted to measure the effects of IRAs on saving through more rigorous econometric analysis. Most of these studies have, with varying degrees of sophistication, examined the underlying correlations between IRA and non-IRA saving activity.

Before describing these studies, it is useful to begin by describing an ideal experiment for assessing the effects of IRAs. The contrast between the ideal data and the available data explains why the measurement of IRA effects has proven so difficult. Imagine that we are given some

large sample of individuals, and that we randomly partition this sample into two subsamples. We treat the individuals in these subsamples exactly the same in all respects (identical initial assets, wages, fringe benefits, working conditions, and so forth), but we permit the individuals in one subsample to contribute to IRAs (the "experimental" group), while withholding this opportunity from the other subsample (the "control" group). In this way, we create exogenous variation in IRA eligibility. We then compare the total saving of individuals in the two subsamples to determine the effects of IRAs.

Unfortunately, between 1982 and 1986, there is no exogenous variation in IRA eligibility. Instead, we observe variations in participation. One could imagine attempting to mimic the ideal experiment by using this variation to identify new "experimental" and "control" groups, in effect asking whether the saving or assets of IRA contributors are higher than, lower than, or the same as the saving or assets of non-contributors. Evidence based on this approach reveals that IRA contributors do not save less in other forms than non-contributors; in fact, they save a good deal more (see e.g. Hubbard, 1984).[3] Unfortunately, this finding tells us very little about the extent to which IRAs displace other saving. Some households save a lot, while some save little. This is presumably attributable to differences in preferences. Since the decision to contribute is endogenous, contributors probably consist of households with stronger preferences for saving. Therefore, one should not be surprised to discover that those who contribute to IRA accounts also save more in other forms than those who choose not to contribute.

In principle, one solution to this problem would be to identify some exogenous variation in IRA contributions that is unrelated to preferences towards saving. One could then use instrumental variables to estimate a specification explaining non-IRA saving or total saving as a function of IRA saving. Since eligibility was universal from 1982 to 1986, a potential source for this variation is difficult to imagine, let alone measure.

Rather than attempt to identify an instrumental variable, the literature has proceeded by reexamining the relation between IRA saving and non-IRA saving, controlling for initial wealth.

This procedure is based on the assumption that two individuals with the same initial wealth must have the same underlying preferences towards saving; thus, the source of the spurious upward bias between IRA saving and total saving is supposedly removed. This approach has been followed in a study by Feenberg and Skinner (1989) and a series of studies by Venti and Wise (1986, 1990, 1991). Analysis of a variety of data sources (including the Michigan Tax Panel, the Survey of Consumer Finances, the Consumer Expenditure Surveys, and the Survey of Income and Program Participation) uniformly demonstrate that total saving is positively correlated with IRA saving, even when one controls for initial wealth. The conditional correlation between IRA saving and non-IRA saving is typically non-negative, which has been taken to indicate that IRA contributions are new wealth.

The central problem with this strategy is that initial wealth may be a relatively poor control for an individual's current underlying disposition toward saving. One problem is that wealth varies for reasons unrelated to tastes for saving (such as the receipt of unexpected inheritances). Another difficulty is that an individual's disposition to save may change through time due to fluctuations in income, household composition, perceived needs, or other factors; thus, the individual's disposition to save during any time period may differ from the dispositions that led to the accumulation of initial wealth at the start of the period. Even if wealth were perfectly correlated with the relevant aspects of tastes, it is well-known that wealth is measured with a great deal of error. Any residual unobserved variation in the current inclination to save that is left after controlling for initial wealth will continue to bias the correlation between IRA saving and non-IRA saving upward: those who, for unobserved reasons, are inclined to save more overall will probably save more in both forms.

The underlying econometric justification for this procedure is also suspect. Suppose it were possible to control perfectly for all aspects of tastes that determine non-IRA saving. Even this would not allow one to calculate the extent to which IRA contributions displace other saving, unless there was some significant exogenous variation in IRA contributions independent of tastes

for saving. But what is the source of this exogenous variation? Since eligibility was universal, it is hard to imagine any significant factor that would have affected IRA saving without also directly affecting non-IRA saving. And if there is no source of exogenous variation, it is difficult to see how the relation of interest could be identified.

In much of their work, Venti and Wise also place additional structure on the data. Specifically, they estimate the parameters of a model in which an individual maximizes a utility function defined over consumption, IRA saving, and non-IRA saving. The specification allows for a range of elasticities of substitution between the two forms of saving. Based on estimates of this model, Venti and Wise conclude that IRA contributions represented new saving, in the sense that they were funded almost entirely by reductions in consumption and income taxes.

The low estimates of the substitution parameter that emerge from estimation of the Venti-Wise model appear to be driven by two considerations. The first consideration is the non-negative correlation (noted above) between IRA saving and non-IRA saving, conditional on initial wealth (which appears in the Venti-Wise model through the budget constraint). For reasons that I have already discussed, this correlation is probably a poor barometer for the true degree of substitutability.

The second consideration has to do with a technical feature of the model. As formulated, the model implies that, if IRA saving and non-IRA saving are perfect substitutes, then no individual would be willing to engage in non-IRA saving until reaching the IRA contribution limit. Since this prediction is manifestly false (many individuals who saved something did not contribute to IRAs), Venti and Wise's estimation strategy automatically guarantees the result that the two forms of saving are imperfect substitutes. This inference is unwarranted. Although it is evident the IRA saving and non-IRA saving must not be perfect substitutes for savers who do not contribute to IRAs (perhaps due to differences in liquidity), it does not follow that these two forms of saving are poor substitutes for individuals who do contribute to IRAs. On the contrary, one could easily imagine that, among IRA contributors, IRAs are quite good substitutes for other

saving. This could occur if, for example, IRA contributors tend to save a lot in all forms, and are therefore relatively unconcerned (on the margin) about liquidity.

Gale and Scholz (1994) estimate an alternative econometric model, in which they permit the parameters of the saving relation to vary according to whether or not an individual is an IRA contributor. This is intended to capture the possibility that those who do not contribute to IRAs may have different attitudes towards IRA and non-IRA saving than those who do contribute. In this way, Gale and Scholz avoid the automatic bias towards low substitution that is present in the analysis of Venti and Wise.

Intuitively, Gale and Scholz identify the degree of substitution between IRA and non-IRA saving as follows. Suppose we measure the marginal propensity to save (out of income) in IRAs ($MPS_{I,N}$), and the marginal propensity to save in other forms ($MPS_{O,N}$) for non-limit contributors, as well as the marginal propensity to save in other forms ($MPS_{O,L}$) for limit contributors. If all IRA saving is new saving, then we should find $MPS_{O,L} = MPS_{O,N}$. On the other hand, if IRA saving simply displaces other saving dollar-for-dollar, we would expect to find $MPS_{O,L} = MPS_{O,N} + MPS_{I,N}$. On the basis of this kind of comparison, Gale and Scholz conclude that a negligible fraction of IRA contributions represent new saving.

The analysis of Gale and Scholz successfully demonstrates that the conclusions of Venti and Wise are highly sensitive to assumptions about the nature and distribution of unobserved preferences. This does not imply, however, that their particular procedure generates reliable estimates of the extent to which IRAs substitute for other forms of saving. On the contrary, this procedure probably has the effect of biasing their findings against the conclusion that IRAs represent new saving. The Gale-Scholz model is identified by the assumption that all IRA contributors have the same preferences towards saving, conditional on a list of covariates, regardless of whether or not they are limit contributors. This homogeneity assumption is distinctly odd, in light of the fact that the their analysis is motivated by the observation that attitudes towards saving differ according to IRA participation status, even when conditioned on

the same list of covariates. It is therefore natural to wonder whether this homogeneity assumption drives their results.

To understand the bias resulting from the Gale-Scholz homogeneity assumption, consider the following illustrative example. Suppose that there are three types of savers, with (respectively) low, medium, and high inclinations to save. Those with greater inclinations to save are assumed to have larger average and marginal saving propensities. Low savers never contribute to IRAs, and are therefore of no further interest to us. As long as moderate savers are not constrained by the IRA contribution limit, they save 5 cents out of each dollar in IRAs, and 5 cents in other forms. If they are constrained by the contribution limit, they still save 5 cents out of each dollar in other forms. As long as high savers are not constrained by the IRA contribution limit, they save 10 cents out of each dollar in IRAs, and 10 cents in other forms. If they are constrained by the contribution limit, they still save 10 cents out of each dollar in other forms. Our final assumption is that all moderate savers end up contributing less than the contribution limit, while all high savers turn out to be limit contributors.

Note that, in this example, all IRA contributions represent new saving. However, applying the Gale-Scholz procedure, one would calculate that $MPS_{O,L} = 0.10 = 0.05 + 0.05 = MPS_{O,N} + MPS_{I,N}$, and infer incorrectly than IRA saving completely displaces other forms of saving. I have constructed this particular example to demonstrate that the bias could be quite large. Obviously, hypothetical examples cannot establish the magnitude of the actual bias. However, the principle (and therefore the direction of the bias) generalizes: heterogeneity among those who contribute to IRAs typically implies that those who contribute more (and who therefore have higher average propensities to save) will also tend to have higher marginal propensities to save. As a result, the data will appear to show that the marginal propensity to save in forms other than IRAs rises as contributions pass the allowable limit. But this is precisely the pattern that Gale and Scholz would interpret as evidence of displacement.

Some authors argue that correlations between IRA saving and non-IRA saving are

particularly informative for new contributors. Using SIPP data from 1984 and 1985, Venti and Wise (1995) demonstrate that the inception of IRA contributions for a household does not coincide with a significant decline in other financial assets. They interpret this to mean that even new contributors engage in very little asset shifting to fund contributions, and that these contributions must therefore represent new saving. Yet the observed patterns do not rule out the possibility that new contributors were simply new savers, and that these new savers would have increased non-IRA savings in the absence of IRAs. Consequently, the evidence is consistent with significant asset shifting.

Attanasio and De Leire (1994) undertake a similar exercise, but suggest that it is appropriate to evaluate the behavior of new contributors treating old contributors as a control group. If new contributions come from consumption, then (it is argued) new contributors should exhibit slower consumption growth, and essentially the same growth in non-IRA assets, as old contributors. In contrast, if new contributions come from saving, then new contributors should exhibit the same growth in consumption, but slower growth in non-IRA assets than old contributors. The authors implement this test using the Consumer Expenditure Surveys, and find the second of these patterns. They conclude that IRA contributions primarily reflect asset reshuffling, rather than new saving.

Unfortunately, there does not appear to be any compelling justification for using old IRA contributors as a control group. Without access to IRAs, new contributors might well have had higher consumption growth and slower growth of non-IRA assets than old contributors.[4] The availability of an IRA might then reduce the consumption growth of new contributors while leaving the growth of their non-IRA assets unaffected, thereby producing the observed pattern. Moreover, even if the control group were appropriate, the Attanasio-De Leire finding would at most establish that each household's initial IRA contributions are funded from asset shifting; it would be entirely consistent with the view that later IRA contributions reflect new saving.

4. Exogenous changes in eligibility. Another possible approach to mimicking the ideal

experiment is to exploit the exogenous variation in IRA eligibility that existed prior to 1982 and after 1986. For example, one could imagine estimating a regression explaining non-IRA saving as a function of IRA contributions using eligibility as an instrument, or directly as a function of eligibility. There are two problems with this suggestion; one is conceptual, the other practical. Conceptually, a problem arises because, in contrast to the ideal experiment, IRA eligibility was non-random. Eligibility was triggered by the absence of pension coverage prior to 1982, and by a combination of pension coverage and AGI after 1986. Since both pension coverage and income are potentially important determinants of household saving, concerns about correlations with underlying preferences are still present. The practical problem arises because, with certain data sources, eligibility is difficult to assess. Information on pension coverage is sometimes unavailable, incomplete or inaccurate, and one must extrapolate AGI from income.

The concern that IRA eligibility (prior to 1982 or after 1986) might have been correlated with preferences towards saving leads to a slightly more sophisticated suggestion. If the heterogeneity in preferences is captured by an individual-specific fixed effect, then it should be possible to eliminate this heterogeneity by differencing saving. One can then relate changes in saving to changes in eligibility, which differed across individuals both in 1982 and 1987. The impact of IRAs is then, in effect, inferred from differences in differences. For example, using panel data that crosses 1982, one attempts to determine whether those who became eligible for IRAs increased their saving by more than those who remained eligible.

This is the general approach taken in Joines and Manegold (1995) and Engen, Gale, and Scholz (1994). Both of these studies make use of the IRS/University of Michigan Tax Panel. Unfortunately, this data set contains no information on pension coverage, and therefore provides no way to measure IRA eligibility prior to 1982. Of course, individuals who contributed to IRAs prior to 1982 must have been eligible. Joines and Manegold therefore propose using this as the control group. By defining the control group in this way, they tend to select individuals who have the highest predispositions to save among the eligible population. To counteract this

selection effect, they use as their experimental group a sample of individuals who also contributed to IRAs (and therefore who also have high predispositions to save), but who began to contribute after 1982.[5] While this experimental group includes some individuals who were eligible prior to 1982, it also includes many individuals who became eligible as of 1982. Therefore, on average, allowable contributions increased by a larger amount for members of the experimental group than for members of the control group. Both studies nevertheless demonstrate that there is relatively little difference between the changes in saving across 1982 for the experimental and control groups. They conclude that IRAs had, at most, a moderate effect on saving (less than 30 cents on the dollar).

One difficulty encountered by Joines-Manegold and Engen-Gale-Scholz is that the IRS/University of Michigan Tax Panel does not contain measures of either saving or wealth. The authors are compelled to impute wealth from dividend and interest income. They then difference estimated wealth to obtain a measure of saving. This variable is the focus of their differences-in-differences analysis. Thus, their key results are based on third differences (twice across time and once across subgroups) of an imputed variable. One must seriously question how much "news" is left over after these operations. Not surprisingly, the key effects are generally estimated with large standard errors, and one typically cannot reject the hypothesis that a large fraction of IRA contributions represented new saving.

The selection criteria used to construct the control subgroup and the experimental subgroup are also potentially problematic. It is doubtful that these groups have comparable characteristics or similar dispositions to save. The differences-in-differences procedure is ostensibly designed to handle this problem, since it removes the fixed effect for each group. However, the validity of this solution depends critically on two assumptions: that tastes enter the saving equation additively, and that tastes do not affect the size of the response to a given change in the policy variable. In this context, the second assumption is objectionable.

To further explore this point, suppose that the saving of group i at time t is given by the

following equation:

$$s_{i,t} = \mu_i + \alpha_t + \eta_i M_{it} ,$$

where μ_i and η_i are fixed group-specific coefficients, α_t is a time effect, and M_{it} is the IRA contribution limit applicable to this group. One would expect μ_i and η_i to be positively correlated, since higher savers are more likely to respond to an increase in the contribution limit. The differences-in-differences estimator is then

$$\Delta s_{et} - \Delta s_{ct} = \left[\eta_e \Delta M_{et} - \eta_c \Delta M_{ct} \right],$$

(where "e" indicates the experimental group, and "c" indexes the control group). Note that one will correctly estimate the effect of the policy change on the experimental group as long as $\eta_e = \eta_c$ (i.e. if there is no heterogeneity in the response to a given change in policy), or if $\Delta M_{ct} = 0$ (i.e. the control group does not experience a change in the policy variable). In this instance, neither condition applies: it is likely that heterogeneity in preferences towards saving (as reflected in η_i) remains, and contribution limits were raised for the control group (albeit to a lesser extent than for the experimental group, so that $\Delta M_{et} > \Delta M_{ct} > 0$).

The resulting bias in the estimates depends on whether the control group is innately more inclined to save or less inclined to save than the experimental group. Suppose for the moment the control group consists of particularly high savers, so that $\eta_c > \eta_e$. Then the sign of the differences-in-differences estimator becomes ambiguous, even if an increase in the contribution limit actually stimulates saving for both groups. To take an example, if a $2,000 increase in the contribution limit induces a $1,000 increase in the average IRA saving of the control group and a $250 increase in the average saving of the experimental group (because the control group largely consists of more highly motivated savers), then a $500 increase in the contribution limit for the control group (e.g. from $1,500 to $2,000) and a $2,000 change in the contribution limit for the experimental group (e.g. from $0 to $2,000) will have the same total effect on saving ($250).

Unfortunately, with the available data, it is impossible to test whether the control group is more or less predisposed to undertake long-term saving than the experimental group. However, the following is suggestive. Prior to 1982, only a tiny fraction of those eligible for IRAs actually made contributions. While these individuals had one characteristic that might be indicative of a predisposition for low saving (no employer pension), they were nevertheless a very highly selected subset of this population. The fact that they both discovered and took advantage of a little-known IRA provision suggests that they may have been exceptionally motivated to save for retirement. In contrast, since a much larger segment of the population contributed to IRAs after 1982, and since IRAs were more widely publicized, the experimental group may be less highly selected. If this is the case, then the differences-in-differences estimator understates the true effect on saving of an increase in the IRA limit. Of course, if the opposite proposition is true, then the differences-in-differences estimator overstates the effect.

As is well-known, the differences-in-differences estimator may go awry for other reasons as well. One obvious possibility is that other changes in the economic environment may have affected the two groups differently. Since the changes in IRA eligibility were accompanied by other significant tax changes, as well as a variety of important macroeconomic developments (including large changes in inflation and interest rates, as well as business cycle effects), attributing the difference-in-difference of saving exclusively to relative changes in IRA eligibility is dicey.

Finally, it is important to realize that, under some of the behavioral alternatives to the LCH, the procedure used by Joines-Manegold and Engen-Gale-Scholz would be incapable of detecting certain kinds of links between IRAs on personal saving. Suppose, for example, that the expansion of the IRA program stimulated saving by enhancing awareness of retirement issues and triggering aggressive promotion of investment vehicles (see the discussion of the evidence on psychological effects, immediately below). If these developments affected members of the control group and the experimental group equally, the differences-in-differences estimator would

falsely indicate no increase in saving.

5. Evidence of psychological effects. The chain of logical reasoning that leads one to doubt the efficacy of IRAs is largely predicated on the view that saving is a consequence of rational and deliberate life cycle planning. It is therefore possible to shed light on the key issue by asking whether other aspects of individuals' responses to IRAs are consistent with the predictions of standard life cycle theory. If they are not consistent, then one should be very cautious about drawing inferences concerning the efficacy of IRAs from anything but the most direct evidence.

The literature identifies a number of patterns in IRA contributions that appear to be anomalous from the perspective of the standard model. The following four are particularly provocative.

First, it is difficult to account for the explosion of IRAs after 1982 and the collapse of IRA contributions after 1986, unless one credits the role of visibility and promotion (Long, 1990, Venti and Wise, 1992). Recall that only 1 percent of taxpayers made contributions to IRAs prior to 1982, despite the fact that roughly half were eligible to contribute up to $1,500. This figure rose to 15 percent by 1986. Recall also that many individuals remained eligible to make deductible IRA contributions after 1986 (those with sufficiently low AGIs, or without pension coverage); moreover, all other individuals could still make non-deductible contributions and benefit from tax-free accumulation. Yet the fraction of taxpayers contributing to IRAs dropped to 4 percent by 1990. IRA contributions followed promotional activity (which exploded after 1982 and contracted after 1986) much more closely than actual economic incentives.[6]

Second, there has been a pronounced tendency for individuals to delay their IRA contributions until the end of a tax year (Summers, 1986). This is puzzling because minimization of tax liabilities requires taxpayers to make these contributions as early as possible. To some extent, the tendency to delay contributions may result from the desire to maintain liquidity throughout the tax year (Engen, Gale, and Scholz, 1994). But, even allowing for the

potential importance of liquidity, it is difficult to explain why more IRA contributors (particularly those with significant non-IRA assets) do not at least make a series of smaller contributions during the course of the tax year (Bernheim, 1994b).

Third, individuals are significantly more likely to make IRA contributions if they owe the IRS money at the end of the tax year. Feenberg and Skinner (1989) interpret this as an indication that, psychologically, individuals would rather write a check to an IRA account than write a somewhat smaller check to the IRS. It is conceivable that this result could reflect spurious correlations of both underwithholding and IRA contributions with some third factor, such as income, tax filing status, or asset holding (Gravelle, 1991). However, the pattern is apparent even when Feenberg and Skinner include plausible controls for these factors.

Fourth, there is considerable evidence of "focal point" saving. Engen, Gale, and Scholz (1994) find that, among those who could have contributed more than $2,000 but who contributed less than the limit, 47 percent contributed exactly $2,000. This finding invites the interpretation that the well-publicized, "officially endorsed" $2,000 figure created a focal target for saving, and that the very existence of this target may have influenced the behavior of many less serious savers (such as those contributing less than the limit).

One alternative explanation for this phenomenon concerns transactions costs. While single-earner married couples could contribute up to $2,250 per year, contributions in excess of $2,000 would have required them to open a second account. A contribution of $250 might seem insufficient to justify the effort. However, it is important to bear in mind that the one-time costs of opening the account must be weighed not against the benefits of a single $250 contribution, but rather against the benefits of a $250 contribution that recurs for many years. Moreover, even among those with a $4,000 limit, 38 percent of those contributing less than the limit contributed exactly $2,000. The transactions cost hypothesis is therefore implausible.

Others have argued that the focal point saving phenomenon results from bargaining among spouses with conflicting objectives (Burman, Cordes, and Ozanne, 1990). Yet it is hard

to see how this would emerge in a formal model of household bargaining, without the introduction of significant transactions costs.

In summary, much of the evidence on IRA contribution patterns remains anomalous. This evidence raises serious concerns about the validity of logical inferences and statistical analyses that are predicated on the assumption that IRA contributions were motivated entirely by rational and deliberate life cycle planning.

B. 401(k)s

Employers were originally authorized to establish 401(k) plans in 1978, but this option remained unpopular until after the Treasury issued clarifying regulations in 1981. In many ways, 401(k)s are similar to IRAs: contributions are tax deductible up to specified limits, the returns to investments are accumulated tax free, and there are restrictions on early withdrawals. There are also a variety of differences. A number of these differences lead one to suspect that 401(k)s might be more effective at stimulating personal saving than IRAs. Since contribution limits are much higher, they are less likely to bind. As a result, 401(k)s can increase the marginal after-tax rate of return for a much larger set of households. This effect is often reinforced through provisions whereby employers match employee contributions. The structure of a 401(k) plan also capitalizes more effectively on the psychology of saving. Since contributions occur through regular payroll deductions rather than through discretionary deposits, 401(k)s are more conducive to the exercise of self-discipline. Higher contribution limits may also provide authoritative validation for higher saving targets. Since 401(k)s are organized around the workplace, they may also create positive spillovers between employees (e.g. through conversations among employees and other "peer group" effects).

From the perspective of econometric modeling, one of the most salient differences between IRAs and 401(k)s is that eligibility for 401(k)s is determined at the level of the employer. This has two implications. First, at all points in time there is substantial variation in

401(k) eligibility across households. Second, at least some of the variation in eligibility (and therefore in contributions) arises from sources that are plausibly exogenous to the individual. These considerations make it easier *in principle* to identify the effects of 401(k)s.

Studying 401(k)s *in practice* is made considerably more difficult by the relative scarcity of good data. For example, none of the available waves of the Survey of Consumer Finances contains a clean measure of 401(k) eligibility. Of the standard public use data sources, only the SIPP contains good information on eligibility, participation, and asset balances for 401(k)s. Unfortunately, the SIPP does not provide longitudinal information that is useful for studying these plans. The literature has therefore treated the SIPP as a series of three cross-sections (1984, 1987, and 1991). An additional limitation of this data is that 401(k) plan balances are not available in 1984. Taken together, these limitations seriously handicap efforts to measure the behavioral effects of 401(k)s. Nevertheless, the literature has developed and explored two separate estimation strategies that attempt to finesse these limitations.

1. Exploiting exogenous variation in eligibility. Imagine for the moment that each firm's decision to offer a 401(k) is completely random. Then 401(k)s would provide the perfect experimental setting for studying the effects of saving incentives. Eligibility is certainly not random, since it is demonstrably correlated with a variety of individual characteristics (such as income). But as long as variation in 401(k) eligibility is orthogonal to the unobserved individual characteristics that determine saving, the experiment is still a clean one.

Poterba, Venti, and Wise (1994, 1995) proceed from the assumption that 401(k) eligibility is exogenous to the process that determines saving. Using the 1987 and 1991 waves of the SIPP, they demonstrate that, controlling for other relevant factors, eligibility is significantly correlated with financial wealth. Indeed, eligibility has very little effect on non-401(k) financial wealth. They interpret this finding as an indication that virtually all 401(k) contributions represented new saving.

The central problem with this procedure is that 401(k) eligibility is probably not

exogenous. On the contrary, there are several reasons to suspect that eligibility would be significantly correlated with the underlying predisposition to save (Bernheim, 1994c, Engen, Gale, and Scholz, 1994). First, employees with tastes for saving probably tend to gravitate towards jobs that provide good pension coverage, including 401(k)s. Second, employers frequently install 401(k) plans as a direct response to expressions of employee interest (Buck Consultants, 1989).

Poterba, Venti, and Wise themselves provide convincing evidence of endogeneity. A careful reading of their tables reveals that differences in median financial assets between eligibles and non-eligibles are often several times as large as 401(k) balances for eligibles (see also Engen, Gale, and Scholz, 1994). Using a more recent cross-sectional household survey sponsored by Merrill Lynch, Bernheim and Garrett (1995) corroborate this finding. Thus, unless one believes that 401(k)s crowd-in other forms of saving at the rate of four or five to one, the evidence must be construed as indicating that eligibility is strongly correlated with the innate inclination to save.

As in the case of IRAs, one could, in principle, attempt to control for the effects of tastes by estimating an equation explaining *flow* saving as a function of 401(k) eligibility, controlling for initial wealth. This approach would be far more coherent in the current setting: since there are identifiable, exogenous sources of variation in 401(k) eligibility, purging 401(k) eligibility of endogenous influences is actually a sensible objective.

While the inclusion of wealth as an explanatory variable would not eliminate bias in the estimated coefficient of 401(k) eligibility, the direction of the bias would no longer be obvious. To understand this point, one must consider the partial correlation between eligibility and flow saving, controlling for initial wealth. Imagine two individuals who are the same in all observable respects (including initial wealth), except that one is eligible for a 401(k), while the other is not. Suppose for the moment that 401(k)s actually stimulate saving to some unknown extent. It is very likely (due to the presence of high serial correlation in eligibility) that the eligible individual was also eligible in past years. Thus, without eligibility, this individual's initial wealth would

have been *lower* than that of the individual who is actually ineligible. The implications of this deduction are striking: under *identical* conditions (including eligibility), the eligible individual would have accumulated less wealth than the ineligible individual. This suggests that the ineligible individual is *more* inclined to save (given the observation that initial wealth is the same). If so, then assuming that 401(k)s stimulate saving, the estimated coefficient of eligibility would be biased downward. The same argument implies that if 401(k)s reduce saving, the estimated coefficient would have an upward bias; however, this is not particularly worrisome, since the primary focus of controversy concerning 401(k)s is the size of the effect, rather than the direction.

In practice, the direction of the bias associated with the proposed estimation strategy is less clear cut. The force of the argument in the preceding paragraph lessens if wealth varies for exogenous reasons (e.g. it is possible that the ineligible individual had the same initial wealth only because he or she received an unexpected inheritance). In the extreme, if measured wealth is pure noise, controlling for wealth obviously cannot purge eligibility of endogenous variation. Nevertheless, evidence of this kind adds one more piece to the puzzle.

Bernheim and Garrett (1995) implement this strategy using the Merrill Lynch household survey data, which contains self-reported measures of flow saving (expressed as fractions of income), as well as categories of wealth. They find that 401(k) eligibility is associated with an increase in the retirement saving rate that exceeds the corresponding increase in the overall rate of saving. Thus, the use of wealth as a control for differences in tastes eliminates the anomalous finding that 401(k)s crowd in other forms of saving at an implausibly high rate. Unfortunately, their study does not permit one to draw reliable inferences concerning the degree of crowding-out.

2. Exploiting transitional effects. A second approach to measuring the effects of 401(k)s does not require one to assume that eligibility is exogenous. Instead, this approach exploits the fact that the legislative authorization for 401(k)s was relatively recent. To understand this

approach, first imagine two idealized worlds, one in which 401(k)s have always been available, and one in which 401(k)s have never been available. Suppose for simplicity that each economy has converged to a steady state with a stable age-wealth profile. This profile may well be higher for the world in which 401(k)s have always been available, but this does not necessarily indicate that 401(k)s stimulate saving, since there may be other differences (such as tastes) between the two worlds. Now imagine a world in which 401(k)s have never been available in the past (so that this economy has also converged to a steady-state age-wealth profile), but where they are established unexpectedly as of some point in time (without any change in tastes). At that point, each individual departs from his or her initial wealth trajectory, and begins to move along some new wealth trajectory. Eventually, the age-wealth profile will converge to a new steady state. But during the transition period, if 401(k)s stimulate saving, this profile should begin to shift upwards relative to the profile from any world in which eligibility is unchanged.

In the ideal implementation of this strategy, one would identify a large group of workers who became eligible for 401(k)s relatively soon after the enabling legislation (say before 1984) and who remained eligible in all subsequent years, as well as a large group of workers who never became eligible for 401(k)s. One would then follow these same individuals through time, estimating age-wealth profiles for each group in each year. The relative amplitudes of these profiles in any particular year would prove nothing, since eligibility may be related to preferences. However, as time passes, the number of years of accumulated eligibility for the first group increases. Therefore, the cross-sectional age-wealth profiles for the eligible group should shift upward relative to the profile of the ineligible group.

Unfortunately, as mentioned above, good panel data on 401(k)s is not available. Poterba, Venti, and Wise (1995) therefore implement this strategy for a series of cross-sections (1984, 1987, 1991) obtained from the SIPP. In each year, they compare the accumulated financial assets of those who are eligible for 401(k)s and those who are not eligible. The data unmistakably show the predicted upward shift in relative financial assets held by those who are eligible for 401(k)s.

Indeed, there is no noticeable decline in the relative level of non-401(k) financial assets held by this group.

This finding supports the hypothesis that individuals funded 401(k) contributions through a combination of reduced taxes and spending, and not by diverting funds that they would have saved in any event.

Of course, Poterba, Venti, and Wise depart from the ideal strategy by using an unrelated sequence of cross-sections. It is important to consider how this deviation is likely to affect their results. If successive cross-sections of eligibles (and ineligibles) are simply random draws from the same population of eligibles (ineligibles), then there is no problem. A problem only arises if the composition of the population of eligibles (ineligibles) changes systematically through time.

It is virtually certain that some compositional changes occurred between the successive surveys used by Poterba, Venti, and Wise, for the simple reason that new workers became eligible. One obvious implication is that, as one moves forward in time by, say, four years, the average length of exposure to 401(k)s within the eligible population increases by less than four years. One can imagine cases in which this could create problem. For example, if 401(k)s pass through a period of sufficiently rapid growth, the average length of eligibility among eligibles could actually decline. However, under more plausible assumptions, the effect would simply be to slow the observed *rate* at which the assets profile of the eligible population shifts relative to the profile of the ineligible population. The underlying logic of the exercise remains valid, as long as newly eligible workers are drawn from the same population as those who have been eligible for longer periods.

A serious issue arises because newly eligible workers are probably systematically different from those who have been eligible for longer periods. As I have already discussed, eligibility for 401(k)s is almost certainly endogenous. It stands to reason that the most motivated "serious" savers would have sought out employers who offered 401(k)s, or encouraged their existing employers to provide 401(k)s, relatively soon after these plans became available. Less

motivated, "occasional" savers would have been less likely to seek out or agitate for 401(k)s, and more likely to drift into these plans slowly through time. Thus, as time passes, the eligible population becomes increasingly skewed towards less motivated savers. Bernheim (1994b) refers to this as the "dilution" effect. It is likely that the dilution effect became more severe after 1986, when more demanding non-discrimination requirements were established for private pensions.

Dilution almost certainly biases the findings of Poterba, Venti, and Wise. However, the direction of the bias is critical. Since dilution creates a *downward* shift in the age-wealth profile, it has the potential to partially offset, completely offset, or even reverse the upward shift due to the behavioral effect of 401(k)s. To see that this is more than a theoretical possibility, one can examine changes through time in the relations between 401(k) eligibility and variables that provide stable proxies for underlying tastes. One plausible proxy for the predisposition to save is ownership of an IRA. It is doubtful that this taste proxy is *stable* for the period of universal IRA eligibility (prior to 1987), since dilution probably affected the set of IRA participants in the same way that it affected the set of 401(k) participants. However, dilution of the IRA population probably declined significantly once eligibility for IRAs was restricted. It is therefore plausible that IRA ownership is a stable taste proxy for the 1987-1991 period. Notably, the fraction of individuals eligible for 401(k)s who owned IRAs declined significantly (relative to ineligibles) between 1987 and 1991; this is a good indication of the dilution effect.[7] Moreover, when one controls for IRA ownership, there is even stronger evidence that the financial asset profiles of eligible workers shift upward after 1987 (Poterba, Venti, and Wise, 1995, 1996, and Engen and Gale, 1995). Thus, Poterba, Venti and Wise may understate the extent to which 401(k)s stimulate saving.

Engen, Gale, and Scholz (1994) criticize Poterba, Venti, and Wise on the grounds that the eligible population differs significantly from the ineligible population. While the conceptual framework of Poterba, Venti, and Wise allows for this possibility, the concern is nevertheless a legitimate one. To the extent the two groups are not comparable, their age-wealth profiles may

have been shifting relative to each other for reasons that are entirely unrelated to 401(k)s. Although this does not provide a reason to believe that the findings of Poterba, Venti, and Wise are biased in any particular direction, it does reduce one's confidence in the validity of their results.

Engen, Gale, and Scholz propose and implement an estimation strategy that uses different experimental and control groups. Specifically, they restrict attention to 401(k) *participants*, and compare them to individuals with IRAs who are ineligible for 401(k)s. The purpose of this strategy is to isolate groups of "high savers" who are more closely comparable than the groups considered by Poterba, Venti, and Wise. Although the authors do not present convincing evidence on the relative degree of comparability, their argument has at least the ring of plausibility.[8] Since eligibility for IRA contributions was restricted in 1986, it is also argued that one would still expect to see an upward shift in the relative assets of 401(k) participants, provided that 401(k)s are effective. On the contrary, the authors find that the total financial assets of 401(k) participants actually declined between 1987 and 1991, whereas the total financial assets of the control populations actually increased. They interpret this as an indication that 401(k)s did not increase personal saving.

It is important to realize, however, that the approach favored by Engen, Gale, and Scholz still suffers from the dilution problem, and therefore is biased against the finding that 401(k)s stimulated personal saving. It is also apparent that, by changing the selection criteria used to define the samples, the authors have probably altered the extent of dilution for the control group, and the extent of reverse dilution for the experimental group (Bernheim, 1994b). Therefore, one should not be too surprised that their findings differ from those of Poterba, Venti, and Wise.

The primary merits and drawbacks of these studies can be summarized succinctly as follows. Poterba, Venti, and Wise estimate that 401(k)s have a large effect on personal saving, and there is at least one good reason to believe that their estimate is biased downwards.[9] Confidence in this result is limited by concerns about the comparability of the experimental and

control groups. There is, however, no reason to believe that non-comparability biases their result in any particular direction, and the results are robust with respect to some efforts to improve comparability (i.e. controlling for IRA status). Engen, Gale and Scholz estimate that 401(k)s have essentially no effect on personal saving. Their experimental and control groups are probably more closely comparable that 401(k) eligibles and ineligibles. However, there is at least one good reason to believe that their estimate is biased downward, and there is no reason to suspect that the magnitude of the bias is similar to that encountered by Poterba, Venti, and Wise.

It is noteworthy that the preceding discussion focuses almost entirely on financial assets. This is a limitation, since it is conceivable that 401(k)s may displace other forms of wealth as well. Engen, Gale, and Scholz present estimates based on broader measures of wealth, and these exhibit similar patterns. In contrast, Poterba, Venti, and Wise focus exclusively on financial wealth. It is therefore natural to wonder whether the inclusion of other forms of wealth would undermine their conclusion.

Engen and Gale (1995) attempt to replicate the analysis of Poterba, Venti, and Wise, while at the same time extending it to broader measures of wealth. Alert to the dilution problem, they attempt to control for unobserved preferences by including a measure of IRA participation status. Their results indicate that mortgages grew and home equity fell in successive cross-sections (1987 and 1991) for the 401(k)-eligible population (both IRA participants and IRA non-participants), resulting in smaller overall wealth growth than for the control groups. They interpret this finding as an indication that 401(k) saving was offset almost completely by larger mortgages.

One difficulty with this interpretation of the Engen-Gale findings is that their results for total wealth are extremely imprecise. Typically, they cannot rule out (at conventional level of confidence) the possibility that 401(k)s contributed significantly to total wealth accumulation. This raises the possibility that their finding might not be robust. In fact, Poterba, Venti, and Wise (1996) find exactly the opposite pattern when comparing 1984 and 1986. It is conceivable that

some of the difficulties here could be linked to data problems, in which case the focus on financial wealth would be justified.

More importantly, Engen and Gale do not "solve" the dilution problem by controlling for IRA status. The trend in the probability that the typical 401(k) worker owns an IRA is properly regarded as a symptom of dilution, rather than as the source of the underlying problem (Bernheim, 1994b). It is highly unlikely that this single binary variable adequately controls for the full variation of preferences towards saving among eligibles and ineligibles. Because of the residual bias, the procedure is only informative as a "one-tailed" test: if the results indicate that 401(k)s stimulate saving *despite* dilution, then one can safely infer something about behavior; however, if the results indicate that 401(k)s do not stimulate saving, one cannot rule out the possibility that dilution has obscured a significant behavioral effect.

Engen and Gale attempt to bolster their contention that no dilution remains after controlling for IRA status by pointing to evidence on home values. Specifically, they note that, conditional on IRA status, home values rose by virtually identical amounts for those eligible for 401(k)s and for those not eligible (in both cases, there was a small increase). It is, however, difficult to understand why this constitutes evidence against dilution. Homes are not like other assets, in that home ownership has an important consumption component, as well as an investment component. Those who are less inclined to save may stretch to buy larger homes, and finance these purchases through greater leverage. Consequently, when Engen and Gale find no evidence of a decline in home values coupled with a 17.9% increase in median leverage for the 401(k) samples, it is entirely natural to construe this as further evidence of dilution.

Engen and Gale also defend their interpretation of the evidence by pointing out that the age-wealth profiles of 401(k)-eligible renters did not shift up relative to those of ineligible renters during the 1980s. Thus, according to Engen and Gale, there is no evidence that 401(k)s contribute to personal saving once one purges the sample of observations for which housing wealth is a potential "contaminant." While the results for renters are thought-provoking, their

proper interpretation is unclear. Renters as a whole are a peculiar group in that they save practically nothing to begin with (Congressional Budget Office, 1993). Those who are eligible for 401(k)s do accumulate significant financial assets (though significantly less than comparable homeowners); however, the median net worth of those who are not eligible for 401(k)s is near zero. These observations have two implications. First, the sample of eligible renters appears to be more highly selected than the sample of eligible homeowners. As a result, eligibility for 401(k)s may be more strongly related to underlying tastes among renters than among homeowners. Sample selection biases and the associated effects of dilution should therefore play out differently in the two samples. It would not be surprising if eligible renters, being more highly selected to begin with, were subject to greater dilution with the passage of time. This would account for the Engen-Gale finding. Second, sample selection issues aside, the absence of significant wealth among ineligible renters can potentially invalidate the methodology used to draw inferences about the effects of 401(k)s. If economic forces were tending to depress saving by renters during the relevant time period, then the absence of a downward shift in the age-wealth profile for eligible renters would indicate that 401(k)s stimulated saving by this group. In theory, the Engen-Gale procedure would detect this by noting a downward shift in the age-wealth profile for ineligible renters. However, in practice, liquidity constraints would have prevented this downward shift from occurring.

It is apparent from this discussion that the evidence presented by Engen and Gale does not discriminate between the hypothesis that 401(k) accumulation displaced net housing wealth (primarily through larger mortgages), and the hypothesis that 401(k) contributions represented new saving. However, other information casts doubt on Engen and Gale's interpretation of this evidence.

To begin with, it is readily apparent that most individuals view and treat housing equity as a highly imperfect substitute for financial assets. In the past, retirees have proven extremely reluctant to downsize their homes or take reverse annuity mortgages for the purposes of using

their home equity to finance normal living expenses during retirement (Venti and Wise, 1989). Among younger workers (who are supposedly doing the substitution), roughly 60 percent say that they view their home equity primarily as a source of financial security, to be used in retirement only in the event of a major emergency (such as entering a nursing home); roughly 23 percent say that they think of it primarily as a method of passing wealth to their children through bequests; and only 14 percent say that they think of it primarily as a store of wealth to finance ordinary living expenses during retirement (Bernheim, 1995a).

Another reason to doubt Engen and Gale's interpretation of the evidence on housing wealth is that 401(k) contributions are largely voluntary. We are asked to imagine that individuals choose to save through 401(k)s, only to undo this saving by borrowing more against their homes. One must wonder why these individuals would voluntarily undertake offsetting transactions.

One possibility is that 401(k) contributors rationally expect to benefit from tax arbitrage, borrowing at the post-tax rate, and investing at the pre-tax rate. However, further reflection suggests that efforts to conduct arbitrage might well prove unprofitable. It is not at all obvious that the appropriately risk-compensated rate of return on 401(k) investments exceeds the after-tax cost of borrowing on mortgages (which, aside from variations in the rate of inflation, are essentially riskless from the point of view of the borrower). Moreover, many strategies for borrowing against one's home (first mortgages, second mortgages, and refinancing) are inherently lumpy, whereas 401(k) contributions are generally made in rather smooth flows. If an individual used one of these devices to fund 401(k) contributions, he or she would be compelled to hold a substantial portion of the mortgage proceeds outside of the 401(k) for a significant period of time, earning the *after-tax* rate of return (which is certainly well below the after-tax cost of borrowing). It might be possible to avoid this problem through a home equity line of credit, which would allow the homeowner to increase the size of his or her mortgage gradually. But the interest rates on home equity lines of credit are much higher than the rates on ordinary

mortgages, and may well exceed the appropriately risk-compensated, pre-tax rates of return on investments, even allowing for deductibility of mortgage payments.

The potential gains from arbitrage activities are certainly enhanced in the presence of employer matching provisions. However, if this is an important motivation for 401(k) contributors, then we would expect to observe a strong positive relation between contributions and matching rates. On the contrary, there is little evidence that such a relation exists (see section III.D).

A second possibility is that 401(k) contributors irrationally expect to benefit from tax arbitrage. While this is certainly conceivable as a matter of principle, the most obvious versions of the irrationality hypothesis are contradicted by the data. Suppose, for example, that individuals fail to understand the risks associated with equity investments, and mistakenly compare the after-tax costs of borrowing with the pre-tax rate of return on the stock market. If this were the case, then one would expect 401(k) contributors to invest their plan balances heavily in equities. On the contrary, participants have tended to invest much more heavily in safe, fixed income securities, particularly during the period in question (see e.g. Arthur D. Little, Inc., 1993).

A third possibility is that myopic behavior leads 401(k) participants to take out larger mortgages. To illustrate this possibility, imagine that 401(k) contributors save less as renters, not because they rationally substitute one form of saving for another, but simply because they have less disposable income. When these individuals begin to think about purchasing a home, they will have less money purchased for a downpayment. Conceivably, they might buy the same house as in the absence of a 401(k), and take out a larger mortgage. This would produce a pattern of 401(k)-mortgage substitution, even if no arbitrage was intended. However, since it is quite common for households to borrow as much as possible against their first homes, one would also expect the reduction in resources available for downpayments to induce many households to purchase smaller homes, or to delay home purchase. One would then observe a downward shift

in home values among eligibles, relative to ineligibles. But Engen and Gale's findings are inconsistent with this prediction.

A final reason to doubt Engen and Gale's interpretation of the evidence on housing wealth concerns the level of mortgage activity. Engen and Gale's figures show that, as of 1991, only a little more than 14 percent of those eligible for 401(k)s had home equity loans. Based on other data in their paper, I would estimate that roughly 12 percent of those eligible for 401(k)s in 1991 extracted equity from their homes through refinancing during the previous four years.[10] According to Gale (private communication), 19% of those eligible for 401(k)s in 1991 purchased primary homes between 1987 and 1991. Given the likelihood of overlap between these groups, it is extremely doubtful that more than 30 to 40 percent of those eligible for 401(k)s behaved in ways that would have permitted them to extract any equity from their homes between 1987 and 1991.[11]

In light of this discussion, it seems unlikely -- though certainly not impossible -- that the Engen-Gale results are attributable to substitution between 401(k)s and home mortgages. A more likely explanation is that 401(k)s increased personal saving significantly, and the apparent increase in mortgage debt reflects the effect of residual dilution. Confidence in this conclusion is, however, qualified by the continuing concern about the comparability of the experimental and control groups.

C. Other evidence

In this section, I briefly touch on some additional sources of evidence on the efficacy of tax-favored savings accounts that do not specifically concern either IRAs or 401(k)s.

1. Comparisons across countries. The generosity of the incentives embodied in tax-favored savings accounts differs significantly across countries. It is doubtful that one could learn much from simple cross-country correlations between tax incentives and saving, since such correlations may exist for spurious reasons (e.g. people who like to save support tax systems that

favor saving). However, if different countries implemented their tax incentives at different points in time, one might hope to learn something about the effect of tax incentives by examining whether the saving rates of different countries converged or diverged when incentives were introduced.

In this spirit, Carroll and Summers (1987) compare historical rates of saving for Canada and the United States. They demonstrate that these rates diverged when Canada expanded its system of Registered Retirement Saving Plans (RRSPs) during the mid-1970s. While this pattern is interesting, an inference of causality requires an unusually large leap of faith, particularly since there are other possible explanations for the increase in Canadian saving during this period. Moreover, the adoption of tax incentives in the U.S. did not result in measurable convergence between the two countries.

2. Comparisons across cohorts. In section B, I discussed the manner in which transitional phenomena generated by the relative novelty of 401(k)s have been used to assess their effects. More generally, one could regard the 1980s as a grand experiment with several different types of tax-favored accounts, and ask whether these accounts had the effect of shifting up the age-wealth profiles of entire cohorts. To take an example, if tax incentives were effective, then the typical individual reaching age 65 in, say, 1991 should have had more wealth than the typical individual reaching retirement in, say, 1984 (due to differences in years of eligibility for tax-favored saving).

Venti and Wise (1993) examine this hypothesis. Their analysis, which primarily relies on the SIPP, documents a substantial upward shift in financial asset profiles. More recent cohorts have greater wealth at the same ages as older cohorts, and the difference is roughly equal to accumulated balances in 401(k)s and IRAs. While these patterns are interesting, it requires an unusually large leap of faith to ascribe all differences in saving between cohorts to tax incentives. The same pattern could emerge if, for example, younger cohorts are wealthier on a lifetime basis.

3. Simulations. The available evidence on 401(k)s and IRAs sheds some light on the

relatively short-run effects of saving incentives. The long-run effects of these programs may differ significantly from their short-run effects. It is conceivable, for example, that contributions to these accounts could increasingly reflect net additions to saving with the passage of time, as more households exhaust their ability to shift assets from other sources.

Unfortunately, since IRAs and 401(k)s were introduced in the relatively recent past, it is impossible to estimate their long-run effects empirically. One alternative is to build simulation models (Engen and Gale, 1993, Engen, Gale, and Scholz, 1994). The available simulations do indicate that tax incentives may have a substantial, positive influence on the long-run capital stock, even when the short-run response is dominated by asset shifting. However, it is important to keep in mind that simulation results do not constitute empirical evidence, and that the simulations are only as good as the assumptions on which they are based. These assumptions generally include rational, farsighted, and deliberate decision-making in a complex environment with considerable uncertainty. As discussed in section II, it is dangerous and potentially misleading to adopt these assumptions unquestioningly, without exploring the implications of alternatives.

V. Indirect third-party responses to tax incentives

Even if the interest elasticity of saving is low and households do not alter their behavior very much as a direct consequence of targeted tax incentives for saving, it might still be possible to influence personal saving through tax policy. In section II.B.3, I argued that non-neutralities in the tax system may encourage various kinds of third-party activities that have the potential to affect the level of personal saving. Specifically, non-neutralities may encourage employers to adopt various kinds of pension plans or to substitute one kind of plan for another, and may influence the activities of employers in the context of these plans. The tax system may also create incentives for the vendors of tax-favored financial vehicles to market and otherwise promote their products. In this section, I briefly discuss the evidence on each of these

possibilities.

A. The size and scope of the pension system

Since pensions provide a tax-favored mechanism for compensating employees, tax policy may have played an important role in stimulating the development of the pension system. To assess the ultimate impact on personal saving, one must answer two questions. First, to what extent is the size and scope of the pension system responsive to changes in tax rates? Second, to what degree does pension saving displace other forms of personal saving? I consider these questions in turn.

1. Incentives for pension saving. It is indisputable that there is a substantial tax incentive for pension formation. Ippolito (1986) estimates that the optimum exploitation of opportunities to defer compensation through pensions can reduce lifetime tax liabilities by 20 to 40 percent. However, this does not imply that the growth of the pension system is exclusively, or even primarily attributable to the tax system. Pensions may enhance the productivity of the work force in a variety of ways. They may bond the workforce against union activity, voluntary job turnover, or poor job performance.[12] Employers may use defined benefit plans to induce a desired pattern of retirement.[13] Mandatory pensions may also provide an effective device for overcoming the problems with adverse selection that characterize the market for private annuities.[14] Thus, it is conceivable that an extensive private pension system would exist even in the absence of tax incentives.

The evidence on the relative importance of these factors is sketchy at best.[15] The existing literature does not allow one reliably to infer the elasticity of pension coverage with respect to the magnitude of the tax breaks granted to pension plans.[16] I am therefore inclined to agree with Parson's (1995) conclusion that the quantitative importance of tax considerations and efficiency considerations remains unknown. However, it is difficult to account for observed behavior unless one attributes an important (though perhaps not preeminent) role to tax policy.

2. Do pensions crowd out other personal saving? The extent to which pensions displace other forms of personal saving probably depends on the characteristics of the pension. For our purposes, it is important to distinguish between employer-controlled pensions that provide the employee with no choice concerning the level of participation, and participant-controlled plans (such as 401(k)s) that permit the employee to determine contributions. I have already discussed the existing evidence on the extent to which contributions to participant-controlled plans crowd out other personal saving (section IV.B). Consequently, I will now focus on employer-controlled plans.

The existing literature contains more than a dozen studies that attempt to measure the degree of substitutability between pensions and other saving. The usual approach is to estimate a cross-sectional relation between either saving or wealth and some measure of pension coverage. The two earliest studies on this topic (Cagan, 1965, and Katona, 1965) conclude that pensions actually crowd *in* other forms of saving. Cagan rationalizes this finding by arguing that pensions induce workers to recognize the need for retirement planning; Cagan suggests that individuals may intensify their efforts to provide adequately for retirement because a pension renders this objective more feasible. Several subsequent studies corroborate the Cagan-Katona finding (Schoeplein, 1970, Green, 1981, Venti and Wise, 1993, and Bernheim and Scholz, 1993a). More commonly, investigators have found either no effect, or a small effect (Munnell, 1974, Kotlikoff, 1979, Blinder, Gordon, and Wise, 1980, King and Dicks-Mireaux, 1982, Diamond and Hausman, 1984, Hubbard, 1986, Wolff, 1988, and Samwick, 1995). Only a few studies have found substantial rates of crowding out (Munnell, 1976, Dicks-Mireaux and King, 1984, Avery, Elliehausen, and Gustafson, 1986, and Gale, 1995), and most of these provide ranges of estimates that include relatively small effects. There is also some evidence that the rate of displacement rises with education (Bernheim and Scholz, 1993b, and Gale, 1995).

While there are many methodological concerns that bear on the reliability (both absolutely and relatively) of these various studies, three issues stand out as particularly salient.

The first concerns the possibility that pension coverage is correlated with underlying tastes for saving. In contrast to the literature on 401(k)s, no existing study has come to grips with this issue. The direction of the resulting bias is ambiguous.[17] The second issue concerns the measurement of compensation. For the most part, the studies listed above control for income, rather than total compensation (which would include the accrual of pension wealth). If the creation of a pension typically entails a shift in the form of compensation rather than incremental compensation, then this practice does not yield the appropriate displacement rate. Bernheim and Scholz (1993a) and Gale (1995) propose different solutions to this problem, and obtain very different results. The final issue concerns the definition of wealth. Although one can point to a number of exceptions, there is some tendency (as in the 401(k) literature) to find higher rates of displacement when one uses a broader measure of wealth. The issues here are similar to those discussed in section IV.B.2.

While the extent of crowding out is therefore not a settled issue, one is hard pressed to find convincing support in any study for the hypothesis that the rate of displacement is dollar-for-dollar. Indeed, there appears to be a significant likelihood that the true offset is much smaller. The importance of this finding becomes obvious when one considers that, between 1980 and 1990, pension assets accounted for significantly more than 100 percent of the real change in national wealth (Shoven, 1991). Thus, the effect of tax incentives on saving through the stimulation (or retardation) of pensions may be substantial, even if the rate of displacement is relatively low. This observation should give pause to those who favor the elimination of existing tax advantages for employer-controlled pensions, as well as those who support the removal of these advantages through the adoption of a consumption tax.

B. Employer-controlled pensions vs. participant-controlled pensions

In evaluating the extent to which 401(k)s contribute to personal saving (section IV.B), we have abstracted from the degree to which these plans substitute for other pensions. If the rate of

substitution is low, then policies that stimulate 401(k)s will tend to increase saving if and only if 401(k) contributions are not fully offset by reductions in non-pension saving. In contrast, if the rate of substitution is high, then policies that stimulate 401(k)s may increase or decrease saving, depending upon whether 401(k) contributions displace non-pension saving at (respectively) a lower or higher rate than other kinds of pensions.

Much has been written about the magnitude and probable causes of the shift from defined benefit to defined contribution pension plans in general, and 401(k)s in particular (see e.g. Parsons, 1995, or Papke, Petersen, and Poterba, 1993, for selective reviews of this literature). The existence of this shift does not, however, establish that 401(k)s have substituted for more traditional plans, since aggregate trends could in principle be driven by changes in the composition and organization of economic activity.

Papke, Petersen, and Poterba (1993) examine data on individual firms, and conclude that wholesale replacement of existing plans (particular defined benefit plans) occurs in a minority of cases. While informative, this evidence does not resolve the central issue, since 401(k)s may displace other pension plans even if they do not directly replace these plans. For example, firms that adopt 401(k)s as supplementary plans may be less inclined to increase, and more inclined to decrease the generosity of other pension plans. The available evidence also indicates that changes in industrial composition and the structure of firms cannot fully account for the aggregate shift to defined contribution plans (see Clark and McDermed, 1990, Gustman and Steinmeier, 1992, and Kruse, 1991). Since the unexplained component of the aggregate shift is large, it is possible that 401(k)s have substituted for other pension plans to a significant degree.

C. Other activities undertaken by employers

Aside from encouraging employers to provide various kinds of pensions, tax policy may also induce employers to engage in other activities that have the potential to influence saving. In some instances, this effect is indirect: by stimulating pensions, tax policy may also encourage

activities that are complementary to pensions. In other cases, subtle features of the tax code may directly affect the activity in question.

Employer-based investment and retirement education is an example of an activity that is complementary to the provision of a pension plan. Tax policies that stimulate pensions in general, and especially participant-controlled plans, may also stimulate complementary educational initiatives (see Bernheim and Garrett, 1995, Bayer, Bernheim, and Scholz, 1996, and Employee Benefit Research Institute, 1995). Subtle features of the tax code, such as non-discrimination requirements, may also encourage employer-based retirement education more directly (in addition to the preceding references, see Garrett, 1995). Generally, the impact of education is not subsumed in estimates of the relation between pensions and saving, since the most of the growth of these offerings post-dates the most commonly used sources of data on household financial behavior.

There are sound reasons to expect that retirement education might have an important effect on household saving. Numerous studies have established that the level of financial literacy among adult Americans is extremely low. This phenomenon is accompanied by an apparently widespread failure to appreciate financial vulnerabilities (Bernheim, 1995b). Although there is little direct evidence on the impact of educational programs, two recent studies conclude that employer-based offerings significantly stimulate both voluntary pension contributions and total household saving (see Bernheim and Garrett, 1995, and Bayer, Bernheim, and Scholz, 1996).

D. Marketing and promotion of financial products

The expansion of IRA eligibility to all taxpayers in 1981 was accompanied by a great deal of media fanfare. Perhaps more importantly, the existence of these retirement saving vehicles created profit opportunities for financial institutions. Although the IRA tax incentive was targeted at households, it generated considerable impetus for private firms to promote saving through a blend of education and pure marketing. Similar phenomena occur in the context other

tax-advantaged savings instruments, such as long-term life insurance policies.

It is natural to wonder whether these promotional activities affect personal saving. Unfortunately, there is virtually no direct evidence on this issue. There are, however, two particularly interesting anecdotes. One concerns the introduction and subsequent scaling-back of IRAs, which I have discussed in section IV.A. The other concerns experience with saving promotion in Japan (Central Council for Saving Promotion, 1981). After World War II the Japanese government launched a national campaign to promote saving. Promotional activities included monthly seminars that extolled the virtues of saving and provided workers with financial guidance, sponsorship of children's banks, the appointment of private citizens as savings promotion leaders, and the extensive dissemination of literature. While the Japanese rate of saving rose precipitously over the relevant time period, other factors were also at work, including the existence of strong tax incentives for saving, as well as various aspects of post-War reconstruction. One can therefore only speculate about the extent to which the increase in saving was attributable to promotion.

VI. A Concluding Lament

As an economist, one cannot review the voluminous literature on saving incentives without being somewhat humbled by the enormous difficulty of learning anything useful about even the most basic economic questions. Having been handed two grand "experiments" with tax incentives during the 1980s, it would seem that we ought to have learned a lot more, and to have achieved greater consensus, than we have. In our defense, it can be said that we have done our best with the information at our disposal. As I have indicated throughout this paper, it is often easy to identify the kinds of data that would have allowed us to answer the pressing policy questions with much greater confidence. Unfortunately, we have had to make due with data that is, at best, a caricature of the ideal. When one thinks of the budgetary costs of tax incentives, and of what is at stake in terms of economic growth and efficiency, it seems a shame that ongoing,

comprehensive, microeconomic data collection has been such a low social priority.

References

Ainslie, George, "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control," *Psychological Bulletin* 82(4), July 1975, 463-496.

Ainslie, George, "A Behavioral Economic Approach to the Defense Mechanisms: Freud's Energy Theory Revisited," *Social Science Information* 21, 1982, 735-779.

Ainslie, George, "Behavioral Economics II: Motivated, Involuntary Behavior," *Social Science Information* 23, 1984, 47-78.

Allen, Steven G., and Robert L. Clark, "Pensions and Firm Performance," in *Human Resources and the Performance of the Firm*, Industrial Relations Research Association Series, 1987.

Allen, Steven G., Robert L. Clark, and Ann McDermed, "Pension Bonding and Lifetime Jobs," *Journal of Human Resources* 28, 1993, 463-481.

Andrews, Emily S., "The Growth and Distribution of 401(k) Plans," in J. Turner and D. Beller, eds., *Trends in Pensions 1992*, Washington, D.C.: U.S. Department of Labor, 1992.

Arthur D. Little, Inc., *America's Retirement Crisis: The Search for Solutions*, Final Report to Oppenheimer Management Corporation, June 1993.

Attanasio, Orazio, and Thomas De Leire, "IRAs and Household Saving Revisited: Some New Evidence," Working Paper No. 4900, National Bureau of Economic Research, October 1994.

Attanasio, Orazio, and Guglielmo Weber, "Consumption Growth, the Interest Rate and Aggregation," *Review of Economic Studies* 60(3), July 1993, 631-649.

Avery, Robert B., Gregory E. Elliehausen, and Thomas A. Gustafson, "Pensions and Social Security in Household Portfolios: Evidence from the 1983 Survey of Consumer Finances," in F. Gerard Adams and Susan M. Wachter, eds., *Savings and Capital Formation*, Lexington, MA: Lexington Books, 1986.

Bayer, Patrick J., B. Douglas Bernheim, and J. Karl Scholz, "The Effects of Financial Education

in the Workplace: Evidence from a Survey of Employers," 1996.

Bernheim, B. Douglas, "Personal Saving, Information, and Economic Literacy: New Directions for Public Policy," in *Tax Policy for Economic Growth in the 1990s*, Washington, D.C.: American Council for Capital Formation, 1994a.

Bernheim, B. Douglas, "Comments and Discussion," *Brookings Papers on Economic Activity*, 1994b, 152-166.

Bernheim, B. Douglas, "Comment on Chapters 4 and 5," in David A. Wise, ed., *Studies in the Economics of Aging*, Chicago: University of Chicago Press, 1994c.

Bernheim, B. Douglas, "The Merrill Lynch Baby Boom Retirement Index: Update '95," mimeo, 1995a.

Bernheim, B. Douglas, "Do Households Appreciate Their Financial Vulnerabilities? An Analysis of Actions, Perceptions, and Public Policy," *Tax Policy and Economic Growth*, Washington, D.C.: American Council for Capital Formation, 1995b, 1-30.

Bernheim, B. Douglas, and Daniel M. Garrett, "The Determinants and Consequences of Financial Education in the Workplace: Evidence from a Survey of Households," mimeo, Stanford University, 1995.

Bernheim, B. Douglas, and John Karl Scholz, "Private Pensions and Household Saving," mimeo, University of Wisconsin, 1993a.

Bernheim, B. Douglas, and John Karl Scholz, "Private Saving and Public Policy," *Tax Policy and the Economy* 7, 1993b, 73-110.

Blinder, Alan S., "Distribution Effects and the Aggregate Consumption Function," *Journal of Political Economy* 83, 1975, 447-475.

Blinder, Alan S., Roger H. Gordon, and Donald E. Wise, "Reconsidering the Work Disincentive Effects of Social Security," *National Tax Journal* 33(4), December 1980, 431-442.

Bloom, David E., and Richard B. Freeman, "The Fall in Private Pension Coverage in the U.S.,"

NBER Working Paper no. 3973, 1992.

Boskin, Michael J., "Taxation, Saving, and the Rate of Interest," *Journal of Political Economy* 86, 1978, S3-S27.

Boskin, Michael J., and Lawrence J. Lau, "Taxation, Social Security and Aggregate Factor Supply in the United States," mimeo, Stanford University, 1978.

Buck Consultants, Inc., "Current 401(k) Plan Practices: A Survey Report," September 1989.

Burkhauser, Richard V., "The Pension Acceptance of Older Workers," *Journal of Human Resources* 14, 1979, 63-72.

Burkhauser, Richard V., "The Early Acceptance of Social Security: An Asset Maximization Approach," *Industrial and Labor Relations Review* 33, July 1980, 484-92.

Burman, Leonard, Joseph Cordes, and Larry Ozanne, "IRAs and National Saving," *National Tax Journal* 43, September 1990, 259-284.

Cagan, Phillip, "The Effect of Pension Plans on Aggregate Saving: Evidence from a Sample Survey," Occasional Paper 95, National Bureau of Economic Research, 1965.

Carroll, Chris, and Lawrence H. Summers, "Why Have Private Savings Rates in the U.S. and Canada Diverged?" *Journal of Monetary Economics* 20, September 1987, 249-80.

Central Council for Savings Promotion, *Savings and Savings Promotion Movement in Japan*, Tokyo: Bank of Japan, 1981.

Clark, Robert L., and Ann McDermed, *The Choice of Pension Plans in a Changing Regulatory Environment*, Washington, D.C.: American Enterprise Institute, 1990.

Congressional Budget Office, *Baby Boomers in Retirement: An Early Perspective*, Washington, D.C.: Government Printing Office, 1993.

Diamond, Peter A. and Jerry A. Hausman, "Individual Retirement and Saving Behavior," *Journal of Public Economics* 23, 1984, 81-114.

Dicks-Mireaux, Louis, and Mervyn King, "Pension Wealth and Household Savings: Tests of

Robustness," *Journal of Public Economics* 23, 1984, 115-139.

Doyle, Robert J., Jr., and Eric T. Johnson, *Readings in Wealth Accumulation Planning*, fourth edition, Bryn Mawr, Pennsylvania: The American College, 1991.

Employee Benefit Research Institute, "Can We Save Enough to Retire? Participant Education in Defined Contribution Plans," *EBRI Issue Brief*, no. 160, April 1995.

Engen, Eric M., "Precautionary Saving and the Structure of Taxation," mimeo, Federal Reserve Board of Governors, June 1994.

Engen, Eric M., and William G. Gale, "IRAs and Saving in a Stochastic Life-Cycle Model," mimeo, Brookings Institution, 1993.

Engen, Eric M., and William G. Gale, "Debt, Taxes and the Effects of 401(k) Plans on Household Wealth Accumulation," mimeo, Federal Reserve Board, 1995.

Engen, Eric M., William G. Gale, and John Karl Scholz, "Do Saving Incentives Work?" *Brookings Papers on Economic Activity*, 1994:1, 85-151.

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

Feldstein, Martin S., "The Welfare Cost of Capital Income Taxation," *Journal of Political Economy* 86(2), Part 2, April 1978, S29-51.

Fields, Gary S., and Olivia S. Mitchell, *Retirement, Pensions, and Social Security*, Cambridge, MA: MIT Press, 1984.

Fisher, Irving, *The Theory of Interest*, London: MacMillan, 1930.

Furnham, Adrian, and Alan Lewis, *The Economic Mind: The Social Psychology of Economic Behavior*, New York: St. Martin's Press, 1986.

Gale, William G., "The Effects of Pensions on Wealth: A Re-evaluation of Theory and Evidence," mimeo, Brookings Institution, 1995.

Gale, William G. and John Karl Scholz, "IRAs and Household Saving," *American Economic*

Review 84, December 1994, 1233-1260.

Garrett, Daniel M., "The Effects of Nondiscrimination Rules on 401(k) Contributions," mimeo, Stanford University, 1995.

Gentry, William M., and Eric Peress, "Taxes and Fringe Benefits Offered by Employers," National Bureau of Economic Research Working Paper No. 4764, June 1994.

Gravelle, Jane G., "Do Individual Retirement Accounts Increase Savings?" *Journal of Economic Perspectives* 5, Spring 1991, 133-148.

Green, Francis, "The Effect of Occupational Pension Schemes on Saving in the United Kingdom: A Test of the Life Cycle Hypothesis," *The Economic Journal* 91, March 1981, 136-144.

Gustman, Alan L., and Thomas L. Steinmeier, "The Stampede Toward Defined Contribution Pension Plans: Fact or Fiction?" *Industrial Relations* 31, 1992, 361-369.

Gustman, Alan L., and Thomas L. Steinmeier, "Pension Incentives and Job Mobility," W.E. Upjohn Institute for Employment Research, 1995.

Hall, Robert E., "Intertemporal Substitution in Consumption," *Journal of Political Economy* 96, April 1988, 339-357.

Hoch, Stephen J., and George F. Lowenstein, "Time-Inconsistent Preferences and Consumer Self-Control," *Journal of Consumer Research* 17(4), March 1991, 492-507.

Howrey, E. Phillip, and Saul H. Hymans, "The Measurement and Determination of Loanable Funds Saving," *Brookings Papers on Economic Activity*, 1978:2, 655-85.

Hubbard, R. Glenn, "Do IRAs and Keoghs Increase Saving?" *National Tax Journal* 37, March 1984, 43-54.

Hubbard, R. Glenn, "Pension Wealth and Individual Saving," *Journal of Money, Credit, and Banking* 18(2), May 1986, 167-178.

Hubbard, R. Glenn, and Jonathan Skinner, "The Effectiveness of Saving Incentives: A Review of the Evidence," mimeo, Columbia University, 1995.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, "Precautionary Saving and Social Insurance," *Journal of Political Economy* 103, April 1995, 360-399.

Ippolito, Richard A., "The Economic Function of Underfunded Pension Plans," *Journal of Law and Economics* 28, 1985, 611-651.

Ippolito, Richard A., *Pensions, Economics and Public Policy*, Homewood, IL: Dow Jones-Irwin, 1986.

Johnson, Alfred P., "Individual Retirement Accounts Help Boost Saving in the U.S.," testimony to the Committee on Finance, U.S. Senate, in Tax Reform Proposals (99-246, part XIII), 1985, 129-149.

Joines, Douglas H., and James G. Manegold, "IRA and Saving: Evidence from a Panel of Taxpayers," mimeo, University of Southern California, February 1995.

Jones, Stephen R. G., *The Economics of Conformism*, Oxford: Basil Blackwell, 1984.

Judd, Kenneth L., "The Welfare Cost of Factor Taxation in a Perfect-Foresight Model," *Journal of Political Economy* 95(4), August 1987, 675-709.

Judd, Kenneth L., "Redistributive Taxation in a Simple Perfect Foresight Model," *Journal of Public Economics* 28(1), October 1988, 59-83.

Katona, George, *Private Pensions and Individual Saving*, University of Michigan Press, 1965.

Katona, George, *Psychological Economics*, Amsterdam: Elsevier, 1975.

King, Mervyn A., and Louis Dicks-Mireaux, "Asset Holdings and the Life Cycle," *Economic Journal* 92, 1982, 247-267.

Kotlikoff, Laurence J., "Testing the Theory of Social Security and Life Cycle Accumulation," *American Economic Review* 69(3), June 1979, 396-410.

Kotlikoff, Laurence J., and Avia Spivak, "The Family as an Incomplete Annuities Market," *Journal of Political Economy* 89, April 1981, 372-91.

Kotlikoff, Laurence J., and David A. Wise, "The Wage Carrot and the Pension Stick: Retirement

Benefits and Labor Force Participation," W.E. Upjohn Institute for Employment Research, 1990.

Kruse, Douglas L., "Pension Substitution in the 1980s: Why the Shift Toward Defined Contribution Pension Plans?" NBER Working Paper, 1991.

Kruse, Douglas L., "Profit Sharing: Does it Make a Difference?" W.E. Upjohn Institute for Employment Research, 1993.

Kusko, Andrea, James M. Poterba, and David W. Wilcox, "Tax Preferred Saving Vehicles and Personal Saving: New Evidence form the Behavior of 12,000 Participants in a 401(k) Plan," mimeo, Federal Reserve Board of Governors, 1993.

Kusko, Andrea, James M. Poterba, and David W. Wilcox, "Employee Decisions With Respect to 401(k) Plans: Evidence from Individual Level Data," Working Paper No. 4635, National Bureau of Economic Research, 1994.

Lazear, Edward P., "Pensions as Severance Pay," in Zvi Bodie, John Shoven, and David Wise, eds., *Financial Aspects of the United States Pension System*, Chicago: University of Chicago Press, 1983.

Lazear, Edward P., and Robert Moore, "Pensions and Mobility," in Zvi Bodie, John Shoven, and David Wise, eds., *Pensions in the U.S. Economy*, Chicago: University of Chicago Press, 1988.

Long, James E., "Marginal Tax Rates and IRA Contributions," *National Tax Journal* 43(2), 1990, 143-53.

Lucas, Robert E., Jr., "Econometric Policy Evaluation: A Critique," in Karl Brunner and Allan H. Meltzer, *The Phillips Curve and Labor Markets*, Carnegie-Rochester Conference Series on Public Policy, vol. 1, suppl., *Journal of Monetary Economics*, 1976.

Maital, Shlomo, "Prometheus Rebound: On Welfare-Improving Constraints," *Eastern Economic Journal* 12(3), July-September 1986, 337-44.

Modigliani, Franco, and Richard Brumberg, "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data," in K.K. Kurihara, ed., *Post Keynsian Economics*, New Brunswick, NJ: Rutgers University Press, 1954.

Munnell, Alicia H., *The Effect of Social Security on Personal Saving*, Cambridge, MA: Ballinger, 1974.

Munnell, Alicia H., "Private Pensions and Saving: New Evidence," *Journal of Political Economy* 84(5), 1976, 1013-1032.

Papke, Leslie E., "Participation in and Contributions to 401(k) Plans: Evidence from Plan Data," NBER Working Paper No. 4199, 1992.

Papke, Leslie, Mitchell Petersen, and James M. Poterba, "Did 401(k) Plans Replace Other Employer Provided Pensions?" Working Paper No. 4501, National Bureau of Economic Research, October 1993.

Parsons, Donald O., "The Employment Relationship: Job Attachment, Work Effort, and the Nature of Contracts," in Orley Ashenfelter and Richard Layard, eds., *Handbook of Labor Economics*, vol. II, Amsterdam: North Holland, 1986.

Parsons, Donald O., "Retirement Age and Retirement Income: The Role of the Firm," mimeo, Ohio State University, 1995.

Poterba, James M., Steven F. Venti, and David A. Wise, "401(k) Plans and Tax-Deferred Saving," NBER Working Paper No. 4181, 1992.

Poterba, James M., Steven F. Venti, and David A. Wise, "401(k) Plans and Tax-Deferred Saving," in David A. Wise, Ed., *Studies in the Economics of Aging*, Chicago: University of Chicago Press, 1994.

Poterba, James M., Steven F. Venti, and David A. Wise, "Do 401(k) Contributions Crowd Out Other Personal Saving?" *Journal of Public Economics* 58, 1995, 1-32.

Poterba, James M., Steven F. Venti, and David A. Wise, Do Retirement Saving Programs

Increase Saving?" *Journal of Economic Perspectives*, 1996, forthcoming.

Quinn, Joseph F., Richard V. Burkhauser, and Daniel A. Myers, "Passing the Torch: The Influence of Economic Incentives on Work and Retirement," W.E. Upjohn Institute, 1990.

Rainwater, Lee, *Behind Ghetto Walls: Black Families in a Federal Slum*, Chicago: Aldine, 1970.

Reagan, Patricia B., and John A. Turner, "The Decline in Marginal Tax Rates During the 1980s Reduced Pension Coverage," mimeo, Ohio State University, June 1995.

Samwick, Andrew A., "The Limited Offset Between Pension Wealth and Other Private Wealth: Implications of Buffer-Stock Saving," mimeo, Dartmouth College, December 1995.

Sandmo, Agnar, "The Effects of Taxation on Savings and Risk Taking," in Alan Auerbach and Martin Feldstein, eds., *Handbook of Public Economics*, vol. I, Amsterdam: North Holland, 1985.

Schelling, Thomas C., "Self-Command in Practice, in Policy, and in a Theory of Rational Choice," *American Economic Review* 74(2), May 1984, 1-11.

Schoeplain, Robert N., "The Effect of Pension Plans on Other Retirement Saving," *Journal of Finance*, 1970, 633-638.

Scitovsky, Tibor, *The Joyless Economy*, Oxford: Oxford University Press, 1976.

Scott, Jason, "The Determinants of Participation in Defined Contribution Pension Plans," mimeo, Stanford University, February 1994.

Shefrin, Hersh M., and Richard H. Thaler, "The Behavioral Life-Cycle Hypothesis," *Economic Inquiry* 26(4), 1988, 609-43.

Shoven, John B., *Return on Investment: Pensions are How America Saves*, Association of Private Pension and Welfare Plans: Washington D.C., September 1991.

Skinner, Jonathan, and Daniel Feenberg, "The Impact of the 1986 Tax Reform on Personal Saving," in Joel Slemrod, ed., *Do Taxes Matter? The effect of the 1986 Tax Reform Act*

on the U.S. Economy, Cambridge, MIT Press, 1989.

Stack, Carol B., *All Our Kin: Strategies for Survival in a Black Community*, New York: Harper & Row, 1974.

Starrett, David A., "Effects of Taxes on Saving," in H. Galper, H. Aaron, and J. Pechman, eds., *Uneasy Compromise: Problems of a Hybrid Income-Consumption Tax*, Washington, D.C.: The Brookings Institution, 1988.

Stock, James H., and David A. Wise, "The Pension Inducement to Retire: an Option Value Analysis," in David A. wise, ed., *Issues in the Economics of Aging*, Chicago: University of Chicago Press, 1990.

Summers, Lawrence H., "Taxation and Capital Accumulation in a Life Cycle Growth Model," *American Economic Review* 71, September 1981, 533-554.

Summers, Lawrence H., "Summers Replies to Galper and Byce on IRAs," *Tax Notes* 31(10), 1986, 1014-1016.

Thaler, Richard H., "Psychology and Savings Policies," *American Economic Review* 84, May 1994, 186-192.

Thaler, Richard H., and Hersh Shefrin, "An Economic Theory of Self-Control," *Journal of Political Economy* 89(2), April 1981, 392-406.

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

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

Venti, Steven F., and David A. Wise, "The Determinants of IRA Contributions and the Effect of Limit Changes," in Zvi Bodie, John B. Shoven, and David A. Wise, eds., *Pensions and the U.S. Economy*, Chicago: University of Chicago Press, 1988.

Venti, Steven F., and David A. Wise, "Aging, Moving, and Housing Wealth," in David A. Wise,

ed., *The Economics of Aging*, Chicago: University of Chicago Press, 1989.

Venti, Steven F., and David A. Wise, "Have IRAs Increased U.S. Saving: Evidence from Consumer Expenditure Surveys," *Quarterly Journal of Economic* 105, August 1990, 661-698.

Venti, Steven F., and David A. Wise, "The Saving Effect of Tax-Deferred Retirement Accounts: Evidence for SIPP," in B. Douglas Bernheim and John B. Shoven, eds., *National Saving and Economic Performance*, Chicago: University of Chicago Press, 1991.

Venti, Steven F., and David A. Wise, "Government Policy and Personal Retirement Saving," *Tax Policy and the Economy* 6, 1992, 1-41.

Venti, Steven F., and David A. Wise, "The Wealth of Cohorts: Retirement Saving and the Changing Assets of Older Americans," Working Paper No. 4600, National Bureau of Economic Research, December 1993.

Venti, Steven F. and David A. Wise, "Individual Response to Retirement Savings Programs: Results from U.S. Panel Data," *Reicherche Economiche*, 1995 (forthcoming).

Whyte, William F., *Street Corner Society*, Chicago: The University of Chicago Press, 1943.

Williamson, Samuel H., "U.S. and Canadian Pensions Before 1930: A Historical Perspective," in John A. Turner and Daniel J. Beller, eds., *Trends in Pensions 1992*, U.S. Department of Labor, Pension and Welfare Benefits Administration, 1992.

Wolff, Edward N., "Social Security, Pensions and the Life Cycle Accumulation of Wealth: Some Empirical Tests," *Annales D'Economie et de Statistique* 9, 1988, 199-226.

Woodbury, Stephen A., and Wei-Jang Huang, "The Tax Treatment of Fringe Benefits," W.E. Upjohn Institute, 1991.

Wright, C., "Saving and the Rate of Interest," in A.C. Harberger and M.J. Bailey, Eds., *The Taxation of Income from Capital*, Washington D.C.: Brookings Institution, 1969.

Footnotes

1. There is considerable evidence that economic decisions in general are strongly affected by peer group effects. See e.g. Whyte (1943), Rainwater (1970), Stack (1974), or Jones (1984).

2. According to Thaler and Shefrin (1981), "[s]imply keeping track seems to act as a tax on any behavior which the planner views as deviant."

3. Hubbard's (1984) data are drawn from the 1979 President's Commission on Pension Policy, and therefore include some non-contributors who were ineligible for IRAs. Thus, the sample selection problem discussed in the text is perhaps less pronounced than for estimates based on data collected between 1982 and 1986.

4. For example, imagine that, absent access to IRAs, new contributors would have consumed all of their disposable incomes (contributing nothing to non-IRA balances), and that their incomes would have been rising more rapidly than the consumption of old contributors.

5. Engen, Gale, and Scholz replicate this procedure, but also estimate a fixed effects model using the full sample, treating all non-contributors prior to 1982 as if they were ineligible. This exaggerates the selection bias discussed below, and thereby increases the likelihood that the resulting estimates will understate the true contribution of IRAs to personal saving. The fact that Engen, Gale, and Scholz significantly overstate the change in the contribution limit for individuals who were eligible for, but did not participate in, IRAs prior to 1982 further exacerbates this bias. To understand this point, imagine for the sake of simplicity that all taxpayers file separately. In that case, the contribution limit for members of the group in question was increased by $500 (from $1,500 to $2,000) in 1982. Engen, Gale, and Scholz falsely assume that the contribution limit increased by $2,000 (from $0 to $2,000) for these individuals. Thus, their incremental saving per dollar increase in the contribution limit is biased downwards by a factor of four. Since this group is large relative to the entire population, the resulting bias is likely to be substantial.

6. Engen, Gale, and Scholz (1994) argue that IRA contributions may have declined after 1986 because of reductions in marginal tax rates and limits on deductibility. But unless one believes that the interest elasticity of saving is enormous, this could not have accounted for the magnitude of the decline in contributions. The authors also attribute the decline in IRA saving to the increased availability of 401(k)s and/or the possible depletion of non-IRA financial assets. To my knowledge, there is no evidence that individual taxpayers replaced IRAs with 401(k)s, or that IRA contributors as a group ran out of shiftable assets. It is also doubtful that either phenomenon can account for the sharpness of the decline in IRA contributions.

7. The fraction of individuals eligible for 401(k)s who owned IRAs rose (relative to ineligibles) between 1984 and 1987. However, this does not imply that "reverse dilution" occurred over this earlier period, for the simple reason that IRA ownership was not a stable taste proxy prior to 1987. IRAs were still universally available during this period, and participation was rising for the population as a whole. Even if the set of individuals eligible for 401(k)s had remained fixed between 1984 and 1987, one would have expected to see IRA ownership rise among eligibles (relative to ineligibles), since those with greater inclinations to save (including 401(k) eligibles) would have been more likely to open new IRA accounts.

8. To be clear, the issue here is not whether observable characteristics are comparable (a point which the authors do address). Rather, it is whether unobserved preferences are similar. While the greater comparability of the experimental and control groups used by Engen, Gale, and Scholz seems intuitive, it does not follow purely as a matter of logic. In essence, the authors censor the distribution of preferences separately for the eligible population and the ineligible population. There is no guarantee that the truncation points are similar, or that the upper tails of the preference distributions for the two samples have similar features. Thus, in principle, censoring could reduce comparability.

9. Engen and Gale (1995) argue that there is also an upward bias because the total compensation of

those with 401(k)s is understated in the data to a greater extent than the total compensation of those without 401(k)s. Even assuming that this is true (which is not at all obvious), the point is not relevant in the current context unless one believes that the authorization of 401(k)s had the effect of increasing the lifetime compensation of those who eventually received 401(k)s.

10. Engen and Gale report that, as of 1989, only 20% of households with mortgages had ever refinanced, and that equity had been extracted in only 57% of these cases. Other figures in their paper point to somewhat higher rates of equity extraction (upwards of 80 percent). As a rough approximation, assume that 20 percent of homeowners covered by 401(k)s in 1991 refinanced their homes between 1987 and 1991 (probably an overestimate), and that 80 percent of these extracted equity. Since roughly three-quarters of those eligible for 401(k)s were homeowners in 1991, this implies that roughly 12 percent extracted equity through refinancing between 1987 and 1991.

11. One might also defend Engen and Gale's interpretation by arguing that, absent 401(k)s, eligible households would have accelerated repayment of their mortgages. However, under this view, the relative increase of mortgage debt among eligible households should be attributable, at least in part, to accelerated repayment by ineligible households (the control group). I doubt that this was the case.

12. See e.g. Ippolito (1985,1986), Parsons (1986,1995), or Williamson (1992).

13. See e.g. Burkhauser (1979,1980), Lazear (1983), Fields and Mitchell (1984), Ippolito (1986), Lazear and Moore (1988), Kotlikoff and Wise (1990), Stock and Wise (1990), and Quinn, Burkhauser and Myers (1990).

14. See Ippolito (1986). Kotlikoff and Spivak (1981) discuss the nature of market failure in private annuity markets.

15. A sampling of pertinent references includes Ippolito (1986), Bloom and Freeman (1992), Reagan and Turner (1995), Kruse (1993), Allen and Clark (1987), Clark and McDermed (1990), and Gustman and Steinmeier (1995).

16. There is, however, some evidence that fringe benefits do respond to variations in state-level

income taxes, particularly for blue collar workers. See Gentry and Peress (1994).

17. It is conceivable that the workers who are most inclined to save, and who have the least problems with self-discipline, sort themselves into jobs that are covered by pension plans with the greatest discretion, such as 401(k)s. Those who are interested in saving, but who have problems with self-discipline, may prefer traditional employer-controlled plans. I know of no evidence that would allow one reliably to distinguish this hypothesis from various other alternatives, and thereby discern the direction of the resulting bias.