International Actors' Willingness to Update: Two Global Field Experiments on Microfinance Institutions

Matthew Brigham, Alex Egbert, Michael Findley, Billy Matthias, Chase Petrey, and Daniel Nielson

Corresponding author: dan_nielson@byu.edu

7 November 2015

Prepared for presentation at the 2015 International Political Economy Society Meeting, 13-14 November 2015, Stanford University, Palo Alto, CA.

Abstract

Leading approaches to international relations emphasize the importance of international actors' ability to acquire new information or learn socially. With a global field experiment on 1,419 micro-finance institutions (MFIs), we test non-governmental organizations' propensity to update their behavior when randomly treated with positive or negative information about their field's current practices. Specifically, we test the effects of scientific findings on MFIs' willingness to learn more about microfinance efficacy and pursue an offered partnership to evaluate their programs. In the positive treatment subjects were randomly assigned to receive a summary of a study by prominent authors finding that microcredit is effective. The negative treatment provided information on research – by the same authors using a very similar design – reporting the *in*effectiveness of microcredit. We compare both conditions to a control in which no studies were cited. In the field experiment the positive treatment elicited twice as many responses as the negative treatment, suggesting significant confirmation bias among microfinance institutions. The results suggest that updating in the face of negative information is difficult for NGOs, which may temper some of the more sanguine conclusions about learning in international relations.

Introduction

Multiple prominent approaches to international relations focus on international actors' ability to update beliefs and learn both from other actors and their environment. Fearon's rationalist explanations for war depend on Bayesian models of updating (1995). Simmons' and collaborators' models of policy diffusion also involve Bayesian updating and other forms of learning as key mechanisms (Elkins and Simmons 2004, Simmons et al. 2006). Risse's (2000) account of communicative action relies on a "logic of arguing" that necessitates that actors be open to persuasion, and likewise Checkel's (2001) framework of social learning demands the ideational flexibility of international actors. While the mechanisms are different across the approaches, these scholars all agree that learning is central to their models of global politics.

Yet the conditions under which international actors learn remain understudied and largely untested with empirical approaches that can reveal causal effects. Moreover, while at least the constructivist approaches of Risse and Checkel invoke psychological mechanisms undergirding the process through which actors learn, the full implications of psychology in international learning have yet to be explored theoretically and especially empirically.

In this article we draw on important ideas from social and cognitive psychology and apply them to the international context. In particular, through a global field experiment on non-governmental organizations we explore the potential of cognitive dissonance and confirmation bias to affect the propensity toward updating of microfinance institutions engaged in efforts at poverty alleviation.

To our knowledge this is the first field experiment to probe causes of updating by international actors generally and by non-governmental organizations in particular. While cognitive dissonance and confirmation bias are well-known phenomena in psychology

experiments with individuals as subjects, prominent international relations theories imply that organizational actors may employ practices that enable smoother updating that is less prone to bias. Thus, this study unpacks key assumptions of models premised on updating in international relations. Through a global field experiment we probe how credible scientific information might affect microfinance institutions' willingness to learn more about and potentially participate in a partnership to evaluate their effectiveness through a randomized impact evaluation.

Randomized evaluation has swept through the international development community, energizing anti-poverty scholarship and practice with the promise of learning the precise causal effects of interventions in foreign aid and private humanitarian efforts. Affiliates of MIT's Jameel Poverty Action Lab, led by economists Abhijit Banerjee and Esther Duflo, have completed and reported the results of 681 randomized control trials (RCTs) in development as of November 2015 (J-PAL 2015). They add more to the list every month. Official development organizations have joined the movement. Indeed, Banerjee and Duflo reported in 2009 that the World Bank had 67 RCTs under way out of a total of 89 program evaluations in the Africa region alone (152).

As field experimenters in anti-poverty and conflict resolution, we celebrate the success of randomized evaluation in motivating large improvements in learning what works in development (see Cohen and Easterly 2009, Banerjee and Duflo 2009). Because anti-poverty programs are interventions by their very nature, evaluators can test their effects rigorously with similar methods to those that have transformed medicine from quackery into a science that saves billions of lives. By assigning interventions to treatment and control groups, researchers can learn the causal effects of the projects and, by replication, accumulate knowledge of effective development practice in which we can place high confidence.

A logical step both in the international relations and the randomized evaluation research programs requires that we rigorously test the willingness of international actors to learn from new knowledge. In the case of anti-poverty programs, as in many human endeavors, development community members have great confidence in their current practices. Their methods make intuitive sense to them, and if their practices are generally followed by many others, the programs may seem "correct," "right," or even "moral" in a normative sense. Contrary to the stance of openness to updating that some prominent international relations models seem to presume, practitioners may resist or ignore evidence that contradicts their prior beliefs. Courting new ideas may defy common sense and feelings of moral obligation. The irony here, of course, is that the goal of the development community is not the perpetuation of current practices but the relief of poverty. Thus, more than in many endeavors, people engaged in anti-poverty efforts ought to be open to information about ways to achieve their goal more effectively. But are they?

We currently have very little information about how open or averse international actors generally, and development organizations in particular, might prove to new knowledge. If the aversion to learning is significant, then prominent IR theories may need to be recalibrated to incorporate less fluid mechanisms for updating. Moreover, learning aversion also suggests that the new wave of development scholarship faces the additional challenge of persuading a resistant target audience of the value of the new knowledge. The present study pursues these questions with a field experiment in which development organizations serve as subjects.

We selected MFIs as subjects both because of the prominence of microfinance's boosters as well as the quality of the randomized control trials evaluating its effectiveness. Cheerleaders for microfinance, such as Nobel Peace Prize Winner and Grameen Bank founder Muhammad

Yunus, have touted small loans to the very poor as the answer to many development problems, including missing labor markets, lack of women's empowerment, limited education opportunities, and poor public health (Yunus 2007).

High-quality randomized evaluations, however, suggest that microcredit can be very helpful in providing capital to entrepreneurs and causing business startups, but it can also induce high indebtedness and may have no effects on women's empowerment, education, or health (Banerjee et al. 2009). Thus, the disconnect between practitioners' beliefs and scholars' current findings creates an opportunity to probe the willingness of development NGOs to update.

We therefore sent sincere offers by email to 1,419 microfinance institutions in Experiment 1 and to 4,375 MFIs in Experiment 2. As affiliates with a development research lab, we are actively seeking partners in many areas of international development with which we might undertake randomized evaluations of their programs. The emails did not offer immediate partnership but instead emphasized current partnership commitments and the need for future funding premised on availability and mutual interest. The emails concluded with an invitation for the MFIs to receive additional information both about studies of microfinance and regarding a potential future partnership with our lab to perform a randomized evaluation. The offer was part of an active effort to recruit potential partners and thus involved no active deception (beyond withholding the knowledge that the organizations were part of a field experiment).

In both experiments we included two treatment conditions and a control. The control condition email introduced our academic organization and offered additional information about randomized evaluation and a potential partnership. The positive condition augmented the control email with a paragraph summarizing the findings from prominent development economists finding positive effects from microfinance. The negative condition also began with the control

email language but added a paragraph summarizing findings from a different study by the same prominent authors finding that a microfinance program produced null effects. The positive condition elicited twice as many requests for additional information as the negative condition in Experiment 1, suggesting significant confirmation bias on the part of microfinance institutions and marking a major challenge for randomized evaluators in persuading development organizations to update their practices. Fearing that the stark wording and projected bias of the treatment language in Experiment 1 may have confounded results, we softened the treatment language in Experiment 2 to make the reports of the research findings more tentative and to suggest that different MFIs may achieve different results than those we relayed. The softened language appeared to neutralize the positive condition, but the negative condition still received significantly fewer invitation acceptances compared to the placebo and also significantly more rejections. The results present revealing evidence on the openness of international actors to new information and thus reflect on important arguments about updating and learning in international relations.

Background and Literature

Several of the most prominent literatures in international relations rest upon the ability of international actors to update based on new information and otherwise learn from counterparts, NGOs, and the global environment. These literatures span the three traditionally dominant approaches to IR: neo-realism, neo-liberalism, and constructivism. While of course the arguments about the specific mechanisms involved in learning vary from author to author, learning proves absolutely central to these leading arguments.

First, Fearon's rationalist models of war are explicitly premised upon Bayesian models of updating (1995). War results when the typical methods of acquiring information about possible opponents are short-circuited by asymmetric information and actor dissembling. The strong implication of the models is that, while war can be fully rational, it is out-of-equilibrium behavior – and, indeed, appropriate updating likely enables most potential belligerents to negotiate solutions to conflict that stop short of violence.

Second, in considering the underlying mechanisms for the diffusion of liberal political and economic practices, Elkins and Simmons (2004) and Simmons, Dobbins, and Garrett (2006) underscore the potential import of learning by international actors. This learning can occur in the standard Bayesian way, or it might also reflect the acquisition of social knowledge or the channeling of new policy ideas through communications networks. Regardless of the specific mechanism involved (which are difficult to tease out through observational techniques but may be more amenable to random assignment in experiments), the key point is that updating either causal beliefs or ultimate goals can have large international relations effects, and Elkins and Simmons (2004) and Elkins, Guzman and Simmons (2006) find evidence consistent with the learning hypothesis in both the diffusion of economic liberalism and the spread of bilateral investment treaties.

Third and finally, constructivists, especially Risse (2000) and Checkel (2000) advance a Habermasian logic of "communicative action" or "argumentative persuasion." In these models, international actors are not merely rational instrumentalists driven by ordered goals or even constructivist norm-abiders acting unconsciously according to taken-for-granted cultural scripts. Rather, international actors – perhaps especially NGOs – actively learn from each other through argument and persuasion where they remain open to updating (see also Risse 2004, Checkel

2001). The openness of NGOs to updating thus forms an important component of communicative action, but its scope and limits remain largely undefined and untested scientifically.

This article thus probes the propensity of international actors and specifically NGOs to update their practices. However, in the substantive domain of anti-poverty programs, the particular history of program evaluation suggests caution here. For decades aid agencies have claimed success rates for all projects ranging from two-thirds to four-fifths (Picciotto 2012, Faiola 2009). Traditional program evaluation involves monitoring the outputs of projects and comparing them to initial goals, which presents a particularly low bar. If program plans state objectives explicitly, say, constructing so many miles of paved road, it should be relatively unproblematic to provide the planned outputs. The new road can be observed and measured accordingly. What is more, the aid agency personnel who produce the monitoring data are often the same people who designed the project in the first place, and both career incentives and confirmation bias likely influence how they report results. Hence, very high success rates for projects naturally follow from such unscientific evaluation.

In the early 2000s, MIT's newly established Jameel Poverty Action Lab (J-PAL) led the charge in arguing that experimental methods provided the most effective way to approach impact evaluation. Esther Duflo, co-founder of J-PAL, stated at a World Bank Conference on evaluation and development effectiveness in 2003 that "Just as randomized trials for pharmaceuticals revolutionized medicine in the 20th century, randomized evaluations have the potential to revolutionize social policy during the 21st" (Duflo and Kremer 2004). Proponents tout the main virtue of randomized evaluations: due to the close collaboration between researchers and practitioners, RCTs allow the estimation of causal effects – the actual impact of projects – that would not otherwise be possible to evaluate (Duflo and Kremer 2004, Banerjee and Duflo 2009).

These claims have proven compelling to many, so randomized field experiments have become a popular tool in development economics research and have found increasing purchase in development practice. As noted above, in the Africa region alone the World Bank in 2009 was performing RCTs on 67 of 89 (or 75 percent of) program evaluations. The Development Impact Evaluation Initiative at the World Bank, which routinely employs randomization, covers 13 percent of the joint IBRD-IDA portfolio of the Bank (Legovini 2010). And this proportion appears to be growing.

In January of 2011 Rajiv Shah, Director of the U.S. Agency for International Development (USAID) announced a major overhaul of the agency's monitoring and evaluation practices. The new policy mandates that all programs be evaluated by third parties reporting directly to USAID (not to project contractors) and requires that all "innovative" programs employing "untested" hypotheses undergo randomized impact evaluation (USAID 2010). These evaluation initiatives by the world's two largest aid organizations suggest that RCTs have broken out of the academic cloister and have captured the attention – and the resources – of important development practitioners.

But randomized evaluation has met with skepticism in the academy. Prominent development economists have questioned both the external validity and theoretical grounding of randomized evaluations (Rodrik 2008, Deaton 2010). Others have openly worried about the perceived high cost of RCTs (Copestake et al. 2009). And yet others point out that RCTs cannot answer many critical questions, including some of the biggest. Writes Avril Subramanian, "What would be the effects of disbursing $1-1.5 billion of foreign aid to Pakistan? RCTs do not, and cannot, have anything to say on the matter – not only because of their narrow focus and applicability, and hence non-generalizability, but also because they cannot speak to

macroeconomic effects. The larger developmental effects of aid may be good or bad but RCTs cannot help us distinguish them" (Subramanian 2011).

Advocates of randomization have generally acknowledged these issues. They have answered that problems of external validity can be addressed through systematic replication of experiments in diverse settings. They have granted that experiments should test discrete causal mechanisms derived from sound theory. They recommend that evaluation costs be built into development projects up front. And they admit that RCTs cannot answer many important questions in development (Karlan 2009). This back and forth between "randomistas" and their critics has proven generally helpful in focusing and refining the practice of randomized evaluation.

The present article, however, addresses an additional – and potentially bigger – problem faced by proponents of randomized evaluation: practitioners' potential unwillingness to accept the results of the studies and update their operations. The topic area of this study, microfinance, perhaps best illustrates the challenges involved in motivating development practitioners to open their minds to scientific findings and change their procedures accordingly. Some of the best designed and most persuasive RCTs in development economics have put microfinance to the test, and the results suggest that microfinance significantly improves entrepreneurs' access to credit and therefore provides an important tool in overcoming poverty (Banerjee et al. 2009, Karlan and Zinman 2010). Even where microfinance fails in its primary goals of income-generation or empowerment of women, it may have ancillary benefits in strengthening community ties, helping borrowers cope with risk, and improving informal credit access (Karlan and Zinman 2011). Scholars performing the studies clearly see microfinance as providing part of the answer to the development puzzle.

But part of the answer is insufficient for the advocates of microcredit. Rather, microcredit has been advanced as a panacea for a panoply of problems in developing countries. Most microfinance institutions organize (predominantly female) borrowers into solidarity groups, which meet together often to repay loans and apply for new financing. Access to small amounts of capital purportedly allows these groups of poor women to invest in their small businesses and generate new sources of income enabling them to lift themselves out of the poverty trap while addressing many other problems of poverty, including poor healthcare, lack of access to education, and discrimination against women. In the thirty-seven years since Bangladeshi economist Muhammad Yunus started the Grameen Bank, thousands of MFIs around the world have been created to join in the effort to alleviate poverty through small loans to the very poor.

In 2006 the Norwegian Nobel Committee awarded the Nobel Peace Prize to Yunus and his Grameen Bank "for their efforts to create economic and social development from below" (Mjøs 2006). In his presentation speech at the Nobel award ceremony, Nobel Committee Chairman Ole Danbolt Mjøs extolled the broad scope of microfinance, which clearly factored into the award decision: "The [female borrower] group meets regularly to sharpen each other's perceptions of borrowing, work, repayment and saving. The members undertake to work for food production, pure drinking water, hygiene, health, family planning, economy, discipline, community and motivation in the group and in their families. The groups form networks with other groups. At the grass-roots level the groups thus help to build up communities." In particular, Mjøs praised Yunus' and Grameen's focus on women: "Micro-credit has proved itself to be a liberating force in societies where women in particular have to struggle against repressive social and economic conditions. Economic growth and political democracy cannot achieve their

full potential unless the female half of humanity on earth contributes on an equal footing with the male" (Mjøs 2006).

Yunus himself has done much to reinforce this impression of the broad impact of microfinance. For example, in Yunus' book, *Banker to the Poor*, he notes that "Grameen is a private-sector self-help bank, and as its members gain personal wealth they acquire water-pumps, latrines, housing, education, access to health care, and so on" (2007, 203). Later, he writes, "Grameen is committed to social objectives: eliminating poverty; providing education, health care, and employment opportunities to the poor; achieving gender equality through the empowerment of women; ensuring the well-being of the elderly" (2007, 209-210). Thus, the claims for the impact of microfinance are quite broad.

As noted above, development economists employing randomized evaluation put these claims of broad scope to the test in a series of studies. The findings were mixed. One study, which we used in our experimental intervention, found strong treatment effects across a wide range of positive outcomes for a microfinance program in South Africa. Access to microcredit caused improvements in economic self-sufficiency, consumption possibilities, and an index measuring subjects' self-reported perceptions of control and positive outlook – including womens' sense of empowerment in their households (Karlan and Zinman 2010). But yet another study employing a similar design by the same authors, which we also used in the experiment, failed to replicate these findings in the Philippines, though as noted it did recover treatment effects for improving community trust, coping with risk, and access to informal credit (Karlan and Zinman 2011). Also as noted, a major study conducted by J-PAL scholars in India found that microcredit improved entrepreneurs' investment in durable goods and that the number of new businesses in treatment neighborhoods increased by one third. This is strong evidence that

microfinance has positive effects. However, the study also showed that microfinance only increased consumption of non-durables (and therefore consumer debt) for people not inclined to business ownership, and it had no effect on health, education, or female empowerment (Banerjee et al. 2009).

After these results first became public, representatives of the six largest MFIs worldwide assumed that all the results would be negative (they were not) and reacted by producing six anecdotes of successful borrowers (Banerjee and Duflo 2011). Brigit Helms, CEO of Unitus, an international MFI, declared in a *Seattle Times* op-ed, "These studies are giving the inaccurate impression that increasing access to basic financial services has no real benefit…. Our worry is that if these studies can't empirically demonstrate significant economic impact in a short time period, the public will be left with the impression that microfinance has no value – especially dangerous at the exact moment microfinance is poised to do more than ever to alleviate global poverty." (Helms 2010). However, the randomized experimental studies do not show negative results, they show mixed results (Banerjee et al. 2009; Dupas and Robinson 2009; Karlan and Zinman 2010; Pitt, Khandker, and Cartwright 2006; Holvoet 2005; Garikipati 2008; Rahman 1999). Taken as a whole, the studies merely suggest that microfinance may be overhyped.

In reaction, the Grameen Bank published an article in which it surveyed the evaluation literature relating to microfinance over the past 20 years. Its conclusions focus on the positive effect that microfinance has and downplays negative findings as limitations to current research methods. The report criticizes RCTs as being under-contextualized short-term evaluations that do not capture the robustness of microfinance's impact (Odell 2010). Rather than accept the evidence showing some positive results while providing a corrective to some of the grander claims of the microfinance movement, microfinance advocates counter-attacked, leaving the

authors of the microfinance studies somewhat baffled. In their general-market book, *Poor Economics*, Banerjee and Duflo describe their experience in the aftermath of their landmark study.

> "As economists, we were quite pleased with these results: The main objective of microfinance seemed to have been achieved. It was not miraculous, but it was working. There needed to be more studies to make sure that this was not some fluke, and it would be important to see how things panned out in the long run, but so far, so good. In our minds, microcredit has earned its rightful place as one of the key instruments in the fight against poverty.
>
> "Interestingly, this is not how the main results played out in the media and the blogosphere. The results were mainly quoted for the negative findings and as proof that microfinance was not what it was made out to be. And though some MFIs accepted the results for what they were (chief among them, Padmaja Reddy [head of the MFI that partnered in the study], who said this was exactly what she had expected, and financed a second wave of the work to study the longer-term impacts), the big international players in microfinance decided to go on the offensive" (Banerjee and Duflo 2011).

Hypothesis

The present study follows up on this anecdote to learn if the reaction Banerjee and Duflo described is systematic and widespread in the microfinance community. Our intuition that MFIs may be slow to update and therefore averse to the possibility of negative results draws on important research from social psychology. Specifically, confirmation bias may encourage MFIs to accept what they already believe and resist what they do not believe in an effort to avoid or resolve cognitive dissonance (Festinger 1957).

Cognitive dissonance is the psychological discomfort an individual feels when presented with information that runs contrary to previously held beliefs (Stone and Cooper 2001; Steele and Liu 1983; Aronson 1969; Festinger 1957). Steele and Liu explain that previously held

beliefs are an individual's "ideal self-image." For this reason, discomfort occurs when presented with information contrary to prior convictions. In the case of MFIs, the "ideal self-image" is that microcredit is a powerful poverty alleviation tool with general effects for a wide range of outcomes.

Confirmation bias clouds the judgment of human beings. Humans are wired to believe chiefly what they want to believe, and what they want to believe rests heavily on priors. For example, in one classic social psychology experiment, both pro-Arab and pro-Israeli citizens interpreted the same media broadcast as being biased against their side (Vallone et al. 1985). In a related experiment, a team of neuroscientists observed the neural responses of Republican and Democratic voters. Each group watched positive and negative campaign ads for candidates George W. Bush and John Kerry before the 2004 presidential election. The neuroscientists found that when subjects confronted a negative campaign ad for their preferred candidate, the region of the brain responsible for reasoning deactivated. When subjects viewed a positive ad, their emotional brains lit up (Westen et al. 2006).

When people are presented with information consistent with prior beliefs, no cognitive dissonance occurs and confirmation bias causes them to readily accept the new information. However, when presented with information inconsistent with prior beliefs, cognitive dissonance occurs and confirmation bias causes them to ignore or minimize the new information. Aronson (1969) describes cognitive dissonance by using Festinger's (1957) real-world example of smoking:

> If a person believes that cigarette smoking causes cancer and simultaneously knows the he himself smoke cigarettes, he experiences dissonance. Assuming that the person would rather not have cancer, his cognition "I smoke cigarettes" is psychologically inconsistent with his cognition "Cigarette smoking produces cancer."

> Perhaps the most efficient way to reduce dissonance in such a situation is to stop smoking. But, as many of us have discovered, this is by no means easy. Thus, a person will usually work on the other cognition. There are several ways in which a person can make cigarette smoking seem less absurd. He might belittle the evidence linking cigarette smoking to cancer ... or he might associate with other cigarette smokers ... or he might smoke filter-tipped cigarettes and delude himself that the filter traps the cancer-producing materials; or he might convince himself that smoking is an important ... activity…. All of these behaviors reduce dissonance. (Aronson 1969).

Although confirmation bias is a well-documented shortcoming in human decision-making, its presence in non-profit organizations, such as MFIs, is not yet known. One might hope that anti-poverty organizations have developed organizational routines to maximize learning and minimize bias. After all, charitable organizations focus on poverty relief as their primary goal, and any information that might help them achieve that objective ought to be privileged.

We fear, however, that organizational routines are created by the same individuals prone to cognitive biases in the first place, so confirmation bias may be built into or even reinforced by organizational structures. Thus, we hypothesize that MFIs confronted with negative evidence of microfinance's effectiveness should be less willing to request scientific material on microfinance effectiveness or pursue additional information about a possible partnership in a randomized evaluation.

Research Design

We executed Experiment 1 in 2011 on 1,419 microfinance institutions worldwide and Experiment 2 in 2012 on 4,375 MFIs globally. Because the experiments were conducted by

email, we carried them out from our organization [ORGANIZATION INFORMATION REMOVED FOR REVIEW] as detailed below.

Subject Pool and Randomization

Thousands of microfinance institutions spend billions of dollars annually and, while they generate significant notice, unfortunately no standard sampling frame exists for such institutions. For Experiment 1 we employed a data source, *Mixmarket.org*, that captures a very large number of MFIs, however. The listings included email addresses and demographic information. But we note that because we are using email to apply our treatments, we are limited to MFIs who have access to the internet. This should bias the sample towards larger, more established groups that likely carry out the majority of microfinance work. For Experiment 2 we made use of the listings and contact information in the *Directory of Development Organizations* and targeted the roughly 4,400 organizations listed as focusing on microfinance in their work.

However, we note that both *Mixmarket.org* and the *Directory of Development Organizations* likely capture MFIs that are already transparent. Because these MFIs have already provided financial information and agreed to be listed, they may not fully represent the many other MFIs of smaller scope and scale, and thus this raises the concern of external validity. The results may therefore only extend to MFIs whose operations are relatively extensive and transparent. The samples we obtained from *Mixmarket.org* and the *Directory of Development Organizations* represent a variety of regions worldwide as Table 1 below demonstrates.

Table 1: Regional Distribution of MFIs

	Experiment 1		Experiment 2	
Region	Freq.	Percent	Freq.	Percent

Africa	290	20	1,142	26
East Asia and Pacific	177	13	818	19
Eastern Europe and Central Asia	294	21	525	12
Latin America and The Caribbean	372	26	809	19
Middle East and North Africa	73	5	116	3
South Asia	213	15	496	11
Western Europe			469	11
Total	1,419	100	4,375	100

Despite some potential selection bias in the samples, randomization should allow us to uncover unbiased causal effects for the contacted organizations. We employed block randomization of subjects using region of the world and size of the MFI loan portfolio for Experiment 1 and world region for Experiment 2 (other information was unavailable in the directory) to demarcate the blocking strata.

Experimental Conditions

Within each stratum we randomly assigned each of the MFI subjects in both experiments to either the control condition or to one of two treatment conditions. The premise of the experiments was to provide information about the effectiveness or ineffectiveness of microfinance and then address whether MFIs are more or less likely to internalize and react to such information. Rather than simply ask them whether they agree with the information, we hoped to elicit a revealed preference through a behavioral response. To accomplish this, we embedded the information about the success or failure of microfinance into an invitation to

consider partnering on an impact evaluation for one of their projects. We expected that if they received positive information, they would be more likely to accept the invitation; if they received negative information, then we expected them to be less likely to accept the invitation.

Email invitations to partner on the evaluation were identical in wording, except for a single paragraph in which we introduce the two information variants for the positive and negative treatments. See the appendix for the complete language included in the experimental conditions for both Experiment 1 and Experiment 2. Since our research lab is actively recruiting partners with which to undertake randomized evaluations, the invitation was sincere and therefore involved no active deception. We have followed up with all organizations that answered our invitation and provided additional information as promised. The emails began with a short introduction to our organization as well as an invitation worded as follows, including a short statement confirming the country in which they operate:

> "We are seeking to assess the interest of qualified microfinance institutions in possible partnerships to perform impact evaluations. We understand that you provide microcredit loans in <country>."

The positive treatment included a priming statement claiming that scientific studies indicate that microfinance is effective:

> "Academic research suggests that microfinance is effective. The results of a recent scientific study show that microcredit loans have a positive effect on economic self-sufficiency and subjective well-being of borrowers, including the decision making power women have in the home (Karlan and Zinman 2009, "Expanding Credit Access," *Review of Financial Studies*). These results are compelling to us, and we wish to learn more so we can further assist those in need."

This statement was designed to signal that our organization subscribes to the idea that microfinance is effective. By citing a published study from prominent authors in the area of microfinance, we also tried to signal that the results and possibly consequences of impact evaluation were not trivial.

The negative treatment is identical to the positive treatment but suggests that microfinance is not effective:

> "Academic research suggests that microfinance is ineffective. The results of a recent scientific study show that microcredit loans have no effect on business growth and subjective well-being, nor are there disproportionate benefits in targeting women with microcredit loans (Karlan and Zinman 2011, "Microcredit in Theory and Practice," *Science*). These results are compelling to us, and we wish to learn more so we can further assist those in need."

The overall email and much of the treatment language is identical across all conditions. We cite the same authors who report different conclusions in two separate studies, which more tightly controls the possible confounding factors between conditions. The control email omitted the critical middle paragraph that contained the treatment language suggesting either the effectiveness or ineffectiveness of microfinance.

For Experiment 2 we retained the same basic design but made some important modifications. Fearing that the long approach email may have depressed response rates in Experiment 1, we shortened the text significantly. More importantly, we worried that the sharply worded language in the treatments for Experiment 1 may have confounded the treatment effects. After all, if receiving the negative treatment, why would a microfinance institution wish to work with a research organization whose director was already asserting that microfinance was ineffective? In Experiment 2 we therefore sought to soften the treatment language to make room

for other outcomes. We did this in two ways. First, rather than assert that microfinance is effective or ineffective, we instead noted that "Academic research suggests that microfinance may be effective" or "may be ineffective." Second, after the paragraph citing the scientific results, we added a key qualifying statement: "These findings are interesting, but microfinance institutions vary, so you may want to know your program's particular results." This was meant to suggest that as an evaluation organization we were open to findings different than those reported in the short citation to prior research. See the appendix for exact wording of the experimental emails.

Additional Email Protocol

First round emails were given the subject "Potential Partnership"; second and third round emails were given the subject "Potential Partnership Reminder". Emails were addressed to the MFI's legal name or abbreviation. We sent all emails from an email address associated with a university [EMAIL REMOVED FOR REVIEW PURPOSES]. By sending it from an educational address, we hoped to increase the validity of our invitation. In addition to the email address, we signed all of the emails in the name of the director of our organization [DIRECTOR NAME REMOVED FOR REVIEW PURPOSES]; we anticipated that the emails were most likely to be opened if the sender was an individual. The emails were signed accordingly with a full electronic signature including the director's professional title, address, and a link to the organization's website (see the Appendix for the full text of the emails).

Because we did not always receive responses to the first inquiry, we followed up only with those providers that did not respond. We followed up only twice with all MFIs before coding them as non-respondents. The reminder emails were prefaced with the following text:

"I sent an email <date email was sent> regarding a potential partnership. Below is a copy

of the email I sent that day. I just wanted to confirm you received the invitation."

To enable more controlled execution of the experiment, we sent all the emails from a proxy server. This approach allowed the timing to be consistent and also enabled us to receive confirmation that all of the emails had been sent. For copies of all the emails see the Appendix.

We conducted a randomization check to consider whether other observable factors were balanced equally across conditions. This check demonstrates that the randomization occurred as expected. Only one of nine variables (age of organization) was related to treatment assignment in Experiment 1, and when this covariate and others are included in robustness checks using regression, the results are substantively similar. See the appendix.

Estimation Strategy

Because we randomized the assignment of experimental conditions, we can employ simple difference-in-means tests. In expectation, all other covariates should be balanced across conditions, thus allowing us to identify treatment effects without bias. Three outcomes are possible: non-response, decline, and accept. Of course, the latter two possibilities are contingent on receiving a response.

Difference-in-means tests should be sufficient to probe whether one or both of the treatments alter the effect relative to the placebo. The randomized design in expectation balances all observable and unobservable covariates. Nevertheless, we estimated a multinomial probit model, separate logit models, as well as a selection model to check robustness. The results of the sensitivity analysis are reported in the appendix. We consider each of the possibilities on the response, decline, and accept outcomes. And finally the selection model allows us to incorporate

response and outcome into the same category. Because we do not have additional information with which to identify the separate stages, we use the model designed by Sartori (2003). As we will highlight below, the results are consistent across model specifications.

Results

Table 2 reports the results of the basic comparisons across treatment and control conditions for the response, decline, and accept outcomes for Experiment 1. We report the numbers of observations in each category, the percentages, as well as p-values indicating statistical significance with significant values in boldface.

We note here that the response rates of 6.06 to 10.44 percent might appear low at first blush. However, these rates of response are roughly equivalent to the response rates public opinion researchers typically achieve in surveys (and higher than many). Given that the emails represented unsolicited "cold calls" from an unknown development lab, the rate of response might have been expected to be even lower. Still, the low response rates may reflect to a degree on the generalizability of the results – particularly compared to what might have been achieved by a better-known development evaluator. Nevertheless, we emphasize here that the baseline propensity to respond, which encompasses all observable and unobservable factors, was balanced across conditions in expectation, so response rates reflect treatment effects.

On basic response rates, there is no statistically meaningful difference between the positive prompt and the control. Although there was a higher response rate based on the positive message (10.44% vs. 8.37%), the result is not statistically significant. Similarly, there is no statistical difference between the negative prompt and the control (6.06% vs. 8.37%). It is interesting that fewer MFIs responded when faced with negative information. The differences

between the treatment and control are not significant statistically, however. Interestingly, there is a strong and statistically meaningful difference between the positive and negative treatments (10.44% and 6.06%; $p = 0.015$).

The next column shows that there were very few MFIs that declined the invitation to receive additional information on a partnership for an impact evaluation. This results holds regardless of the condition (control = 4, positive = 3, negative = 5). None of these differences across experimental conditions even come close to standard levels of statistical significance. The results that begin to emerge from the response category are thus not reflected in the decline outcome, but rather in the acceptance.

The numbers and proportions of acceptances do appear to change in response to the experimental condition. The difference between the positive treatment and control (9.81% vs. 7.53%) is not statistically significant ($p = 0.210$), but the differences between the negative treatment and control (4.98% vs. 7.53%; $p = 0.107$), as well as the positive and negative treatments (9.81% vs. 4.98%; $p = 0.005$), are significantly below or close to conventional levels.

Table 2: Contingency Table of Outcomes across Experiment 1 Conditions

Condition	*N*	Response	Decline	Accept
Control	478	40	4	36
Proportion		8.37%	0.84%	7.53%
Positive	479	50	3	47
Proportion		10.44%	0.63%	9.81%
P-value vs. Control		0.273	0.703	0.210
Negative	462	28	5	23
Proportion		6.06%	1.08%	4.98%
P-value vs. Control		0.173	0.700	0.107
P-value vs. Control		**0.015**	0.447	**0.005**
Total	1,419	118	12	106
		8.32%	0.85%	7.47%

In Experiment 2 with the softened language the results display some key differences from Experiment 1, as reported in Table 3. Response rates were a bit lower than in the first experiment, and we no longer see significant differences between the positive and negative conditions. However, while the proportion responding the invitation emails was comparable between the control and treatments, and there is no statistically significant difference between the positive treatment and control in either the rate of acceptance of rejection of the invitation, there are small but statistically significant differences between the control condition and the negative treatment in both the proportion of the MFIs that declined and accepted the invitation for an evaluation. The proportion declining increased from 3.5 in control to 4.9 percent for the negative treatment (p = .067) and the proportion accepting the invitation decreased from 3.1 percent in control to 2.1 percent in the negative treatment (p = .064). While these are small differences in

terms of percentage points, they are potentially large substantive alterations from baseline, representing changes of roughly one third.

The deliberate changes in wording between the two experiments appear to have resulted in substantive changes in the findings. While the intent of the revisions was to soften the potentially harsh and pre-judgmental language in the negative condition, it appears that the alterations to the experimental conditions mostly served to neutralize the positive condition. While the softer wording does not seem to have made MFIs in the negative condition more likely to accept the invitation for a randomized evaluation (or made a rejection of the offer less likely), it does appear to have made MFIs in the positive condition less inclined toward acceptance than in Experiment 1. We are tempted to conclude that MFIs assigned to the positive condition in Experiment 2 were prompted by the new language to consider the possibility that the results, while positive in the evaluation we cited, may not be equally positive for their organization. Both experiments thus suggest evidence of confirmation bias among MFIs.

Table 3: Contingency Table of Outcomes across Experiment 2 Conditions

Condition	*N*	Response	Decline	Accept
Control	1443	95	51	44
Proportion		6.58%	3.53%	3.05%
Positive	1463	95	62	33
Proportion		6.49%	4.24%	2.26%
P-value vs. Control		0.921	0.327	0.183
Negative	1469	101	72	29
Proportion		6.88%	4.9%	1.97%
P-value vs. Control		0.753	**0.067**	**0.064**
P-value vs. Positive		0.679	0.390	0.596
Total	4375	291	185	106

To check robustness, we estimated separate multinomial probit models for each of the experimental conditions. Tables 4 and 5 displays these results. The findings confirm what we learn in the basic difference-in-means tests showing that receiving the negative prompt makes MFIs on average less likely to request additional information on the offered partnership for an impact evaluation than when receiving the control. It also shows a very strong difference between the positive and negative conditions in Experiment 1.

Table 4: Multinomial Probit Table of Outcomes across Experiment 1 Conditions

Treatments	Response	Decline	Accept	*N*
Positive	Base	-0.104	0.198	945
	Base	(0.368)	(0.164)	
Constant	Base	-3.076***	-1.953***	
	Base	(0.287)	(0.143)	
Negative	Base	0.064	**-0.329***	930
	Base	(0.338)	**(0.186)**	
Constant	Base	-3.119***	-1.944***	
	Base	(0.293)	(0.147)	
Pos. vs. Neg.	Base	-0.167	**0.522*****	929
	Base	(0.354)	**(0.179)**	
Constant	Base	-3.125***	-2.365***	
	Base	(0.260)	(0.155)	

Robust standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1
These models include the age variable, but the results are not reported here. It is insignificant in all of the regression models.

Table 5: Multinomial Probit Table of Outcomes across Experiment 2 Conditions

Treatments	Response	Decline	Accept	*N*
Positive	Base	0.101	-0.164	2906
	Base	(0.118)	(0.134)	
Constant	Base	-2.489***	-2.571***	
	Base	(0.086)	(0.090)	
Negative	Base	0.191*	-0.227*	2912
	Base	(0.116)	(0.137)	
Constant	Base	-2.489***	-2.571***	
	Base	(0.086)	(0.090)	
Pos. vs. Neg.	Base	-0.091	0.063	2932
	Base	(0.112)	(0.143)	
Constant	Base	-2.297***	-2.798***	
	Base	(0.077)	(0.103)	

Robust standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1

We also considered the comparisons as a set of logit models on the outcome variables separately. Like the multinomial model, we compared the negative prompt to control, positive prompt to control, and positive prompt to negative prompt, but in the basic logit models we conduct each of these regressions separately. Tables 6 and 7 display the results of these analyses.

Table 6: Experiment 1 Logit Results for Accept, Reject, Response

Variables	***Response***	***Response***	***Response***
Positive vs. Control	0.236		
	(0.225)		
Negative vs. Control		-0.407	
		(0.263)	
Positive vs. Negative			**0.632****
			(0.249)
Age	-0.203	-0.246	0.131
	(0.232)	(0.269)	(0.244)
Constant	-2.296***	-2.278***	-2.817***
	(0.198)	(0.205)	(0.218)
	Accept	***Accept***	***Accept***
Variables	***Offer***	***Offer***	***Offer***
Positive vs. Control	0.284		
	(0.234)		
Negative vs. Control		**-0.508***	
		(0.284)	
Positive vs. Negative			**0.780*****
			(0.267)
Age	-0.166	-0.235	0.142
	(0.240)	(0.287)	(0.257)
Constant	-2.425***	-2.397***	-3.038***
	(0.207)	(0.216)	(0.239)
	Decline	***Decline***	***Decline***
Variables	***Offer***	***Offer***	***Offer***
Positive vs. Control	-0.323		
	(0.775)		
Negative vs. Control		0.236	
		(0.690)	
Positive vs. Negative			-0.548
			(0.735)
Age	-0.593	-0.288	0.028
	(0.849)	(0.727)	(0.735)
Constant	-4.541***	-4.646***	-4.514***
	(0.589)	(0.613)	(0.512)
Observations	945	930	929

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 7: Experiment 2 Logit Results for Accept, Reject, Response

Variables	***Response***	***Response***	***Response***
Positive vs. Control	-0.015		
	(0.150)		
Negative vs. Control		0.047	
		(0.148)	
Positive vs. Negative			-0.061
			(0.148)
Constant	-2.653***	-2.653***	-2.606***
	(0.106)	(0.106)	(0.103)
Variables	***Accept***	***Accept***	***Accept***
Positive vs. Control	-0.310		
	(0.233)		
Negative vs. Control		-0.446*	
		(0.242)	
Positive vs. Negative			0.136
			(0.257)
Constant	-3.459***	-3.459***	-3.905***
	(0.153)	(0.153)	(0.186)
Variables	***Decline***	***Decline***	***Decline***
Positive vs. Control	0.189		
	(0.193)		
Negative vs. Control		0.341*	
		(0.187)	
Positive vs. Control			-0.152
			(0.177)
Constant	-3.307***	-3.307***	-2.965***
	(0.143)	(0.143)	(0.121)
Observations	2906	2912	2932

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The results are consistent with the previous findings reported in Tables 2 and 3. The significant results are again bolded and show robustness across specifications. As one final

check, we consider the possibility that accepting or declining is a two-stage process. We therefore estimate a selection model in which assignment to treatment is the independent variable generating response as well as outcome (decline or accept). Because we do not have enough information to satisfy an exclusion restriction, we use the selection model developed by Sartori (2003). These results (reported in Tables 8 and 9) are largely consistent with those reported earlier. The effect of the negative condition on acceptance is in the same direction but not quite significant at conventional levels ($p = 0.107$) for Experiment 1. That result is consistent with the difference-in-means results reported above. The results for Experiment 2 in Table 9 are consistent with the results above.

Table 8: Experiment 1 Selection Model of Response and Accept/Decline Outcome

Treatments	Response	Accept	Resp. Constant	Out. Constant	N
Pos. v. Control	0.124	0.145	-1.381***	-1.437***	957
	(0.113)	(0.115)	(0.082)	(0.085)	
Neg. v. Control	-0.169	-0.210	-1.381***	-1.437***	940
	(0.124)	(0.130)	(0.082)	(0.085)	
Pos. v. Neg.	0.293**	0.355***	-1.550***	-1.647***	941
	(0.120)	(0.126)	(0.092)	(0.098)	

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 9: Experiment 2 Selection Model of Response and Accept/Decline Outcome

Treatments	Response	Accept	Resp. Constant	Out. Constant	N
Pos. v. Control	-0.007	-0.130	-1.508***	-1.874***	2906
	(0.072)	(0.098)	(0.051)	(0.066)	
Neg. v. Control	0.022	-0.186*	-1.508***	-1.874***	2912
	(0.071)	(0.100)	(0.051)	(0.066)	
Pos. v. Neg.	-0.029	0.056	-1.485***	-2.059***	2932
	(0.071)	(0.105)	(0.050)	(0.076)	

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Conclusion

We hypothesized that MFIs would be less willing to accept our invitation in response to the negative treatment emails than to the control or the positive treatment. Although we do not explicitly test the causal mechanisms in this study, the results are at least consistent with the conjecture that organizations engage in significant confirmation bias when confronted with new information about their field.

Giving MFIs information on the effectiveness of microfinance could have reinforced their belief in the industry, raised questions about MFI efficacy, or had no effect. Since we assume MFIs believe in their cause, providing an MFI with unqualified positive scientific information on microfinance (as in Experiment 1) appears to reinforce previously held beliefs. The MFIs on average seemed to engage in confirmation bias by agreeing with the content of the new information; they were more likely to respond favorably to receiving additional information on a possible impact evaluation partnership. However, when the positive findings were qualified and MFIs were prompted with the possibility that their organization may achieve disparate results, the positive boost toward accepting the invitation seems to disappear in Experiment 2. However, the MFIs that received negative scientific information on microfinance in both experiments appeared to react consistently with expectations of confirmation bias. The evidence is consistent with the proposition that the information that microfinance is ineffective ran contrary to the MFIs staff members' previously held beliefs in a way that induced cognitive dissonance and either lower inclination to accept the invitation (in Experiment 1) or both higher

propensity to decline the invitation coupled with reduced probability of acceptance (in Experiment 2).

When representatives of an organization experience cognitive dissonance, they could either be open to updating their methods or they could rationalize their organizational behavior. The results of this study suggest that MFI staffers assigned to the negative treatment may be significantly less interested in updating. It seems that MFI representatives in the negative condition were prone to reject invitations associated with the information causing dissonance.

While we are fascinated by the results, we are also disappointed by their implications. Multiple studies in social psychology show that humans are susceptible to confirmation bias (Westen et al. 2006; Vallone et al. 1985). However, we hoped that MFIs' organizational structure would transcend this human tendency, especially given that MFIs' chief purpose is poverty alleviation. We also hoped that the MFIs that received the negative treatment would have more of a desire to at least explore the idea of an impact evaluation. We thought that if MFIs were shown some scientific evidence suggesting that current methods may not be effective, they would want to discover if their specific practices could be improved.

A more optimistic interpretation, on the other hand, would point to the five percent of subjects in the negative treatment condition in Experiment 1 and the three percent in Experiment 2 that accepted the invitation for additional information about a partnership to perform a randomized evaluation. They accepted the invitation despite the fact that they received information suggesting that a negative result might be found. The invitation may have also signaled that the researchers proposing the partnership may have themselves been biased against microfinance. Yet a non-trivial share of MFIs was still willing to work with the academic team to

learn their own organizations' effectiveness. This provides some grounds for optimism about the willingness of some development organizations to update.

However, if these results extend to learning in development more generally, on balance they are not good news. With the recent evaluation revolution in development, there is substantial hope that practices will be updated based on the findings and that development activities will subsequently become more effective. But a missing step has been overlooked between the execution of impact evaluations and the planning of new interventions: the willingness of organizations to update based on scientific information has been assumed and not established. If organizations continue to seek confirmation of priors, then moving from evaluation to better interventions may take much longer than expected.

Of course, further research is necessary to determine whether other NGOs and development organizations behave consistently with MFIs. This industry may be unique. We suspect, however, that the extensive findings from social psychology and neuroscience on confirmation bias will extend to additional organizations involved in poverty relief. But additional research will need to establish the scope of the problem. On balance, however, it appears that even if all of the other stipulated problems with randomized evaluations can be addressed (and we happen to believe they can be), the willingness of organizations to update based on the findings from RCTs may still attenuate the effectiveness of field experiments in development. Future research should therefore also explore the conditions that enable organizational openness to new information and willingness to update established practices accordingly.

References

A short history of Grameen Bank. Grameen Communications. http://www.grameen-info.org/index.php?option=com_content&task=view&id=19&Itemid=114 (accessed November 12, 2011).

Armendáriz, Beatriz, and Jonathan Morduch. 2010. Measuring impacts. In *Economics of Microfinance*, 267. Cambridge, Massachusetts: MIT Press.

Aronson, Elliot. 1969. The theory of cognitive dissonance: A current perspective. In *Advances in experimental social psychology* 4, no. 2-32. New York: Academic Press Inc.

Banerjee, Abhijit V., and Esther Duflo. 2009. The Experimental Approach to Development Economics. *Annual Review of Economics* 1, 1: 151-178.

Banerjee, Abhijit V., and Esther Duflo. 2010. Giving credit where it is due. *Journal of Economic Perspectives* 24, no. 3: 61-80.

Banerjee, Abhijit V., and Esther Duflo. 2011. *Poor economics: A radical rethinking of the way to fight global poverty.* New York: Public Affairs.

Banerjee, Abhijit V., Esther Duflo, Rachel Glennerster, and Cynthia Kinnan. 2009. The Miracle of Microfinance?: Evidence from a Randomized Evaluation. Massachusetts Institute of Technology.

Checkel, Jeffrey T. 2001. Why Comply? Social Learning and European Identity Change. *International Organization* 55 (3): 553-588.

Checkel, Jeffrey T. 2005. International Institutions and Socialization in Europe: Introduction and Framework. *International Organization* 59 (4): 801-826.

Copestake, James, Nathanael Goldberg and Dean Karlan. 2009. Randomized control trials are the best way to measure impact of microfinance programs and improve microfinance product designs. *Enterprise Development and Microfinance* 20, no. 3 (September): 167-176.

Elkins, Zachary and Beth Simmons. 2004. The Globalization of Liberalization: Policy Diffusion in the International Political Economy. *American Political Science Review* 98: 171-190.

Elkins, Zachary, Andrew T. Guzman, and Beth Simmons. 2006. Competing for Capital: The Diffusion of Bilateral Investment Treaties 1960–2000. *International Organization* 60 (4): 811-846.

Deaton, Angus S. 2010. Instruments, Randomization, and Learning about Development. *Journal of Economic Literature* 48(2): 424-55.

Duflo, Esther and Michael Kremer. 2004. "Use of Randomization in the Evaluation of Development Effectiveness." In *Evaluating Development Effectiveness*. Edited by George K. Pitman, Osvaldo N. Feinstein, and Gregory K, Ingram. New York: Transaction Publishers.

Faiola, Anthony. 2009. "Some World Bank Health Programs Ineffective, Report Says." Washington Post 1 May 2009, accessed 15 March 2013 athttp://articles.washingtonpost.com/2009-05-01/business/36842038_1_world-bank-emergency-programs-nutrition-and-population.

Fearon, James. 1995. Rationalist Explanations for War. *International Organization* 49 (3): 379-414.

Festinger, Leon. 1957. *A theory of cognitive dissonance.* Stanford, CA: Stanford University Press.

Garikipati, Supriya. 2008. The impact of lending to women on household vulnerability and women's empowerment: Evidence from India. *World Development* 36, no. 12:2620-2642.

Hansen, Henrik and Finn Tarp. "Aid Effectiveness Disputed." *Journal of International Development* 12, 3: 375-398.

Helms, Brigit. 2010. Microfinancing changes lives around the world—Measurably. *Seattle Times*. 7 April 2010. Accessed 17 March 2013 at http://seattletimes.com/html/opinion/2011545639_guest08helms.html.

Holvoet, Nathalie. 2005. The impact of microfinance on decision-making agency: Evidence from south india. *Development and Change* 36, no. 1:75-102.

Karlan, Dean S. 2009. Cairo Evaluation Clinic: Thoughts on Randomized Trials for Evaluation of Development (June 26). Yale University Economic Growth Center Discussion Paper No. 973; Yale Economics Department Working Paper No. 65. Accessed 17 March 2013 at http://ssrn.com/abstract=1426130

Karlan, Dean, and Jonathan Zinman. 2011. Microcredit in theory and practice: Using randomized credit scoring for impact evaluation. *Science* 332, no. 6035:1278-1284.

Karlan, Dean and Jonathan Zinman. 2010. Expanding credit access: Using randomized supply decisions to estimate the impacts. *The Review of Financial Studies* 23, no. 1:433-464.

Legovini, Arianna. 2010. "Development Impact Evaluation Initiative: A World Bank-Wide Strategic Approach to Enhance Developmental Effectiveness." Washington, D.C.: World Bank. Accessed 17 March 2013 at http://siteresources.worldbank.org/INTDEVIMPEVAINI/Resources/3998199-1286546178578/7465778129130657202 8/Legovini_dime_paper_ext.pdf.

Lord, Charles G., Lee Ross, and Mark R. Lepper. 1979. Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology* 37, no. 11:2098-2109.

Morduch, Jonathan. 1998. Does Microfinance Really Help the Poor: New Evidence from Flagship Programs in Bangledesh. Department of Economics and HIID Harvard University and Hoover Institution Stanford University.

Mosley, Paul. 1986. “Aid Effectiveness: The Micro-Macro Paradox.” *The Institute of Development Studies Bulletin* 17, 2 (April): 22-27.

Picciotto, Robert. 2012. Experimentalism and development evaluation: Will the bubble burst? *Evaluation* 18, no. 2 (April): 213-229.

Pitt, Mark M., Shahidur R. Khandker, and Jennifer Cartwright. 2006. Empowering women with micro finance: Evidence from Bangladesh. *Economic Development and Cultural Change* 54, no. 4:791-831.

Rahman, Aminur. 1999. Microcredit initiatives for equitable and sustainable development: Who pays? *World Development* 27, no. 1:67-82.

Rajan, Raghuram G. and Arvind Subramanian. 2008. "Aid and Growth: What Does the Cross-Country Evidence Really Show?," *The Review of Economics and Statistics* 90, 4: 643-665.

Ravallion, Martin. 2001. The mystery of the vanishing benefits: An introduction to impact evaluation. *World Bank Econ Review* 15, no. 1: 115-140.

Risse, Thomas. "Let's Argue!": Communicative Action in World Politics." *International Organization* 54 (1): 1-40.

Risse, Thomas. 2004. "Global Governance and Communicative Action." *Government and Opposition* 39 (2): 288-313.

Rodrik, Dani. 2008. “The New Development Economics: We Shall Experiment, but How Shall We Learn?” Harvard Kennedy School Working Paper RWP08-055.

Sartori, Anne E. 2003. An Estimator for Some Binary-Outcome Selection Models without Exclusion Restrictions. *Political Analysis* 11 (2): 111-38.

Simmons, Beth, Frank Dobbin, and Geoffrey Garrett. 2006. Introduction: International Diffusion of Liberalism. *International Organization* 60 (4): 781-810.

Steele, Claude M., and Thomas J. Liu. 1983. Dissonance processes as self-affirmation. *Journal of Personality and Social Psychology* 45, no. 1:5-19.

Stone, Jeff, and Joel Cooper. 2001. A self-standards model of cognitive dissonance. *Journal of Experimental Social Psychology* 37, no. 3.

Subramanian, Arvind. 2011. “Nicholas Kristof and Aid.” Washington, D.C.: Center for Global Development. Accessed 17 March 2013 at http://blogs.cgdev.org/globaldevelopment/2011/05/nicholas-kristof-and-aid.php.

Vallone, Robert P., Lee Ross, and Mark R. Lepper. 1985. The hostile media phenomenon: Biased perception and perceptions of media bias in coverage of the Beirut massacre. *Journal of Personality and Social Psychology* 49, no. 3:577-585.

United States Agency for International Development. 2011. *USAID Evaluation Policy*. Washington, D.C.: USAID. Accessed 15 March 2013 at http://www.usaid.gov/sites/default/files/documents/1868/USAIDEvaluationPolicy.pdf.

Westen, Drew, Pavel S. Blagov, Keith Harenski, Clint Kilts, and Stephan Hamann. 2006. Neural bases of motivated reasoning: An fMRI study of emotional constraints on partisan political judgment in the 2004 U.S. presidential election. *Journal of Cognitive Neuroscience* 18, no. 11:1947-1958.

Wesselink, Ben. 2010. Directory of development organizations. http://www.devdir.org/ (accessed February 15, 2012).

Yunus, Muhammad. 2007. Banker To The Poor: Micro-lending and the Battle Against World Poverty. New York: Perseus Books Group.

Appendix: Treatment Language & Robustness Checks

Experiment 1 Control Email

<MFI Name>,

I am contacting you as director of the [ORGANIZATIONAL NAME OMITTED FOR REVIEW PURPOSES]. Founded in 2008, we study the relationship between politics and economics with a special focus on global development, including the impact microfinance institutions have on the poor. We are seeking to assess the interest of qualified microfinance institutions in possible partnerships to perform impact evaluations. We understand that you provide microcredit loans in <country>.

As I am sure you understand, in order to improve MFI processes we must carefully evaluate impact. This is best accomplished through scientific evaluations using random assignment. Should grant funding, balance of prior commitments, and mutual interest allow, would your organization be interested in receiving more information about potentially partnering with [NAME REMOVED] on a future impact evaluation?

Please understand that this is not an invitation for immediate partnership. We have several other commitments to partners currently and thus can pursue only a few new joint projects going forward – and those will, of course, depend on future grant funding. But we are hoping to gauge your possible interest.

Due to numerous research commitments, we would prefer to communicate – at least through this initial phase – through email. In order that we can keep better track of your response, please reply directly to this email. We hope to hear from you soon.

Thank you very much for attention to this inquiry.

Sincerely,

[NAME OMITTED FOR REVIEW PURPOSES]

Experiment 1 Negative Treatment Email

<MFI Name>,

I am contacting you as director of the [ORGANIZATIONAL NAME OMITTED FOR REVIEW PURPOSES]. Founded in 2008, we study the relationship between politics and economics with a special focus on global development, including the impact microfinance institutions have on the poor. We are seeking to assess the interest of qualified microfinance institutions in possible partnerships to perform impact evaluations. We understand that you provide microcredit loans in <country>.

Academic research suggests that microfinance is ineffective. The results of a recent scientific study show that microcredit loans have no effect on business growth and subjective well-being, nor are there disproportionate benefits in targeting women with microcredit loans (Karlan and Zinman 2011, "Microcredit in Theory and Practice," *Science*). These results are compelling to us, and we wish to learn more so we can further assist those in need.

As I am sure you understand, in order to improve MFI processes we must carefully evaluate impact. This is best accomplished through scientific evaluations using random assignment. Should grant funding, balance of prior commitments, and mutual interest allow, would your organization be interested in receiving more information about potentially partnering with [NAME REMOVED] on a future impact evaluation?

Please understand that this is not an invitation for immediate partnership. We have several other commitments to partners currently and thus can pursue only a few new joint projects going forward – and those will, of course, depend on future grant funding. But we are hoping to gauge your possible interest.

Due to numerous research commitments, we would prefer to communicate – at least through this initial phase – through email. In order that we can keep better track of your response, please reply directly to this email. We hope to hear from you soon.

Thank you very much for attention to this inquiry.

Sincerely,

[NAME OMITTED FOR REVIEW PURPOSES]

Experiment 1 Positive Treatment Email

<MFI Name>,

I am contacting you as director of the [ORGANIZATIONAL NAME OMITTED FOR REVIEW PURPOSES]. Founded in 2008, we study the relationship between politics and economics with a special focus on global development, including the impact microfinance institutions have on the poor. We are seeking to assess the interest of qualified microfinance institutions in possible partnerships to perform impact evaluations. We understand that you provide microcredit loans in <country>.

Academic research suggests that microfinance is effective. The results of a recent scientific study show that microcredit loans have a positive effect on economic self-sufficiency and subjective well-being of borrowers, including the decision making power women have in the home (Karlan and Zinman 2010, "Expanding Credit Access," *Review of Financial Studies*). These results are compelling to us, and we wish to learn more so we can further assist those in need.

As I am sure you understand, in order to improve MFI processes we must carefully evaluate impact. This is best accomplished through scientific evaluations using random assignment. Should grant funding, balance of prior commitments, and mutual interest allow, would your organization be interested in receiving more information about potentially partnering with [NAME REMOVED] on a future impact evaluation?

Please understand that this is not an invitation for immediate partnership. We have several other commitments to partners currently and thus can pursue only a few new joint projects going forward – and those will, of course, depend on future grant funding. But we are hoping to gauge your possible interest.

Due to numerous research commitments, we would prefer to communicate – at least through this initial phase – through email. In order that we can keep better track of your response, please reply directly to this email. We hope to hear from you soon.

Thank you very much for attention to this inquiry.

Sincerely,

[NAME OMITTED FOR REVIEW PURPOSES]

Experiment 2 Control Email

<MFI Name>,

I am contacting you as director of [ORGANIZATION NAME REMOVED FOR REVIEW PURPOSES]. We want to identify microfinance institutions in <country> that may be interested in participating in an impact evaluation using random assignment.

Would your organization be interested in receiving more information about potentially partnering with PEDL on a future impact evaluation?

Please understand that this is not an invitation for immediate partnership, which would require funding and mutual availability. We are, however, hoping to identify interested organizations.

We hope to hear from you soon.

Sincerely,
[NAME OMITTED FOR ANONYMOUS REVIEW]

Experiment 2 Negative Treatment Email

<MFI Name>,

I am contacting you as director of the [ORGANIZATION NAME REMOVED FOR REVIEW PURPOSES]. We want to identify microfinance institutions in <country> that may be interested in participating in an impact evaluation using random assignment.

Credible academic research suggests that microfinance may be ineffective. A recent scientific study shows that microcredit loans have no effect on economic self-sufficiency, subjective well-being, or women's empowerment (Karlan and Zinman 2011, "Microcredit in Theory and Practice," Science). These findings are interesting, but microfinance institutions vary, so you may want to know your program's particular results.

Would your organization be interested in receiving more information about potentially partnering with PEDL on a future impact evaluation?

Please understand that this is not an invitation for immediate partnership, which would require funding and mutual availability. We are, however, hoping to identify interested organizations.

We hope to hear from you soon.

Sincerely,

[NAME OMITTED FOR ANONYMOUS REVIEW]

Experiment 2 Positive Treatment Email

<MFI Name>,

I am contacting you as director of [ORGANIZATION NAME REMOVED FOR REVIEW PURPOSES]. We want to identify microfinance institutions in <country> that may be interested in participating in an impact evaluation using random assignment.

Credible academic research suggests that microfinance may be effective. A recent scientific study shows that microcredit loans have a positive effect on economic self-sufficiency, subjective well-being, and women's empowerment (Karlan and Zinman 2010, "Expanding Credit Access," Review of Financial Studies). These findings are interesting, but microfinance institutions vary, so you may want to know your program's particular results.

Would your organization be interested in receiving more information about potentially partnering with PEDL on a future impact evaluation?

Please understand that this is not an invitation for immediate partnership, which would require funding and mutual availability. We are, however, hoping to identify interested organizations.

We hope to hear from you soon.

Sincerely,
[NAME OMITTED FOR ANONYMOUS REVIEW]

Experiment 2 Control Email

<MFI Name>,

I am contacting you as director of [ORGANIZATION NAME REMOVED FOR REVIEW PURPOSES]. We want to identify microfinance institutions in <country> that may be interested in participating in an impact evaluation using random assignment.

Would your organization be interested in receiving more information about potentially partnering with PEDL on a future impact evaluation?

Please understand that this is not an invitation for immediate partnership, which would require funding and mutual availability. We are, however, hoping to identify interested organizations.

We hope to hear from you soon.

Sincerely,
[NAME OMITTED FOR ANONYMOUS REVIEW]

Robustness Checks

In addition to the main difference-in-means analysis reported above, we conducted a series of robustness checks using different estimation strategies. The multinomial probit model sets response as the base category and then estimates the likelihood of declining or accepting the invitation. The logit models set up a series of dichotomies between the treatment 1 and control, treatment 2, and control, and treatment 1 vs. treatment 2. We consider each of the possibilities on the response, decline, and accept outcomes. And finally the selection model allows us to incorporate response and outcome into the same category. Because we do not have additional information with which to identify the separate stages, we use the model designed by Sartori (2003). As we will highlight below, the results are consistent across model specifications.

First, we estimated separate multinomial probit models for each of the experimental conditions. Table A1 displays these results. The findings confirm what we learn in the basic difference-in-means tests showing that receiving the negative prompt makes MFIs on average less likely to request additional information on the offered partnership for an impact evaluation than when receiving the placebo ($p < 0.1$). It also shows a very strong difference between the positive and negative conditions ($p < 0.01$).

Table A1: Multinomial Probit Table of Outcomes across Conditions

Treatments	Response	Decline	Accept	*N*
Positive	Base	-0.104	0.198	945
	Base	(0.368)	(0.164)	
Constant	Base	-3.076***	-1.953***	
	Base	(0.287)	(0.143)	
Negative	Base	0.064	**-0.329***	930
	Base	(0.338)	**(0.186)**	
Constant	Base	-3.119***	-1.944***	
	Base	(0.293)	(0.147)	
Pos. vs. Neg.	Base	-0.167	**0.522*****	929
	Base	(0.354)	**(0.179)**	
Constant	Base	-3.125***	-2.365***	
	Base	(0.260)	(0.155)	

Robust standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1
These models include the age variable, but the results are not reported here. It is insignificant in all of the regression models.

We also considered the comparisons as a set of logit models on the outcome variables separately. Like the multinomial model, we compared the negative prompt to placebo, positive prompt to placebo, and positive prompt to negative prompt, but in the basic logit models we conduct each of these regressions separately. Table 4 displays the results of these analyses.

Table A2: Logit Results for Accept, Reject, Response

Variables	***Response***	***Response***	***Response***
Positive vs. Placebo	0.236		
	(0.225)		
Negative vs. Placebo		-0.407	
		(0.263)	
Positive vs. Negative			**0.632****
			(0.249)
Age	-0.203	-0.246	0.131
	(0.232)	(0.269)	(0.244)
Constant	-2.296***	-2.278***	-2.817***
	(0.198)	(0.205)	(0.218)
	Accept Offer	***Accept Offer***	***Accept Offer***
Variables			
Positive vs. Placebo	0.284		
	(0.234)		
Negative vs. Placebo		**-0.508***	
		(0.284)	
Positive vs. Negative			**0.780*****
			(0.267)
Age	-0.166	-0.235	0.142
	(0.240)	(0.287)	(0.257)
Constant	-2.425***	-2.397***	-3.038***
	(0.207)	(0.216)	(0.239)
	Decline Offer	***Decline Offer***	***Decline Offer***
Variables			
Positive vs. Placebo	-0.323		
	(0.775)		
Negative vs. Placebo		0.236	
		(0.690)	
Positive vs. Negative			-0.548
			(0.735)
Age	-0.593	-0.288	0.028
	(0.849)	(0.727)	(0.735)
Constant	-4.541***	-4.646***	-4.514***
	(0.589)	(0.613)	(0.512)
Observations	945	930	929

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The results are consistent with the previous findings reported in Tables 2 and 3. The significant results are again bolded and show robustness across specifications. As one final check, we consider the possibility that accepting or declining is a two-stage process. We

therefore estimate a selection model in which assignment to treatment is the independent variable generating response as well as outcome (decline or accept). Because we do not have enough information to satisfy an exclusion restriction, we use the selection model developed by Sartori (2003). These results (reported in Table A3) are largely consistent with those reported earlier. The effect of the negative condition on acceptance is in the same direction but only marginally significant ($p = 0.107$). That result is consistent with the difference-in-means results reported above.

Table A3: Selection Model of Response and Accept/Decline Outcome

Treatments	Response	Accept	Resp. Constant	Comp. Constant	N
Pos. v. Placebo	0.124	0.145	-1.381***	-1.437***	957
	(0.113)	(0.115)	(0.082)	(0.085)	
Neg. v. Placebo	-0.169	-0.210	-1.381***	-1.437***	940
	(0.124)	(0.130)	(0.082)	(0.085)	
Pos. v. Neg.	0.293**	0.355***	-1.550***	-1.647***	941
	(0.120)	(0.126)	(0.092)	(0.098)	

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1