EDUCATIONAL SAVING INCENTIVES FOR LOW-INCOME FAMILIES: EXPERIMENTAL EVIDENCE FROM THE MICHIGAN SEED PROGRAM

Gary V. Engelhardt and *Alissa L. Dubnicki,* Syracuse University
Ellen L. Marks and *Bryan B. Rhodes,* RTI International

INTRODUCTION

There has been long-standing interest among academics and policy makers in the ability of public policies to increase saving for low-income American families. Unfortunately, traditional tax subsidies to saving, such as 401(k)s, IRAs, Coverdells, and 529 plans, that allow tax-deductible contributions and accrual at the pre-tax rate of return, provide very limited incentives to save for such families who face low marginal tax rates. Efforts to promote saving are further hampered by this population's relative unfamiliarity with financial institutions. In fact, some estimates suggest that more than 25 percent of lower-income families are "unbanked."

Recent initiatives to stimulate saving among low-income families have been structured either to provide matching contributions to raise the net return to saving, as in the Retirement Saver's Credit, or a mix of matching contributions, account provision as a means of financial institution access, and case management services, as in Individual Development Accounts (IDA). However, empirical findings on the impact of these initiatives thus far have been mixed. Duflo et al. (2006) found that matching had a sizeable impact on IRA contributions in a field experiment run through H&R Block offices in St. Louis, MO. Mills et al. (2008) found that IDAs raised homeownership substantially after four years in a field experiment in Tulsa, OK. However, subsequently, Grinstein-Weiss et al. (2011) found that these effects dissipated almost fully after 10 years.

This paper adds to the small, but growing, literature that uses field experiments to analyze saving behavior (e.g., Duflo & Saez, 2003; Ashraf et al., 2006; Duflo et al., 2006; Mills et al., 2008; Saez, 2009; Grinstein-Weiss et al., 2011). It reports evidence from the Michigan Saving for Education, Entrepreneurship and Downpayment (SEED) Program, a group randomized trial on the effects of matching contributions, account provision, and case management on saving for children's post-secondary education among low-income families.

There are three primary findings. First, four years after the intervention, or when the children were roughly 8 years old, 67 percent of treatment-group families had accepted an offer of a $1,000 deposit and opened a 529 plan. In comparison, about 8 percent of the control group had independently opened a 529 plan. Therefore, the SEED initiative raised 529 account ownership substantially, by 59 percentage points. Second, 22 percent of treatment-group families contributed their own funds. Among these, the average family contribution was $440. Finally, for every dollar of 529 plan saving, total saving for college rose by 25-45 cents. Therefore, 529 plan saving resulted in 55-75 percent crowd-out of other saving.

BACKGROUND ON THE MICHIGAN SEED EXPERIMENT

SEED is the acronym for a national initiative known as Saving for Education, Entrepreneurship and Downpayment (SEED), funded by the Ford Foundation and other philanthropic organizations to develop and test approaches for building assets for children from low-income families. The initiative featured 12 community-based programs. This study focuses on the Michigan SEED Program operated by the Oakland Livingston Human Service Agency (OLHSA) in Pontiac, MI. Importantly, this is the only program of the 12 in the initiative that had a random assignment experimental design. OLHSA is a community-action agency that provides services to low-income families in the greater Pontiac metropolitan area.

The program described in this paper was designed by Michael Sherraden and Margaret Clancy at the Center for Social Development, Washington University at St. Louis; and Deb Adams and Sandy Beverly at the University of Kansas School of Social Work. We thank the Ford Foundation for funding, and seminar and conference participants at Syracuse University, Federal Reserve Bank of Cleveland, APPAM, and the National Tax Association for helpful comments. All errors and opinions are those of the authors and should not be taken to represent the views of any of the organizations with which they are affiliated.

The experiment was conducted at 14 OLHSA Head Start centers, as described in more detail in Marks et al., 2009a, 2009b. Recruitment took place in summer 2004 and participants took a baseline survey in fall 2004. For families with unmarried heads of households, the respondent was the biological parent or legal guardian of the focal child; for married couples consisting of both biological parents, the respondent was the female. Throughout the remainder of the paper, the term "parent" refers to the parent or guardian who was the survey respondent. Within each family, the experiment targeted one child, referred to as the "focal" child. This child had to have been enrolled in Head Start. For families with multiple children in Head Start, the youngest was designated the focal child; for families with twins in Head Start, one was chosen at random.

After the baseline survey was completed, the Head Start centers were matched in pairs according to poverty rates, racial and ethnic composition, and the proportion of one-parent families, and one center of each pair was randomly assigned to the treatment group. Hence, there were seven centers each in the treatment and control groups, respectively. Then, families of focal children attending centers in the treatment group were informed of their eligibility for and the structure of the SEED program. The control-group families were told they were not eligible.

Columns 1 and 2 of table 1 show baseline means of key variables used in the analysis for the treatment- and control-group families, respectively. Column 3 shows the difference between the treatment and control groups at baseline. One concern in analyzing data from a grouped randomized trial, such as MI SEED in which randomization was at the level of the Head Start center, is that standard cluster techniques for the calculation of standard errors may suffer from downward finite-sample bias from too few clusters (centers). With just 14 Head Start centers in MI SEED, this is a bona fide concern. Therefore, to address this, we follow Angrist and Lavy (2009) and calculate all standard errors in the paper using Bell and McCaffrey's (2002) Biased Reduced Linearization (BRL) estimator designed to mitigate the finite-sample bias from a small number of clusters. Therefore, in column 3, BRL standard errors are used to perform the tests for statistical significance.

Overall, statistically significant differences in baseline characteristics of treatment- and control-group families were about as frequent as would be expected based on chance alone. However, relative to controls, treatment-group families were less likely to have been employed (at the 10 percent level of significance). There also were some noticeable differences by race and marital status between the groups, although they were not statistically significant. There were very few differences in asset holdings between the groups at baseline.

A follow-up survey of all families was conducted in fall 2008, about four years later. A total of 686 families (or 86 percent) completed the follow-up survey, 86 of which had item nonresponse in the saving outcomes used below, which yielded an analysis sample of 600 families, 302 in the treatment group and 298 in the control group.

THE STRUCTURE OF THE MICHIGAN 529 PLAN

In a 529 plan there are two key parties: the beneficiary and the owner of the account. The beneficiary is a child for whom the funds are intended to be used for qualified educational expenses. The owner is typically an adult who makes a contribution to the plan, such as a parent, grandparent, other relative, or friend, but also can be a legal entity, such as a non-profit. A child may be named a beneficiary in multiple accounts; an owner may have accounts for multiple beneficiaries.

Control of the plan rests with the owner, who can re-designate the beneficiary or withdraw the funds at any time. The former can be done without tax consequences as long as the new beneficiary is in the same family as the original one when the account was established. If the new beneficiary is from a different family, then the re-designation constitutes a non-qualified withdrawal. In terms of liquidity, all non-qualified withdrawals result in a penalty of 10 percent on the earnings portion, and the earnings portion of the withdrawal is treated as ordinary income for the owner's federal tax purposes.

Contributions to a 529 are not tax deductible for the purposes of determining the owner's federal taxable income, but in Michigan they are tax deductible for determining state taxable income up to an annual limit ($5,000 for single individuals and $10,000 for married couples, regardless of the owner's income level). Earnings on all deposits are tax free at both the federal and state level, as long as the funds are withdrawn and used for qualified educational purposes. In Michigan, these qualified purposes include tuition, books, supplies,

Table 1
Sample Means for Baseline Economic and Demographic Characteristics of Families and Focal Children for Treatment- and Control-Group Families, Respectively

	(1)	*(2)*	*(3)*
Explanatory Variable	*Treatment-Group Families*	*Control-Group Families*	*Difference (T-C)*
Received TANF	0.272	0.295	-0.023
Received Food Stamps	0.570	0.597	-0.027
Parent/Guardian Covered by Health Insurance	0.709	0.755	-0.046
Focal Child Covered by Health Insurance	0.927	0.946	-0.019
Income (in $1000)	21.62	20.12	1.500
Log Tax Price	-0.0419	-0.0415	-0.0004
Employed	0.487	0.554	-0.067*
Age	30.54	30.26	0.280
High School Graduate	0.305	0.315	-0.010
More than High School	0.394	0.419	-0.025
Married	0.394	0.315	0.079
Widowed or Divorced	0.182	0.195	-0.013
Number of Siblings for the Focal Child	1.500	1.490	0.010
Black	0.348	0.527	-0.179
Other Race	0.136	0.074	0.062
English Spoken at Home	0.877	0.94	-0.063
Had Bank Account	0.722	0.748	-0.026
Had an Automobile	0.891	0.876	0.015
Had Financial Assets	0.252	0.272	-0.020
Had a Credit Card	0.311	0.302	0.009
Homeowner	0.298	0.275	0.023
Focal Child was a Boy	0.483	0.443	0.040
Focal Child Could Count Higher than Twenty	0.119	0.094	0.025
Focal Child Knew All Letters in Alphabet	0.159	0.134	0.025
Focal Child's Health was Very Good/Excellent	0.964	0.943	0.021
Focal Child Had Special Needs	0.199	0.158	0.041
Sample Size	302	298	—

Note: Statistical significance is indicated as follows: * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$. Statistical significance was determined by a two-tailed *t*-test using Bell and MacCaffrey's (2002) heteroskedasticity consistent BRL standard errors.

required fees, and some room and board costs at an in-state or out-of-state public or private university or trade school.

In general, the main benefit of a 529 plan is the accrual of tax free earnings on plan contributions. However, for the purposes of the Michigan SEED Program, the 529 platform was chosen to give these families access to a well-established, reputable saving vehicle, specifically targeted toward educational saving – factors emphasized to the treatment group through interactions with case managers. For those families in the treatment group who took up the offer to participate, the SEED program consisted of two legally separate 529 plans. Specifically, those in the treatment group were offered an $800 initial deposit into a Michigan state 529 plan, invested in the owner's chosen investment option, naming the focal child as the beneficiary and the parent as the owner. This is the parent-owned 529 plan.

In addition, contributions to 529 plans in Michigan may be eligible for a matching contribution from the state, known as a "state matching grant," in the first year that the beneficiary is enrolled in the plan. The match rate is 33 percent on every dollar contributed up to a $600 contribution threshold, for a maximum state match amount of $200, after which this match is exhausted. Eligibility is determined by the characteristics of the beneficiary, who must be under seven years old, reside in a household with annual income of $80,000 or less, and be a Michigan resident.

Under SEED, the initial deposit of $800 *automatically* qualified as a "contribution" for the purposes of the income-tested 33 percent state match. Since the initial deposit exceeded the maximum contribution limit for the state match, SEED treatment-group children received a combined initial deposit of $1,000: $800 initial deposit plus the maximum state match of $200. For administrative purposes the state match was placed into a second, custodial 529 plan, naming the child as the beneficiary and the SEED program as the owner, and automatically invested in the guaranteed option.

Savings for college also could have been accumulated for the focal child in treatment-group families through contributions by the parent into the parent-owned plan. SEED matched these contributions dollar-for-dollar on the first $1,200 of own contributions, after which the match was exhausted. The $1,200 limit applied to all parent contributions pooled across the first three years of the demonstration. The SEED matching contributions were placed in the custodial plan, outside of the control of the parent, and automatically invested in the guaranteed option. Contributions by other adults were not eligible for the SEED match.

The saving opportunities differed for control-group families. As the Michigan 529 plan was open to the general public, they were free to open a college savings plan and make their own contributions, but they were not eligible for the $800 deposit or SEED matching contributions. For qualified families, the first $600 of their own contributions was eligible for the 33 percent state match, up to the $200 match cap. The treatment also included the offer of some financial education and case management, but participation in which was voluntary. OLHSA held several events, such as family education nights, at each treatment Head Start center. These were not offered to control-group families.

IMPACTS ON 529 PLAN OWNERSHIP AND COLLEGE SAVING

In the first row of table 2, just over 8 percent of focal children in control-group families had any 529 plan naming the focal child as a beneficiary (either a parent-owned or custodial plan) by the end of the demonstration. This represents the counterfactual: in the absence of the treatment, about 1 in 12 children would have had a 529 plan anyway.

The second row in table 2 shows that 67 percent of the MI SEED treatment-group families had any 529 plan naming the focal child as a beneficiary by the end of the demonstration. The third row shows the simple difference in 529 plan ownership between the treatment and control groups. The average impact of the treatment was to raise ownership of 529 plans by 58.5 percentage points. This is an economically substantial effect. The associated BRL standard error is 7.4 percentage points. Therefore, the simple difference in 529 plan ownership between the treatment and control groups of 58.5 percentage points is a statistically significant treatment effect at conventional significance levels.

Table 3 gives basic statistics on 529 saving for the focal children at the end of the demonstration. Column 1 shows the average 529 balance then. For the treatment group, the balance is the sum of the balances in the parent-owned and custodial accounts. This average was $912 for the treatment group, of which $98 came from own contributions into the parent-owned plan (column 2). In fact, 22 percent of treatment-group families contributed

Table 2
Sample Means for 529 Plan Ownership for the Focal Child, by Treatment Status, Standard Errors in Parentheses

Sample	*(1) Focal Child has a 529 Plan*
Control-Group Families	0.084
Treatment-Group Families	0.669
Difference (T – C)	0.585 (0.074)

Note: Standard errors using Bell and MacCaffrey's (2002) heteroscedasticity-consistent BRL technique are reported in parentheses.

Table 3
Sample Means for 529 Contributions and Plan Balances for the Focal Child, by Treatment Status, Standard Errors in Parentheses, All Amounts in 2008 Dollars

Sample	*(1) Total 529 Balance for Focal Child, All Families*	*(2) Amount of Own Contribution, All Families*	*(3) Family Made Own Contribution to 529 Plan*	*(4) Amount of Own Contribution, Contributors*
Control-Group Families	288	—	—	—
Treatment-Group Families	912	98	0.223	440
Difference (T – C)	624 (158)	—	—	—

Note: The means in columns 2-4 were calculated for the treatment-group families only, as all contributions to 529 plans for control-group families were by definition "own" contributions. Standard errors using Bell and MacCaffrey's (2002) heteroscedasticity-consistent BRL technique are reported in parentheses.

to a 529 plan (column 3), and the average balance from these contributions at the end of the demonstration was $440 (column 4), equivalent to a monthly contribution of $9.16 over the course of the experiment.

The table also shows that the average 529 balance for the focal child was $288 for the control group. The difference between the mean treatment- and control-group balances was $624 (the bottom row of column 1), with a standard error of $158. This is the estimated treatment effect on 529 savings.

To account for baseline differences between the two groups, the first row of column 1 in table 4 shows parameter estimates from the following econometric model,

$$(1) \quad S_{ij}^{529} = \mu + \alpha T_j + \zeta \mathbf{X}_i + \chi \mathbf{C}_i + v_{ij},$$

where μ is a constant, and the dependent variable, S^{529}, is the 529 balance for the focal child at the end of the demonstration. The parameter α measures the treatment effect on 529 balances for the focal child, conditional on baseline parent, family, and child characteristics, **X** and **C**. The ordinary least squares (OLS) estimate of this treatment effect, adjusting for covariates, (labeled as the "first-stage estimate" in the table) was to raise 529 balances by $617 dollars. With a standard error of $318, this estimate is statistically significant at conventional levels.

The next question is to what extent did the treatment increase in overall saving for college for the focal child, including saving in non-529 assets. The second row of column 2 in table 4 shows parameter estimates from the following econometric model,

$$(2) \quad S_{ij}^{Total} = \gamma + \rho T_j + \varphi \mathbf{X}_i + \xi \mathbf{C}_i + \upsilon_{ij},$$

where γ is a constant, and the dependent variable, S^{Total}, is the total saving in all forms for education for the focal child at the end of the demonstration. The parameter ρ measures the treatment effect on total saving for the focal child, conditional on baseline parent, family, and child characteristics, **X** and **C**. The OLS estimate of this treatment effect, adjusting for covariates (labeled as the "reduced-form estimate" in the table), was to raise overall saving by $156 dollars. Note that this is far below the estimated treatment effect on 529 saving in column 1, suggesting substantial crowd-out. Indeed, with a standard error of $202, this overall saving estimate is not statistically different than zero at conventional levels.

The third row of the table combines the reduced-form and first-stage estimates to form an instrumental variable estimate of crowd-out. Formally, parameters from the following econometric model were estimated,

$$(3) \quad S_i^{Total} = \kappa + \beta S_i^{529} + \psi \mathbf{X}_i + \varsigma \mathbf{C}_i + u_i,$$

where κ is a constant. The parameter β measures the extent to which a dollar in the form of a 529 plan raised total saving for the focal child. In turn, $1 - \beta$ measures the extent to which a dollar in the 529 plan was offset or crowded out by reductions in saving for the focal child in other (non-529) forms.

Table 4
Parameter Estimates of the Effect of 529 Plan Saving on Total Saving for the Focal Child, Standard Errors in Parentheses

	(1)	*(2)*
	Dependent Variable	
Explanatory Variable	*Total 529 Plan Balance for Focal Child*	*Total Saving for Focal Child*
First-Stage OLS Estimate		
Assigned to Treatment Group	617 (318)	—
Reduced-Form OLS Estimate		
Assigned to Treatment Group	—	156 (202)
Instrumental Variable Estimate		
Total 529 Plan Balance for Focal Child	—	0.253 (0.329)
Instrumental Variable Tobit Estimate		
Total 529 Plan Balance for Focal Child	—	0.454 (0.230)
OLS Estimate for Control Group Only		
Total 529 Plan Balance for Focal Child	—	0.733 (0.020)
Ordinary Tobit Estimate for Control Group Only		
Total 529 Plan Balance for Focal Child	—	0.782 (0.019)
Number of Families	600	600

Note: Each cell of the table shows a parameter estimate from a different specification. The first through fourth rows show estimates based on the pooled sample of treatment- and control-group families 4 years after randomization. Each specification controls for parent, family, and child characteristics at baseline, as described in the text. The instrument for the IV and IV Tobit estimates is randomly assigned treatment status. To illustrate the bias from not instrumenting, the fifth and sixth rows shows the OLS and Tobit estimates based on the sub-sample of control group observations 4 years after randomization. This is what would be estimated using just observational data. Standard errors using Bell and MacCaffrey's (2002) heteroscedasticity-consistent BRL technique are reported in parentheses. There are 298 observations in the control group.

The fundamental problem in estimating β using OLS is that the estimator will be biased from unobserved heterogeneity if families who otherwise would have saved for their children's education self-selected into participation in saving via 529 plans. In this case, the bias would be upward, suggesting little crowd-out. Therefore, in (3), the randomly assigned treatment status, T, is used as an instrument for S^{529}, and the parameters are estimated using the instrumental variable (IV) estimator. The IV estimate of $\hat{\beta} = 0.253$ indicates that for every dollar of 529 saving, overall saving rose 25 cents, or 75 percent crowd-out. The t-statistic for the test of the null hypothesis of no crowd-out ($\beta = 1$) is -2.27, so that we can reject the null of no crowd-out (in favor of the alternative of some crowd-out) at conventional significance levels. However, we cannot reject the null of full crowd-out ($\beta = 0$).

About 45 percent of the sample reported at follow up having zero saving for the focal child's education. Because of this and the heavy reliance on the Tobit estimator in past studies of saving behavior, the fourth row of the table presents the IV Tobit estimates of β, allowing for censoring at zero. The IV Tobit estimate of $\hat{\beta} = 0.454$ indicates that for every dollar of 529 saving, overall saving rose 45 cents, or 55 percent crowd-out. The *t*-statistic for the test of the null hypothesis of no crowd-out ($\beta = 1$) is -2.37, so we can reject the null of no crowd-out (in favor of the alternative of some crowd-out) at conventional significance levels. In addition, the *t*-statistic for the test of the null of full crowd-out ($\beta = 0$) is -1.97. Thus, we can reject this null (in favor of the alternative of less than full crowd-out) at conventional significance levels.

The final two rows of column 2 in table 4 illustrate the bias from unobserved heterogeneity if families who otherwise would have saved for their children's education self-selected into participation in 529 plans. In particular, they show the OLS and ordinary Tobit estimates of β in (3), respectively, using just the subsample of control-group families. In the absence of a randomized controlled trial, this is what would be estimated using observational data. The IV estimate indicates that each dollar of 529 savings raises total college savings by 73 cents $\hat{\beta}^{OLS} = (0.733)$, or 529 savings results in 27 percent crowd-out. The ordinary Tobit estimate is similar, $\hat{\beta}^{OLS} = 0.782$, or 22 percent crowd-out. A comparison of these estimates to their IV counterparts in the table indicates substantial upward bias. The ordinary and IV estimates are statistically different from one another at the 5 percent level based on a Hausman test.

CAVEATS AND DISCUSSION

There are a number of caveats to these findings. First, this paper represents a positive, not a normative analysis, of an existing program. There is no discussion of the conditions under which, in general, it is optimal or desirable to offer low-income families institutional access, lump-sum, and match-based incentives to save for education, etc., or at what age those should commence, especially for families likely to have full financial need in the absence of the SEED intervention. It may be that the intervention would have been better targeted when the focal children were older. Second, the extent to which changes in educational saving associated with incentives, like matching, access to financial institutions, and the other elements of the SEED treatment, emerge in the longer run – especially when the children reach ages at which family saving for college is more prevalent – is an open question. Third, and related, it is difficult to measure the early effects of this intervention on the children's actual educational outcomes. The length of the experiment does not allow for an examination of any long-term impact, such as increased college enrollment, that would result from greater savings for education. Fourth, because the experiment was not designed with multiple, differentiated treatment groups, it is not possible to quantitatively isolate the specific reasons for low (or no) saving, nor the independent impacts of matching, account provision, and case management, respectively. Finally, although novel, the experiment covered a relatively small number of Head Start centers and, therefore, a small sample of low-income families in and around Pontiac, MI.

References

Angrist, Joshua and Victor Lavy. The Effects of High School Achievement Awards: Evidence from a Randomized Trial. *The American Economic Review* 99 (2009): 1384-1414.

Ashraf, Nava, Dean Karlan, and Wesley Yin. Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines. *Quarterly Journal of Economics* 121 (2006): 635-672.

Bell, Robert M. and Daniel F. McCaffrey. Bias Reduction in Standard Errors for Linear Regressions with Multi-stage Samples. *Survey Methodology* 28 (2002): 169-181.

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

Duflo, Esther and Emmanuel Saez. The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment. *Quarterly Journal of Economics* 118 (2003): 815-842.

Grinstein-Weiss, Michal, Michael Sherraden, William Gale, William M. Rohe, Mark Schreiner, and Clinton Key. Ten-Year Impacts of Individual Development

Accounts on Homeownership: Evidence from a Randomized Experiment. Chapel Hill, NC: University of North Carolina, 2011. Mimeo.

Marks, Ellen L., Bryan B. Rhodes, Gary V. Engelhardt, Scott Scheffler, and Ina F. Wallace. Building Assets: An Impact Evaluation of the MI SEED Children's Savings Program. Research Triangle Park, NC: RTI International, 2009a.

Marks, Ellen L., Bryan B. Rhodes, Jennifer Wheeler-Brooks, and Deborah Adams. A Process Study of the SEED Community Partners Initiative. Research Triangle Park, NC: RTI International, 2009b.

Mills, Gregory, William G. Gale, Rhiannon Patterson, Gary V. Engelhardt, Michael Eriksen, and Emil Apostolov. Effects of Individual Development Accounts on Asset Purchases and Saving Behavior: Evidence from a Controlled Experiment. *Journal of Public Economics* 92 (2008): 1509-1530.

Saez, Emmanuel. Details Matter: The Impact of Presentation and Information on the Take-up of Financial Incentives for Retirement. *American Economic Journal: Economic Policy* 1 (2009): 204-228.