National Tax Journal, June 2012, 65 (2), 283–312

THE EFFECT OF RECENT TAX CHANGES ON TAX-PREFERRED SAVING BEHAVIOR

Bradley T. Heim and Ithai Z. Lurie

This paper estimates the extent to which changes in tax policy induce changes in contributions to tax-preferred savings accounts using a panel of tax returns from 1999–2005 that spans the tax changes enacted in the Economic Growth and Tax Relief Reconciliation Act of 2001 and the Jobs and Growth Tax Relief Reconciliation Act of 2003. The results suggest a statistically significant, though small, response of taxpayers to changes in the after-tax price of contributing on the extensive margin, but not on the intensive margin. There is also some suggestive evidence that increases in after-tax incomes may have increased the probability of contributing to tax-preferred savings accounts and the amount contributed, though these results are not robust to alternative specifications.

Keywords: retirement savings, individual taxation

JEL Codes: H24, D91

I. INTRODUCTION

Since the creation of Individual Retirement Accounts (IRAs) in 1974, and the issuance of 401(k)-clarifying regulations in 1981, tax-preferred accounts have become an important vehicle for retirement saving in the United States. Copeland (2009) estimates that in 2007 assets in IRAs, self-employed plans, and employer-based defined contribution accounts (including 401(k), 403(b), 457(b), SIMPLE, SEPs, and other plans) amounted to almost $8.5 trillion, far exceeding the $2.7 trillion in assets held in defined benefit plans. In addition, 30.6 percent of all families owned some form of IRA — including regular, rollover, and Roth IRAs — or self-employed retirement account, and 33.5 percent of families participated in an employer-based defined contribution plan.

Bradley T. Heim: School of Public and Environmental Affairs, Indiana University, Bloomington, IN, USA (heimb@indiana.edu)

Ithai Z. Lurie: Office of Tax Analysis, U.S. Department of the Treasury, Washington, DC, USA (ithai.lurie@treasury.gov)

In traditional IRAs,[1] employer-based accounts, and accounts for the self-employed, taxpayers exclude contributions to the plan when calculating taxable income and plan assets grow tax free, but distributions from the plan (including contributions and account earnings) are taxable.[2] For Roth IRAs, contributions are not deductible, but account earnings accumulate tax free and distributions are not taxable. As a result, these accounts provide benefits both in the form of the tax-free buildup of assets, and in the possibility of lower overall taxes, depending on how the tax rate at the time of contribution differs from the tax rate at the time of withdrawal.

In recent years, reductions in marginal tax rates, the introduction of the Saver's Credit, and a change in the way the Earned Income Tax Credit (EITC) is calculated have altered the benefits of contributing to tax-preferred accounts. This paper seeks to answer two questions about these changes. First, did these policy changes induce more or less participation in retirement savings plans? Second, among those who contribute to tax-preferred plans, did these changes in tax policy induce individuals to increase or decrease contributions to those plans?

Knowing the answers to these questions is important for at least two reasons. First, knowledge of the extent to which changes in tax law encourage or discourage the use of tax-preferred savings accounts is needed to accurately gauge the revenue impact of tax changes, since an increase in tax-preferred contributions will cause the present value of income tax revenue to decline (though the timing differs depending on the type of account). Second, if tax changes lead to changes in tax-preferred retirement savings, then policymakers can use tax policy as an additional instrument to affect the level of savings in retirement accounts.

The effect of the availability of tax-preferred savings plans on the savings rates of households has been widely debated in the literature. The difficulties in estimating the effects of these plans on savings are twofold. First, individuals may decrease savings held elsewhere or increase debt at the same time that they increase contributions to tax-preferred accounts, so that an individual's entire balance sheet must be observed to be able to infer whether any net increase of savings has occurred when a contribution is made. Second, eligibility for plans is not randomly assigned, but rather is correlated with income or with the individual's choice of employer, which can lead to biased estimates. Several papers (surveyed in Poterba, Venti, and Wise, 1996) found that these plans tend to increase savings among plan participants, whereas other authors (surveyed in Engen, Gale, and Scholz, 1996) have found that most of the contributions to these plans consist of amounts that would have been saved in any case. However, recent results in Engen and Gale (2000) and Chernozhukov and Hansen (2004) suggest that 401(k) accounts do increase wealth among those with low earnings or assets.

Due to the difficulties in examining the effect of tax-preferred plans on overall savings, much of the literature has focused on savings in tax-preferred accounts. A number of

[1] Throughout the paper, the term "traditional IRA" is used as a shorthand for a traditional IRA to which a taxpayer makes a deductible contribution. Taxpayers who are not eligible to make deductible contributions to a traditional IRA may make non-deductible contributions.

[2] If a withdrawal is made before a certain age (generally age 59½), the withdrawal is subject to an additional 10 percent penalty. However, there are several exceptions to this rule.

studies — including Papke (1995), Bassett, Fleming, and Rodrigues (1998), Clark and Schieber (1998), Kusko, Poterba, and Wilcox (1998), Springstead and Wilson (2000), and Munnell, Sunden, and Taylor (2003) — have examined how participation in these accounts varies with a wide variety of individual and plan characteristics, and have generally found that the probability of participating in a tax-preferred retirement savings plan increases with income, age, job tenure, the existence of an employer match, and if the savings plan is the sole retirement plan for the household. A recent set of papers, including Madrian and Shea (2001), Choi, Laibson, and Madrian (2004), Choi et al. (2004) and others have found that default policies of the retirement plan — for example, whether workers are automatically enrolled in the plan, or whether they have to affirmatively choose to contribute to the plan — significantly affect whether a worker contributes to a tax-preferred account.

However, little attention has been paid to how participation in and contributions to tax-preferred savings plans vary with the tax environment. A handful of studies have examined whether marginal tax rates affect the amount contributed to different types of tax-preferred accounts, but these studies have come to conflicting conclusions. Venti and Wise (1988) found no relationship between marginal tax rates and IRA saving, while Feenberg and Skinner (1989) and Long (1990) found that higher marginal tax rates increased contributions to IRAs.

More recently, Joulfaian and Richardson (2001) use a 1996 cross-section of tax returns and find that higher marginal tax rates tend to increase the probability of participation in any type of tax-preferred retirement savings plan. Power and Rider (2002) use a panel of tax returns to examine the effect that tax-based savings incentives have on the self-employed, and estimate that the elasticity of the amount of contributions to the tax price of contributing is around –2. They also find a statistically significant response on the extensive margin, though it is unclear how applicable these estimates are to employees of firms. Cunningham and Englehardt (2002) estimate a Tobit model of contributions as a function of the value of contributing to a 401(k) plan relative to the value of contributing to an IRA using W-2 records from households in the Health and Retirement Survey; their estimated coefficients on the relative value variable are statistically insignificant in all specifications. Finally, Englehardt and Kumar (2007) estimate a structural model of contributions to 401(k)-type plans using data from the Health and Retirement Study, and find large and statistically significant effects of marginal tax rates on both the decision to contribute and the amount contributed.

However, the existing literature examining the responsiveness of tax-preferred retirement savings has generally used cross-sectional data, which implies that unobserved characteristics that are correlated both with the contribution decision and with the after-tax prices of contributing and incomes could bias the results. Further, the one paper that did use panel data (Power and Rider, 2002) focused exclusively on the self-employed. In addition, no paper has examined the effect of recent tax changes. Finally, with the exception of Joulfaian and Richardson (2001), all of the papers have examined a particular type of retirement account in isolation. Thus, the effect that taxes have on all types of tax-preferred accounts when unobserved individual characteristics are properly accounted for is an open question.

This paper uses data from a 1999–2005 panel of tax returns to examine whether the propensity of taxpayers to contribute to tax-preferred savings accounts, as well as the amount contributed, changed in response to changes in tax policy over these years. This paper is the first to use panel data to estimate the effects of taxes on contributions to all types of tax-preferred retirement accounts, including employer-based accounts, traditional and Roth IRAs, and self-employed plans. Because a panel is utilized, it is possible to control for time-invariant unobserved characteristics of taxpayers that might be correlated with both their after-tax cost of contributing to an account and their propensity to contribute. In addition, it is possible to instrument for changes in taxes and after-tax incomes with synthetic changes in these variables, thereby identifying tax price and income effects solely from the changes in tax laws that occurred during this time period.

The results suggest a statistically significant response of taxpayers to changes in the after-tax price of contributing on the extensive margin, but no statistically significant response on the intensive margin. The estimates imply that, by lowering the after-tax price of contributing, the 2001 and 2003 tax changes increased the probability of contributing to a tax-preferred account by between 0.23 and 0.71 percent on average, with low income taxpayers being more likely to contribute and high income taxpayers being less likely. There is also some suggestive evidence that tax induced increases in after-tax incomes may have increased the probability of contributing and increased the amount that was contributed, though these results are not robust to alternative specifications.

The paper proceeds as follows. In Section II, the policy environment is reviewed. Section III describes the estimation strategy and the data used, and Section IV presents the results. Section V concludes.

II. Policy Environment

Between 1999 and 2005, two major federal tax laws were passed that affected the benefit taxpayers receive from contributing to a tax-preferred savings account: the Economic Growth and Tax Relief Reconciliation Act of 2001 (EGTRRA) and the Jobs and Growth Tax Relief Reconciliation Act of 2003 (JGTRRA).[3] Prior to the passage of EGTRRA, the federal tax rate structure consisted of five brackets, ranging from 15 to 39.6 percent. After the passage of JGTRRA, however, ordinary marginal tax rates decreased, ranging from 10 to 35 percent.[4] In addition, JGTRRA reduced capital gains tax rates, and lowered statutory rates on "qualified" dividends down to the level of the capital gains rate.

[3] In addition, the Working Families Tax Relief Act was passed in 2004. This act, however, consisted mainly of extensions of previously passed provisions, and so did not impact marginal tax rates directly.

[4] Because the EGTRRA marginal rate cuts were to be phased in gradually between 2001 and 2006, taxpayers may have had some incentive to shift contributions to 2001 and 2002 to take advantage of lower tax prices due to higher marginal tax rates. This would tend to bias estimates upward. However, since the yearly reductions were not sizable, and since the passage of JGTRRA accelerated the scheduled reductions to 2003, the scope for such shifting was limited, and so the extent of this bias is likely to be small. Ideally, one could estimate a specification that accounts for past and anticipated future tax rates (as was done in Bakija and Heim (2011)), but the short length of the Edited Panel precludes this.

Starting in 2002, EGTRRA also established a non-refundable tax-credit for low and moderate income filers known as the Saver's Credit.[5] The Saver's Credit provides a tax credit for contributions to qualified retirement savings accounts, including both employer-based plans and IRAs. The credit is figured as a percentage (either 50, 20, or 10 percent) of contributions up to $2,000, where the applicable rate depends on the adjusted gross income and filing status of the taxpayer.

In addition, EGTRRA changed the way contributions to employer-based accounts are treated when figuring the EITC. Prior to 2002, a taxpayer added pre-tax salary reductions (including contributions to employer-based plans) to the wages and salaries reported on the taxpayer's Form 1040 when calculating earned income, so that contributing a dollar of earnings to a tax-preferred account had no effect on a taxpayer's EITC. Starting in 2002, however, these amounts were not added back when figuring earned income for the EITC, so an additional dollar of earnings contributed to an employer-based account may affect the taxpayer's EITC.[6] This policy change has the effect of lowering the after-tax price of contributing to an employer-based account for those in the phase-out range of the EITC.[7]

If ordinary marginal income tax rates at the time of withdrawal are expected to be the same as at the time of contribution, the marginal tax rate changes (as well as the capital gains and "qualified" dividend tax rate reductions) would be expected to affect savings in tax-preferred accounts through their impact on the after-tax rate of return in ordinary accounts. Theoretically, an increase in the after-tax rate of return on savings will have two offsetting effects, a substitution effect, which increases savings, and an income effect (or a wealth effect in a two or more period model (Summers, 1981)), which decreases savings. Though the overall effect is ambiguous, Summers (1981, 1984) argues that increases in the after-tax rate of return should result in an increased level of savings. This, in turn, suggests that the existence of a tax-preferred savings vehicle, in which assets accumulate at the pre-tax rate of return, should increase savings (primarily in the tax-preferred account).[8] However, a reduction in the marginal tax rate and capital gains and dividend rate reductions will increase the after-tax rate of return for ordinary accounts, decreasing the tax advantage of saving in a tax-preferred account relative to an ordinary account, and would be expected to lead to decreased amount of savings in tax-preferred accounts. In addition, depending on the taxpayer's expectations about

[5] The Saver's Credit was initially scheduled to expire at the end of 2006, but was made permanent by the Pension Protection Act of 2006.

[6] This occurs as long as a taxpayer's earned income is greater than modified adjusted gross income (AGI) (defined as AGI plus tax exempt interest), which is most often the case. For example, consider a taxpayer with one child in the phase-out range of the EITC, who faces a phase-out rate of 15.98 percent. Prior to 2002, if they earned an additional dollar, regardless of whether they took it as income or contributed it to an employer-based account, their EITC would decrease by 15.98 cents. Starting in 2002, however, if they took the dollar as income, their EITC would decrease by 15.98 cents, but if they contributed the dollar to an employer-based account, their EITC would not change.

[7] It also has the effect of increasing the price of contributing for those in the phase-in range of the EITC, but the proportion of that group contributing to employer-based accounts is likely to be very small.

[8] As noted above, empirical studies on this topic have come to conflicting results.

future tax rates, a decrease in marginal tax rates today may increase the rate of return in a Roth IRA relative to a traditional IRA or an employer-based account, which could result in shifting between types of tax-preferred accounts.

In contrast to changes in marginal tax rates, the introduction of the Saver's Credit and the EITC change do not affect the after-tax rate of return in an ordinary account. They do, however, reduce the after-tax price of saving a dollar in tax-preferred accounts, thereby increasing the tax advantage of saving in a tax-preferred account relative to an ordinary account. As a result, they would be expected to lead to increased contributions to tax-preferred accounts because of a reduction in the after-tax price of contributing to such an account.

Marginal tax rate decreases will also result in higher after-tax incomes. Though the income effect of a change in the after-tax rate of return, as noted above, is negative, much previous empirical has found that higher incomes are associated with increased savings overall, and in tax-preferred retirement accounts in particular. In what follows, the effect of after-tax income on contributions will be referred to as the "after-tax income effect."

During the sample period, several other non-tax related policy changes occurred which may also affect the propensity of taxpayers to contribute to tax-preferred accounts.[9] First, in 1997, Simple IRA plans were introduced, and the number of firms offering plans and number of participants grew steadily over the subsequent years. Second, in 2001, participation in the federal Thrift Savings Plan was extended to members of the uniformed services. Third, the share of firms implementing automatic enrollment features in their employer-sponsored plans increased from 7.4 percent in 2002 to 16.9 percent in 2005. Finally, the number of 401(k) and other defined contribution plans increased, increasing the ability of employees to save in an employer-sponsored plan. In the estimation that follows, year fixed effects will be included to control for these other policy changes.

III. ESTIMATION METHOD AND DATA

A. Estimation Method

One difficulty in estimating the effect of tax policy on contributions to tax-preferred accounts is that the benefits from contributing, and thus the decision to contribute, depend on expectations about variables (such future tax rates, expected returns, and expected time until withdrawal) that are unobserved by the econometrician. In addition, the tax benefit from contributing to a traditional IRA, employer-based plan, or self-employed plan may differ from the benefit from contributing to a Roth IRA if marginal tax rates change.

Following much of the previous literature, to keep the estimation specification parsimonious and to reflect the parameters over which policymakers have control, the

[9] We thank Peter Brady for providing information on these other policy changes. For more detail about these changes, see Brady and Holden (2006), Holden, Brady, and Hadley (2006), and "Thrift Savings Plan," Federal Retirement Thrift Investment Board, http://www.tsp.gov/features/def_ch2-certain-other-employees.html.

after-tax price of contributing a dollar to a tax-deferred account[10] is used as a proxy for the benefits from contributing to any type of tax-preferred account.[11] [12] To derive this variable, consider a dollar contributed to a tax-deferred account prior to 2002. Assuming that the taxpayer can exclude or deduct the amount of the contribution from taxable income,[13] the taxpayer's income taxes would decrease by the amount of their marginal regular income tax rate. In those years, however, a taxpayer's EITC amount would be unaffected by the contribution, and the Saver's Credit was not available. Starting in 2002, in addition to decreasing taxes owed by the taxpayer's marginal regular income tax rate, the contribution would not be counted as income for EITC purposes, so that contributing a dollar would change an EITC eligible taxpayer's taxes owed by their marginal EITC rate. In addition, if a taxpayer were eligible for the Saver's credit, contributing a dollar would decrease their taxes owed by their Saver's Credit rate. Formally, the after-tax price of contributing to a tax-deferred account is given by[14]

$$(1) \quad P = \begin{cases} 1 - \tau_I & \text{if} \quad year < 2002 \\ 1 - \tau_I + \tau_{EITC} - \tau_{SAV} & \text{if} \quad year \geq 2002, \end{cases}$$

where τ_I denotes the individual's marginal income tax rate (including federal and state taxes),[15] τ_{EITC} is the combined marginal federal and state EITC rate,[16] and τ_{SAV} is the marginal Saver's Credit rate.[17]

---

[10] The after-tax price of contributing to a tax-deferred account is used rather than the after-tax price of contributing to the type of account that was actually contributed to — which for a Roth account would be 1 — to avoid introducing additional endogeneity of the after-tax price to contributing behavior.

[11] In addition to reflecting the cost in after-tax dollars of contributing a dollar to a tax-deferred account, the after-tax price of contributing (as defined below, but ignoring the EITC and Saver's Credit) is indirectly related to the relative after-tax rates of return in tax-preferred compared to ordinary accounts. If $(1 - t_I)$ is lower, the after-tax rate of return in an ordinary accounts taxed at ordinary rates is lower, and so the relative return in a tax-preferred account (either traditional or Roth) is higher.

[12] Since expectations about future rate changes are not observed in the tax data, it is not possible for us to control for these separately. If some changes in contemporaneous tax rates are correlated with changes in expectations about future tax rates (for example if marginal rate reductions for high-income taxpayers are correlated with the expectation that rates will go up in the future for these taxpayers), the estimates will attribute the effects of both changes to the change in the contemporaneous rate.

[13] As robustness checks, specifications that exclude those who are likely to be ineligible to make a contribution to a tax-preferred account, or for whom the contribution limit is likely to be binding, were estimated.

[14] For IRA contributions, this equation only applies to deductible contributions. For non-deductible contributions (which are not included in the measure of contributions used in this study), the after-tax price is simply 1.

[15] Burman, Gale, and Weiner (2001) found that the marginal effective tax rate for a Roth IRA would have been 11 percentage points higher on average than that for a traditional IRA or 401(k)-type plan during the 1982–1995 period. As a robustness check, a specification was tried in which the actual marginal tax rate was used if the taxpayer contributed to a Roth IRA, and 10 percentage points were subtracted off of the marginal tax rate if the taxpayer did not. The intensive margin results were similar to those presented below. The extensive margin results displayed more significant tax effects, but this is likely due to the fact that this procedure introduces a negative correlation between the contributing and after-tax price.

[16] As specified in this equation, τ_{EITC} is positive in the phase-in range and negative in the phase-out range.

[17] Marginal tax rates were calculated using tax calculators provided by Jon Bakija. Documentation for these calculators is detailed in Bakija (2009).

The taxpayer's contribution decision is assumed to also be a function of their after-tax income, denoted as *Y*. The measure of income that is used for this variable includes the taxpayer's income before any retirement contributions were made less the income taxes that applied to this amount of income.

The taxpayer's retirement account contribution decision, then, is modeled as consisting of two parts. First, the taxpayer chooses whether to contribute to an account or not. Second, if they choose to contribute they decide how much to contribute. Hence, the two choices can be represented by the following equations:

$$(2) \quad C_{it}^{*} = \beta_1 \ln\left(P_{it}^{0}\right) + \beta_2 \ln\left(Y_{it}\right) + \beta_3 X_{it} + \gamma_t + \nu_i + \varepsilon_{1it},$$

$$(3) \quad \ln\left(C_{it}\right) = \alpha_1 \ln\left(P_{it}\right) + \alpha_2 \ln\left(Y_{it}\right) + \alpha_3 X_{it} + \gamma_t + \nu_i + \varepsilon_{2it}, \text{ conditional on } C_{it}^{*} = 1,$$

where C_{it}^{*} is one if the taxpayer makes a contribution to one or more tax-preferred accounts and zero otherwise, and C_{it} is the total amount contributed. In these equations, X_{it} is a vector of the taxpayer's characteristics that vary over time, γ_t is a year fixed effect, ν_i is a taxpayer specific effect, and ε_{it} is the error term.[18]

The decision of whether or not to contribute is assumed to be a function of the after-tax price of the first dollar contributed to a retirement account,[19] denoted by P_{it}^{0}, while the decision of how much to contribute is assumed to be a function of the marginal or last dollar contributed to a retirement account, denoted by P_{it}. Since Koenig and Harvey (2005) found that 34 percent of those eligible to claim the Saver's Credit did not do so, to account for imperfect take-up, the Saver's Credit rate is incorporated in the first dollar tax price for any taxpayer that would be eligible to claim the credit it if they contributed to a tax-preferred account, but is only incorporated in the last dollar tax price if the taxpayer actually claimed the Saver's Credit in that year. Finally, both the extensive and intensive margins are assumed to be functions of the same pre-contribution after-tax income variable.

Included in X_{it} are age and age squared (because contributions have been found to significantly vary with the age and/or tenure of an employee), the number of children and number of children away from home (because the presence of children, and particularly children in college, may affect a taxpayer's ability to save for retirement), and

[18] Note that, unlike in Madrian and Shea (2001) and others, the effect of plan defaults on contribution behavior is not explicitly accounted for since such information is not available in tax data. If these plan defaults were the primary driver of contributing behavior, the estimation in this paper would find correspondingly small effect of taxes. To the extent that plan characteristics are constant for an individual over time, the individual fixed effects will control for these. In addition, if changes in plan characteristics (for example the institution of auto-enrollment) are orthogonal to the tax changes, the estimates presented here will not be biased by individuals responding to these changes.

[19] We also tried a specification in which the decision of whether to contribute was a function of the last dollar tax price. When this was done, the results were qualitatively similar to the first dollar specification, though the magnitudes of the tax price coefficients were substantially larger.

indicator variables for being married, the primary filer being female, being an itemizer, and the region of the country in which the taxpayer resides.

In addition, year fixed effects are included to account for changes in the policy environment that affected the availability and desirability of contributing to tax-preferred accounts, including changes in the capital gains and qualified dividend tax rates, nationwide trends in the availability of employer-based plans, and nationwide trends in match rates.[20]

Because panel data are used in this study, first differencing or including taxpayer fixed effects can be utilized to control for unobservable characteristics that are correlated with both the tax price of contributing and the propensity to contribute or the amount contributed. For example, suppose that some people tend to have a greater preference toward saving, and that this group also tends to have higher income, which in turn implies a lower after-tax price of contributing. This would lead to a spurious positive correlation between saving and income, and a negative correlation between savings and the after-tax price of contributing. First differencing or including taxpayer fixed effects controls for these characteristics, so that any effect of tax prices or income comes from variation across time for individual taxpayers. However, transitions in marital status between years may be correlated both with changes in marginal tax rates and changes in the propensity to contribute to tax-preferred accounts. As a result, a taxpayer observed after a marital status change is treated as a different taxpayer than the one observed before the change. In addition, the difference between transition years is not included in the first-differenced estimation sample.

The advantage of using first differencing instead of taxpayer fixed effects is that, if the change in the tax price or after-tax income is endogenous (as is argued below), one can use as instruments the changes in these variables that would have occurred if the taxpayer's income had remained the same (in real terms) between the two years of the difference.[21] However, using first differencing has a disadvantage that the right hand side variable is only non-zero in a year in which there was a tax change. As a result, if behavior is sticky or changes slowly, the estimates of the effect of changes in taxes could be downward biased. Estimating a fixed effects specification, on the other hand, will yield estimated coefficients that reflect behavior that changes slowly (since variables will be above or below their within-individual mean when later changes in behavior materialize), but requires an alternative instrumenting strategy. Since neither specification is clearly better, results for both specifications are presented below.

In this study, instrumenting for the after-tax price and income variables is necessary for three reasons. First, because tax prices are a nonlinear function of income in the current

[20] To examine the impact of including year fixed effects to account for these factors, we also estimated some specifications that excluded year effects. The finding of statistically significant effects of the after-tax price and after-tax income on the extensive margin was robust to this change in the fixed effects specification, though the after-tax price coefficient drops to statistical insignificance in the first differencing specification. On the intensive margin, all after-tax price and income effects were statistically insignificant when year effects are omitted, and the after-tax income effects reversed sign.

[21] A similar instrumental variable strategy was used in Gruber and Saez (2002), among others.

year, the contemporaneous tax price coefficient may still pick up nonlinear effects of income. So, in the main specifications that follow, when the data are first differenced, the difference in contemporaneous tax prices is instrumented with a difference in synthetic tax prices, where the synthetic tax prices are calculated by applying the tax law from each year of the difference to income and other tax relevant variables from the first year of the difference.[22] Alternatively, when fixed effects are included, the contemporaneous tax price is instrumented with a synthetic tax price that is calculated by applying the tax law from the current year to income and other tax relevant variables from the first year that the taxpayer is observed with the current year's filing status.[23]

Second, in the intensive margin (amount contributed) specifications, the last dollar tax price is directly a function of the amount contributed to tax-preferred accounts. Initially, then, the last dollar tax price is instrumented using the contemporaneous first dollar tax price as an instrument. In the main specification, however, the contemporaneous last dollar tax price is instrumented with synthetic first-dollar tax price instruments that are calculated in the manner described above.

Finally, without instrumenting for after-tax income, the coefficient on this variable will be identified both by variation in incomes due to changes in tax laws and by variation in incomes due to changes in pre-tax income. Since the focus of this paper is on the effect of tax rates on retirement savings, using both sources of variation will be problematic if year to year changes in pre-tax incomes have a different effect on savings (perhaps because they are viewed as transitory) than changes in incomes due to tax law changes (perhaps because they are viewed as being more persistent). So, to identify the effect of tax laws per se on retirement savings, in the first differencing specification the first difference in after-tax income is instrumented with a difference in synthetic after-tax incomes, calculated by applying the tax law from each year of the difference to the pre-contribution income from the first year of the difference. In the fixed effects specification, after-tax income is instrumented using synthetic after-tax income, calculated by applying tax law from the current year to the pre-contribution income from the first year the taxpayer is observed with the current year's filing status. When specified in this manner, any within taxpayer variation in the instruments comes solely from changes in tax law across years.

Ideally, one would control for the unobserved propensity to contribute, control for the endogeneity of the after-tax price, and instrument for changes in after-tax incomes, while estimating both equations simultaneously and explicitly controlling for selection bias (for example using a Heckman two-step procedure). However, the tax data contain

[22] For example, for the 1999–2000 difference, variables from 1999 are used to calculate the tax rates in 1999, then inflated by the CPI to 2000 levels and used to calculate the tax rates that would have applied to the taxpayer in 2000 if their income and other nominal variables had only increased with inflation.

[23] For example, for a taxpayer observed in 2002 that was first observed with a given marital status in 1999, variables from 1999 are inflated by the CPI to 2002 levels and used to calculate the tax rates that would have applied to the taxpayer in 2002 if their income and other nominal variables had only increased with inflation. In this specification, the first year of each taxpayer-filing status combination is dropped, since no instrument can be calculated for that observation.

no variables that could plausibly be included in the selection equation and excluded from the contribution equation, so the selection correction term could only be identified by making restrictive functional form assumptions. In addition, such a specification does not readily incorporate first differencing, individual fixed effects, or instrumental variables. Since the results are markedly different when either first differencing or individual fixed effects are utilized, we view controlling for unobserved characteristics as the more important issue, and thus estimate the two equations separately.

B. Data

The data used in this study come from a seven-year panel of tax returns known as the 1999–2005 Edited Panel.[24] Data from tax returns has a number of advantages for this type of study. First, because contribution information comes from information reports that are sent to the IRS, the exact amount of contributions made is observed. Second, contributions to all types of accounts are observed, including to employer-based accounts, IRAs (both traditional and Roth), and self-employed plans. Finally, because all of the taxpayer's tax-relevant information is observed, it is possible to accurately calculate the taxpayer's marginal tax rate, Saver's Credit rate, EITC rate, and tax-preferred account contribution limits.[25]

The full Edited Panel consists of a random sample known as the Continuous Work History Subsample (CWHS), which contains a random sample of taxpayers for whom the primary filer's social security number ended in one of five four-digit combinations, and a high income oversample. For the estimation, the panel is cut to include only the CWHS subsample, which includes 384,517 returns from 63,321 different taxpayers over the seven years. Of these taxpayers, 69 percent are in the sample all seven years.

The sample was also cut to include only returns where both taxpayers (primary and secondary) are age 25 to 55, to focus on taxpayers who are in their prime earning years. This age cut eliminates those who might be finishing their schooling, and also eliminates those who would be eligible to withdraw from their tax-preferred savings account without penalty because they had reached the threshold age.

Amounts contributed to tax-preferred accounts come from three sources. First, deductible contributions to traditional IRA accounts and self-employed plans come from the relevant lines of the taxpayer's Form 1040.[26] Second, contributions to Roth IRAs are reported on Form 5498. Finally, contributions to employer-based retirement accounts come from the relevant boxes on the W-2 form(s) that were matched to the sample of

[24] For more information on the 1999 Edited Panel, see Weber and Bryant (2005).

[25] A downside of tax data, however, is that only contributions are observed, and not account balances (unlike the Survey of Consumer Finances). Thus, it is not possible to assess whether any changes in contributions were associated with an increase in savings in tax-preferred accounts or overall savings. The availability or amount of an employer match is also not observed.

[26] The relevant lines from Form 1040 include line 23 in 1999–2001, line 24 in 2002–2003, line 25 in 2004, and line 32 in 2005.

tax returns.[27] There were, however, a small fraction of observations for which wages and salaries were reported on Form 1040, but no corresponding W-2 form existed in the W-2 dataset. Because it is not possible to observe whether contributions to an employer-based retirement account were made for these taxpayers, they are eliminated from the estimation sample.

Unfortunately, the W-2 form does not contain detailed information about the taxpayer's employer-based plan, including information on the presence and amount of an employer match.[28] The lack of detail about match rates in an employer's plans is offset somewhat by the availability of panel data, since individual and time fixed effects can be included to control for unchanging match levels at the individual level and nationwide trends in match rates. To the extent that changes in match rates are correlated with changes in tax law at the within-individual level (for example, if individuals with larger increases in the after-tax price of contributions tended to have also experienced larger decreases in match rates), the regression results may be biased, though this bias is likely to be smaller than that found when using a cross-section.[29]

The dataset also does not contain an indicator for eligibility to participate in an employer-based retirement savings plan. To deal with this issue, we assume that all taxpayers are eligible to contribute to one of the types of tax-preferred plans, and include all observations in the sample. The robustness of the results to this assumption is examined below.

Sample statistics are presented in Table 1. Throughout the paper, all dollar values have been deflated to 2000 dollars using the consumer price index for all urban consumers (CPI-U). The first and second columns present means and standard deviations of the relevant variables for all observations, the third and fourth columns present these statistics for all contributors, and the fifth and sixth columns present these statistics for all noncontributors. The full sample consists of 192,293 observations, whereas the contributor subsample consists of 92,951 observations. The mean after-tax income among contributors is $71,474, which is roughly double the mean among those who do not contribute, and the mean first dollar tax price is lower for those who contribute than those who do not. Contributors also tend to be more likely to be married, older, and to itemize deductions.

Figure 1 presents the time trend of the first dollar tax price for all observations and for subgroups divided by income (whether income modally fell below $50,000 across all years of the panel, modally fell between $50,000 and $100,000, or modally was above $100,000). In order to remove variation due to changes in sample composition across years, for this figure the sample was cut to include only taxpayers who were present in the sample all seven years. In this figure, the three policy changes that

[27] The relevant boxes from the W-2 forms include box 13 in 1999–2000 and box 12 in 2001–2005.

[28] Employer matching contributions are widespread, though the magnitude of the match varies. For example, Holden and VanDerhei (2001) report that, in data from 1999, 91 percent of participants in 401(k) plans were offered an employer matching contribution. Under the most common offering, the employer matched 50 percent of up to 6 percent of the employee's salary (for a maximum employer contribution of 3 percent of salary).

[29] Higher income taxpayers (whose tax price is lower) are likely to work at firms that provide a larger match, possibly yielding a spurious negative correlation between contributing and the tax price in the cross-section.

Table 1

Sample Statistics

	Eligible to Contribute		Contributors		Non-contributors	
	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation
Income variables						
Indicator variable for making a contribution	0.483	0.500	1.000	0.000	0.000	0.000
Total contributions ($)	2,090	3,884	4,324	4,643	0	0
Last dollar tax price	0.780	0.149	0.741	0.098	0.816	0.177
First dollar tax price	0.778	0.150	0.736	0.100	0.816	0.177
After-tax income ($)	52,769	128,162	71,474	171,938	35,268	59,158
Demographic variables						
Married	0.500	0.500	0.619	0.486	0.389	0.487
Age	39.532	8.257	40.639	8.036	38.496	8.327
Age squared	1,631	654.6	1,716	647.1	1,551	651.6
Sex of primary filer (1 = female)	0.291	0.454	0.231	0.422	0.347	0.476
Itemizer	0.474	0.499	0.634	0.482	0.324	0.468
Number of children	1.048	1.126	1.048	1.116	1.048	1.136
Number of children away from home	0.010	0.102	0.010	0.101	0.010	0.102
Census division						
New England	0.051	0.221	0.060	0.238	0.043	0.203
Mid-Atlantic	0.142	0.349	0.144	0.351	0.139	0.346
East North Central	0.163	0.369	0.174	0.379	0.153	0.360
West North Central	0.068	0.251	0.078	0.268	0.058	0.234
South Atlantic	0.189	0.392	0.183	0.387	0.195	0.396
East South Central	0.060	0.238	0.053	0.224	0.067	0.251
West South Central	0.105	0.307	0.094	0.291	0.116	0.321
Mountain	0.063	0.242	0.063	0.243	0.062	0.241
Pacific	0.155	0.362	0.148	0.355	0.162	0.368
Year						
1999	0.155	0.362	0.145	0.353	0.164	0.370
2000	0.150	0.357	0.149	0.356	0.151	0.358
2001	0.147	0.354	0.148	0.355	0.146	0.353
2002	0.139	0.346	0.142	0.349	0.136	0.343
2003	0.135	0.342	0.137	0.344	0.133	0.340
2004	0.132	0.338	0.136	0.343	0.128	0.334
2005	0.142	0.349	0.143	0.350	0.142	0.349
Number of observations	192,293		92,951		99,342	

Notes: Data are from the 1999–2005 Edited Panel of tax returns. All dollar values have been deflated to 2000 dollars using the CPI-U.

Figure 1

Means of Tax Price Variables, by Year

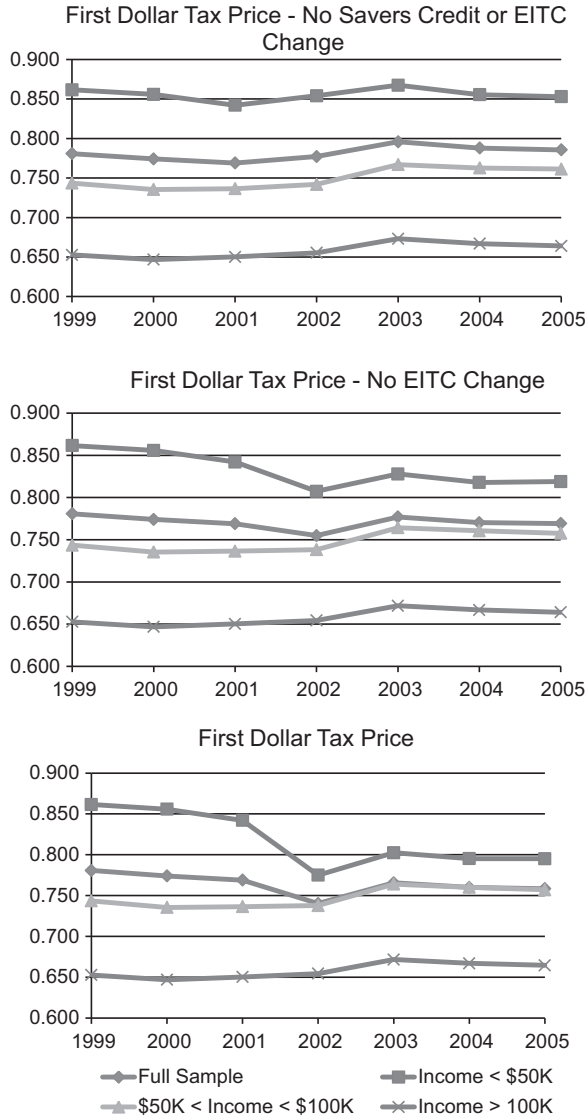

Notes: Data are from the 1999–2005 Edited Panel of tax returns. The sample was cut to include only taxpayers who were present in the sample all seven years, and who were imputed as being eligible to contribute.

occurred during the sample period — marginal tax rate changes, the introduction of the Saver's Credit, and the exclusion of employer-based retirement plan contributions from earnings for EITC purposes — are added sequentially, so that it is possible to examine how the different policies contribute to the identifying variation in the tax price.

The first panel presents what the first dollar tax price trends would have been had marginal tax rates changed in isolation. For the full sample, the mean tax price would have increased by about 0.02 over the sample period, driven mainly by the marginal tax rate reductions in 2001 and 2003 for the middle and upper income groups, while the tax price for the lowest income group stayed relatively constant. The middle panel then adds the effect of the Saver's Credit. As expected, this addition primarily affects the tax price of the lowest income group, which now displays a decreasing trend due to a decline in the tax price after the introduction of the Saver's Credit in 2002. The tax price for the middle group also declines somewhat starting in 2002, while the trend for the upper income group is unchanged. As a result, the overall trend now decreases. The bottom panel adds the effect of the EITC. Compared to the previous panel, the tax price of the lowest and middle groups again declines starting in 2002, while the tax price of the highest income group is unchanged.

The bottom panel of this figure presents the actual trend in the first dollar tax price. Among all members of the sample in this panel, the mean first dollar tax price drops slightly between 1999 and 2001, drops a larger amount in 2002, then increases in 2003 to its former level before decreasing slightly in 2004 and 2005. Breaking this trend out by income group, for the lowest group, the tax price declines slightly from 1999–2001, decreases substantially in 2002 due to the introduction of the Saver's Credit and the change in the EITC, and then increases in 2003 due to the JGTRRA marginal tax rate reductions. For the middle group, the tax price declines a small amount in 2002 due to the Saver's Credit and the EITC, and then increases in 2003 due to the JGTRRA marginal rate reductions, while for the upper income group, the tax price stays relatively constant from 1999–2002, and then increases in 2003 due to the JGTRRA marginal tax rate reductions.

Figure 2 presents the time trends of the contribution variables for the full sample and the same three subgroups. In the top panel, levels of participation in tax-preferred retirement savings accounts are seen to increase across the years of the sample, with a jump between 1999 and 2000, and a steady increase thereafter. It appears that this increase is centered among the lowest two income groups, with increases of 25 percent for the lowest income group and 9 percent for the middle income group. The probability of contributing for the highest income group increases in 2000, but then stays relatively stable through the rest of the period.

The bottom panel presents mean contributions (in deflated 2000 dollars) among contributors. Overall, contributions increased steadily from 1999–2005. The increases are spread across all groups, with contributions from 1999–2005 increasing among the lowest group by 20 percent , by 31 percent for the middle group, and by 41 percent for the highest group.

Figure 2

Means of Contribution Variables, by Year

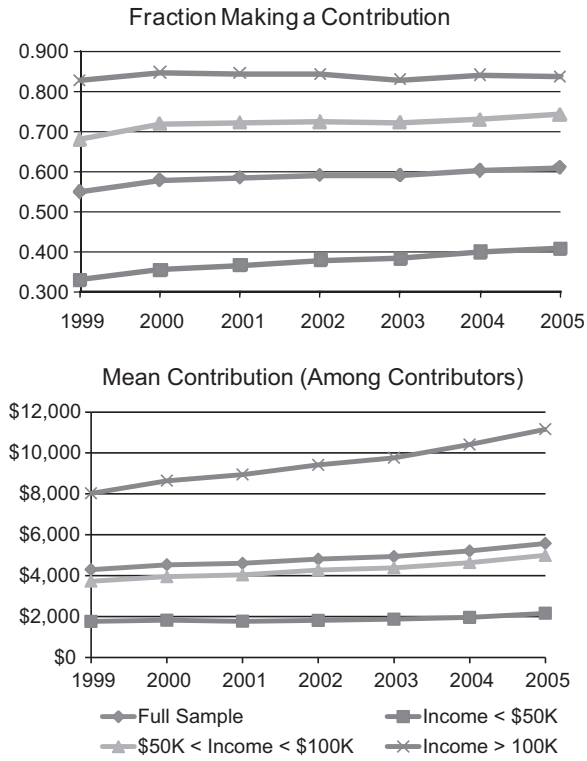

Notes: Data are from the 1999–2005 Edited Panel of tax returns. The sample was cut to include only taxpayers who were present in the sample all seven years, and who were imputed as being eligible to contribute. All dollar values have been deflated to 2000 dollars using the CPI-U.

Looking across these trends, it is hard to discern any obvious negative relationship between the price of contributing to a tax-preferred account and either the propensity to contribute or the level of contributing. However, these aggregate trends could be confounded by changes in incomes and the macroeconomy that affect contributions to retirement accounts, and they do not exploit variation across individuals and across time within an individual taxpayer.

The instrumental variables regressions thus will help to discern whether any relationship between these variables arises under a more comprehensive identification strategy. To illustrate the identifying variation used in the empirical specification, in Figure 3, the sample each year was divided into quintiles based on the difference in the synthetic after-tax price. The means of the differences in the synthetic after-tax price were then calculated for each quintile in each year. In this figure, it is clear that during

Figure 3

Mean of Instrument for First Difference of Log Tax Price, by Quintile

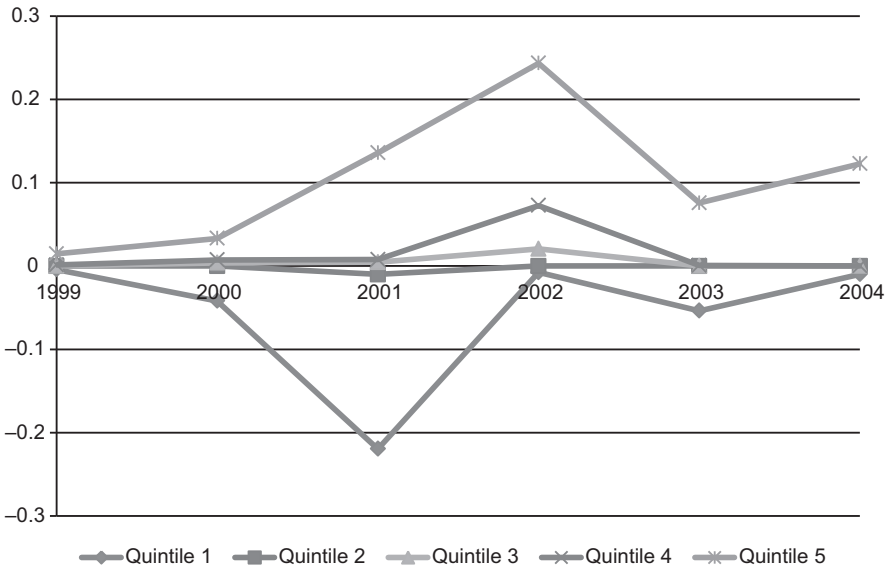

Notes: Data are from the 1999–2005 Edited Panel of tax returns.

most of the sample period there was considerable variation in yearly differences in the synthetic tax prices. In 2001, when the spread was the greatest, the lowest quintile had a mean log difference of –0.22, while the highest quintile had a mean log difference of 0.14. Other years also exhibit substantial spreads across quintiles, suggesting that there exists sufficient variation to identify the effect of these changes on contributions to tax-preferred retirement plans.

IV. RESULTS

A. Extensive Margin

Table 2 presents the results from estimating the responsiveness to recent tax changes on the extensive margin — the decision of whether to contribute — using a linear probability model. Since the tax price and net-of-tax income are entered in the regressions in logs, the estimated coefficients can be interpreted as semi-elasticities.[30]

In column 1, (2) is estimated without first differencing, fixed effects, or instrumenting for the after-tax price or income. In this specification, both the after-tax price and

[30] That is, the coefficients reflect the percentage point change in the probability of contributing that would result from a 1 percent change in the tax price or after-tax income.

Table 2

Extensive Margin Estimation Results — Full Sample

	Contributor to Any Plan				
	Regression in Levels (1)	First Difference (2)	Fixed Effects (3)	First Difference, Base Year Instrument (4)	Fixed Effects, First Year Instrument (5)
Ln(first dollar tax price)	–0.021***	0.000	0.004	–0.046***	–0.136***
	(0.007)	(0.006)	(0.006)	(0.017)	(0.029)
Ln(after-tax income)	0.250***	0.097***	0.132***	0.054	0.258**
	(0.003)	(0.004)	(0.003)	(0.059)	(0.119)
Married	0.056***				
	(0.005)				
Age	0.006***				
	(0.002)				
Age-squared	–0.007***	–0.019***	–0.013***	–0.021***	–0.005
	(0.002)	(0.006)	(0.003)	(0.008)	(0.007)
Sex of primary filer	0.057***				
	(0.004)				
Itemizer	0.070***	–0.003	0.009**	0.003	–0.012
	(0.004)	(0.004)	(0.004)	(0.009)	(0.021)
Number of children	–0.044***	–0.004	–0.007***	–0.001	–0.008
	(0.002)	(0.002)	(0.002)	(0.003)	(0.005)
Number of children away from home	–0.014	–0.016	–0.002	–0.017	–0.005
	(0.015)	(0.013)	(0.014)	(0.013)	(0.014)
First difference		yes		yes	
Individual fixed effects			yes		yes
Observations	192,293	132,296	147,013	132,271	146,795

Notes: Data are from the 1999–2005 Edited Panel of tax returns. Robust standard errors are in parentheses. Asterisks denote significance at the 10% (*), 5% (**), and 1% (***) levels. In column (4), the change in the log first dollar tax price is instrumented with the change in the log first dollar tax price calculated using income variables from the first year of the two year difference (inflated by the CPI for the tax calculation in the later year), and the change in the log of after-tax income is instrumented using the change in the log of income calculated using income variables from the first year of the two year difference (inflated by the CPI for the tax calculation in the later year). In column (5), the log first dollar tax price is instrumented with the log first dollar tax price calculated using income variables from the first year that the tax unit is observed with that marital status (where the income variables are inflated by the CPI for the tax calculation in the later year), and the log of after-tax income is instrumented using the log of after-tax income calculated using income variables from the first year that the tax unit is observed with that marital status (where the income variables are inflated by the CPI for the tax calculation in the later year). All specifications include a year effects, region effects, and a constant term.

after-tax income variables are highly statistically significant, with an after-tax price coefficient of –0.021 and an after-tax income coefficient of 0.250. The demographic characteristics generally enter with the expected sign, with married taxpayers, older taxpayers, and itemizers more likely to contribute, and taxpayers with more children less likely to contribute.

However, if people who have an unobserved preference toward savings in general also tend to have higher income (and thus a lower tax price due to the progressivity of the income tax), these coefficients will both be biased upward (in absolute value). In column 2, the data are first differenced in an attempt to control for the taxpayer-specific propensity to contribute, while in column 3 fixed effects are added. When this is done, the coefficient on the tax price drops to a statistically insignificant 0.000 in the first differencing specification and 0.004 in the fixed effects specification, suggesting that unobserved propensity to contribute to tax-preferred accounts was biasing the after-tax price effect upward in column 1. The after-tax income coefficient also drops to (a still statistically significant) 0.097 in the first differencing specification and 0.132 in the fixed effects specification.[31] However, it is important to note that, in these specifications, identification of the income effect comes from year to year fluctuations in income. As a result, the after-tax income coefficient represents the effect of transitory changes in after-tax income on the decision to contribute, so that the decrease in the after-tax income effect is reflecting a combination of both controlling for taxpayer fixed effects and a switch to estimating a transitory after-tax income effect.[32]

To deal with the potential endogeneity of marginal tax rates, and to remove the effects of transitory income shocks, the specifications in columns 4 and 5 implement the instrumental variable strategies described above using the synthetic first dollar tax price and after-tax income.[33] When this is done, the coefficient on the tax price increases to –0.046 in the first differencing specification, and to –0.136 in the fixed effects specification, and is again highly statistically significant. To put the estimated effect in context, from 1999–2005 the mean after-tax price — among those in the sample in 1999 — decreased by 2.9 percent. Given the coefficient of –0.046 (or –0.136), these results suggest that the tax changes during this period increased the probability of contributing by 0.13 percentage points (or 0.39 percentage points) which, compared to a baseline contribution rate of 54.9 percent in 1999, corresponds to about a 0.23 percent (or 0.71 percent) increase. These average effects, however, mask heterogeneity across income levels in the direction of the effect. As can be seen in Figure 1, during the sample period, the mean tax price among taxpayers with income less than $50,000 decreased

[31] Estimating a fixed-effects logit specification without instrumenting for the price and income variables yielded larger and statistically significant effects for both the after-tax price and after-tax income variables.

[32] For more detailed discussion on estimating the effect of transitory and permanent income on contributions to tax-preferred retirement accounts, see Heim and Lurie (2010).

[33] Results from the first stage regressions for these specifications are available from the authors. The F statistics on the excluded instruments were generally well in excess of 10, suggesting that weak instruments are not a concern. The only exception was in the first differenced specification, in which case the F statistic was 8.64 for the excluded instruments in the ln(after-tax income) first stage.

by 7.7 percent, while it increased by 1.7 percent for taxpayers with income greater than $100,000. These coefficients imply that the probability of contributing for the lowest income group increased by 0.35 (or 1.05) percentage points, but that it decreased by 0.08 (or 0.23) percentage points for the higher income group. In any case, it appears that the changes in the after-tax price had a statistically significant, though small, effect on the decision to contribute.

The after-tax income coefficient, on the other hand, declines to 0.054 in the first differencing specification and is now statistically insignificant, while it increases to 0.258 in the fixed effects specification. This latter coefficient provides some suggestive evidence that higher after-tax incomes resulting from decreases in taxes may have led to an increased rate of contributions to tax-preferred accounts, though the result is not robust.[34]

Overall, the results in this table suggest that changes in the after-tax price of contributing had a statistically significant, though relatively small, effect on the probability of contributing to a tax-preferred retirement account. The evidence on changes in after-tax incomes, however, is mixed, with one explanation for the discrepancy between the two specifications being that taxpayers take time to adjust their contributing behavior in response to tax-induced changes in incomes.

The specifications above assumed that all taxpayers are eligible to contribute to one of the types of tax-preferred plans, which is clearly a simplification. As a robustness check, we ran a specification that included only taxpayers who either (1) are eligible to contribute to a traditional or Roth IRA or self-employed plan,[35] or (2) are working in a particular year for an employer at which a contribution to an employer-based plan was made in at least one year of the panel.[36] [37] Using this imputation eliminated 904 observations from the estimation sample of 131,367 when the data are first differenced and the base year instrument is used.

Results from these specifications are presented in column 1 of Table 3. Results from the first difference specification are presented in the top panel, while results from the fixed effects specification are presented in the bottom panel. Comparing these results to those in Table 2 demonstrates that the change in sample does not qualitatively affect the results.

---

[34] A possible concern with the specifications in columns 4 and 5 is that the estimated tax price effect may be biased if income is mean-reverting while taxes change to a greater extent for high-income taxpayers, since contributions may be falling due to mean reversion while increasing (or decreasing) due a drop (or increase) in the tax price. To account for this, following Gruber and Saez (2002), a spline in the log of after-tax income from the base year was included. These results suggested that mean reversion did not appreciably bias the results in columns 4 and 5, as the coefficient on the tax price declines only slightly to –0.037 in the first differences specification and to –0.099 in the fixed effects specification.

[35] Eligibility is inferred using information from Form 1040 and from W-2 forms.

[36] For example, if a taxpayer was observed working for the same employer in all years of the panel and contributed to an employer-based plan in 2000, they would be considered eligible to contribute in all years of the panel.

[37] This definition will correctly pick up taxpayers who were eligible in all years they worked for a particular employer, and who contributed at least once. However, it will not pick up workers who were eligible in all years but never contributed, and it will erroneously include taxpayers who stayed at the same employer, but whose eligibility changed over time.

Table 3

Extensive Margin Estimation Results — Sample of Taxpayers Imputed as Eligible to Contribute, by Type of Contribution

	Contributor to Any Plan (1)	Contributor to Employer-Based Plan (2)	Contributor to Traditional IRA (3)	Contributor to Roth IRA (4)	Contributor to Self-Employed Plan (5)
First difference specifications					
Ln(first dollar tax price)	–0.047***	–0.112**	0.001	–0.010	0.004
	(0.017)	(0.046)	(0.008)	(0.007)	(0.012)
Ln(after-tax income)	0.045	0.232	0.009	0.004	–0.037
	(0.059)	(0.202)	(0.028)	(0.033)	(0.047)
Observations	131,367	74,815	95,150	119,814	9,463
Fixed effects specifications					
Ln(first dollar tax price)	–0.133***	–0.392***	–0.103***	–0.139***	0.058
	(0.029)	(0.097)	(0.034)	(0.030)	(0.111)
Ln(after-tax income)	0.240**	0.391*	0.131	0.167*	0.369
	(0.116)	(0.236)	(0.117)	(0.101)	(0.262)
Observations	146,051	82,951	114,115	135,666	15,744

Notes: Data are from the 1999–2005 Edited Panel of tax returns. Robust standard errors are in parentheses. Asterisks denote significance at the 10% (*), 5% (**), and 1% (***) levels. In the top panel, the change in the log first dollar tax price is instrumented with the change in the log first dollar tax price calculated using income variables from the first year of the two year difference (inflated by the CPI for the tax calculation in the later year), and the change in the log of after-tax income is instrumented using the change in the log of income calculated using income variables from the first year of the two year difference (inflated by the CPI for the tax calculation in the later year). In the bottom panel, the log first dollar tax price is instrumented with the log first dollar tax price calculated using income variables from the first year that the tax unit is observed with that marital status (where the income variables are inflated by the CPI for the tax calculation in the later year), and the log of after-tax income is instrumented using the log of after-tax income calculated using income variables from the first year that the tax unit is observed with that marital status (where the income variables are inflated by the CPI for the tax calculation in the later year). All specifications include the demographic characteristics in Table 2, region effects, year effects, and a constant term.

An advantage of using this finer cut is that it allows for the estimation of responsiveness on the extensive margin for each type of plan in isolation (traditional IRA, Roth IRA, employer-based plan, or self-employed plan). In columns 2 through 5 of Table 3, the estimation sample includes only those imputed as eligible to contribute to each of the four types of tax-preferred retirement savings plans. When this is done, the after-tax price is consistently estimated to have a statistically significant effect on contributions only for contributions to employer-based plans, with an estimated coefficient of –0.112

in the first differencing specification and an estimated coefficient of –0.392 in the fixed effects specification. In addition, the after-tax income effect is estimated to be largest for those imputed as eligible to contribute to an employer-based plan (though the coefficient is only statistically significant in the fixed effects specification), suggesting that the contributing behavior of those eligible for employer-based plans is more responsive on the extensive margin than the behavior of taxpayers eligible for other types of plans.

B. Intensive Margin

Table 4 presents the results of estimating the effect that recent tax changes have had on tax-preferred savings on the intensive margin — the decision of how much to contribute. For this specification, the dependent variable is the log of contributions, and so any taxpayer-year in which a contribution was not made is excluded from the estimation sample. Thus, estimated coefficients should be interpreted as the effect of after-tax prices and income on the amount contributed among those contributing a positive amount. Since both the amounts contributed and the price and income variables are in logs, the coefficients on the price and income variables can be interpreted as the elasticity of contributions to after-tax prices and after-tax incomes.

In column 1, when the data are not first-differenced and individual fixed effects are omitted, the coefficient on the after-tax price estimated to be –0.489, and the coefficient on after-tax income estimated to be 1.081; both coefficients are statistically significant.[38] The coefficients on the demographic characteristics generally have the expected sign, as itemizers are estimated to contribute more, and those with more children and with children away from home estimated to contribute less. Interestingly, married couples are estimated to contribute less, and the estimated effect is strongly statistically significant.

However, not controlling for the unobserved propensity to save could bias both the price and income coefficients upward in absolute value. In column 2, the data are first differenced, while fixed effects are included in column 3. When this is done, the after-tax price coefficient drops to a statistically insignificant –0.012 in the first differencing specification and to a statistically insignificant 0.015 in the fixed effects specification. As noted above, the after-tax income coefficient now reflects the effect of transitory changes in income, but that coefficient declines as well to 0.581 and 0.656.

Since tax prices are a nonlinear function of income in the current year, the tax price in this specification may still be biased by picking up nonlinear effects of income. In addition, the identifying variation for the income coefficient includes income changes

[38] Estimating a Tobit specification with a lower limit at zero, no upper limit, and instrumenting for the last dollar tax price with the first-dollar tax price yielded a smaller after-tax price effect, but a larger after-tax income effect. Estimating a Tobit specification with a lower limit at zero and an upper limit at the log of the contribution limit but with no instrumenting for the last dollar tax price, yielded larger after-tax price and after-tax income effects. However, neither of these specifications control for unobserved individual characteristics, which (as shown below) has a substantial impact on the estimated coefficients.

Table 4

Intensive Margin Estimation Results

	Regression in Levels (1)	First Difference (2)	Fixed Effects (3)	First Difference, Base Year Instrument (4)	Fixed Effects, First Year Instrument (5)
Ln(last dollar tax price)	–0.489***	–0.012	0.015	0.047	–0.098
	(0.056)	(0.050)	(0.039)	(0.105)	(0.277)
Ln(after-tax income)	1.081***	0.581***	0.656***	0.225*	0.385
	(0.020)	(0.023)	(0.013)	(0.126)	(0.503)
Married	–0.159***				
	(0.016)				
Age	0.007				
	(0.007)				
Age-squared	0.005	–0.077***	0.001	–0.104***	–0.017
	(0.008)	(0.016)	(0.011)	(0.019)	(0.034)
Sex of primary filer	–0.120***				
	(0.015)				
Itemizer	0.141***	0.001	–0.002	0.030*	0.028
	(0.013)	(0.012)	(0.011)	(0.016)	(0.056)
Number of children	–0.126***	–0.032***	–0.043***	–0.026***	–0.037***
	(0.006)	(0.008)	(0.007)	(0.008)	(0.013)
Number of children away from home	–0.101**	–0.038	0.025	–0.033	0.025
	(0.044)	(0.039)	(0.039)	(0.040)	(0.040)
First difference		yes		yes	
Fixed effects			yes		yes
First dollar after-tax price instrument	yes	yes	yes		
First dollar synthetic after-tax price and income instrument				yes	yes
Observations	92,951	60,302	73,918	60,296	73,827

Notes: Data are from the 1999–2005 Edited Panel of tax returns. Robust standard errors are in parentheses. Asterisks denote significance at the 10% (*), 5% (**), and 1% (***) levels. In column (2), the change in the log last dollar tax price is instrumented with the change in the log first dollar tax price. In column (3), the log last dollar tax price is instrumented with the log first dollar tax price. Refer to the notes from Table 2 for columns (4) and (5). All specifications include region effects, year effects, and a constant term.

not caused by changes in taxes. In columns 4 and 5, the change in the last dollar tax price is instrumented with the change in the synthetic first dollar tax price, and the change in after-tax income is instrumented with the change in synthetic after-tax income.[39] In these specifications, the coefficient on the after-tax price of 0.047 reflects a change in sign to 0.047 in the first differencing specification and is –0.098 in the fixed effects specification, though both of these coefficients are still statistically insignificant. The coefficient on the after-tax income variable is marginally statistically significant in the first differencing specification, with a coefficient of 0.225. To put this estimate in context, between 1999 and 2005, the mean after-tax income increased by 3.3 percent, implying that the tax policy changes during this period increased the amount contributed by 0.74 percent by increasing after-tax incomes. The estimate in the fixed effects specification is somewhat larger, at 0.385, but due to a large standard error is statistically insignificant.[40] [41]

Overall, it appears that the changes in the after-tax price of contributing had little effect on the amount contributed during the sample period. There is, however, suggestive evidence that tax-policy induced increases in after-tax incomes may have led to increased amounts contributed to tax-preferred retirement savings accounts.[42]

To examine whether contributing behavior differs depending on the type of accounts to which the taxpayer is contributing, we estimate regressions in which the left hand side variable is changed to the amount of contributions to a particular type of account. The results are provided in Table 5. Results from the first differencing specifications are presented in the top panel, while results from the fixed effects specifications are presented in the bottom panel. For comparison, column 1 repeats the estimated coefficients from the base specification in columns 4 and 5 of Table 4, where the dependent variable was

[39] Results from the first stage regressions for these specifications are available from the authors. The F statistics on the excluded instruments were well in excess of 10, suggesting that weak instruments are not a concern.

[40] We also estimated a specification that included a 10-piece spline in the log of after-tax income to control for mean reversion of income. In this case, the tax price coefficients were statistically insignificant in both the first differencing and the fixed effects specifications. The after-tax income variable, on the other hand, decreased and became statistically insignificant in the first differencing specification, while it increased and became statistically significant in the fixed effects specification.

[41] One further concern with these specifications arises because there are annual limits on the amount that may be contributed to tax-preferred accounts. As a result, including in the sample taxpayers whose contributions are at or near these limits may bias the coefficients in this regression, since these taxpayers could only respond to the change in tax policy by increasing their contributions up to the limit that applies in the next year, if that limit changed. As a robustness check, a specification was tried in which any taxpayer who was observed to be contributing an amount within $100 of their limit was dropped from the sample. When this was done, the coefficient on the after-tax price declined slightly in absolute value to 0.022 in the first differencing specification and increased to –0.201 (though the coefficient was still statistically insignificant), while the coefficients on after-tax income were marginally statistically significant 0.405 in the first differencing specification and 0.304 in the fixed effects specification.

[42] Specifications using a balanced panel were also estimated, but the results did not change markedly from the base specification.

Table 5

Intensive Margin Estimation Results — by Type of Contribution

	All Contributions (1)	Employer-Based Plan Contributions (2)	Traditional IRA Contributions (3)	Roth IRA Contributions (4)	Self-Employed Plan Contributions (5)
First difference specifications					
Ln(last dollar tax price)	0.047	0.039	0.456	–0.158	–1.330
	(0.105)	(0.106)	(0.322)	(0.450)	(4.468)
Ln(after-tax income)	0.225*	0.233*	–0.342	0.359	–2.056
	(0.126)	(0.136)	(0.480)	(0.428)	(3.860)
Observations	60,296	55,484	2,554	5,545	744
Fixed effects specifications					
Ln(last dollar tax price)	–0.098	–0.288	1.389	1.372	–14.549
	(0.277)	(0.256)	(1.310)	(4.883)	(21.840)
Ln(after-tax income)	0.385	0.214	–0.223	0.848	–4.843
	(0.503)	(0.522)	(1.308)	(5.018)	(7.518)
Observations	73,827	68,116	4,399	8,305	1,114

Notes: Data are from the 1999–2005 Edited Panel of tax returns. Robust standard errors are in parentheses. Asterisks denote significance at the 10% (*), 5% (**), and 1% (***) levels. Refer to notes from Table 3 for the treatment of panels. All specifications include the demographic characteristics in Table 4, region effects, year effects, and a constant term.

contributions to all types of tax-preferred accounts. In column 2, the dependent variable includes only contributions to employer-based retirement accounts, the dependent variable in column 3 includes deductible contributions to traditional IRAs, the dependent variable in column 4 includes contributions to Roth IRAs, and the dependent variable in column 5 includes contributions to self-employed plans. Across all columns, the coefficients on the after-tax price variables are statistically insignificant. In addition, only for employer-based plans is the effect of after-tax income on contributions estimated to be statistically significant, and even then only marginally in the first differencing specification, with a coefficient of 0.233.

C. Discussion

The results imply statistically significant but relatively small effects of the after-tax price of contributing on the extensive margin, and little effect of the after-tax price of

contributing on the intensive margin. A possible reason for the difference between intensive and extensive margin behaviors may be found in the literature that examines the effects of tax-preferred account plan characteristics on contributions to those accounts, particularly since the decision to contribute to these types of accounts tends to be the most responsive. It could be that changes in taxes (or possibly in after-tax incomes) induce taxpayers to create an account, but once created, they simply contribute some default amount, or the amount that maximizes their match.

In any case, the extensive margin results are broadly consistent with Venti and Wise (1988), Joulfaian and Richardson (2001), Power and Rider (2002), and Englehardt and Kumar (2007), who each find statistically significant effects of the marginal tax rates on the decision to contribute. It is not possible to compare the magnitudes of the coefficients estimated here to those in Joulfaian and Rider (2001), as they do not report the semi-elasticities that their estimates would imply. However, the estimated effects in the latter two papers are substantially larger than those found here, with implied semi-elasticities of –0.250 in Power and Rider (2002) and –0.828 in Englehardt and Kumar (2007). In addition, both of these papers find statistically significant effects of marginal tax rates on the intensive margin, which were not found here. Differences in samples likely go a long way in explaining these differences, as Power and Rider (2007) focus on self-employed retirement savings, and Englehardt and Kumar (2007) focus on 401(k) savings of individuals who are between age 51 and 61 and are nearing retirement, and both of these groups are likely to be more responsive to incentives than taxpayers in general.[43]

The intensive margin results are, however, inconsistent with Feenberg and Skinner's (1989) and Long's (1990) examinations of IRAs. It should be noted, though, that these papers used cross-sectional variation to identify their tax prices coefficients. In this study, when the tax price coefficient was identified by treating the data as pooled cross-sections, statistically significant and larger effects of the tax price on the intensive margin were also found. This suggests that being able to control for individual fixed effects and instrumenting for the tax price and after-tax income are largely responsible for the differences between this study and those previous studies.

Finally, previous studies, including Feenberg and Skinner (1989), Venti and Wise (1988), and Joulfaian and Richardson (2001) tended to find significant effects of income. This study also found some statistically significant effects of after-tax income on both margins, though these results were not robust to changes in the estimating specification.

V. CONCLUSION

This paper estimates the responsiveness of contributions to tax-preferred retirement accounts (including employer-based plans, IRAs, and self-employed plans) to recent tax changes, including the marginal tax rate reductions that were part of EGTRRA in

[43] Indeed, when Power and Rider (2002) include in their sample those who have some self-employment and some wage income, the semi-elasticity drops to –0.005.

2001 and JGTRRA in 2003, the introduction of the Saver's Credit, and the change in the EITC treatment of contributions to employer-based plans.

The results suggest that taxpayers respond to changes in the after-tax price of contributing on the extensive margin, but not on the intensive margin. The estimates imply that, by lowering the average after-tax price of contributing and increasing after-tax incomes, the tax changes between 1999 and 2005 increased the probability of contributing to a tax-preferred retirement account by about 0.23 to 0.71 percent, on average, with low income taxpayers being more likely to contribute and high income taxpayers being less likely to respond to the tax changes. There is also some suggestive evidence that changes in after-tax incomes led to increases in the probability of contributing to a tax-preferred account and in the amount contributed, though the results are not robust to changes in the estimating specification.

This study thus suggests that changes in tax rates may indeed have a role in increasing contributions to tax-preferred retirement accounts. Unfortunately, the tax data do not include the more complete information on a taxpayer's finances that would be required to estimate whether overall savings increased as a result of the tax changes. As a result, whether these increased contributions translate to increased savings rates is still an open question.

ACKNOWLEDGEMENTS

We would like to thank Pete Brady, Seth Giertz, and conference participants at the 101st Annual Conference on Taxation of the National Tax Association in 2008 and the Annual Meeting of the American Economic Association in 2009, the editors Bill Gentry and George Zodrow, and anonymous referees for helpful comments and suggestions. We also would like to thank Jon Bakija for providing his comprehensive tax calculator. The views expressed are those of the authors and are not necessarily those of the U.S. Department of the Treasury.

REFERENCES

Bakija, Jon., 2009. "Documentation for a Comprehensive Historical U.S. Federal and State Income Tax Calculator Program." Working paper. Williams College, Williamstown, MA.

Bakija, Jon, and Bradley T. Heim, 2011. "How Does Charitable Giving Respond to Incentives and Income? New Estimates from Panel Data." *National Tax Journal* 64 (2), 615–650.

Bassett, William F., Michael J. Fleming, and Anthony P. Rodrigues, 1998. "How Workers Use 401(k) Plans: The Participation, Contribution, and Withdrawal Decisions." *National Tax Journal* 51 (2), 263–289.

Brady, Peter, and Sarah Holden, 2006. "The U.S. Retirement Market, 2005." *Research Fundamentals* 15 (5), 1–12.

Burman, Leonard E., William G. Gale, and David Weiner, 2001. "The Taxation of Retirement Savings: Choosing Between Front-Loaded and Back-Loaded Options." *National Tax Journal* 54 (3), 689–702.

Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick, 2004. "For Better or For Worse: Default Effects and 401(k) Savings Behavior." In Wise, David (ed.), *Perspectives on the Economics of Aging*, 81–126. University of Chicago Press, Chicago, IL.

Choi, James, David Laibson, and Brigitte Madrian, 2004. "Plan Design and 401(k) Savings Outcomes." *National Tax Journal* 57 (2), 275–298.

Chernozhukov, Victor, and Christian Hansen, 2004. "The Effect of 401(k) Participation on the Wealth Distribution: An Instrumental Variable Quantile Regression Analysis." *The Review of Economics and Statistics* 86 (3), 735–751.

Clark, Robert, and Sylvester J. Schieber, 1998. "Factors Affecting Participation Rates and Contribution Levels in 401(k) Plans." In Mitchell, Olivia S., and Sylvester J. Schieber (eds.), *Living with Defined Contribution Pensions: Remaking Responsibility for Retirement*, 69–97. University of Pennsylvania Press, Philadelphia, PA.

Copeland, Craig, 2009. "Individual Account Retirement Plans: An Analysis of the 2007 Survey of Consumer Finances, with Market Adjustments to June 2009." EBRI Issue Brief No. 333. Employee Benefit Research Institute, Washington, DC.

Cunningham, Christopher R., and Gary V. Englehardt, 2002. "Federal Tax Policy, Employer Matching, and 401(k) Saving: Evidence from HRS W-2 Records." *National Tax Journal* 55 (3), 617–645.

Engen, Eric M., and William G. Gale, 2000. "The Effects of 401(k) Plans on Household Wealth: Differences Across Earnings Groups." NBER Working Paper 8032. National Bureau of Economic Research, Cambridge, MA.

Engen, Eric M., William G. Gale, and John Karl Scholz, 1996. "The Illusory Effects of Saving Incentives on Saving." *Journal of Economic Perspectives* 10 (4), 113–138.

Englehardt, Gary V., and Anil Kumar, 2007. "Employer Matching and 401(k) Saving: Evidence from the Health and Retirement Study." *Journal of Public Economics* 91 (10), 1920–1943.

Feenberg, Daniel, and Jonathan Skinner, 1989. "Sources of IRA Saving." In Summers, Lawrence H. (ed.), *Tax Policy and the Economy, Volume 3*, 25–46. MIT Press, Cambridge, MA.

Gruber, Jon, and Emmanuel Saez, 2002. "The Elasticity of Taxable Income: Evidence and Implications." *Journal of Public Economics* 84 (1), 1–32.

Heim, Bradley T., and Ithai Z. Lurie, 2010. "Taxes, Income, and Retirement Savings: Differences by Permanent and Transitory Income." Working Paper. Indiana University, Bloomington, IN.

Holden, Sarah, Peter Brady, and Michael Hadley, 2006. "401(k) Plans: A 25-Year Retrospective." *Research Perspective* 12 (2), 1–40.

Holden, Sarah, and Jack VanDerhei, 2001. "Contribution Behavior of 401(k) Plan Participants." EBRI Issue Brief Number 238. Employment Benefit Research Institute, Washington, DC.

Joulfaian, David, and David Richardson, 2001. "Who Takes Advantage of Tax-preferred Retirement Saving Programs? Evidence from the Federal Income Tax Data." *National Tax Journal* 54 (3), 669–688.

Koenig, Gary, and Robert Harvey, 2005. "Utilization of the Savers Credit: An Analysis of the First Year." *National Tax Journal* 58 (4), 787–806.

Kusko, Andrea L., James M. Poterba, and David W. Wilcox, 1998. "Employee Decisions With Respect to 401(k) Plans." In Mitchell, Olivia, and Sylvester J. Schieber (eds.), *Living with Defined Contribution Pensions: Remaking Responsibility for Retirement*, 98–112. University of Philadelphia Press, Philadelphia, PA.

Long, James E., 1990. "Marginal Tax Rates and IRA Contributions." *National Tax Journal* 43 (2), 143–153.

Madrian, Brigitte C., and Dennis F. Shea, 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics* 116 (4), 1149–1187.

Munnell, Alicia H., Annika Sunden, and Catherine Taylor, 2003. "What Determines 401(k) Participation and Contributions?" *Social Security Bulletin* 64 (3), 64–75.

Papke, Leslie E., 1995. "Participation in and Contributions to 401(k) Pension Plans: Evidence from Plan Data." *Journal of Human Resources* 30 (2), 311–325.

Poterba, James M., Steven F. Venti, and David A. Wise, 1996. "How Retirement Saving Programs Increase Savings." *Journal of Economic Perspectives* 10 (4), 91–112.

Power, Laura, and Mark Rider, 2002. "The Effect of Tax-Based Savings Incentives on the Self-Employed." *Journal of Public Economics* 85 (1), 33–52.

Springstead, Glenn R., and Theresa M. Wilson, 2000. "Participation in Voluntary Individual Savings Accounts: An Analysis of IRAs, 401(k)s, and the TSP." *Social Security Bulletin* 63 (1), 34–39.

Summers, Lawrence H., 1981. "Capital Taxation and Accumulation in a Life Cycle Growth Model." *American Economic Review* 71 (4), 533–544.

Summers, Lawrence H., 1984. "The After-Tax Rate of Return Affects Private Savings." *American Economic Review* 74 (2), 249–253.

Venti, Steven F., and David A., Wise, 1988. "The Determinants of IRA Contributions and the Effect of Limit Changes." In Bodie, Zvi, John B. Shoven, and David A. Wise (eds.), *Pensions in the U.S. Economy*, 9–52. University of Chicago Press, Chicago, IL.

Weber, Michael, and Victoria L. Bryant, 2005. "The 1999 Individual Income Tax Return Edited Panel." In *2005 Proceedings of the American Statistical Association*, 159–166. American Statistical Association, Alexandria, VA.