SAVINGS INCENTIVES FOR LOW- AND MODERATE-INCOME FAMILIES IN THE UNITED STATES: WHY IS THE SAVER'S CREDIT NOT MORE EFFECTIVE?

Esther Duflo
MIT

William Gale
Brookings Institution and Retirement Security Project

Jeffrey Liebman
Harvard University Kennedy School of Government

Peter Orszag
Brookings Institution and Retirement Security Project

Emmanuel Saez
University of California, Berkeley

Abstract
This paper uses data from the largest tax preparer in the United States to estimate the impact of the "saver's credit," a US federal program providing financial incentives to encourage retirement savings, on the decision to contribute to an IRA. It finds significant, but very modest, effects. This is contrasted with results from a field experiment showing much larger impacts of clearly presented matching incentives. Various explanations are discussed for why the saver's credit is not more effective. (JEL: H00, H31, C93, D14)

1. Introduction

A significant share of low- and middle-income American families appear to be saving little, either for retirement or for any other purpose. Families with incomes below $40,000 have low rates of coverage under employer-provided pensions, are extremely unlikely to contribute to individual retirement accounts (IRAs), and in 2001 had median net financial wealth outside of retirement accounts of just $2,200.[1]

Acknowledgments: We thank Brigitte Madrian and numerous seminar participants for comments and discussion. E. Duflo, J. Liebman, and E. Saez are research associates at NBER.

E-mail addresses: Duflo: eduflo@mit.edu; Gale: wgale@brookings.edu; Liebman: jeffrey_liebman@harvard.edu; Orszag: porszag@brookings.edu; Saez: saez@econ.berkeley.edu

1. See Burman et al. (2004) for data on defined contribution pension coverage rates by income group. Calculations from the 2001 Survey of Consumer Finances (SCF) imply that only one quarter of households with income below $40,000 have defined benefit coverage. Burman et al. (Appendix Table 6) report that among households with cash income below $40,000, less than 2.1% contributed to either a Roth or traditional IRA in 2004. Median net financial wealth in the text is calculated from the 2001 SCF.

Journal of the European Economic Association April–May 2007 5(2–3):647–661

Researchers and policymakers have long considered ways to raise saving among these families. Financial incentives, in the forms of tax deductions or matching of employees contributions in 401(k) plans have been in existence for some time but have not reached many low- and moderate-income families, in part because the value of the tax deduction is lower for these families and in part because they are less likely to be working in a firm offering such plan.

The Saver's Credit is an attempt to fill this gap. Enacted in 2001, the credit provides a federal income tax reduction of up to 50% of funds contributed to a 401(k) or IRA by qualified filers. This paper uses a rich data set (covering the clients of the largest tax preparer in the country, H&R Block) to provide estimates of the impact of the Saver's Credit on contribution to an IRA. It shows that, despite its generosity, the Saver's Credit has had a very small impact on this decision.

This is in contrast with experimental evidence showing a substantial impact of clearly presented matching incentives, evaluated by the authors in a randomized experiment conducted by H&R Block (Duflo et al. 2006). This paper argues that the small impact of the credit may not be due to lack of responsiveness to incentives but that the credit as enacted may not accomplish the full potential of matching, both due to some essential features (low threshold for eligibility and nonrefundability) and to design features that may have seemed innocuous at the outset but turn out to have large unintended consequences on the impact of the program. This study, therefore, provides some useful lessons about how to reform the Saver's Credit or design similar programs in other countries.

2. Evaluating the Impact of the Saver's Credit Program

2.1. The Saver's Credit Program

The Saver's Credit provides, over certain income ranges, matching incentives for low- and middle-income tax filers. It was first implemented in tax year 2002 (for tax returns filed in 2003) and is scheduled to expire after 2006 (tax returns filed in 2007 will be the last to benefit from the program unless it is extended). The Saver's Credit is a nonrefundable tax credit on the first $2,000 (for each spouse) contributed to IRAs (Roth and traditional) or voluntary pension plans (Keogh, 401(k), 403(b), SIMPLE IRA, etc.).[2] As shown in Table 1, the credit rate decreases with adjusted gross incomes (AGI) and is zero above an AGI threshold that depends on filing status. The credit rate is 50% at the bottom, 20% in a narrow AGI band, and then 10% for a relatively broad range.

Note that a credit at rate t is economically equivalent to a match rate on the contribution of $t/(1-t)$. For example, a tax filer facing the 50% credit rate and

2. Those contributions are netted of any withdrawals made during the last three years to limit gaming possibilities.

TABLE 1. Saver's credit parameters.

Credit rate, % t (1)	Equivalent Match rate, % $t/(1-t)$ (2)	Married filing jointly AGI range, $ (3)	Head of household AGI range, $ (4)	Single and others AGI range, $ (5)
50	100	0–30,000	0–22,500	0–15,000
20	25	30,001–32,500	22,501–24,375	15,001–16,250
10	11	32,501–50,000	24,376–37,500	16,251–25,000
0	0	50,001+	37,501+	25,001+

Notes: The Saver's Credit is a nonrefundable federal income tax credit proportional to the sum of all elective retirement contributions (all IRAs, 401(k)s, etc.) up to $2,000 of contributions ($2,000 for each spouse for married taxpayers). Full-time students, individuals claimed as dependents by other taxpayers, and individuals aged under 18 are not eligible. Withdrawals from IRAs and 401(k)s (within the last three years) are netted out of annual retirement contributions to estimate eligible contributions. As shown in columns 3 to 5, the credit rate varies by AGI range and marital status. The bracket length for heads of household and singles are 75% and 50% of the bracket length for married, respectively. A credit rate of t (column 1) is equivalent to a match rate of $t/(1-t)$ (column 2). AGI used to compute the credit rate t is net of most retirement contributions with the exception of Roth IRAs and is therefore endogenous. The Savers' Credit is nonrefundable and can be used only to offset tax liability net of other non-refundable credits. As a result, the Saver's Credit is independent of the EITC but interacts with the partially refundable child tax credit. A taxpayer benefits from the Saver's Credit if its income tax liability net of non-refundable credits (excluding the child tax credit) and net of (full child tax credit − child tax credit, which can be recovered with additional refundable children tax credit) is positive. Taxpayers might have positive Saver's Credit and not benefit from it because the Saver's Credit might crowd out the nonrefundable child tax credit one for one.

See IRS Form 8880 and IRS Publication 590 for more details.

contributing $1,000 would receive a $500 tax credit, so that her out-of-pocket cost for a $1,000 contribution is only $500, which is effectively a 100% match rate. Therefore, the Saver's Credit generates effective match rates of 100%, 25%, and 11%. The higher rate, in particular, thus provides strong incentives (as large as those provided in many employer provided 401(k)).

Because the Saver's Credit is nonrefundable, however, many low-income tax filers who would qualify for the highest credit rate on the basis of their AGI benefit from the credit only to a very limited extent, or even not at all, because they have little or no tax liability due to standard or itemized deductions, personal exemptions, and use of other non-refundable credits (in particular, the Child Tax Credit).[3] The data allow us to define precisely whether a tax filer benefits or could potentially benefit from the Saver's Credit. In what follows, we call those filers who are or could be benefiting "eligible" filers. Those who could not benefit are called "ineligible."

At first glance, comparing retirement savings contributions on each side of the AGI boundary points defining the credit rate brackets might seem like a promising way to analyze the effects of the credit (or match equivalent) rates, because there is a discontinuity in the match rate at that point. AGI, however, is net of retirement contributions (with the exception of the Roth IRA). As a result, even if gross income (defined here as AGI with all tax-deductible retirement

3. The Saver's Credit is determined before refundable credits. Therefore, the Earned Income Tax Credit and the refundable portion of the Child Tax Credit do not reduce the Saver's Credit.

contributions added back) is exogenous, AGI itself is not, because rational tax filers just above the boundary have incentives to contribute to tax-deductible IRAs in order to fall below the boundary and benefit from having a higher credit rate that would apply to all of their contributions.

An implication, however, is that we should expect an abnormally large number of taxpayers bunching at AGIs just below the boundary point. Symmetrically, we should expect to see few taxpayers just above the boundary point, because increasing contributions even a little would increase the credit by a discrete percentage (and the higher credit rate would apply to the entire contribution, not just the marginal contribution). Finding bunching in the share of filers who contribute among filers just to the left of the boundary point (and a lower contribution rate to the right of the boundary) would thus constitute convincing evidence that individuals understand and respond to the incentives provided by the Saver's Credit.

In practice, tax filers may not be aware of the precise dollar amount of their annual incomes and pension contributions.[4] Furthermore, it might be difficult for tax filers to tailor their pension contributions (often specified as a percentage of salary) during the year so that their AGI falls precisely below the boundary points. However, they can adjust their AGI by purchasing an IRA after having completed their taxes, when their exact AGI amount is revealed, and they can estimate whether a higher contribution would make them eligible for a higher rate.

In particular, H&R Block offers an "express-IRA" (X-IRA) product. The X-IRA allows the taxpayer to open an account and make IRA contributions at the time of tax preparation and to fund those contributions with part or all of their federal income tax refunds or from other sources. In effect, the X-IRA allows the client to "split" their anticipated refund between contributions to a retirement account and other uses. This instrument should allow a taxpayer to fine-tune their AGI, so that if taxpayers were well informed and responsive to the credit, we should see bunching to the left of the AGI thresholds, and "holes" to the right.

2.2. Graphical Evidence around the Saver's Credit Cliffs

To investigate responses to the Saver's Credit, we use national H&R Block retail tax return data for tax season 2005. The data consist of a 100% sample of tax returns with positive X-IRA contributions (about 180,000 returns) and a 9%

4. Indeed, annual salary contracts do not run in general from January 1 to December 31. Actual payment dates for work in December might fall in January. There are a number of additional factors, such as Social Security and Medicare taxes, and pretax parking and health care contributions, which make it difficult to evaluate precisely taxable wages and salaries before the W2 form arrives in January. At that time, it is too late to modify employer pension plan contributions.

random sample selected based on nine two-digit endings of the Social Security number of the primary taxpayer (about 1,400,000 returns). In what follows, the data are weighted by the inverse of the sampling probabilities. Because the minimum X-IRA contribution is $300, we focus on the sub-sample of taxpayers with earnings (wages and salary, alimony, and self-employment income) above $300. We also exclude tax filers aged less than 18 (as of January 1, 2005) or claimed on someone else's return because such taxpayers are ineligible to claim the Saver's Credit. We multiply the incomes of single and married filing separately tax filers by 2 and incomes of heads of households by 1.3333 so that the boundary points, where the Saver's Credit rate changes, are aligned for all types of tax filers. We call these income figures "normalized AGI" (see Table 2).

Figure 1 plots the percentage of tax units making an X-IRA contribution by $250 bands of normalized AGI, for normalized AGI ranging between $20,001 to $60,000. The sample is further limited to "eligible" tax filers defined as those for whom an X-IRA contribution (real or potential) would trigger a positive Saver's Credit not offset by reductions in other credits.[5] The three Saver's Credit cliffs, where the implicit match falls from 100% to 25%, 25% to 11%, and 11% to 0%, respectively, are indicated by vertical lines. The figure shows clear evidence of a behavioral response of X-IRA contributions to the Saver's Credit: There are three visible spikes in the fraction of taxpayers contributing just below each of the three cliffs.

To further test the hypothesis that the spike is due to the Saver's Credit, we compare "eligible" tax units from Figure 1 to "inelegible" tax units whose X-IRA contribution would not trigger a Saver's Credit because they have no income tax liability (net of other credits) to offset. Ineligible tax units have more children on average, are less likely to be single, and more likely to itemize deductions, each of which reduces tax liability. Figure 2 shows the likelihood of being an X-IRA contributor for eligibles (dark shading) and ineligibles (light shading) around the first two cliffs of the Saver's Credit.[6] The figure appears consistent with the tax explanation. As we saw in Figure 1, the spikes are clearly visible for eligible tax units, and the contribution rates are higher in the 100% match bracket than in the 25% and 11% match brackets. In contrast, there are no spikes for the ineligibles and the contribution rate is about the same across the 100%, 25%, and 11% match brackets.

Figures 1 and 2 also show that, while there is a drop in the fraction of contributors to the right of the threshold, a relatively large fraction of taxpayers is contributing in those income ranges. In principle, many of the filers whose income

5. Because of the interaction with the Child Tax Credit, about 3% of eligible tax filers do not actually benefit from the Saver's Credit even if their return shows a positive Saver's Credit (because the Saver's Credit is offset one for one by a reduction in the Child Tax Credit).

6. Almost all taxpayers around the third cliff have positive tax liability and hence the group of ineligibles is too small in that case.

TABLE 2. The effects of the Saver's Credit on X-IRA behavior.

	Levels					Differences		
Match rate	100%	25%	11%	11%	0%	100% vs. 25%	25% vs. 11%	11% vs. 0%
Credit rate	50%	20%	10%	10%	0%	50% vs. 20%	20% vs. 10%	10% vs. 0%
Normalized AGI	$27.5K–$30K	$30K–$32.5K	$32.5K–$35K	$45K–$50K	$50K–$55K	$27.5K–$32.5K	$30K–$35K	$45K–$55K
						(1) − (2)	(2) − (3)	(4) − (5)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Savers' Credit eligible taxpayers								
X-IRA take-up rate (percent)	3.28 (0.083)	1.87 (0.063)	1.55 (0.058)	1.31 (0.043)	0.89 (0.040)	1.41 (0.104)	0.32 (0.086)	0.42 (0.058)
X-IRA average contributions (unconditional)	$19.2 (0.67)	$9.8 (0.48)	$7.9 (0.45)	$9.0 (0.47)	$6.4 (0.45)	$9.5 (0.82)	$1.8 (0.66)	$2.6 (0.65)
X-IRA average contributions (conditional on take-up)	$587 (4.8)	$522 (6.2)	$513 (7.0)	$690 (8.9)	$722 (12.7)	$65 (7.9)	$9 (9.4)	−$32 (15.1)
Sample size	45,747	45,721	44,508	71,052	61,388	91,468	90,229	132,440
B. Savers' Credit ineligible taxpayers (no tax liability)								
X-IRA take-up rate (percent)	2.09 (0.110)	2.00 (0.145)	1.86 (1.640)			0.09 (0.182)	0.14 (0.220)	
X-IRA average contributions (unconditional)	$9.5 (0.64)	$9.4 (0.90)	$9.0 (1.05)			$0.1 (1.09)	$0.4 (1.38)	

X-IRA average contributions (conditional on take-up)	$455	$472	$485	−$17	−$13
	(6.3)	(9.7)	(12.0)	(11.2)	(15.4)
Sample size	17,041	9,301	6,831	26,342	16,132
C. Difference between eligible and ineligible taxpayers (A−B)				**Difference-in-differences**	
X-IRA take-up rate (percent)	1.19	−0.14	−0.31	1.33	0.18
	(0.148)	(0.156)	(0.164)	(0.220)	(0.228)
X-IRA average contributions (percent)	$9.7	$0.3	−$1.1	$9.4	$1.4
	(1.13)	(1.16)	(1.23)	(1.67)	(1.70)
X-IRA average contributions (percent)	$132	$51	$28	$81	$23
	(10.0)	(14.2)	(17.3)	(17.2)	(22.3)
Sample size	62,788	55,022	51,339	117,810	106,361

Notes: This table shows X-IRA take-up rates, average contributions (unconditional including zeros and conditional on take-up, excluding zeros) for five AGI groups as well as differences and difference-in-differences. Standard errors are reported in parentheses. The sample consists of 100% of H&R Block retail returns with X-IRAs, and a 9% random sample of all other returns. All results are re-weighted. Taxpayers aged below 18 or who are claimed as dependents on someone else's return are excluded (do not qualify for the Savers' Credit). Taxpayers with less than $300 in earnings are excluded (cannot contribute to an X-IRA with minimum $300 contribution). First, groups are defined relative to normalized AGI (Normalized AGI = 100% of AGI for married taxpayers, 133.33% for heads of household, and 200% for singles and others). Columns 1 to 5 display the levels for those five groups and columns 6 to 8 the differences across groups. Second, groups are defined by Savers' Credit eligibility (panel (A) versus panel (B)).

Group A are the eligible, defined as taxpayers whose X-IRA contributions (actual or potential) would benefit from the savers' credit. Group B are the ineligible, defined as taxpayers whose X-IRA contributions (actual or potential) would not benefit from the savers' credit because their tax liability net of other non refundable tax credits (excluding the child tax credit) and net of (full tax credit less potentially refundable additional tax credit) is zero or negative. There is no group B for columns 4 and 5 because almost all taxpayers in that AGI range have positive tax liability. Panel (C) displays the difference between group A and group B. The bottom right panel displays the difference-in-differences.

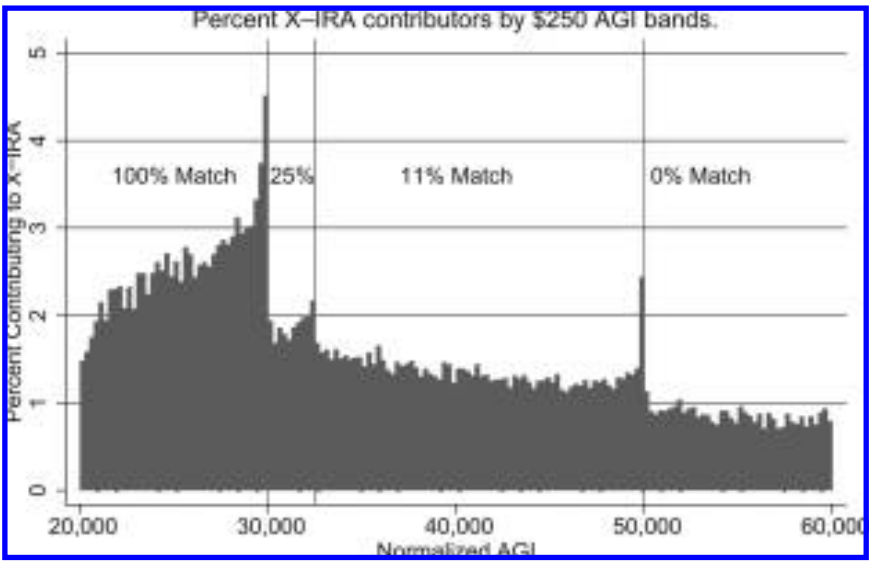

FIGURE 1. Percent X-IRA contributors by $250 AGI bands.

is just above the threshold could significantly increase their tax refund by contributing more to the X-IRA and hence possibly increase both their X-IRA savings and their after-tax current income net of X-IRA contributions. In particular, taxpayers with 401(k) contributions just above the 100% match cliff have very strong

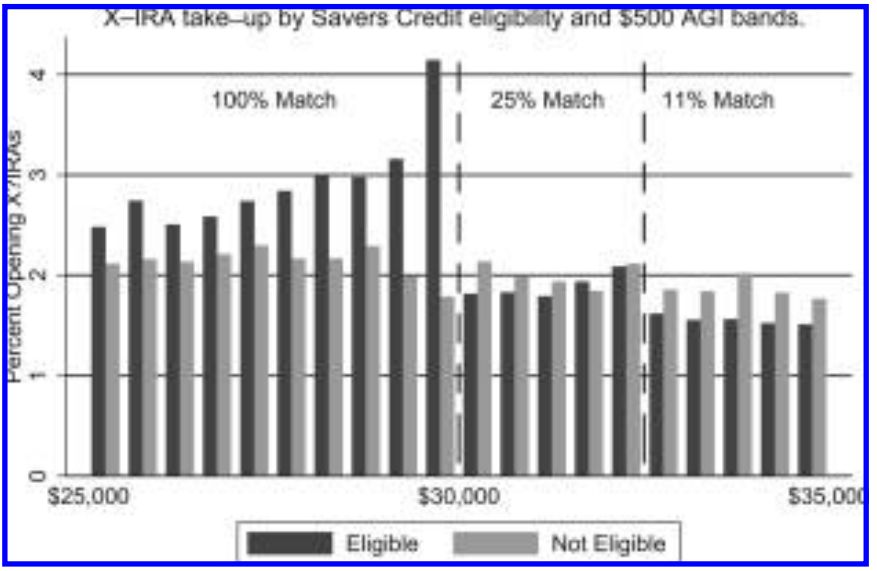

FIGURE 2. X-IRA take-up by Savers Credit eligibility and $500 AGI bands.

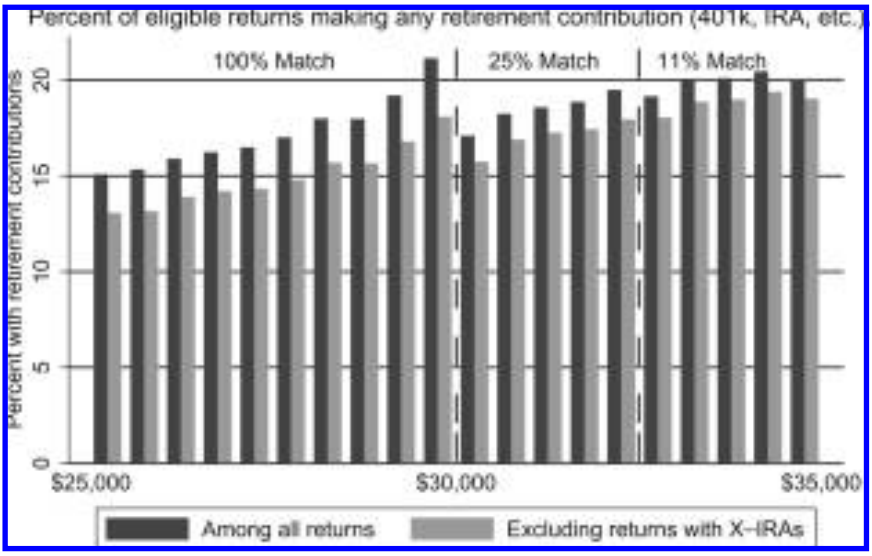

FIGURE 3. Percent of eligible returns making any retirement contribution (401k, IRA, etc.).

incentives to make potentially modest X-IRA contributions in order to increase the credit rate on their preexisting retirement contributions. In some of those cases, an X-IRA contribution can even increase the tax refund net of the contribution. However, the X-IRA take-up rate among taxpayers facing this situation is only around 6.2%, showing that the vast majority of those fortunate taxpayers actually fail to exploit this opportunity to be paid for saving.

Figure 3 plots the percentage of returns with positive retirement contributions (all IRA types, and other retirement contributions such as 401(k)s, etc.) among all those potentially eligible for the Saver's Credit and then excluding those who made an X-IRA contribution. The proportion with any retirement contributions (such as 401(k)s, other IRAs, etc.) is much higher than the fraction with X-IRAs (from Figure 1) and increases from about 15% to about 20% as (normalized) AGI rises from $25,000 to $35,000. There is a moderate spike visible at the $30,000 cliff. However, most of this spike goes away when X-IRA contributors are excluded, and there is no strong evidence that the 100% match rate induces more tax filers to make retirement contributions. As we discussed, it is not surprising that tax filers cannot fine tune the Saver's Credit optimization before tax preparation.[7]

7. Comparing 401(k) and IRA participation rates among low and moderate income earners across years (before and after the Saver's Credit was implemented) could provide suggestive evidence on the effects of the Saver's Credit on overall retirement contributions. The evidence would only be suggestive because it could be confounded by nonsavers' credit-related time trends.

This suggests that the responses may be even lower for taxpayers who do not work with a tax preparer that can offer an opportunity to split the refund.

The lack of utilization of the credit is not confined to the extensive margin; most of those who contribute are far from contributing to the extent that will maximize their refund. Define the fraction of credit used by an X-IRA contributor as the ratio of the actual X-IRA contribution to the amount of potential contribution (before the X-IRA contribution) that would exhaust the Saver's Credit. This maximum amount is attained either when tax liability is exhausted or when total retirement contributions reach $2,000 (per spouse in case of married filers). Given the significant effective match rate, we would expect to observe a significant fraction of tax filers contributing up to this maximum and having a fraction of credit used equal to 100% if households were optimizing their contributions. Less than 3% of those taxpayers are exactly maximizing their Saver's Credit. The vast majority of tax filers actually contribute less than the maximum amount eligible for the credit. This is another important difference between the Saver's Credit and our experiment, where about 50% of those who contributed were bunching at the $1,000 contribution level, which maximized their match amount.

2.3. Estimates of the Saver's Credit Effects on X-IRA Take-up and Amounts

The graphical analysis from Figures 1 and 2 suggests that simple comparisons could be made to estimate the effects of the changes in the match rates generated by the Saver's Credit using the fact that both the AGI and the amount of other credits determine whether a particular tax filer may benefit from the Saver's Credit.

We consider five normalized AGI groups of interest: (1) AGI between $27,501 and $30,000 (entitled to a 100% match rate), (2) AGI between $30,001 and $32,500 (entitled to a 25% match rate), (3) AGI between $32,501 and $35,000 (entitled to an 11% match rate), (4) AGI between $45,001 and $50,000 (entitled to an 11% match rate), and (5) AGI between $50,001 and $55,000 (entitled to a 0% match rate). We consider relatively narrow AGI groups around the cliffs in order to compare tax filers with relatively similar income levels. Columns 1 to 5 in panel (A) of Table 2 display the X-IRA take-up rate, average X-IRA contributions, and conditional X-IRA contributions for those five groups. Finally, columns 6 to 8 display the differences between the 100% and the 25% groups, between the 25% and the 11% groups, and between the 11% and 0% groups, respectively. All standard errors are obtained from a weighted OLS regression.

These differences are very precisely estimated, but remarkably small. Going from a 25% match rate to a 100% match rate increases take up by 1.4 percentage points, from 1.9% to 3.3%. Going from 11% to 25% match rates and from 0% to 11% match rates increases take-up much more modestly,

by 0.3 and 0.4 percentage points respectively. We also find similarly small but significant increases in amounts contributed (unconditional). Amounts contributed conditional on contributions increase modestly (by $65) around the 25% to 100% match rate jump but not around the other two smaller match rate jumps.[8]

Furthermore, the simple difference estimates we obtain in Table 2 are likely to overestimate the effect of the match rate due to the artificial "piling up" effect below the cliff (those who manipulate their AGI to be low the cliff are probably more likely to save, and to save more). If we exclude boundary tax filers with AGI between $29,501 and $30,500, the difference falls further from 1.4 to 1.0 (0.11).

Another reason why the comparisons from panel (A) may be an overestimate of the effect of the Saver's Credit is that the differences in take-up rates might not be due to the causal effects of differences in the Saver's Credit rate but rather the fact that tax filers with different AGIs also have different propensities to save. A simple but admittedly imperfect way to control for this is to consider the same AGI groups for tax filers ineligible for the Saver's Credit because they lack income tax liability before refundable credits. Such an analysis of ineligible control groups is shown in panel (B) of Table 2 for the bottom three groups.[9] For those groups, there are much smaller (and insignificant) differences in X-IRA take-up rates and contribution amounts, suggesting that the differences from panel A were indeed caused by the Saver's Credit.

Panel (C) displays the difference between panel (A) (the eligible) and panel (B) (the ineligible). If those two groups were identical in their savings tastes, such differences could also potentially capture the Saver's Credit effect. Those differences are even smaller than the results in panel (A). In column 1, the difference in contribution rates is just 1.2 percentage points for the 100% match rate and very close to zero and insignificant for the 11% or 25% match groups. These negligible effects could be due either to the fact that the Saver's Credit has almost no impact (and the differences in panel (A) were biased upward), or that eligible and non-eligible are systematically different. Indeed, eligible tax units tend to have fewer dependent children and are more likely to be single, and hence perhaps have a lower taste for savings. Differences in amounts are equally small and barely significant.

Finally, columns 6 and 7 in panel (C) display the differences-in-differences estimates. These are unbiased estimates of the differential impact of the 100% rate versus 25% rate and the 25% versus the 11% match rate under the assumption that, absent the Saver's Credit, the difference in X-IRA behavior between eligible and

8. The X-IRA amounts contributed are implicitly inclusive of the equivalent match rate because the Saver's Credit is formally a credit rebate instead of a match.

9. Almost all tax filers around the third cliff of the Saver's Credit (above which the effective credit rate falls to zero) have positive tax liability and are therefore eligible, making it impossible to obtain control groups large enough for statistical analysis.

ineligible filers would be the same across the two AGI groups. The difference-in-differences estimates for the 100% versus 25% match comparison are quite close to the simple difference estimates from panel (A). The difference-in-differences estimates for the 25% versus 11% match comparison also remains of similar magnitude but becomes insignificant, again suggesting that the causal effect of the Saver's Credit is at best very small.

3. Contrasting the Saver's Credit to a Randomized Experiment on Savings Incentives

All these results suggest statistically significant, but modest effects of the Saver's Credit on take-up and even more modest effects on amounts contributed. The fact that these effects are present suggest that at least some of the taxpayers are aware of the programs. This modest effect could be attributed to a low elasticity of taxpayers with respect to incentives.

However, the results of an experiment we conducted in collaboration with H&R Block in St. Louis in the 2005 tax season (Duflo et al., 2006) suggests this is not the entire explanation. The experiment was run in 60 H&R Block tax preparation offices in the St. Louis metropolitan area from March 5 to April 5, 2005. Like this analysis of the saver's credit, the experiment was also built around the Express IRA (X-IRA) product.

Each client preparing a tax return in one of the 60 offices during the period was randomly assigned to one of three match rates for X-IRA contributions: zero (the control group), 20%, or 50%. In the two "match" groups, any funds deposited by the taxpayer in the X-IRA were matched by H&R Block up to $1,000, a limit that applied separately for each spouse for married tax filers. Each client, including those in the control group, received a waiver of the $15 set-up fee for opening an X-IRA. The minimum X-IRA contribution of $300 was maintained in all groups.

The results from this experiment show that match rates can have large effects on IRA participation and contributions. Take-up rates were 3%, 8%, and 14%, respectively, for the control group, the 20% match group, and the 50% match group. Conditional on take-up, average contribution levels (excluding the match) were $765, $1,100, and $1,110, respectively. With the match included, average IRA deposits were $765, $1,280, and $1,590 respectively, among contributors. Average IRA deposits (including noncontributors, excluding the match) with the 20% and 50% matches were 4 and 7 times higher, respectively, than with no match.

These results indicate considerably higher elasticities than for the saver's credit. Even the differences we found in the saver's credit analysis around the largest jump (from the 25% to 100% match rates) are much smaller than the difference we found between our 20% and 50% experimental match rates, where

we found a 6.3 percentage point difference in take-up and a $310 difference in amounts contributed (conditional on take-up and inclusive of match). Even though the differences in match rates are smaller in our experiment, and even though the effects of the Saver's Credit we estimate here are upper bounds, the effect of a 50% match rate in our experiment on take-up and contribution rates are respectively 4.5 and 4.8 times larger than what we estimate for the Saver's Credit. In the next section, we consider a number of explanations for this result.

4. Why is the Saver's Credit not More Effective?

A first explanation could be that the population on which we estimated the impact of the Saver's Credit is different from the population on which we estimated the impact of the experiment; for example, those who file early in the season may be more likely to be liquidity constrained than those who file in the mid-season as in the experiment. However, if we restrict our analysis of the Saver's Credit to those filing between March 5 and April 5, and compare take-up in the 100% match rate versus 25% match rate groups, we still obtain an increase in take-up of only 1.33 (s.e. 0.29) percentage points (relative to 1.29 percentage points for all filers), leaving the puzzle almost entirely intact. This suggests that this is not the entire explanation.

A second possible explanation is that the experiments presented the incentives in the form of a match, rather than a credit. Although a match is economically equivalent to a credit, this may not be evident to the individuals, who may find that a 50% match looks "larger" than an equivalent 33% credit. Previous experiments have indeed shown that, in the context of charitable giving, the match presentation generates higher take-up than the credit presentation.[10] This hypothesis will be tested in future experiments by randomly assigning taxpayers in a group of offices to one of three possible groups: control, 50% match, or 33% credit. To make the offer economically equivalent, the minimum contribution in the 33% credit group will be $450 (and $300 in the 50% match group).

A third possible explanation is that the saver's credit has complex rules (due to the combination of the fact that it is not refundable and it has low exemption thresholds). This has two consequences: First, taxpayers may not realize, even at the time of tax preparation, that they are eligible for the credit. Even tax preparers may not be completely on top of the rules, and until 2005 the H&R Block tax preparation software did not have a way to systematically flag the taxpayer's option. H&R Block has now changed their Saver's Credit screen to make it more explicit to taxpayers that they qualify for the credit if they open an X-IRA. Second,

10. See, for example, Davis, and Millner (2005), Davis, Millner, and Reilly (2005), and Eckel and Grossman (2003, 2005).

this complication makes it difficult to plan ahead: Because taxpayers do not know whether or not they would be eligible, they have little rationale to set money aside for contributing to an IRA at the time of tax preparation, should they find out they are eligible. Thus, even if it appears at the time of tax preparation that they could benefit from the credit, they may not be willing or able to do it. Focus groups we conducted after the 2005 experiments suggested that a number of H&R Block clients who decided not to take advantage of the match had already "spent" their refund (i.e., borrowed against it) or had made plans on how to use it. They had no slack left for investing in an X-IRA. To further test whether the ability to plan ahead leads to an increase in the take up of a match offer, an experiment can be set up where taxpayers can be warned in advance that a match will be available.

5. Conclusion: Reforming the Saver's Credit to Increase its Effectiveness

These results suggests a number of avenues for reforming the current Saver's Credit. First, the credit could be replaced by an equivalent matching incentive. Although it is more complex for the government to match retirement contributions than to provide a credit, our results show that this is has important implications for its success. Second, the credit could be made refundable and the complex graduated structure should be replaced by a single salient rate (such as a 50% match) applying to retirement contributions up to some level (such as $2,000 of contributions per tax filer); this maximum amount (rather than the credit rate) could be linearly phased out above a specified AGI threshold. Those changes would make the program easier to understand and more predictable. Third, the IRS could offer split refund options to allow tax filers to direct tax refunds into retirement savings accounts so that tax filers can take advantage of the annual lump-sum nature of tax refunds to save. Finally, to reduce the uncertainty and allow taxpayers to plan ahead, the eligibility could be based on the previous year's AGI.

It is worth pointing out, however, than even with these reforms the saver's credit is unlikely to cause a dramatic increase in retirement contributions. Even well-presented matches in 2005 produced relatively modest increase in IRA contribution, especially compared to default enrollment.

References

Burman, Leonard E., William G. Gale, Matthew Hall, and Peter R. Orszag (2004). "Distributional Effects of Defined Contribution Plans and Individual Retirement Arrangements." *National Tax Journal*, 57(3), 671–701.

Davis, Douglas D., and Edward L. Millner (2005). "Rebates, Matches and Consumer Behavior." *Southern Economic Journal*, 72, 410–422.

Davis, Douglas D., Edward L. Millner, and Robert J. Reilly (2005). "Subsidy Schemes and Charitable Contributions: A Closer Look." *Experimental Economics*, 8, 85–106.

Duflo, Esther, William Gale, Jeffrey Liebman, Peter Orszag, and Emmanuel Saez (2006). "Saving Incentives for Low- and Middle-Income Families: Evidence from a Field Experiment with H&R Block." *Quarterly Journal of Economics*, 121(4), 1311–1346.

Eckel, Catherine, and Philip J. Grossman (2003). "Rebate versus Matching: Does How We Subsidize Charitable Contributions Matter?" *Journal of Public Economics*, 87(3–4), 681–701.

Eckel, Catherine, and Philip J. Grossman (2005). "Subsidizing Charitable Contributions: A Field Test Comparing Matching and Rebate Subsidies." Working paper, Virginia Tech.