

Impact of Credit: How to Measure Impact, and Improve Operations Too

Nathanael Goldberg & Dean Karlan

January 2008

Contributions to this research made by a member of The Financial Access Initiative and Innovations for Poverty Action.

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, 2nd Floor
New York, NY 10012-9604

T: 212.998.7523
F: 212.995.4162
E: contact@financialaccess.org

www.financialaccess.org

Innovations for Poverty Action applies rigorous research techniques to develop and test solutions to real-world problems faced by the poor in developing countries.

Innovations for Poverty Action
85 Willow St, Building B, 2nd Floor
New Haven, CT 06511

T: 203.772.2216
F: 203.772.2428
E: contact@poverty-action.org

www.poverty-action.org

As the microfinance industry attracts greater media attention, donations and now investments with each passing year, a single provocative fact becomes increasingly surprising: the lack of clear evidence demonstrating whether it works. Undoubtedly the microfinance industry has been spectacularly successful at reaching millions of poor households with credit and savings services. And those clients appear to be satisfied customers: many return. Yet in many markets, many have access but do not avail themselves of the services. And in developed markets, regulators often take proactive steps to restrict high-interest loans under the premise that they can do harm.

The basic premise of microfinance, that credit extended to the poor for investing in entrepreneurial activities increases their welfare of their households, remains untested by rigorous scientific standards. Given the massive flow of money pouring into microfinance, policymakers, donors and investors should demand to have clearer, more decisive evidence in favor of microfinance. Resources are scarce, after all, and every dollar invested in microfinance is a dollar not donated to or invested in health or education or other projects aimed at alleviating poverty.

Why Should We Measure Impact?

All impact evaluations attempt to answer the same question: “How are the lives of the participants different than they would have been had the program, product, service or policy not been implemented?” This is a powerful question and a challenging one to answer well. The first part of it is straightforward: how did the lives of participants change? To do that, one merely needs to follow individuals, survey them before and after the program, and observe how their lives changed. The second part is the tough part. *How would their lives have changed had the program not existed?* The difference between the outcomes of the participants and this *counterfactual* situation is the impact of the program. To measure the counterfactual one needs a control group. In this focus note we will discuss the relative merits of different approaches to measuring the counterfactual.

Measuring impact can also help improve operations. Getting the answer right can provide practitioners, donors, investors, and policymakers with critical information about what types of services are effective, helping them to design good programs and allocate scarce resources toward interventions that work. This is true both for impact evaluations of an entire program (e.g., testing the impact of expanding access to credit), and impact evaluations of program innovations (e.g., testing the impact of one loan product versus another loan product, or of providing add-on services such as entrepreneurship or health training).

In this sense an impact evaluation is akin to good market and client research. By learning more about the impact on clients, or particular types of clients, one can design better products and processes. This is not at all limited to non-profit organizations, or even firms with an explicit social mission. For-profit firms can and should invest in learning how best to have a positive impact on their clients. By improving client loyalty and wealth, an institution is likely to keep its clients longer and provide them with the resources to use a wider range of services, thus improving profitability.

Thus, impact evaluations are not simply about measuring whether a given program is having a positive effect on participants. The microfinance industry, like many industries, is often driven by a set of assumptions about the “right way” to do business. For example, until fairly recently group liability was widely considered the “best practice” way to offer collateral-free loans

because the group guarantee was essential to ensure repayment. It took a maverick lender (ASA, Bangladesh) to show other MFIs it was possible to offer collateral-free individual loans to the poor. Impact evaluations could have simplified the process, in a transparent and controlled way, by individually testing assumptions about how to offer financial services to the poor. Examples of processes successfully evaluated this way include changing repayment schedules from weekly to monthly, offering door-to-door deposit collections, and targeting “ultra-poor” households which have been excluded from credit.

Lastly, even financially self-sufficient financial institutions often receive indirect subsidies in the form of soft loans or free technical assistance from donor agencies. Therefore, it is reasonable to ask whether these subsidies are justified relative to the next best alternative use of these public funds. Donor agencies have helped create national credit bureaus and worked with governments to adopt sound regulatory policies for microfinance. What is the return on these investments? Impact evaluations allow program managers and policymakers to compare the cost of improving families’ well-being through microfinance to the cost of achieving the same impact through other interventions.

Bottom line: impact evaluations if done well are about the future, not the past. They should not merely satisfy donors about past performance, but rather guide future decisions by financial institutions, policymakers, donors and investors.

What Should We Measure?

There is no single definition of household welfare or poverty. We typically think of impact as steps, following the chain of events. The first question is how the immediate cash flows in the household or enterprise changed. Was more money invested in durable investment goods, working capital, consumption, household durables, or health care or other such emergencies? If money was invested in the enterprise, what change in profits did this lead to? What is the net impact? If there was an increase in enterprise profits, how did the additional profits end up getting spent? Were the funds reinvested, or did consumption rise in the family? If money was spent to absorb a shock, did the family meet the need, or did they have to sell productive assets?

Then, on the non-economic outcomes, each of the above steps could lead to changes in less tangible outcomes, such as female empowerment, sentiments of self-worth, lower (or higher) stress or levels of depression. Qualitative research is helpful for understanding the local context in order to construct appropriate measures of these types of outcomes. Like income, these generally rely on self-reported measures, but here they will reflect the self-perception of respondents. In some cases proxies can reduce the ambiguity inherent in qualitative measures. For instance, instead of asking clients how healthy they are, researchers could ask them how many days of work they missed or how many visits to the doctor they made in the previous month. For empowerment, researchers might ask women the extent to which important decisions about family planning, or household purchases, are aligned with their preferred choices.

Given the importance of health and education, it is typically important to measure changes in usage of health services, nutrition, disease and education. There are two potential channels for observing changes in health and education outcomes: the loan funds may have been used directly to cover health, nutrition or education expenditures. Or, the funds may have been invested in an enterprise, which led to higher enterprise income and then higher health, nutrition or education expenditures. Specifically regarding education, an important question remains unanswered: if the investment in the enterprise requires labor as well, does the family shift labor *into* the enterprise, potentially thus *reducing* education investments in their children? This potential adverse

consequence is potentially very important to watch for, as it would undermine the long-term development goals.

The evaluation process can be used to provide evidence to practitioners and policymakers on the most basic—and often, the most important—questions about how credit is used. Many MFIs assume (or certainly claim) their loans are used for entrepreneurial investments. Some MFIs attempt to ensure this by targeting only prospective clients with active businesses. But this does not tell them anything about what borrowers actually spend their loans on. Other MFIs may attempt to enforce this more aggressively by demanding to see receipts for capital investments from clients, or verifying the purchases through spot checks in the field. But even this cannot determine for certain how loan proceeds were truly used. Money, of course, is fungible, and an entrepreneurial loan made to a borrower who was already going to make an investment anyway (out of other sources of funds) could easily be used for consumption. A rigorous evaluation can look much deeper by collecting data on business investments, monthly cash flows, profits, and household consumption and assets.

The Challenge of Rigorous Impact Evaluation

Now we return to the question stated above, “How are the lives of the participants different than they would have been had the program, product, service or policy not been implemented?” Without measuring the counterfactual—how participants’ lives *would have* been had they *not* participated—“impact” research tells us little about the true impact of the program. Even if participants experience an improvement in welfare, many factors could contribute to the changes in their outcomes. For instance, participants’ income could increase, but this could be due to general economic changes in the region (referred to as an “omitted variable” bias), or changes may simply be a consequence of entrepreneurial participants and their preexisting trend towards better outcomes (referred to as a “selection bias”).

Measuring the counterfactual is the key to determining the impact of the program, and how well the design allows the researcher to measure it is one critical difference between a reliable and an unreliable evaluation. The typical strategy for creating the counterfactual is to compare clients to a group of non-clients, and to attribute any differences between the groups to the impact of the program. This will only work if the two groups were otherwise similar at the outset of the program. Otherwise the comparison merely may capture differences between the groups unrelated to how well the program works (i.e., a “selection bias”).

Identifying suitable comparison groups can be particularly difficult with credit evaluations because of the problem of selection bias. What kind of person chooses to participate in a microfinance program (and thus, to become an entrepreneur)? If microfinance clients have a special determination and ability to improve their welfare, then comparing their outcomes to non-clients (without this drive) will overstate the impact of the credit program. The extent to which this increases (or decreases) the estimate of the program impact is the self-selection bias. If microfinance clients are especially driven, it could be that they would have improved their welfare just as much even without borrowing. What makes selection bias so problematic is that it is driven by presumably unobservable characteristics. Factors like “entrepreneurial ability” or “drive” are difficult to measure, and therefore we do not have a suitable way of correcting for these differences in the impact analysis. We won’t know for sure how serious a problem this is until rigorous impact assessments of credit are completed, side-by-side with alternative approaches, and the results can be compared.

Coleman (1999) shows how serious a concern selection bias can be in a study of microfinance borrowers in northern Thailand. By forming a group of prospective clients who signed up a year

in advance to participate in two village banks, Coleman was able to create a comparison group mostly free of selection bias: both the borrowers and the non-borrowers had selected into the program, at the same point in time. Coleman then generated two estimates of the impact of the program, an unbiased estimate using the clients who signed up in advance as the comparison group; and a “naïve” estimate using a group of seemingly similar non-participants (like in typical non-rigorous evaluations). Comparing his unbiased impact estimate to the estimate he would have calculated had he naïvely compared program participants to a group of non-participants Coleman finds the “naïve” estimate substantially overstated the gains from participation on several outcomes (especially women’s landholding) because participants turned out to be initially wealthier than non-participants. Unobservable differences between participants and non-participants have significant effects on the measures of some common indicators of microfinance impact, including women’s business revenue and female-owned business assets.

A related pitfall is bias from program placement, in which outcomes in program villages are compared to outcomes in non-program villages. The problem with this method is that programs choose where they operate for a reason. They may target the poorest villages, for instance, or they may start cautiously with better-off clients before expanding their outreach. The bias from non-random program placement, therefore, can go either way, depending on whether the evaluation compares program villages to non-program villages that may be (even unobservably) better or worse off.

These problems can be solved with randomized controlled trials (RCTs). By randomly assigning clients into the treatment and control group at the outset, RCTs ensure that both groups are similar across observable and unobservable dimensions. This eliminates both the omitted variable bias and the selection bias noted above, as long as there is a sufficiently large sample size. Thus the only difference between the groups, on average, is whether or not they are assigned to participate in the program. This is the critical link in allowing researchers to determine whether the program is the *causal factor* in creating impact: since in an RCT there are no other differences between the groups, it has to be the program that causes any changes.

However, while RCTs can be powerful, they are not always easy to implement. Since clients must be assigned into treatment and control groups before they receive services, RCTs require careful advance planning. Researchers cannot simply collect data after the fact. Randomizing operations requires substantial ongoing cooperation and patience from the lender to adhere to the research design. In the best cases the study will take advantage of limited ability by the MFI to expand, but after some time the program may encounter complaints from field staff as they pass over control villages they might find convenient to serve.

Another important barrier is scale: to establish statistical significance of results (that is, to rule out any improvement in outcomes of the treatment group compared to the control group as being due to mere chance) researchers require a minimum sample size. This is true of any quantitative study, not just RCTs, but it is true that for a randomized evaluation one needs to have a large enough coverage area to create both a treatment and control group, not just a treatment group. Thus, depending on the intervention being evaluated and the design of the program, this can require in some cases thousands of clients to participate in the study, or over one or two hundred villages if the program is randomizing its entry into villages. Small MFIs cannot expand into a few hundred villages in short time periods, and thus randomized evaluations may be unrealistic for these programs. Randomizing loans at the individual level is more realistic for smaller MFIs, for example using credit-scoring techniques discussed below, but this then requires employing particular processes for marketing and approval of loans which may be incongruent with the

MFI's desired policies. Bottom line: even though randomized evaluations provide the cleanest evidence, some situations simply will not allow for them to be conducted.

Two Randomized Evaluation Designs

Here we describe two methods for conducting a randomized evaluation of credit. First, we describe a randomized program placement approach, and second, we describe a randomized individual-lending approach using credit scoring.

Randomizing program placement means that the MFI lays out a list of desirable villages, communities, or urban markets, and then randomizes which to enter and which not to (or to delay entering until the end of the study). This method has the most promise for measuring important concepts such as spillovers to non-participants within a village, but also is the most challenging in terms of cost and sample size. A typical take-up rate for microcredit programs is quite low (this is in its own right an important concern, but another topic). But since the evaluator does not know who within the control villages would have said yes and who would have said no, the evaluation is at the village level, on everyone in the villages (or on everyone in the villages who meet some objective criteria that is not a function of the credit offer process, such as having a microenterprise and being female). Otherwise there would be selection bias: the evaluation cannot compare just those in treatment villages who joined the program to everyone in the control villages, since this would mean there was a self-selection process in the treatment villages that is likely correlated with outcomes, whereas no such self-selection occurred in the control groups. We have typically found that one needs a minimum of at least 100 villages or markets, ideally at least 200, in order to conduct such a study without any individual household-level randomization. Researchers at Innovations for Poverty Action and Jameel Poverty Action Lab are in the midst of several such evaluations.

A second method randomizes at the individual level. This could be done in any setting in which the lender has less money to lend than loan applicants are seeking. Credit scoring can easily incorporate a randomized component in order to facilitate this process. Karlan and Zinman (2006b) employed such a design with an individual consumer-lending program in South Africa. Applicants with credit scores just below the normal cutoff for approval were randomly selected to be reconsidered for loans, or to go into the control group and remain rejected. Despite the lack of any entrepreneurial targeting, the loans led to considerable improvement in income: those randomly selected to be approved for loans were more likely to be employed, less likely to be below the poverty line, less likely to have family members go to bed hungry, and had better credit scores. The evaluation approach also provides operational benefits: the lender learns about its optimal loan-approval rule. The authors find the marginal applicants—those who would have been rejected by the lender—to be profitable clients, despite their higher default rates. Note that this approach need not employ credit scoring *per se*. As long as there are more applicants than there is money to lend, the final decision can be done via lottery in order to create treatment and control groups.

Again, both of these methods require utmost cooperation from the financial institutions, significant enough sample sizes to draw statistically sound results, prior planning by the MFI as to its intended growth or loan-approval process, and of course, significant enough resources to conduct independent surveys to measure outcomes.

The Opportunities for Rigorous Impact Evaluation to Help Improve Programs

Evaluation for Product Innovation

High-quality impact evaluations can allow us to benchmark the social performance of different MFIs and identify program features that excel, or fail, in improving clients' welfare. Exploring why top-performing programs have the impact they do can help policymakers develop and disseminate best practice policies for MFIs to adopt. Impact evaluations may be able to provide insight into which specific practices can cause programs to fail or succeed. In a similar manner, evaluations can provide important information to practitioners and policymakers about the types of products and services that work best for particular types of clients.

This stands in sharp contrast to the way most MFIs test new products or procedures. Many MFIs test new product designs by allowing a few volunteer clients or areas to use a new lending product. Alternatively, an MFI can implement a change throughout an entire branch of its operations. These approaches are risky for two reasons. First, such approaches do not help establish whether the innovation or change *causes* an improvement for the institution (or the client). There may be something about the particular group of clients selected for the trial (MFIs often try out new products on their best clients) that can make it appear more successful than it might be on the rest of their clients; or there might be something about the particular branch chosen for the trial, such as especially experienced staff. Establishing a causal link between the innovation and its impact should be important not only for the MFI implementing the change, but also for policymakers and other MFIs that want to know whether they should implement similar changes.

Second, it can be risky to test new products or processes by trying them out on too few clients or too many clients. Too few, and the MFI risks making inferences based on insufficient information: the outcomes of the test sample may be due to simple chance rather than the innovation itself. Too many clients (e.g., every client in a single branch, or the entire MFI) and the MFI risks rolling out a potentially ineffective product to thousands of clients. Rigorous evaluations optimally balance these risks by piloting innovations on a calculated number of participants: the smallest number of clients that will allow researchers to establish a causal link between the innovation and the change in outcome.

Piloting program innovations is a situation in which impact evaluations, especially randomized control trials, are a win-win proposition: less risky (and hence less costly in the long run) from a business and operations perspective, and optimal from a public goods perspective, in that the lessons learned from establishing these causal links can be disseminated to other MFIs. Impact evaluations are not just about the impact on clients. Program innovations—even those specifically designed to increase client welfare—can have a big effect on the MFI's bottom line, through repayment rates, client retention, and loan and savings balances. Other innovations, such as incentive payments for loan officers, may be targeted at MFI staff.

There are many examples of randomized control trials used to evaluate the impact of a microfinance product or process innovation on the institution. To illustrate the types of questions, we now discuss eight different randomized evaluations that have helped to shed insight into product innovations, and in some cases challenge conventional wisdoms in microfinance.

1. Gine and Yang (2007) conducted a study which randomized whether farmers borrowing for the purchase of higher-risk hybrid seeds were required to purchase weather insurance along with the loan. Though the policy was priced at actuarially fair rates, and rainfall represents the dominant source of risk, take-up was much lower among those offered insurance with the loan (17.6 percent take-up vs. 33% percent among the uninsured group). This may be due to the simple fact that insurance is a more complicated and less familiar product: take-up among the insured group was positively correlated with farmer

education levels. This suggests it may be important to design appropriate marketing or educational interventions for insurance products.

2. Many microfinance institutions grapple with whether to include business training or not in their program. In a randomized evaluation of credit with business education, compared to credit only, FINCA Peru tested the training provided by Freedom from Hunger and Atinchik (Karlan and Valdivia 2006), and found considerable improvements in both institutional outcomes (client repayment and retention) as well as business income for clients (and particularly strong improvements in *bad*-month income; thus the training helped women diversify and thus avoid declines in income during tough months).
3. Third, de Janvry, McIntosh and Sadoulet (2007) investigated the effects of the introduction of a credit bureau for microfinance borrowers in Guatemala. While they could not randomize which clients were subject to the bureau they could choose clients to *educate* about it. With some randomly selected borrowers they conducted an intervention explaining to clients how reporting their repayment performance to other institutions would affect their future options: a bad credit record would limit their ability to borrow while a good record would increase their ability to borrow from commercial banks at lower rates. They find borrowers who are made aware of the credit bureau are less likely to default, but also 10 percent more likely to leave the MFI for another lender.
4. In the Philippines, researchers measured the impact of a new commitment savings product (a specialized savings account with which the client set a savings goal; her money could not be withdrawn until she reached her goal), as well as an accompanying deposit collection service, and compared the savings balances of clients who receive it to clients who already had traditional savings accounts (Ashraf, Karlan and Yin 2006a; Ashraf, Karlan and Yin 2006b; Ashraf, Karlan and Yin 2006c). 28 percent of those offered the account signed up for it, and the treatment group (i.e., everyone offered the product, not just those who took-up) increased their savings held in the bank by over 80 percent relative to the control group.
5. The same evaluation process can be used to provide evidence to practitioners and policymakers on the most basic—and often, the most important—questions about how credit is used, and how it can work best. Much of this work tests the assumptions made by the microfinance industry about how microcredit “should” be done. One especially prominent and contentious issue concerns interest rates: some practitioners and policymakers believe interest rates are too high, while others argue a purely market-based approach is best: firms should generate as much profit as possible and let competition drive prices down for customers. Much of the acrimony around this issue can be traced to a lack of basic evidence: (1) what is the demand for credit at different interest rates? (2) what are microfinance clients’ returns to the capital they borrow?

Generating answers to these questions is extremely difficult outside of a rigorous evaluation. Think about simply asking potential borrowers how much they are willing to pay for credit, and how much they are able to pay. Are they likely to understate the answer? Are they likely to estimate well? The only way to get a sure answer is to go out and make real offers and see if people take them up. In another study with the same consumer credit lender in South Africa mentioned above we did just that: we measured clients’ sensitivity to interest rates by mailing out over 50,000 credit offers with randomly selected interest rates. The results: borrowers were less sensitive to changes in price than expected, a finding which suggests the poor have limited outside options for access to

- finance (Karlan and Zinman 2006c; Karlan and Zinman 2006a). At the same time as we were testing sensitivities, we tested the effectiveness of different marketing approaches on the likelihood that individuals borrowed. We find that some costless marketing approaches, such as presenting only one rather than several loans or including a woman's photo on the mailer, were as effective at increasing demand as dropping the interest rate as much as four percentage points per *month* from an average rate across the sample of 7.9 percent (Bertrand, Karlan, Mullainathan, Shafir and Zinman 2005).
6. Sixth, and closely connected to the pricing issue just discussed: if proceeds are in fact invested in enterprises, how much are clients actually earning on the additional capital in order to justify the high interest rates? If they can make anything more than a 100 percent return on the investment they make with the money they borrow from an MFI then it will be worthwhile for them to pay anything up to 100 percent interest (keeping the difference as profit). But returns to capital are hard to measure because of selection bias; if we look only at returns to microloans we are measuring the profitability of a select group: those who have decided they can earn sufficient returns to make borrowing worthwhile. Likely this group has greater entrepreneurial ability or investment options. This distinction becomes important when practitioners and policymakers explore options for expanding access to finance: what will be the impact of credit on those not yet reached? De Mel, McKenzie et al. (2007a) solved this problem by randomly selecting entrepreneurs and simply *giving* them working capital (either in cash or in a productive asset, again randomly selected). Among entrepreneurs in Sri Lanka they find the average return on the additional capital was 5.7 percent per month. This is substantially higher than the market interest rate, suggesting clients may be able to afford even very high rates. However, in related work (2007b) they find markedly heterogeneous returns: females had close to zero returns on capital on average. Some of this may have been due to occupational choice and traditional gender roles: in industries in which both men and women are owners (i.e., not traditionally male nor female occupations), women have average returns of 3.1 percent, compared to 8.9 percent for men. The clear lesson for practitioners: even if entrepreneurs *on average* can pay high rates, there may be important welfare effects for different types of clients.
 7. Conventional wisdom in microfinance says that group lending helps generate higher repayment rates. However, some believe that it thwarts growth (because many potential clients do not want the tension of the group process) and does not in fact do better than a bank can do on its own. In a study in the Philippines, a rural bank put this foundational question to the test by removing group liability from randomly selected clients, keeping every other aspect of the loan contract the same. They found no measurable drop in repayment, while client retention and new client enrollment increased in the individual-liability centers (Giné and Karlan 2006).
 8. In another study, an MFI in India switched randomly selected borrowers from a weekly payment schedule to a monthly schedule (Field and Pande 2007). Most MFIs use weekly payments because it keeps payment amounts low and weekly meetings enforce discipline among borrowers but they make for high labor costs. Field and Pande found no drop in repayment after the switch.

Returning to the point we made at the opening of this focus note, there is no more basic question about microfinance than *does it work?* Researchers from the Financial Access Initiative, Innovations for Poverty Action, and the Jameel Poverty Action Lab are finally addressing this core question with rigorous evidence through RCTs in urban slums of India, rural areas of

Morocco, the Philippines, and elsewhere. In each of these studies the impact of credit will be cleanly evaluated by randomizing access to credit across entire villages (or slums) through a phase-in design: villages selected to receive credit after a certain period of time will constitute the control group. By comparing entire villages to each other—those who choose to borrow along with those who do not—there is no selection bias. Better yet, this type of evaluation can also capture impact spillovers: increased economic activity or changes in empowerment among non-borrowers from the presence of the microfinance program. The findings from these evaluations will provide much-needed evidence on the value proposition of microfinance for lifting the poor out of poverty.

Conclusion

Although microfinance has much promise, the evidence to date on its impact on well-being and poverty is extremely limited. The obstacles to quality evaluations outlined here make this surprising fact somewhat less surprising. There are no easy answers. Yet it is also clear that rigorous evaluations are both necessary and possible. With the size of global investment in microfinance—and the hopes pinned on its success—there is a mandate to learn how well microfinance works, and how it works best. Scientific rigor can be applied to measuring not only the impact of microfinance, but to answering critical operational questions as well. The ideal impact studies, in fact, do both. By creating more win-win opportunities, where impact research also answers key management questions, microfinance hopefully can push forward into new frontiers. Despite the popularity of microfinance, the majority of the “targeted” population still is not reached, and more rigorous evidence on impacts in general, as well as relative impacts of specific products and processes on clients and institutions, should be able to help donors, investors, policymakers and practitioners break through and deepen the quantity and quality of access to finance around the world.

In this sense good evaluation is about helping practitioners and funders allocate *future* resources most effectively. Programs that want immediate feedback, or cannot meet the requirements for conducting randomized trials, while they will find it difficult to establish definitely what is their impact on their clients, can still improve their operations through monitoring tools. Focus groups, client exit surveys, activity-based costing, etc. can all be useful tools for generating insight into the program’s operations. Nonetheless, as rigorous evaluations grow in popularity, microfinance managers will increasingly turn to findings from RCTs (including studies conducted with other MFIs) to inform their operational decisions for their programs.

REFERENCES

- Ashraf, N., D. Karlan and W. Yin (2006a). "Deposit Collectors." Advances in Economic Analysis & Policy 6(2): Article 5.
- Ashraf, N., D. Karlan and W. Yin (2006b). "Female Empowerment: Further Evidence from a Commitment Savings Product in the Philippines." Yale University Economic Growth Center Discussion Paper 939.
- Ashraf, N., D. Karlan and W. Yin (2006c). "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines." Quarterly Journal of Economics 121(2): 673-697.
- Bertrand, M., D. Karlan, S. Mullainathan, E. Shafir and J. Zinman (2005). What's Psychology Worth? A Field Experiment in Consumer Credit Market. Yale University Economic Growth Center Discussion Paper. 918.
- Coleman, B. (1999). "The Impact of Group Lending in Northeast Thailand." Journal of Development Economics 60: 105-141.
- de Janvry, A., C. McIntosh and E. Sadoulet (2007). "The Supply and Demand Side Impacts of Credit Market Information." working paper.
- de Mel, S., D. McKenzie and C. Woodruff (2007a). "Returns to Capital in Microenterprises: Evidence from a Field Experiment." World Bank Policy Research Working Paper.
- de Mel, S., D. McKenzie and C. Woodruff (2007b). "Who does Microfinance Fail to Reach? Experimental Evidence on Gender and Microenterprise Returns." working paper.
- Field, E. and R. Pande (2007). "Repayment Frequency and Default in Micro-Finance: Evidence from India." working paper.
- Giné, X. and D. Karlan (2006). "Group versus Individual Liability: Evidence from a Field Experiment in the Philippines." Yale University Economic Growth Center working paper 940.
- Giné, X. and D. Yang (2007). "Insurance, Credit, and Technology Adoption: Field Experimental Evidence from Malawi." World Bank Policy Research Working Paper.
- Karlan, D. and M. Valdivia (2006). "Teaching Entrepreneurship: Impact of Business Training on Microfinance Institutions and Clients." Yale University Economic Growth Center working paper.
- Karlan, D. and J. Zinman (2006a). "Credit Elasticities in Less Developed Economies: Implications for Microfinance." working paper.
- Karlan, D. and J. Zinman (2006b). "Expanding Credit Access: Using Randomized Supply Decisions to Estimate the Impacts." working paper.
- Karlan, D. and J. Zinman (2006c). "Observing Unobservables: Identifying Information Asymmetries with a Consumer Credit Field Experiment." working paper.