



FINANCIAL ACCESS INITIATIVE

Contributions to this research report made by a member of The Financial Access Initiative research group.

## Framing Note No. 1

### The Impacts of Savings

Dean Karlan

Financial Access Initiative

January 2008

The Financial Access Initiative is a consortium of researchers at New York University, Harvard, Yale and Innovations for Poverty Action.

NYU Wagner Graduate School  
295 Lafayette Street, 2<sup>nd</sup> Floor  
New York, NY 10012-9604

T: 212.998.7536  
F: 212.995.4162  
E: [contact@financialaccess.org](mailto:contact@financialaccess.org)

[www.financialaccess.org](http://www.financialaccess.org)

Ask for a definition of “microfinance” and most will say something about small loans. Why not savings? There is a broadly held assumption that the poor save all that they can already, and that they can better use an infusion of outside capital to break the cycle of poverty. There’s some truth to that. But three observations suggest that this is not entirely correct, and that efforts to spur savings can help the poor.

The first observation is that poor households *do* have discretionary spending. Most people that are classified as “poor” according to official statistics are not living hand to mouth. They are often living in dire circumstances, without security, and without luxuries. But they are not without choices. Banerjee and Duflo (2007) document large budget shares spent on expenses like social and religious events that are in part discretionary, giving potential scope for saving within budgets.

Second, financial institutions appear to have ways to stimulate higher savings. Recent experience shows that simple changes to savings account policies, designs, and even marketing can encourage the poor to save more.

Third, there remains a huge gap between the cost of debt and the return on savings—and thus big potential impacts to saving. Individuals who “save up” sufficiently can “earn” the interest *not* paid on debt. In short, if the poor can use their own capital rather than borrowing from moneylenders the effective net gain (i.e., the implicit return on savings) could be as high as 100 or 200 percent per year.

Building on these stylized observations, this focus note will address two broad questions, one methodological and one substantive: (1) Why is measuring the impact of savings hard? (2) What rigorous evidence do we have to date about the impact of savings, and of specific savings products?

### **Why is Measuring the Impact of Savings Hard?**

In order to measure the impact of savings we must first understand what it means for poor families *who save regularly* to have their own source of funds to draw from. This is relatively simple to document: we can describe how important it is for families to have a secure and stable place to store their money, and we can, for example, collect data on how frequently households use their savings to reduce hungry periods for their families or pay for medical expenses. These are useful first observations, but the analysis is limited to today’s savers and their existing savings practices. It gives little guidance on the effect of expanding the extent of saving or introducing innovations.

Understanding the effect of savings programs on household welfare and on financial institutions requires more than tracking and describing outcomes. Impact evaluations are needed to answer the critical question: “How are the lives of the participants different than they would have been had the program, product, service or policy not been implemented?” The second part of the question, what would have happened in the absence of the program, i.e. the *counterfactual*, is key to understanding the true impact of the program. Researchers generally create the counterfactual by comparing program participants to a control group who do not get the program. In the case of savings, a program that does not increase participants’ savings balances *by more than the control group* will show no impact (or even a negative impact).

Having a control group is thus critical, as is making sure the control group is chosen well.

### ***Statistical issues***

Efforts to measure the impact of savings can be seriously affected by statistical biases, especially *selection bias*. The selection problem occurs when we form the control group: if we simply look at a population and compare savers to non-savers we risk comparing *two different types of people*. Savers may be better planners, or more educated; they may have more disposable income, or they may be more risk averse. These differences can be important enough to skew impact studies, since what would seem like an impact of saving might instead be an impact of a general propensity to plan well. The root problem is that we can easily control for, say, age in statistical analysis, but something like the ability to plan ahead is difficult to quantify, and thus difficult to properly take into account.

An example of this difficulty is evident if we compare health outcomes, for instance, between savers and non-savers. We might find a big difference between the two groups, but that difference might not be due to savings itself but to correlated factors such as planning and risk aversion. There might even be *reverse causality*: healthy people might be able to earn more, or they might spend less on hospital visits, increasing savings balances.

Experience shows that the best way to deal with this problem is with randomized control trials (RCTs). By randomly assigning potential participants from the same sample population into treatment and control groups at the outset of the program, RCTs ensure that the treatment and control groups are, on average, similar *except for participation in the program*. This allows researchers to establish causality: if there is any difference between the groups after participation in the program it will be due to the program and not any confounding factors. Comparing the two groups thus gives a clean, unbiased measure of the welfare gain to access.

### ***Measurement problems***

In practice matters are a bit more complicated. One of the biggest challenges to measuring the impact of savings programs is that savings are not easy to capture in household surveys. Four issues in particular contribute to this difficulty: size, timing, diffusion, and confounding of other services.

Unlike credit inflows, which can be sizable relative to household income, savings flows can be quite small, and balances accumulate slowly. The most recent edition of the *Microbanking Bulletin* shows an average loan balance of \$456 per borrower among clients of leading MFIs, while average deposits are \$251. Since less can be purchased with smaller amounts of money, the expected impact on outcomes such as enterprise investment, education and health likewise are smaller. When conducting empirical research, the smaller the expected effect the larger the required sample size needed to study impacts.

Savings are also more difficult to measure than credit because the timing of the change in behavior and outcome is less obvious. For many households, savings develop slowly through a small reduction in consumption that accumulates over time. At some point the household will have built up enough savings to protect itself from shocks (such as sickness or unemployment), to pay school fees, or to start a business. But from the outside that point is hard to predict. Researchers therefore need to measure savings balances at multiple points in time, often over several years, a requirement which can be challenging and expensive, or measure the household's propensity to make important investments or absorb shocks.

The third issue, diffusion, is perhaps the most important. Poor households save in many different ways. In addition to opening an account with a formal savings institution they may save cash at home, with deposit collectors, within a savings club (such as a rotating savings and credit association, or ROSCA), or by lending to family members. They may also save in non-cash assets such as jewelry or livestock. An evaluation that failed to capture the total savings held in these different savings vehicles, as well as the shifts in savings from one vehicle to the other, could understate

savings rates, or overstate savings growth, potentially leading to incorrect conclusions about the impact of the program. Specifically, an evaluation which focused on just one savings channel may incorrectly conclude that net savings goes up, whereas in fact the only impact was a shift from one type of savings to another (we will revisit this issue below when discussing the expansion of the Pahnal savings accounts in Mexico).

Accurate measurement of the impact of *savings per se*, not *banking services more generally*, can also be confounded by the fact that few institutions offer *only* savings. This problem can be solved by offering savings in a controlled enough fashion that one can study the impact of it, unconfounded by other financial services. This does *not* mean one must stop other services. To the contrary, what a firm must do is control the marketing or ease of access of the savings component to some individuals randomly chosen, and not others. Everyone can have access if they seek it out, and everyone can have access to the other financial services. This allows the firm to measure the marginal impact of the savings component of their services.

### ***Context matters***

In research there are two very different concepts used to gauge the extent to which others can rely on the results from a study. The first is *internal validity*: How rigorous was the study? How well did the study answer the research question within the particular population it considered? A carefully executed RCT is perfect for ensuring internal validity. The second concept is *external validity*. How applicable are the results from the study to other settings? If researchers show that a program works in one location, can they be sure it will be effective in another? There is no easy way to determine this. This is an issue with *all* evaluations, whether randomized or not randomized, and whether quantitative or qualitative. Even if researchers design the best possible study and nail the measurement issues discussed above, lessons from one randomized trial may not necessarily transfer to entirely different contexts. To escape this issue, one must have a theory as to why and how something works, an assessment of the local environmental, cultural and economic contexts, and last but not least, a replication of clean, randomized evaluations in order to test the theory and establish confidently whether, and when and where, something works.

### **What Have We Learned about the Impact of Access?**

Not surprisingly, evidence is limited on the impacts of efforts to increase access to savings and to improve the quality of existing savings accounts. Two types of studies have been conducted to date: studies that generally measure the impact of access to savings; and studies that hone in on specific design features to understand, conditional on access, the qualities that generate higher savings levels.

We are unaware of any studies to date that have truly rigorous evidence on the impact of savings, though two studies stand out and are worthy of mention despite issues which leave us unsatisfied.

First, a study in India by Robin Burgess and Rohini Pande (2005) examines the impact of a general expansion of *banking services*. This study takes advantage of a short-lived change in Indian banking law which required that banks open more branches in poor, rural areas. When the law was later rescinded, the before-and-after effect of availability could be measured. The study finds a noticeable improvement in poverty reduction at the time the law was enacted, and a clear reversion when it was repealed. Although the study provides evidence on the overall impact of access to a bank, it does not provide separate analysis for access to saving accounts versus credit. To do that, we would need to examine an expansion of savings that was not simultaneously accompanied by an expansion of credit.

For that dynamic, we turn to an analysis of the expansion of Patronato del Ahorro Nacional (Pahnal), a government financial institution in Mexico, (Aportela 1999). Starting in 1993, Pahnal decided to expand its savings operations through the postal office network. This approach is similar to the postal

bank network found in countries such as Germany, South Africa and Japan. By using the postal offices, Pahnal was able to serve people who were otherwise hard to reach, without incurring the fixed cost of opening a vast branch network. The study is thus relevant to innovations such as mobile phone banking, which aim to expand financial access without bank branches.

As part of this expansion, Pahnal offered two savings options: (1) A fixed term instrument (12, 24 or 36 months) with compulsory monthly deposits of five dollars, and a prohibition on withdrawals before maturity; and (2) A liquid savings account with a minimum balance of just over five dollars and no fees, but a lower interest rate than the fixed term. Both accounts also had a monthly lottery component. Aportela found substantial use of the new accounts but is unable to rule out substitution from other savings vehicles. Moreover, the study does not allow for identification of which of the two products, or within either product which particular feature, led to the impact. The study does allow measurement of the “kitchen-sink” impact, which is important in its own right.

The key with any impact study, as discussed above, is to answer the following question: “How did the lives of those given access change relative to how their lives would have changed had they not been provided access?” The ideal methodology for answering the question of how the lives of those given access changed relative to how their lives would have otherwise been is a randomized evaluation. When randomization is not possible, evaluators search for scenarios that function *as if* they were random. This is called a “natural experiment” (rather than a “controlled experiment”). Aportela uses the partial (i.e., not all at once) expansion as a sort of natural experiment. For this approach to work, researchers must verify that preexisting savings rates or poverty rates are not correlated with the expansion path. It is also helpful to have information on *why* and *how* Pahnal decided to expand as it did, in order to make sure there was not a senior manager advocating to expand specifically into areas where unmet demand had been identified. In this case there was no evidence of any such correlation. The expansion appears to have more closely followed Pahnal’s operational convenience i.e., proximity to its prior branches. An unavoidable issue remains as to whether Pahnal’s prior branches were located strategically so as to create a preexisting trend, and a resulting selection bias. With that one caution, the analysis proceeds.

The analysis uses data from the 1992 and 1994 Mexican Household’s Income and Expenditure surveys. The advantage of using these data, rather than client data, is that there is no individual selection bias. The analysis compares the change in outcomes for those in communities that received the Pahnal expansion to the change in outcomes for those in communities that did not receive the Pahnal expansion.

The impact on formal sector savings balances was noticeable: the average savings rate increased by 3 to 5 percentage points, and for low-income individuals it increased at a higher rate of 5.7 to 8 percentage points. The emphasis of the analysis in the study is on formal sector savings volume, broken down by different income levels. The attempt to measure aggregate savings, including informal savings, fails to reach decisive conclusions because of a lack of statistical precision. Understanding whether this was a shift from informal to formal, or an increase in net savings, is of course important. Regardless, we tend to think that the formalization of savings implies safer savings for individuals. Regarding income, the study reveals that the impact on income was positive as well (although not statistically significant).

### **How Does Quality Matter? How are Saving Decisions Made?**

Isolating the difference in outcomes caused by variations in products or services offered can help practitioners optimize their products and services. Randomized trials are especially effective here because they can enable management to keep everything the same across clients except for one particular product or policy change.

One example of this process is a study we conducted of a “commitment savings product” offered by a rural bank in the Philippines (Ashraf, Karlan and Yin 2006). The bank created a product, called SEED, designed to help clients save by locking away their money until they reached a self-specified savings goal. The goal could be a pre-set savings amount (e.g., enough money for a new roof), or a future date (e.g., in time for village festivals). The clients had complete flexibility to choose which of these restrictions they would like on their account. Once the decision was made it could not be changed, however, and SEED clients could not withdraw funds from the account until they met their chosen goal.

The bank was primarily interested in increasing savings balances among its clients. To evaluate the impact of the product the bank implemented an RCT where it randomly assigned about 1,800 individuals to either receive an offer to open the SEED account (the treatment group) or not (the control group). Of those offered the account, 28 percent accepted and opened the account (note that 100 percent take-up would have been a failure, as a commitment savings account, paying the same interest as a normal savings account, is only ideal for those with either self or familial control issues to overcome). We found a strong effect on savings: after twelve months, average balances increased by 80 percent in the treatment group compared to the control group. Those in the treatment group who actually opened the commitment-savings product increased savings after one year by 337 percent. When we followed up after the initial program period (after two and a half years), however, the differences between the treatment and control groups were smaller and no longer significant statistically. The bank provided no further encouragement, reminder or marketing, beyond the initial field visit. Further study is needed to determine whether continued marketing would have maintained higher savings levels, or whether the commitment device was instead “unraveled” by competing pressures.

As part of the study we also measured whether the spousal control channel was important. We found a significant increase in women’s decision-making power within the household, and an increase of the purchase of household-oriented durable goods. We can attribute this change causally to having access to this specialized account, which likely helped women gain control over household assets.

Hence, we learn from this that liquidity matters, but more liquidity is not necessarily better. Some have a specific preference for *illiquidity*, for either self-control or spousal/familial control motivations.

Careful product testing has also shown us that timing matters. In a series of controlled field experiments in Busia, Kenya, in which farmers were given the opportunity to put proceeds from their harvest directly into a commitment savings account dedicated to fertilizer use, we learn about the importance of the *timing* of opportunities to save (Duflo, Kremer and Robinson 2006). The experiments were conducted over a four-year period, with treatments varying in different seasons. One of the beauties of this series of experiments is the ability to tease apart the precise mechanism needed to generate the result (as compared to, for example, the SEED experiment above, in which more of a “kitchen-sink” approach is taken). A second striking characteristic of these experiments is that they focus on the decision to buy fertilizer (and the decision to save in order to buy fertilizer). Hence, whereas in many cases the impact of savings is diffuse, here there is a simple and concrete outcome that can be tracked as well: use of fertilizer on one’s farm.<sup>1</sup>

---

<sup>1</sup> Fertilizer use has huge consequences for farm yield, and thus household income. In the context studied in Busia, Kenya, fertilizer use leads to a 36 percent increase in returns over a season, or 69.5 percent on an annualized basis (Duflo, Kremer and Robinson 2008).

The key finding from this study is that when farmers were given the opportunity to put proceeds from their harvest immediately into a commitment-savings account dedicated to fertilizer use, they were more likely to use fertilizer in the next planting season (56 percent versus 17 percent who were offered fertilizer at planting time). Even farmers who were offered the account but were given a period of a few days before they had to pay were much less likely to do so (30 percent). Interestingly, these individuals *said* they wanted to buy the fertilizer, but a few days later when the payment was due they did not complete the purchase. This group was correspondingly less likely to use fertilizer months later. Farmers who were told about the fertilizer but not given the opportunity to pre-pay into the commitment-savings account behaved exactly like the control group in terms of fertilizer use (even though they were given valuable information about the potential impact of using fertilizer, and the control group was not).

As a final way of comparing the relative importance of timing versus subsidy, the researchers offered one treatment group a 50 percent subsidy on their fertilizer (but at the time of planting), and compared that to those offered the pre-purchased full-price fertilizer in the commitment savings account at the time of harvest. The take-up rate was the same. In other words, getting the timing right by giving the farmer access to a commitment savings account at the moment she is flush with cash had just as big an effect on fertilizer use as subsidizing the fertilizer use by 50 percent (and the increase in yield on fertilizer is increasing in usage throughout the range tested, albeit not linearly). Thus improving the ability to save can have very big potential impacts on farm income, household income, and, by implication, the various outcomes that tend to result from increased income (such as improved education, health, and nutrition).

These two cases offer good examples of how financial institutions can use costless mechanisms (like illiquidity in the case of SEED, or optimal timing as in the fertilizer experiments) to generate major increases in savings, and also increase the likelihood that farmers make profitable investments.

## **Conclusion**

The potential gains to better ways to save are, in theory, huge. Informal savings options are plentiful but insecure, while formal savings products as currently offered can be made much better. This is likely due both to the price and quality of the service. In developed countries, we have seen a proliferation of innovations designed to guide people towards saving more. Poor people in developing countries are no less in need of such guidance. The core idea is that although choice is good, how one frames choices and designs the architecture of the decision-making process can have just as much an effect on outcomes as the choices themselves. The good news is that there is typically no “default” method of offering financial services, and designing the right defaults and ways of framing choices can have large potential impacts.

As with any program, of course, resources are scarce. Should donors focus more attention on savings, and less on education, health, credit, or insurance? Only proper impact evaluations can help determine the most effective place for donor money. With the attention of the microfinance industry on credit, we have yet to see clean evaluations of savings initiatives, but that is beginning to change. The careful introduction of ideas from behavioral economics and psychology, and the adaptation of new savings product ideas from developed countries breed promise for global improvements in the quality of, and access to, formal sector savings.

## REFERENCES

- Aportela, F. (1999). "Effects of Financial Access on Savings by Low-Income People." MIT Department of Economics Dissertation Chapter 1.
- Ashraf, N., D. Karlan and W. Yin (2006). "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines." Quarterly Journal of Economics 121(2): 673-697.
- Banerjee, A. and E. Duflo (2007). "The Economic Lives of the Poor." Journal of Economic Perspectives 21(1): 141-167.
- Burgess, R. and R. Pande (2005). "Do Rural Banks Matter? Evidence from the Indian Social Banking Experiment." American Economic Review 95(3): 780-795.
- Duflo, E., M. Kremer and J. Robinson (2006). "Understanding Technology Adoption: Fertilizer in Western Kenya Evidence from Field Experiments." working paper.
- Duflo, E., M. Kremer and J. Robinson (2008). "How High are Rates of Return to Fertilizer: Evidence from Field Experiments in Kenya." American Economic Review Papers and Proceedings forthcoming.