

How Much do Existing Borrowers Value Microfinance? Evidence from an Experiment on Bundling Microcredit and Insurance

By ABHIJIT BANERJEE[†], ESTHER DUFLO[†] and RICHARD HORNBECK[‡]

[†]MIT [‡]University of Chicago

Final version received 26 March 2018.

Randomized controlled trials have found only modest effects of microfinance, but these studies focus on new clients. Existing estimates may thus understate ongoing gains for more experienced borrowers and the longer-run potential of microfinance. We estimate impacts of microfinance on *experienced* borrowers, using an episode when a microfinance institution modestly increased existing clients' fees in randomly selected villages (in exchange for a mandatory health insurance policy that turned out to be useless). This increase in fees led to a 22 percentage point decline in loan renewal in treatment villages (95% confidence interval: 16 to 27), compared to control villages where the policy was not introduced. Using this randomly generated variation in microfinance participation among *experienced* borrowers, we find impacts of microfinance that are strikingly similar to previous estimates for new clients: neither business outcomes nor household consumption were affected, on average. Also, consistent with prior studies, we find significant impacts on business outcomes among clients who had started their businesses before microfinance entered the village (0.06 standard deviation decline in an index of business outcomes from the loss of microfinance, 95% confidence interval: -0.002 to -0.12). However, despite these measured losses, these clients were just as willing to give up microfinance.

INTRODUCTION

Several randomized evaluations of microfinance, from various settings and countries, have not found that microfinance loans have a strong positive impact on household consumption, income or social outcomes, such as female empowerment, health and education (Crepon *et al.* 2015; Augsburg *et al.* 2015; Tarozzi *et al.* 2015; Angelucci *et al.* 2015; Attanasio *et al.* 2015; Banerjee *et al.* 2015b).¹ These studies estimate the impacts of extending microfinance loans to new clients, using a research design that takes advantage of the randomized expansion phase of microfinance. Microfinance institutions usually resist the idea of denying credit to anyone who wants it once they are officially open for business in the area, and of course these institutions would be hesitant to randomly stop providing credit to existing clients. As a consequence, however, an important caveat for interpreting these previous results is that estimated impacts of extending microfinance to new clients may understate the impacts of microfinance on inframarginal experienced borrowers (Banerjee *et al.* 2015c). Early adopters of microfinance might be expected to have the highest returns to microfinance loans (Wydick 2016). Consistent with this possibility, the above studies estimate somewhat greater impacts in settings that are less saturated with microfinance before the randomized evaluation (Wydick 2016). Further, if experienced clients are systematically better at using microfinance than new clients—perhaps because it takes time to determine how best to use the credit—then the previous results could be misleadingly pessimistic about the potential long-run positive impacts from microfinance. Breza and Kinnan (2016) estimate impacts from the loss of microfinance at the district level, using the sudden retrenchment of microfinance in Andhra Pradesh, India. In contrast to the modest impacts on borrowers implied by the above randomized evaluations, they find large negative impacts on district-level

outcomes. This result could be due to larger impacts of microfinance on existing clients or, in their preferred interpretation, due to general equilibrium effects on both borrowers and non-borrowers (see also Buera *et al.* 2014).

To shed light on this question, this paper examines the (partial equilibrium) impact of microfinance on *experienced* microfinance borrowers. The paper takes advantage of an episode when existing microfinance clients, in randomly chosen villages in India, became obliged to purchase a health insurance policy on renewal of their microfinance loan. As a consequence, many long-established clients left the microfinance institution and stopped borrowing. This episode therefore provides an unusual opportunity to examine whether the existing literature on new clients applies to a very different population of experienced microfinance borrowers. Further, this focus on experienced clients allows us to test a conjecture of the existing microfinance literature: that among those who joined microfinance to finance a business, the treatment effect of microfinance is very different for clients who joined to finance an existing business as compared to the treatment effect for those who joined to finance a new business.

In 2007, in rural Karnataka and Andhra Pradesh, SKS Microfinance was one of India's leading microfinance organizations, and they began requiring all new and renewing clients to purchase a health insurance policy that provided coverage for catastrophic incidents, hospitalization and maternal care. At the beginning of this initiative, for two districts in Northern Karnataka,² we coordinated with SKS to exclude randomly some villages from the health insurance expansion to enable the evaluation of this health insurance product. SKS had been operating in those districts for over two years, and microfinance was a known product in the area. We collected data at baseline (before the introduction of the health insurance requirement), at endline, and at regular intervals on a randomly selected sample of existing SKS clients in 101 'treatment' villages (in which clients became required to purchase insurance) and 100 'control' villages (in which clients were not required to purchase insurance).³

To the surprise of SKS, the insurance product was extremely unpopular. There were anecdotal accounts of client complaints from the beginning. In the course of events, the insurance scheme was never properly implemented, the relationship between SKS and the third-party insurer (ICICI-Lombard) soured, and eventually the purchase of the insurance policy was made voluntary and then discontinued.

Our first result is that the insurance requirement and the associated fee led to a large decline in loan renewal rates. Administrative data show that loan renewal rates declined by 22 percentage points (or 30%), with a standard error of 2.9 percentage points, in treatment villages compared to control villages where 72% of clients renewed. Self-reported data from clients suggest that few of those who left SKS obtained microfinance loans from other organizations, even in villages where they were available, so this led to a net decline of participation in microfinance. The effect is large: the health insurance policy was inexpensive (525 rupees, compared to an average renewal loan size of around 10,300 rupees in control villages) and could be rolled into the loan so that it essentially represented a 5 percentage point increase in the interest rate on a base annual percentage rate of 24%. Even if the clients assigned zero value to the insurance product—probably rightly—this fee represented only a modest increase in borrowing costs. Before interest rates were capped by the Reserve Bank, microfinance institutions in India often charged rates in excess of 30%. The implied price elasticity of microfinance participation (1.4), however, is comparable to the participation elasticity estimates from Karlan and Zinman (2016) using experimental variation in the interest rate in Mexico.

Given the failures in the implementation of the insurance scheme, detailed further below, we find unsurprisingly a precisely estimated but very small impact of providing health insurance on utilization of healthcare, healthcare spending, and the financing of healthcare. We therefore treat this as a pure microfinance experiment, where the increased cost of borrowing in some villages generated random variation in the continued use of microfinance loans.

Consistent with the previous literature on microfinance, we find little impact on income, consumption, social outcomes, and whether SKS clients continued to own a business or started a new business. These results are not simply the mechanical consequence of the fact that these borrowers chose to stop borrowing and so, by revealed preference, had little to gain from microfinance. While the net welfare gain from a microfinance loan may well be small for those who stopped borrowing (more on this later), we might still expect changes in business activity, income, and/or consumption. The clients borrowing money are either investing the funds or consuming them, prior to repaying the loan, so decreased borrowing should be reflected in the nature and timing of investment and/or consumption. For business owners specifically, the loss of the loan should imply lower investment, lower revenues and lower profits (gross of interest payments).⁴

Consistent with this prediction for business owners, we do see a decline in the scale of businesses and a significant decline in an index of business outcomes (that includes profit, sales, employment and assets). Interestingly, this decline is concentrated entirely among the 80% of business owners at baseline who had a business before SKS started lending in their village. By contrast, we find no effect of losing microfinance on those businesses at baseline that were begun after SKS started lending in their village. Indeed, for this latter group of ‘new’ entrepreneurs, leaving microfinance had a substantial *positive* effect on their household consumption and particularly their non-durables consumption (e.g. food). This result is similar to findings from several of the previous studies of microfinance, which also estimate positive impacts of microfinance on businesses that started before microfinance was available (Banerjee *et al.* 2015b; Crepon *et al.* 2015; Augsburg *et al.* 2015) and little or no effect on the general population. Banerjee *et al.* (2015a) also examine businesses that were started after microfinance was launched in treatment and control areas in India, and find no impact of microfinance on their business outcomes. However, this previous result combines a possible negative selection effect (i.e. businesses started with microfinance money may be less productive businesses) with the potentially positive effect of additional borrowing, so the net difference may be zero even if the loan by itself has a positive effect. By contrast, our estimate compares businesses started after microfinance, in both treatment and control areas, so the selection is exactly the same. The absence of a treatment effect on businesses launched after microfinance became available is confirmation that additional borrowing does not make these households’ businesses more productive.

The remainder of the paper proceeds as follows. In Section I, we lay out a simple model that highlights why the impact of losing access to microfinance may vary systematically across households with older or newer businesses. In Section II, we describe the empirical setting. In Section III, we describe the experimental research design and data collection. We lay out the empirical methodology in Section IV. We report the results in Section V, which we interpret further in Section VI. We conclude in Section VII.

I. POTENTIAL IMPACTS OF EXIT FROM MICROFINANCE

In this section we sketch a simple model of consumption and investment choice for households that are credit constrained. The purpose is to highlight the different responses

to a tightening of the credit constraint among households who have been in business for some time, as compared to those households who are relatively new to business. An important advantage of our empirical setting is the opportunity to differentiate between these two groups.

Model setup

Each consumer ‘lives’ for two periods. At the beginning of each period, the consumer can spend money on two goods that we call consumption goods and capital goods. Consumption, denoted by c , is fully divisible and purchases are consumed during the period in which they are bought. The per-period utility function is given by $u(c)$, which we assume is defined only on the positive orthant, with $u'(c) \rightarrow \infty$ as $c \rightarrow 0$. There is no discounting.

The capital good comes embodied in two available technologies. Technology 1 is purely linear: the consumer can invest an amount k in technology 1 in any period and get a return ak at the end of that period. Technology 2 involves a one-time fixed cost f to be paid in period 1. If the consumer pays that fixed cost, and invests an amount k in technology 2 in period 1 or period 2, then the consumer receives a return $Ak > ak$ at the end of that period.

Given the increasing returns of technology 2, the model is well-defined only if there is a credit constraint. Assume that the consumer starts period 1 with wealth zero but can borrow up to a limit b at the beginning of each period, at interest rate $r < a$, which is repaid at the end of the period. At the end of the first period, the consumer can also save and therefore begins the next period with wealth $w \geq 0$.

Analysis of the model

Given the above model setup, the consumer must decide whether to invest in technology 1 or technology 2. In the former case, the utility will be

$$v(b) = \max_w u\left(b\left(1 - \frac{r}{a}\right) - \frac{w}{a}\right) + u\left(b\left(1 - \frac{r}{a}\right) + w\right).$$

In the latter case, the utility will be

$$V(b) = \max_w u\left(b\left(1 - \frac{r}{A}\right) - \frac{w}{A} - f\right) + u\left(b\left(1 - \frac{r}{A}\right) + w\right).$$

We can verify the following.

Lemma 1 Regardless of the choice of technology (technology 1 or technology 2), it is optimal to set $w = 0$.

Using this lemma, it follows that for a large enough fixed cost (f)—i.e. for $b(1-r/A)-f$ close enough to zero—the marginal value of credit to a consumer investing in technology 2 is greater than the marginal value of credit to a consumer investing in technology 1 (i.e. $V'(b) > v'(b)$). This gives our first result.

Result 1. Following a reduction in borrowing capacity (b), consumers in the first period will tend to switch from technology 2 to technology 1 (as long as the fixed cost f is large enough). When this happens, the consumer's first-period consumption will increase and second-period consumption will decrease.

However, not all consumers will make this switch from technology 2 to technology 1. Those consumers that start with a high enough b will continue to invest in technology 2 and just decrease their first-period consumption. Likewise, consumers that start with a low b , and therefore were never going to choose technology 2, will decrease their first-period consumption.

Result 2. Following a reduction in borrowing capacity (b), there will be a drop in first-period consumption for consumers with either very low or very high initial values of b .

This result suggests what we might expect for consumers who started their businesses after SKS began lending in their village. Period 1 consumption may go up or down on average for these consumers, depending on the initial distribution of b . If we allowed their initial wealth to vary, then the average impact would also depend on the initial distribution of wealth.

For the households that already had a business before SKS started lending in their village, we assume that the decrease in borrowing capacity occurs in the second period of their life. For these households, the only effect of decreased borrowing capacity is a reduction in k and a resulting decline in second-period revenue and consumption. The reduction in revenue and consumption will be larger for those households that had adopted technology 2.

Result 3. A reduction in borrowing capacity (b) leads to reduction in investment, business earnings and consumption for consumers that are in the second period of their lives.

Discussion of the model

This model reflects, of course, a great deal of simplification. By ending in period 2, the model rules out the possibility that consumers in the second period may cut back their consumption to rebuild their capital stock. This effect would depress their consumption even further. On the other hand, consumers may borrow to buy indivisible consumption goods. In this case, dropping microfinance might lead them to substitute divisible consumption for non-divisible consumption. While our data do not directly distinguish between these two categories, it is reasonable to assume that durable goods are more indivisible and non-durable goods are more divisible. Weddings and other celebrations are an important exception, however, tending to be somewhat indivisible without being durable.

The model also abstracts from selection effects. It is plausible that those who started a business before SKS began lending are, on average, more productive and/or more committed to their business. In this case, a reduction in borrowing capacity may have a greater effect on these older businesses because the marginal product of capital is higher for those businesses (even if there is no indivisibility in production investment). However, if there is no indivisibility, then there should not be increased consumption in either period for clients who experience a reduction in borrowing capacity.

II. THE CONTEXT: BUNDLING OF CREDIT AND INSURANCE

In 2006, SKS Microfinance decided that it should offer health insurance to its clients. At that time, SKS was the largest microfinance institution in India and sought to leverage its

administrative advantage in dealing with low-income clients spread across rural areas of India. While ICICI-Lombard would provide the back-end insurance, SKS would administer enrolment and the initial processing of claims.

In June 2007, SKS began requiring loan clients to purchase health insurance across most of their area of operation. We persuaded them to use the expansion as an opportunity to conduct a randomized evaluation of the insurance product in 201 candidate villages with SKS presence in two districts of Northern Karnataka.⁵ In 100 randomly selected villages (the control group), SKS continued with business as usual. In the remaining 101 villages (the treatment group), insurance subscription would become mandatory for clients at the time of loan renewal. The typical health insurance policy cost 525 rupees (approximately \$13 at 2007 exchange rates), which was loaded into the amount of the loan and paid in weekly instalments along with the loan payments. By way of comparison, the average renewal loan size was 10,300 rupees.⁶ The insurance premium thus represented a 5 percentage point increase in the interest rate, which was roughly a 24% annual percentage rate at the time. The health insurance policy was intended to be actuarially fair, though SKS was prepared to lose money initially on administrative costs.

The launch of the insurance product did not go smoothly. SKS initially planned to make the purchase of insurance mandatory for all existing clients. Faced with rebellion by its clients, SKS decided to make the insurance product mandatory only for new clients and for existing clients when renewing their loans. Still, discontent with the policy and resulting client drop-out led SKS to make the insurance voluntary starting in October 2008. This unilateral change to the insurance product, and anecdotal accounts of adverse selection and outright fraud, led to a breakdown of relations between SKS and ICICI-Lombard, and insurance enrolment was discontinued in March 2009. Thus, by the time of our endline survey, clients had become free to rejoin SKS without purchasing the insurance policy.

As it turned out, SKS clients were correct in not wanting to purchase this particular health insurance policy. In principle, the policy covered hospitalization and maternity expenses, and clients had the option of going to approved health facilities to get cashless treatment or paying out-of-pocket for treatment at other facilities and submitting a claim for reimbursement. In practice, however, the implementation was badly managed by the partnership of SKS and ICICI-Lombard. Reimbursement claims were difficult for clients to file, and submitted reimbursement claims often went unprocessed. In an attempt to deal with this problem, the focus of the programme was shifted to upfront cashless treatment, but the number of hospitals that were networked for this service was inadequate, and many SKS clients did not receive the required insurance card for cashless treatment. As a result, the cashless approach was also ineffective. Below, we show that obtaining insurance had no meaningful impact on the way SKS clients handled major health incidents or on their health status and expenditures.

III. RANDOMIZATION AND DATA COLLECTION

SKS Microfinance originally identified 201 villages where it had been offering microfinance and where SKS was interested in evaluating its new health insurance product. SKS operations were organized by centre, with multiple centres in a village. To minimize the risk of spillovers between treatment areas and control areas, centres were grouped by village such that all centres in close proximity would receive the same treatment/control status.

In December 2006, using SKS's list of villages, our research team randomly selected 101 villages in which SKS would pilot the health insurance product. The remaining 100 villages formed the control group, in which SKS did not offer health insurance (although some clients had insurance through other sources). The randomization was performed by the Principal Investigators using the Stata random number generator, after stratification by branch and number of microfinance clients.⁷ The stratification ensured an even geographic distribution of treatment villages and control villages, as well as a similar number of clients in treatment and control.

SKS introduced the insurance requirement on a rolling basis, whereby the first village was reached in June 2007 and the last village was reached in November 2007. Once insurance was introduced in a village, its purchase became mandatory on loan renewal for all microfinance clients within the village.

We draw on four sources of data for the analysis.

First, we collected detailed baseline data from a random sample of SKS client households: 29 households per village, on average, in all treatment and control villages. We collected data from December 2006 to March 2007, and the survey instruments and data are available for download.⁸ A household survey module was administered to the household head in sampled households, and an adult module was administered to each adult found in the household.⁹ The household survey measured a number of household characteristics: household composition, economic status and assets, means of livelihood, and household expenses. The adult survey covered the adult's means of livelihood, income, educational background, expenses, health status, and medical treatment patterns. For rarer health incidents, the household survey covered the household's experience with major health incidents in the previous year: all incidents in which a household member died, gave birth, experienced an injury or illness that prevented them from performing their normal daily activities for more than a week, had any other health problem that required hospitalization, or otherwise spent more than 300 rupees (\$7) to treat a health incident. For each of these health incidents, the survey records basic information on its type, the way it was handled, and how the household paid for its associated expenses.

In the baseline data, we see similar client characteristics in treatment and control villages (Table 1). For the subsample of clients who report owning a business at the time of the baseline survey, panel A reports average business outcomes over the previous year. Following Kling *et al.* (2007), we also pool these four outcomes into a single index of business performance.¹⁰

Second, we collected similar survey data at endline, which came after insurance enrolment had been discontinued and clients had the opportunity to rejoin SKS without purchasing insurance. From 2009 to 2010, approximately two years after clients had faced enrolment decisions, we collected detailed data on the same households. Of the baseline households surveyed, only 1.3% were not found for the endline survey, and this attrition was not differential by treatment status.¹¹

Third, we draw on administrative data provided by SKS, which can be merged to our detailed surveys through SKS's client identification numbers. The SKS administrative data come in two main forms. First, SKS provided loan histories for its entire client base in our research areas, including when clients took out past loans and the amounts received. This gives us detailed information on clients' previous loan activities, as well as the ability to calculate the effect of the requirement to purchase health insurance on loan renewal. In a previous paper (Banerjee *et al.* 2014), we combined these data with our baseline sample to show that there was no adverse selection in client sign-up: less healthy

TABLE 1
 BASELINE CLIENT CHARACTERISTICS, BY TREATMENT AND CONTROL VILLAGES

	All villages (1)	Treatment villages (2)	Control villages (3)	Difference: (2)–(3) (4)	Number of clients (5)
<i>Panel A: SKS client businesses</i>					
Expenditures on assets, previous year	4568 [22,929]	4707 [16,951]	4410 [28,211]	347 (1095)	2118
Expenditures on workers, previous year	1560 [8118]	1700 [9266]	1401 [6579]	313 (382)	2112
Total sales, previous year	36,339 [65,433]	35,386 [56,969]	37,411 [73,823]	–1270 (3558)	1968
Total profits, previous year	14,558 [23,351]	13,536 [22,683]	15,761 [24,074]	–2404 (1587)	1581
Index of business outcomes	–0.005 [0.656]	–0.01 [0.633]	0.001 [0.681]	–0.007 (0.034)	2136
<i>Panel B: SKS loan activity</i>					
SKS loan, at time of baseline survey	0.893 [0.309]	0.895 [0.307]	0.892 [0.31]	–0.002 (0.024)	5366
SKS loan amount, at time of baseline survey	7619 [2625]	7604 [2634]	7635 [2615]	–43 (160)	4794

Notes

Column(1) reports average household characteristics from the baseline survey, with standard deviations reported in brackets. Columns (2) and (3) report average characteristics for households in randomly-assigned treatment villages and control villages, respectively. Column (4) reports the estimated difference between treatment and control households, controlling for the randomization stratification groups (SKS branch and above/below median number of clients within branch). Robust standard errors clustered by village are reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, 10% level, respectively.

households were not disproportionately likely to renew their loan in treatment areas compared to control areas. Second, SKS maintained a database of everyone who was enrolled in insurance, and all insurance claims requested and processed. This database provides information on who used a cashless facility and who received reimbursement for health expenses at other facilities.

Finally, to identify the effects of relatively uncommon major health shocks, we collected detailed data on health incidents and the way households handled them through a ‘Major Health Events Survey’. A ‘major’ health incident is defined to be any health incident that substantially disrupted a person’s ability to perform normal daily activities for more than one week.¹² This survey was conducted on a continuous basis, from April 2008 to December 2009, and covers 25,000 major incidents that happened to 7000 unique households. The survey was conducted in two stages.

In the first stage, a survey monitor accompanied the SKS loan officer to multiple centre meetings and asked the clients about any major health incidents in their household. At the centre meeting, the surveyor recorded the name of the person who was affected, the category of health problem (sickness, accident, birth, other), the relationship

between the affected person and the head of household, and whether the person went to a hospital.¹³

In the second stage, the full survey was conducted with the SKS client who had been identified at the first stage and was generally conducted in the presence of the person affected by the health incident. The full survey began with verification of the information collected at the microfinance centre meeting, and included a brief description of the incident, when it began, and the timing of treatment received. The person also categorized the seriousness of the incident, along with the length of time for which it caused an inability to perform normal daily activities. The person provided a list of symptoms, which allows us to further characterize the seriousness of the problem. The surveyor then collected information on all health providers the person visited, along with basic information about the provider, what treatment was received and at what cost, and the amount of lost income for this person and family caregivers resulting from this episode. For expenses incurred, the person was asked about how they were covered, including by saving, borrowing or the sale of assets. Information was also collected on whether and how this person used insurance and other finances to pay for the treatment expenses as well as the person's expectations for receiving reimbursement.

IV. METHODOLOGY

The empirical analysis compares client outcomes in treatment villages to client outcomes in control villages. We focus on existing clients who had loans by June 2007, the date of the roll-out of the health insurance requirement in the first village in the sample. The roll-out then took place across different villages from June to November.

For each client i in village v and randomization strata s , we regress each outcome (Y) on an indicator variable for treatment village (T) and randomization strata fixed effects (α):

$$Y_{ivs} = \beta T_v + \alpha_s + \varepsilon_{ivs}.$$

The coefficient of interest β indicates the average impact from the requirement to purchase health insurance. For all regressions, the standard errors are clustered by village to adjust for local geographic correlation.

We begin by considering impacts on SKS clients' loan renewal decisions using administrative data from SKS. Given the troubled implementation of the health insurance product, we then verify the expected absence of meaningful impacts on healthcare utilization and healthcare expenses using both the major health incident survey and the endline survey. Whatever impact we find on other outcomes is therefore presumably unrelated to the provision of health insurance.

To avoid the potential for specification search when estimating impacts on household outcomes and business outcomes, we follow the analysis template that was adopted by the randomized evaluations of microfinance in the 2015 Microfinance issue of the *American Economic Journal: Applied Economics*. Following this template, we classify the outcomes into three categories: consumption effects, business effects and social effects.¹⁴ To avoid misleading inference due to multiple inference, we compute an index of outcomes for each category and regress that index on treatment (Kling *et al.* 2007). Further, we verify the estimated p -value on the business outcome index using a Hochberg

correction for multiple hypothesis testing across total consumption and an index of social effects (Hochberg 1988).¹⁵

We also separately report estimated impacts on business outcomes for the entire sample and for households that had a business at the time of our baseline, following previous literature. Further, we split the sample of businesses that existed at the time of the baseline survey into two subgroups: those businesses that were started before SKS began lending in the village, and those businesses that were started after SKS began lending in the village (but before our baseline survey). The idea is to look for heterogeneous treatment effects due to indivisible production investments and/or due to differential self-selection into entrepreneurship. Previous studies of microfinance (Banerjee *et al.* 2015a,b,c) have attempted to get at this distinction by separately estimating treatment effects for those households that started their businesses before the introduction of microfinance and for those households that started their businesses after the introduction of microfinance. This separation is imperfect, however, as those households that start businesses after microfinance are necessarily not the same in treatment and control areas because of differential self-selection, and therefore the estimated effect of microfinance on these businesses is potentially biased downwards. A nice feature of our experiment is that it allows us to separate the treatment effect from the selection effect because the businesses that we compare are from the same pool of post-microfinance firms in both treatment and control areas, unlike in previous work where the control firms are necessarily formed before the availability of microfinance.

V. RESULTS

Impacts on loan renewal

The requirement to purchase health insurance substantially lowered SKS clients' loan renewal rates. Column (1) of Table 2 reports that clients in treatment villages were 22 percentage points (or 30%) less likely to take out an annual loan within one year after the pilot began. Specifically, clients were less likely to take out a new loan between 7 June 2007 and 3 July 2008.¹⁶ Since the roll-out took place between June and November, these estimates are intent-to-treat estimates because not all clients in treatment villages faced the health insurance requirement. We estimate that 73% of clients in treatment villages renewing during the experimental period actually faced the health insurance requirement in order to renew, so these intent-to-treat estimates might be scaled up by a factor of 1.37 to get a sense of the magnitude of the impact on those facing the requirement to buy insurance.¹⁷

The estimates indicate that clients left SKS Microfinance in large numbers due to a modest increase in fees, and perhaps also due to a loss of trust in SKS from their shifting product offering and problematic implementation of the health insurance policy. Trust and future expectations of reliable loans would be particularly important if clients maintained their existing loan in part to secure future loans. Loan renewal could also have declined if clients started defaulting on loan payments and SKS then denied new loans, though loan denial was not the focus of SKS as it was shedding clients. Missed loan payments were rare and were not differentially affected by the increase in fees associated with the treatment.¹⁸

Interestingly, this difference in loan renewal persisted after the health insurance requirement had been eliminated. At the time of the endline survey, SKS clients in treatment villages remained substantially less likely to have an SKS loan. Based on administrative data, clients in treatment village were 16 percentage points (30%) less

TABLE 2
ESTIMATED IMPACTS OF TREATMENT ON LOAN RENEWAL (EXTENSIVE MARGIN)

	First year after treatment:		At time of endline survey:	
	Administrative data		Self-reported	Other MFI loan
	(1)	(2)	(3)	(4)
<i>Panel A: Loan renewal, full sample</i>				
Treatment	-0.220*** (0.029)	-0.162*** (0.029)	-0.076*** (0.024)	0.006 (0.004)
Control group mean	0.724	0.541	0.717	0.011
Number of clients	5353	5353	5219	5353
<i>Panel B: Loan renewal, business owners at baseline</i>				
Treatment	-0.255*** (0.034)	-0.220*** (0.033)	-0.054** (0.027)	0.005 (0.006)
Control group mean	0.778	0.596	0.722	0.014
Number of clients	2149	2149	2102	2149
<i>Panel C: Loan renewal, business owners at baseline, business started before SKS entry</i>				
Treatment	-0.257*** (0.037)	-0.219*** (0.035)	-0.051* (0.030)	-0.001 (0.006)
Control group mean	0.787	0.609	0.721	0.016
Number of clients	1715	1715	1676	1715
<i>Panel D: Loan renewal, business owners at baseline and endline, business started after SKS entry</i>				
Treatment	-0.217*** (0.055)	-0.187*** (0.057)	-0.057 (0.048)	0.026* (0.014)
Control group mean	0.731	0.527	0.717	0.006
Number of clients	392	392	384	392
<i>Panel E: Loan renewal, non business owners at baseline</i>				
Treatment	-0.201*** (0.033)	-0.133*** (0.032)	-0.091*** (0.028)	0.005 (0.005)
Control group mean	0.687	0.510	0.713	0.009
Number of clients	3071	3071	2984	3071

Notes

Column (1) reports the estimated impact of treatment (imposing the insurance requirement) on whether clients took out a new SKS loan in the first year after treatment (by the end of June 2008, for the SKS clients in our baseline and endline surveys who had an annual loan prior to June 2007). Column (2) reports the impact on whether baseline SKS clients had a loan at the time of the endline survey, continuing to use SKS administrative data, whereas column (3) uses clients' self-reported loan data. Column (4) reports the impact on whether clients self-report having a microfinance loan from a non-SKS MFI, at the time of the endline survey. Panel B restricts the sample to business owners at baseline, panel C restricts the sample to business owners at baseline whose businesses started before the entry of SKS in the business owner's village (using the date of the first loan as reported in the administrative data), and panel D restricts the sample to business owners at baseline whose business started after the entry of SKS. Panel E restricts the sample to non business owners at baseline. All regressions control for the randomization stratification groups (SKS branch and above/below median number of clients within branch), and robust standard errors clustered by village are reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, 10% level, respectively.

likely to have an outstanding SKS loan (column (2) of Table 2). This difference is smaller based on self-reported data (column (3)), which may reflect measurement error because many clients report having an SKS loan when these do not appear in the administrative data. While both our survey data and the administrative data may contain errors, we suspect that the administrative data are more accurate because they reflect the firm's centralized records related to funds disbursed. Average loan renewal rates should also

decline over time, as previous clients naturally drop out from SKS, and the self-reported mean renewal rate in control villages is higher than would be expected.

Panel B of Table 2 reports corresponding estimates when restricting the sample to clients who report owning a business in the baseline survey. Of particular use later, panels C and D split this sample into those baseline business owners who report their business starting before SKS entered the village (panel C) and those baseline business owners who report their business starting after SKS entered the village (panel D). Panel E reports estimates for households that did not include a surveyed business at baseline.

Table 3 reports the accompanying declines in SKS loan sizes, where a non-renewing client's loan size is set to zero. Outstanding loan sizes decline, mostly due to changes on the extensive margin (i.e. whether or not someone has a loan).

TABLE 3
ESTIMATED IMPACTS OF TREATMENT ON LOAN RENEWAL (EXTENSIVE AND INTENSIVE MARGINS)

	First year	At time of endline survey:		
	after treatment:			
	Administrative data		Self-reported	Other MFI loan
	(1)	(2)	(3)	(4)
<i>Panel F: Loan amount, full sample</i>				
Treatment	-2075*** (340)	-1995*** (438)	—	72 (45)
Control group mean	7485	7195		63
Number of clients	5352	5353		5353
<i>Panel G: Loan amount, business owners at baseline</i>				
Treatment	-2713*** (418)	-2869*** (516)	—	110 (75)
Control group mean	8505	8343		64
Number of clients	2148	2149		2149
<i>Panel H: Loan amount, business owners at baseline, business started before SKS entry</i>				
Treatment	-2843*** (445)	-2724*** (562)	—	97 (96)
Control group mean	8669	8505		78
Number of clients	1714	1715		1715
<i>Panel I: Loan amount, business owners at baseline, business started after SKS entry</i>				
Treatment	-1902*** (664)	-2868*** (837)	—	191* (113)
Control group mean	7725	7461		0
Number of clients	392	392		392
<i>Panel J: Loan amount, non business owners at baseline</i>				
Treatment	-1728*** (361)	-1552*** (464)	—	29 (44)
Control group mean	6802	6513		65
Number of clients	3071	3071		3071

Notes

The reported estimates correspond to those reported in Table 2, but report impacts on the amount of the loan (including zeros for non-renewal). All regressions control for the randomization stratification groups (SKS branch and above/below median number of clients within branch), and robust standard errors clustered by village are reported in parentheses.

***, **, * denote statistical significance at the 1%, 5%, 10% level, respectively.

Some of the villages where the experiment took place had another microfinance organization, so part of the flight from SKS may have been compensated for by increased borrowing from another provider. Column (4) in each of Tables 2 and 3 reports the estimated impact of treatment on whether the household reports having a loan from another microfinance institution (MFI). With the caveat that these data are self-reported and may underestimate actual borrowing from other MFIs, we find little impact of the treatment on borrowing from these alternative sources. In general, there is very little reported borrowing from other MFIs at endline by current SKS clients (1.0%) or former SKS clients (1.6%).

The characteristics of those clients who leave SKS are discussed in a previous paper (Banerjee *et al.* 2014). In general, we found that clients who drop out are similar to those who remain. In particular, we found no evidence of adverse selection based on health characteristics, and that extends to health incidents that are fairly predictable (e.g. propensity to have a child). We also found little difference in the households' economic characteristics, including the propensity to own a business.

Impacts on health status and health expenditures

Table 4 reports impacts on insurance usage following surveyed 'major health incidents'. People in treatment villages are 51 percentage points more likely to report having health insurance at the time of the health incident (panel A, column (1)). However, they are only 0.3 percentage points more likely to receive insurance benefits (column (2)). This number includes the use of both a cashless facility and reimbursements, either of which taken separately shows very small increases (columns (3) and (4)). The major health incident survey was generally conducted shortly after the incident, however, and in 4.5% of the cases, responders say that they expect to receive reimbursement (column (5)). While these differences are statistically significant due to the large sample size and near absence of insurance in control villages, the magnitudes are all very small.

We see the same pattern when we group the health incident survey data by client (panel B of Table 4). People in treatment villages were 68 percentage points more likely to report ever having insurance for a major health incident and more likely to report ever having received insurance benefits (1%), or expecting reimbursement (9%), but the magnitudes remain small.

For this sample of clients who report a major health incident, we can also use administrative claims data to see whether they appear to have ever used insurance (panel C of Table 4). Column (1) reports that 84% appear in administrative data as being enrolled in the insurance programme at any point.¹⁹ In terms of these clients receiving insurance benefits at any time, 7.5% receive some benefit, of which 2.6% used a cashless facility and 5.3% received some reimbursement. This number falls between the rate of reimbursement observed in panel B (1%) and the rate of expected reimbursement (9%). Since these numbers are *conditional* on an eligible health incident occurring, these usage numbers are very low.

Given that insurance benefits were rarely used, it is unsurprising that we see no meaningful difference in how households responded to a major health incident (Table 5). Following one of these health incidents, there is no significant impact on whether the person stayed overnight in a hospital, the total cost of healthcare (including lost income), or the financing of associated costs. The point estimates and standard errors are small, suggesting that the lack of a significant finding is not driven by noise.

In the endline survey, there is also no meaningful impact on clients' health and their healthcare usage in the previous year (Table 6). Specifically, we find no impact on

TABLE 4
ESTIMATED IMPACTS OF TREATMENT ON INSURANCE USAGE

	Have insurance (1)	Received insurance benefits (2)	Used cashless facility (3)	Received reimbursement (4)	Expect to receive reimbursement (5)
<i>Panel A: For each major health incident (self-reported)</i>					
Treatment	0.510*** (0.035)	0.0030*** (0.0006)	0.0016*** (0.0004)	0.0018*** (0.0005)	0.045*** (0.012)
Control group mean	0.057	0.0003	0.0001	0.0002	0.012
Number of incidents	25,072	25,072	25,072	25,072	25,072
<i>Panel B: For each client (self-reported)</i>					
Treatment	0.683*** (0.0402)	0.0106*** (0.0019)	0.0055*** (0.0014)	0.0062*** (0.0016)	0.091*** (0.0231)
Control group mean	0.0573	0.0003	0.0001	0.0002	0.0119
Number of clients	6941	6941	6941	6941	6941
<i>Panel C: For each client (administrative data)</i>					
Treatment	0.842*** (0.018)	0.075*** (0.007)	0.026*** (0.004)	0.053*** (0.007)	
Control group mean	0	0	0	0	
Number of clients	6941	6941	6941	6941	

Notes

For each 'major health incident' surveyed, panel A reports the estimated impact of treatment (imposing the insurance requirement) on clients' self-reported insurance usage for that incident: whether clients self-report the affected person having insurance at the time of the incident (column (1)); whether clients self-report having received any insurance benefits from that incident, through the affected person either using a cashless facility or receiving reimbursement (column (2)); whether clients self-report the affected person used a cashless facility (column (3)) or they received reimbursement (column (4)); or whether clients self-report expecting to receive reimbursement from the insurance policy (column (5)). Panel B aggregates the 'major health incidents' by client, indicating the impact of treatment on whether clients self-report any affected person in their household having: had insurance for any surveyed incident (column (1)); received insurance benefits for any surveyed incident (column (2)); used a cashless facility for any surveyed incident (column (3)); received reimbursement for any surveyed incident (column (4)); or expect to receive reimbursement for any surveyed incident (column (5)). Panel C uses administrative claims data, merged to clients ever surveyed on a major health incident, to report the impact of treatment on whether the client or persons covered under the client's insurance policy: were ever enrolled in insurance (column (1)); ever received insurance benefits (column (2)); ever used a cashless facility (column (3)); or ever received reimbursement (column (4)). All regressions control for the randomization stratification groups (SKS branch and above/below median number of clients within branch), and robust standard errors clustered by village are reported in parentheses.

***, **, * denote statistical significance at the 1%, 5%, 10% level, respectively.

healthcare expenditures (column (1)), whether clients borrowed for healthcare expenses (column (2)), how much they borrowed for healthcare expenses (column (3)), the number of serious health incidents (column (4)), or the probability of staying overnight in a hospital (column (5)). The absence of impacts on healthcare utilization parallels recent estimates from Nicaragua (Thornton *et al.* 2010). There is also no impact on the ability of individuals to perform basic activities in daily life (column (6)).²⁰ Curiously, households have significantly *worse* self-reported health (column (7)), which may reflect

TABLE 5
ESTIMATED IMPACTS OF TREATMENT ON MAJOR HEALTH INCIDENTS

	Overnight hospitalization (1)	Health incident expenses (2)	Borrowed from any MFI (3)	Borrowed from family/friends (4)	Borrowed from moneylender (5)
Treatment	0.014 (0.011)	-80 (151)	0.002 (0.005)	0.021 (0.029)	0.009 (0.031)
Control group mean	0.067	1662	0.022	0.437	0.288
Number of clients	25,072	25,072	25,072	25,072	25,072

Notes

Each column reports the estimated impact of treatment (imposing the insurance requirement) on the indicated outcome variable from the 'major health incidents' survey. The outcome variables are: whether the affected person stayed overnight in a hospital for that health incident (column (1)); total health expenses for that incident, including lost income (column (2)); and whether these expenses were partly paid by borrowing from a microfinance organization (column (3)), family or friends (column (4)), or a moneylender (column (5)). All regressions control for the randomization stratification groups (SKS branch and above/below median number of clients within branch), and robust standard errors clustered by village are reported in parentheses.

***, **, * denote statistical significance at the 1%, 5%, 10% level, respectively.

an insurance information campaign by SKS leading clients to focus more on catastrophic health incidents or their health more generally.²¹ Overall life satisfaction, however, is not substantially affected (column (8)).

The health insurance product had no direct effect on the impacts that it sought to achieve: health status, healthcare usage, and the financing of healthcare expenditures. While many people in these areas did pay the health insurance premium and enrol in health insurance, very few received insurance benefits following major health incidents. For whatever reason (failure to communicate to households, failure of SKS field officers to effectively intermediate between the clients and ICICI-Lombard, clients' lack of understanding, etc.), the product turned out to be mostly useless, and anecdotal evidence suggests that clients found this out fairly quickly. The requirement to purchase insurance did inadvertently lead to a substantial decline in microfinance borrowing, however, and the following subsections explore how this impacted households.

Impacts on client businesses

Table 7 reports the impact on clients' businesses from imposing the requirement to purchase health insurance. For the full sample of clients (panel A) or the sample of clients who owned a business at baseline (panel B), there was no substantial or statistically significant impact on whether they owned a business at endline (column (1)). Columns (2)–(5) report impacts on endline business outcomes for those with businesses at baseline, including zeros for those who do not report owning a business at endline.²² All the point estimates suggest that these clients invested less in their businesses and generated less profit, though only expenditure on workers is individually statistically significant. Column (6) reports the estimated impact on an index of business outcomes, drawing on the outcomes in columns (2)–(5), which is negative and statistically significant at the 10% level, though not with a Hochberg correction for multiple hypothesis testing across all three categories of household outcomes (business, consumption and social outcomes).

The point estimates imply a substantial decline in business scale, which is somewhat masked by the substantial churn in businesses. Only 32% of all self-reported business

TABLE 6
ESTIMATED IMPACTS OF TREATMENT ON HEALTH OUTCOMES AT ENDLINE

	Health expenses (1)	Borrowed for health expenses (2)	Amount borrowed (3)	Number of health incidents (4)	Overnight hospitalization (5)	ADL index (6)	Self-reported health (7)	Overall life satisfaction (8)
Treatment	-18 (277)	0.009 (0.018)	-234 (236)	0.020 (0.034)	0.011 (0.015)	-0.010 (0.013)	-0.208*** (0.065)	-0.005 (0.020)
Control group mean	2747	0.337	1462	1.691	0.271	0.004	6.865	3.556
Number of clients	5353	5353	5353	5353	5353	5353	5353	5353

Notes

Each column reports the estimated impact of treatment (imposing the insurance requirement) on the indicated outcome variable from the endline survey. The outcome variables are: expenses on healthcare for major incidents over the previous year (column (1)); whether the household reports having borrowed money to pay some of those health expenses (column (2)); the amount borrowed to pay health expenses (column (3)); the number of major health incidents experienced by the household over the previous year (column (4)); whether a household member was hospitalized overnight in the previous year (column (5)); an index reflecting adults' self-reported ability to perform 15 typical daily activities, averaged across adults in the household (column (6)); a self-reported index of health, averaged across adults in the household (column (7)); and a self-reported index of overall life satisfaction, averaged across adults in the household (column (8)). All regressions control for the randomization stratification groups (SKS branch and above/below median number of clients within branch), and robust standard errors clustered by village are reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, 10% level, respectively.

TABLE 7
ESTIMATED IMPACTS OF TREATMENT ON BUSINESS OUTCOMES AT ENDLINE

	Business outcomes, summing over the previous year (in rupees):					
	Owns business (1)	Spent on assets (2)	Spent on workers (3)	Total sales (4)	Total profits (5)	Index of (2)–(5) (6)
<i>Panel A: Full sample</i>						
Treatment	0.008 (0.014)	–278 (171)	–333** (159)	–1359 (1258)	–783 (701)	–0.038* (0.020)
Control group mean	0.177	437	507	12,142	6120	0
Number of clients	5353	5353	5353	5353	5353	5353
Hochberg <i>p</i> -value						0.191
<i>Panel B: Business owners at baseline</i>						
Treatment	0.006 (0.025)	–308 (192)	–537* (311)	–2554 (2335)	–1470 (1269)	–0.050* (0.028)
Control group mean	0.316	503	973	21,681	11030	0
Number of clients	2149	2149	2149	2149	2149	2149
Hochberg <i>p</i> -value						0.213
<i>Panel C: Business owners at baseline, business started before SKS entry</i>						
Treatment	–0.008 (0.028)	–408** (204)	–658* (373)	–3706 (2777)	–1963 (1528)	–0.063** (0.031)
Control group mean	0.335	532	1141	23,599	11709	0
Number of clients	1715	1715	1715	1715	1715	1715
Hochberg <i>p</i> -value						0.134
<i>Panel D: Business owners at baseline, business started after SKS entry</i>						
Treatment	0.065 (0.041)	191 (399)	–66 (197)	1707 (3038)	–282 (1721)	0.018 (0.060)
Control group mean	0.222	410	234	12,860	7951	0
Number of clients	392	392	392	392	392	392
Hochberg <i>p</i> -value						0.760
<i>Panel E: Non business owners at baseline</i>						
Treatment	–0.003 (0.012)	–231 (249)	–187* (102)	–1111 (892)	–677 (435)	–0.045* (0.024)
Control group mean	0.094	410	217	5788	2981	0
Number of clients	3071	3071	3071	3071	3071	3071
Hochberg <i>p</i> -value						0.174

Notes

Each column reports the estimated impact of treatment (imposing the insurance requirement) on the indicated outcome variable from the endline survey. Panel B limits the sample to business owners at baseline, panel C restricts the sample to business owners at baseline whose businesses started before the entry of SKS in the business owner's village (using the date of the first loan as reported in the administrative data), panel D restricts the sample to business owners at baseline whose business started after the entry of SKS, and panel E restricts the sample to non business owners at baseline. The outcome variables are: whether the household earns money from owning a business at endline (column (1)); total amount spent on business assets, over the previous year (column (2)); total amount spent on hiring labour for the business, over the previous year (column (3)); total sales over the previous year (column (4)); and a direct measure of self-reported total profits over the previous year (column (5)). When there is no business reported at endline (in panels A and B), zero values are assigned for the outcome variables in columns (2)–(5). In column (6), the outcome variable is an index reflecting the equal-weighted average of the component variables in columns (2)–(5) (each normalized to have mean 0 and standard deviation 1). All regressions control for the randomization stratification groups (SKS branch and above/below median number of clients within branch), and robust standard errors clustered by village are reported in parentheses.

The Hochberg *p*-value reflects the statistical significance of the treatment effect on the index, adjusting for multiple hypothesis testing across the three summary outcomes (business, consumption, social impacts) in Tables 7, 9 and 11. ***, **, * denote statistical significance at the 1%, 5%, 10% level, respectively.

owners at baseline continue to own a business at endline, so we have many zeros in the data.²³ For example, the point estimates in panel B of Table 7 imply a reduction of 55% in expenditures on workers for existing businesses, a 61% in reduction in asset expenditure, and 12% reduction in sales.

In panels C and D of Table 7, we separate the existing businesses according to whether they are reported to have been started before SKS began lending in the village (panel C) or started after SKS began lending in the village (panel D).

The negative impact of losing microfinance access is driven entirely by the older businesses (panel C of Table 7). These older businesses average approximately twice the sales as those businesses that were started after SKS began lending in the village. For these older businesses, the effects are even larger: we find a reduction of 82% in asset expenditure and 58% in expenditure on workers (significant at the 10% and 5% level,

TABLE 8
BASELINE DIFFERENCES BETWEEN OLDER BUSINESSES AND NEWER BUSINESSES

	Business owners at baseline		
	Business started before SKS entry (1)	Business started after SKS entry (2)	Difference: (2)–(1) (3)
Expenditure on assets, previous year	3804 [20,906] 1703	7896 [30,337] 390	3992** (1726) 2093
Expenditure on workers, previous year	1702 [8761] 1696	964 [4705] 390	–680* (349) 2086
Total sales, previous year	39,327 [68,785] 1574	22,442 [43,145] 363	–17,161*** (3324) 1937
Total profits, previous year	16,177 [24,690] 1272	7401 [14,919] 287	–8830*** (1307) 1559
Index of business outcomes	0.017 [0.681] 1709	–0.081 [0.578] 392	–0.098** (0.042) 2101
SKS loan amount, at time of baseline survey	8148 [2422] 1596	8023 [2726] 354	–123 (169) 1950
Total annual consumption	56,288 [69,372] 1715	49,140 [31,398] 393	–7435*** (2191) 2108
Business income as fraction of consumption (business income = total profits)	0.33 [0.56] 1271	0.16 [0.30] 287	–0.17*** (0.03) 1558

Notes

Columns (1) and (2) report average characteristics at baseline for businesses that opened before SKS began operations in the village (column (1)) and businesses that opened after SKS began operations in the village (column (2)), with standard deviations reported in brackets. Column (3) reports the difference (column (2) – column (1)). Robust standard errors clustered by village are reported in parentheses.

***, **, * denote statistical significance at the 1%, 5%, 10% level, respectively.

TABLE 9
ESTIMATED IMPACTS OF TREATMENT ON ANNUAL CONSUMPTION AND HOUSEHOLD LABOUR INCOME

	Consumption categories								
	Total (1)	Durables (2)	Non-durables (3)	Food (4)	Health (5)	Education (6)	Temptation goods (7)	Festivals and celebrations (8)	Other labour income (9)
<i>Panel A: Full sample</i>									
Treatment	1957 (1206)	-202 (127)	2159* (1170)	1667* (851)	-13 (46)	-91 (102)	152** (73)	-312 (561)	-344 (1035)
Control group mean	50,728	1175	49,552	28,969	1092	1241	1345	9694	28,253
Number of clients	5353	5353	5353	5353	5353	5353	5353	5353	5353
Hochberg <i>p</i> -value	0.212								
<i>Panel B: Business owners at baseline</i>									
Treatment	2458 (1671)	-313 (218)	2771* (1614)	1340 (1084)	-17 (51)	-101 (155)	170* (92)	-33 (701)	-1138 (1523)
Control group mean	51,624	1329	50,295	29,729	1093	1262	1321	9608	26,386
Number of clients	2149	2149	2149	2149	2149	2149	2149	2149	2149
Hochberg <i>p</i> -value	0.286								
<i>Panel C: Business owners at baseline, business started before SKS entry</i>									
Treatment	478 (1650)	-555** (241)	1033 (1588)	-65 (1007)	20 (54)	-59 (121)	160 (100)	67 (775)	-687 (1670)
Control group mean	51,935	1369	50,566	30,033	1070	1188	1339	9381	25,836
Number of clients	1715	1715	1715	1715	1715	1715	1715	1715	1715
Hochberg <i>p</i> -value	0.772								
<i>Panel D: Business owners at baseline, business started after SKS entry</i>									
Treatment	9516*** (3071)	605 (463)	8911*** (3011)	5718*** (2085)	-52 (110)	-190 (623)	331* (169)	-21 (1968)	-3397 (3664)
Control group mean	49,822	1188	48,634	28,350	1130	1691	1089	10,913	29,463
Number of clients	392	392	392	392	392	392	392	392	392
Hochberg <i>p</i> -value	0.007								

TABLE 9
CONTINUED

	Consumption categories								
	Total (1)	Durables (2)	Non-durables (3)	Food (4)	Health (5)	Education (6)	Temptation goods (7)	Festivals and celebrations (8)	Other labour income (9)
<i>Panel E: Non business owners at baseline</i>									
Treatment	955 (1341)	-220 (155)	1174 (1305)	1594* (958)	-30 (61)	-122 (129)	171* (102)	-823 (759)	514 (1251)
Control group mean	50,473	1097	49,376	28,572	1091	1250	1349	9792	29,734
Number of clients	3071	3071	3071	3071	3071	3071	3071	3071	3071
Hochberg <i>p</i> -value	0.693								

Notes

Each column reports the estimated impact of treatment (imposing the insurance requirement) on the indicated outcome variable from the endline survey. Panel B limits the sample to business owners at baseline, panel C restricts the sample to business owners at baseline whose businesses started before the entry of SKS in the business owner's village (using the date of the first loan as reported in the administrative data), panel D restricts the sample to business owners at baseline whose business started after the entry of SKS, and panel E restricts the sample to non business owners at baseline. The outcome variables reflect total household consumption over the previous year, where columns (4)–(7) are monthly measures multiplied by 12. Other labour income in column (9) sums earnings from salaried jobs and daily labour. All regressions control for the randomization stratification groups (SKS branch and above/below median number of clients within branch), and robust standard errors clustered by village are reported in parentheses.

The Hochberg *p*-value reflects the statistical significance of the treatment effect on total consumption, adjusting for multiple hypothesis testing across the three summary outcomes (business, consumption, social impacts) in Tables 7, 9 and 11.

***, **, * denote statistical significance at the 1%, 5%, 10% level, respectively.

TABLE 10
ESTIMATED IMPACTS OF TREATMENT ON ANNUAL CONSUMPTION (PER CAPITA, ADULT EQUIVALENT)

	Consumption categories							
	Total (1)	Durables (2)	Non-durables (3)	Food (4)	Health (5)	Education (6)	Temptation goods (7)	Festivals and celebrations (8)
<i>Panel A: Full sample</i>								
Treatment	3 (268)	-58* (33)	61 (256)	111 (181)	-14 (12)	-24 (22)	23 (18)	-152 (106)
Control group mean	11,365	261	11,105	6508	249	258	299	2105
Number of clients	5353	5353	5353	5353	5353	5353	5353	5353
Hochberg <i>p</i> -value	0.991							
<i>Panel B: Business owners at baseline</i>								
Treatment	177 (338)	-105 (68)	282 (317)	127 (217)	-14 (12)	-17 (36)	14 (21)	-113 (130)
Control group mean	11,345	309v	11,037	244	248	296	2062	2149
Number of clients	2149	2149	2149	2149	2149	2149	2149	2149
Hochberg <i>p</i> -value	0.601							
<i>Panel C: Business owners at baseline, business started before SKS entry</i>								
Treatment	-164 (330)	-159** (79)	-5 (306)	-142 (193)	-5 (12)	-2 (19)	14 (24)	-65 (148)
Control group mean	11,383	327	11,056	6573	239	228	297	2008
Number of clients	1715	1715	1715	1715	1715	1715	1715	1715
Hochberg <i>p</i> -value	0.620							
<i>Panel D: Business owners at baseline, business started after SKS entry</i>								
Treatment	1260** (590)	102 (85)	1158* (587)	854* (448)	-33 (23)	-54 (160)	49 (38)	-290 (334)
Control group mean	11,257	233	11,024	6417	257	364	264	2385
Number of clients	392	392	392	392	392	392	392	392
Hochberg <i>p</i> -value	0.103							

TABLE 10
CONTINUED

	Consumption categories							
	Total (1)	Durables (2)	Non-durables (3)	Food (4)	Health (5)	Education (6)	Temptation goods (7)	Festivals and celebrations (8)
<i>Panel E: Non business owners at baseline</i>								
Treatment	-219 (304)	-48 (34)	-171 (295)	59 (207)	-19 (15)	-38 (26)	36 (26)	-253* (152)
Control group mean	11,395	233	11,162	6488	252	268	297	2140
Number of clients	3071	3071	3071	3071	3071	3071	3071	3071
Hochberg <i>p</i> -value	0.693							

Notes

Each column reports the estimated impact of treatment (imposing the insurance requirement) on the indicated outcome variable from the endline survey. Panel B limits the sample to business owners at baseline, panel C restricts the sample to business owners at baseline whose businesses started before the entry of SKS in the business owner's village (using the date of the first loan as reported in the administrative data), panel D restricts the sample to business owners at baseline whose business started after the entry of SKS, and panel E restricts the sample to non business owners at baseline. The outcome variables reflect per capita household consumption over the previous year, where columns (4)–(7) are monthly measures multiplied by 12. Per capita consumption is calculated per adult equivalent, following the conversion to adult equivalents used by Townsend (1994) for rural Andhra Pradesh and Maharashtra (the weights are: 1.0 for adult males, 0.9 for adult females, 0.94 for males and 0.83 for females aged 13–18, 0.67 for children aged 7–12, 0.52 for children aged 4–6, 0.32 for toddlers aged 1–3, and 0.05 for infants). All regressions control for the randomization stratification groups (SKS branch and above/below median number of clients within branch), and robust standard errors clustered by village are reported in parentheses. The Hochberg *p*-value reflects the statistical significance of the treatment effect on total consumption, adjusting for multiple hypothesis testing across the three summary outcomes (business, consumption, social impacts) in Tables 7, 10 and 11.

***, **, * denote statistical significance at the 1%, 5%, 10% level, respectively.

respectively), and a 0.063 standard deviation decline in the index of business outcomes that is significant at the 5% level (and 13% with the Hochberg correction).

In panel D of Table 7, we find no impact on businesses that were started more recently, after SKS had entered the village. This is a smaller sample of clients, and the standard errors are correspondingly larger, but the point estimates are also small and often have the opposite sign. Thus it appears that either the microfinance funds were not invested in the business (consistent with the absence of a decline in asset expenditure), or the marginal product of capital for these businesses is close to zero.

Table 8 explores these differences further, comparing characteristics in the baseline survey of older businesses (started before SKS began lending in the village) and newer businesses (started after SKS began lending in the village). At the time of the baseline survey, the older businesses were approximately twice as large as the newer businesses. The older businesses also had greater profits, and greater profits as a share of household consumption, along with lower expenditure on assets.²⁴

Panel E of Table 7 reports estimated impacts on business outcomes for clients that did not own a business at baseline. We see negative impacts on these clients' business

TABLE 11
ESTIMATED IMPACTS OF TREATMENT ON SOCIAL OUTCOMES

	Share of children in school, aged 5–15		Share of teenagers in school, aged 16–20		Index of women's attitudes (5)	Index of columns (1)–(5) (6)
	Girls (1)	Boys (2)	Girls (3)	Boys (4)		
<i>Panel A: Full sample</i>						
Treatment	0.009 (0.019)	−0.028* (0.015)	0.002 (0.020)	−0.02 (0.024)	−0.075** (0.037)	−0.021 (0.030)
Control group mean	0.68	0.732	0.166	0.279	0.034	−0.042
Number of clients	3001	3155	1677	1627	1444	4761
Hochberg <i>p</i> -value						0.986
<i>Panel B: Business owners at baseline</i>						
Treatment	−0.039 (0.029)	−0.002 (0.023)	0.001 (0.030)	−0.029 (0.037)	−0.088 (0.054)	−0.031 (0.042)
Control group mean	0.705	0.715	0.168	0.312	0.026	−0.035
Number of clients	1201	1222	657	630	561	1906
Hochberg <i>p</i> -value						0.601
<i>Panel C: Business owners at baseline, business started before SKS entry</i>						
Treatment	−0.040 (0.030)	−0.019 (0.025)	−0.009 (0.035)	−0.001 (0.040)	−0.081 (0.064)	−0.042 (0.047)
Control group mean	0.700	0.720	0.174	0.308	0.022	−0.032
Number of clients	963	972	505	507	444	1517
Hochberg <i>p</i> -value						0.620
<i>Panel D: Business owners at baseline, business started after SKS entry</i>						
Treatment	−0.035 (0.064)	0.091* (0.055)	0.064 (0.062)	−0.107 (0.099)	−0.095 (0.108)	0.052 (0.086)
Control group mean	0.716	0.672	0.115	0.380	−0.014	−0.076
Number of clients	216	226	134	107	101	349
Hochberg <i>p</i> -value						0.760

TABLE 11
CONTINUED

	Share of children in school, aged 5–15		Share of teenagers in school, aged 16–20		Index of women's attitudes	Index of columns (1)–(5)
	Girls (1)	Boys (2)	Girls (3)	Boys (4)	(5)	(6)
<i>Panel E: Non business owners at baseline</i>						
Treatment	0.043* (0.022)	−0.047** (0.019)	−0.006 (0.026)	−0.025 (0.028)	−0.056 (0.048)	−0.015 (0.039)
Control group mean	0.664	0.740	0.171	0.263	0.035	−0.047
Number of clients	1719	1857	980	970	857	2739
Hochberg <i>p</i> -value						0.693

Notes

Each column reports the impact of treatment (imposing the insurance requirement) on the indicated outcome variable from the endline survey. Panel B limits the sample to business owners at baseline, panel C restricts the sample to business owners at baseline whose businesses started before the entry of SKS in the business owner's village (using the date of the first loan as reported in the administrative data), panel D restricts the sample to business owners at baseline whose business started after the entry of SKS, and panel E restricts the sample to non business owners at baseline. In columns (1)–(4), the outcome variables are the shares of household children that are in school (by age and gender). In column (5), the outcome variable is an index of adolescent girls' self-reported attitudes concerning: whether men should be more educated than women; whether men should eat before women; the ideal age of marriage for women; whether women should have children immediately after marriage; and the ideal number of children. The index in column (5) reflects an equal-weighted average across responses to each question, after the responses are normalized to have mean 0, standard deviation 1, and the sign of the response oriented toward a more positive number reflecting more 'progressive attitudes' (e.g. for the questions above: 'No', 'No', older ages, 'No', fewer children). In column (6), the outcome variable is an index reflecting the equal-weighted average of the component variables in columns (1)–(5) (each normalized to have mean 0, standard deviation 1). In columns (1) and (2), the sample is restricted to households with children between the ages of 5 and 15. In columns (3) and (4), the sample is restricted to households with children between the ages of 16 and 20. In column (5), the sample is restricted to households with girls between the ages of 14 and 19. All regressions control for the randomization stratification groups (SKS branch and above/below median number of clients within branch), and robust standard errors clustered by village are reported in parentheses.

The Hochberg *p*-value reflects the statistical significance of the treatment effect on the index in column (6), adjusting for multiple hypothesis testing across the three summary outcomes (business, consumption, social impacts) in Tables 7, 9 and 11.

***, **, * denote statistical significance at the 1%, 5%, 10% level, respectively.

outcomes, which are smaller in magnitude than the effect for businesses started before SKS began lending in the village. About 9% of these clients have a business at endline, so in terms of a scaled magnitude, the impact per active business is comparable to the impact for the businesses started before SKS began lending in the village. These businesses (panel E) started after SKS became less attractive, so they may be more similar to those businesses that were started before SKS began lending in the village.

Impacts on other household outcomes

Papers discussing recent randomized evaluations of microfinance have also found some impacts on household businesses, but little impact of gaining access to microfinance on overall consumption and various social outcomes. We analyse data for these other outcomes, following the same template as those papers, and similarly find little average impact from our 'reverse' experiment (the loss of microfinance). We then extend this

analysis to consider differential impacts for households with businesses that were started before SKS began lending in the village and households with businesses that were started after SKS began lending in the village.

Table 9 reports estimated impacts on household consumption at endline.²⁵ For the sample as a whole (panel A) or the sample of baseline business owners (panel B), we see little impact on households' total annualized consumption or on households' consumption within particular categories. There is some decline in durable goods consumption, which aligns with other estimates on the impacts of gaining access to microfinance (with the opposite sign, as expected).²⁶ There is also some evidence of increased non-durable goods consumption, which we explore further below.

Focusing on business owners at baseline, a very interesting pattern emerges when we separate the older businesses that started before SKS began lending in the village and those newer businesses that started after SKS began lending in the village. Households that own older businesses spend significantly less on durable goods when they lose access to microfinance, and we estimated also that they reduced business asset expenses. Households that own newer businesses *increase* their total consumption substantially when they lose access to microfinance, and this estimate is statistically significant at the 1% level even after correcting for multiple hypothesis testing. This consumption increase is focused on non-durables (food, in particular). This same pattern emerges in per capita terms, adjusting for the number of adult equivalent household members (Table 10).

Finally, focusing on the households that did not have a business at baseline (of which the vast majority still do not have a business), we find significant increases in food and temptation good consumption when losing access to microfinance, and an insignificant decline in durables consumption and celebrations.

Column (9) of Table 9 reports no increase in the household's adult labour income from salaried jobs or day labour, which was reported in the adult modules of the household survey. Households do not appear to have compensated for declining business income by earning additional income from sources outside the business. Indeed, all the point estimates are negative. Thus we do not see an increase in another source of revenue that might compensate for the loss of business income among older business owners. For more recent businesses, this raises the question of how the consumption increase could have been financed. We return to this question in Section VI.

Table 11 shows no systematic impacts on social outcomes, consistent with findings across other studies of microfinance. There is little change in the share of younger or older children in school, and while there is some impact on adolescent girls' reported 'progressive' attitudes,²⁷ an index across all outcomes in this category shows a statistically insignificant impact.

VI. INTERPRETATION

For our sample of experienced microfinance borrowers, most of the estimated average impacts are in line with the existing randomized evaluations of the impact of microfinance on new borrowers.

For non-business-owning households, there is little systematic impact of microfinance on household consumption. Microfinance may encourage these households to borrow more to finance some form of consumption, and then cut back other forms of consumption to pay for the loan. We find that decreased microfinance borrowing causes some increase in these households' food consumption and temptation goods

consumption, which suggests the loan may provide a disciplining device. There is some question about why overall consumption is unaffected by microfinance, as these households pay interest on the microfinance loan, though perhaps they are switching away from more expensive loans (e.g. from a moneylender) or working a little more in a manner that we do not detect in the data. Notably, the effect on durable goods consumption is negative, but insignificant and not very large in magnitude (compared to the size of the loan), which suggests that they are not using the money to finance additional durable goods purchases that might be lumpier.

For business-owning households, we find results that are consistent with much of the previous literature on microfinance and consistent with the simple model described in Section I. As predicted for those households in period 2 of their 'life' in the model, a reduction in microfinance borrowing results in older business owners investing less in working capital (payments to workers) as well as fixed capital (business assets). We also see a large reduction in business profits, as the implied treatment effect on treated households is nearly 10,000 rupees per business per year, though this effect is imprecisely estimated because of the substantial churn in business ownership and the large number of zeros. These effects are perhaps reflected in lower durables consumption, although we do not see a decline in non-durables consumption. For newer business owners, corresponding to those who are in period 1 of their 'life' in the model, we see the predicted increase in consumption from decreased microfinance borrowing as these households switch away from making a lumpy investment. The increase in consumption is substantial, at over 9500 rupees per household. If we assume that this 9500 rupee increase is driven entirely by the 22% of households that stopped borrowing due to the health insurance requirement, this implies an increase in consumption of more than 43,000 rupees per year for those households that stopped borrowing. This effect is much larger than the 12,800 rupees that an average borrowing household repays in a year (10,300 rupees in principal and 2500 rupees in interest). This result suggests that these households were borrowing to buy an indivisible item that cost much more than the amount of the loan, and that they were paying for the item by cutting consumption as well as by borrowing.

However, we do not estimate the predicted corresponding decline in the acquisition of business assets among treatment households (which is why the control households were borrowing in the first place). One way to square this with the model is to think of the first period in the model as consisting of multiple years. At the beginning of this period, the household takes the microfinance loan but also borrows from other sources (some of which are more expensive) to finance a single investment. In each subsequent year, all part of the first period in our model, the household then both saves and takes additional microfinance loans to repay the expensive loans that it is carrying until its obligations are discharged. When that is done, the household uses the microfinance loan to pay for additional working capital for their business (i.e. what business owners in the second period do in our model). What happens in our model is that a certain fraction of households that were at the cusp of making this investment decide to not make the investment after all, due to the imposition of this new borrowing cost from the requirement to purchase health insurance. Therefore they do not make the large financial commitment and do not buy the asset. This means that their consumption can be higher because they are not servicing these loans.

A similar story can be told about these households borrowing from a moneylender to finance an expensive wedding, and then using microfinance borrowing and savings to finance the moneylender loan over multiple years. The

loss of microfinance makes people unwilling to have such expensive weddings, which is why they are now richer. This would also help to explain why we see no impact on their business earnings after they gave up their microfinance borrowing, though it might also reflect borrowers paying down their loans before they build up enough working capital to make proper use of their assets.

VII. CONCLUSION

The attempt by SKS Microfinance to bundle health insurance with microfinance was clearly a failure. The requirement to purchase insurance led to a substantial 22 percentage point decline in clients' loan renewal (or 30%), which led to abrupt changes in the programme rules that further undermined the insurance product and led to its rapid demise.

However, this failure reveals something very interesting about the nature of microfinance: clients' apparent willingness to abandon their microfinance relationship suggests that many microfinance clients receive little surplus from their borrowing relationship. SKS certainly did not anticipate this phenomenon. While SKS did worry that clients would not understand the value of the insurance product that they were getting, at least initially, SKS was convinced and that the benefits their clients derived from microfinance were sufficiently large and that those benefits would prevent any significant loss of demand for their bundled product. This expectation, in fact, is why SKS took on the risk to introduce a new health insurance product.

One reason for this misperception might be the finding—suggested by casual observation and now supported by more research—that what we have called 'experienced' businesses do suffer financially when they lose microfinance. These 'experienced' businesses expand through microfinance borrowing, but if these business impacts do not imply large gains in welfare, this then would explain SKS's mistaken expectations. Indeed, the absence of large gains in welfare from expansion of microfinance businesses would explain the mistaken expectations of many microfinance organizations worldwide, and their supporters. Many remain convinced that microfinance can: (1) cause substantial improvements in business outcomes, and *hence* (2) cause substantial increases in welfare.

Our results show that the first part of the above statement is partially correct: as with other microfinance institutions that have been evaluated, SKS loans do contribute to better reported business outcomes (although only for the 'experienced' businesses that were started before SKS began lending in their village).

The second part of the above statement appears to be wrong, even for the subset of households that experience financial benefits from microfinance: a large share of borrowers are fairly willing to give up microfinance, and even those households whose businesses do appear to benefit from microfinance are similarly willing to give up microfinance. That is, these households act as if they receive very little in terms of welfare. Perhaps even households with expanding businesses receive little welfare gain because expanding these businesses has many non-financial costs, such as long hours worked by the owners, stress from business uncertainty, or other factors.

Our results are discouraging regarding microfinance businesses as an engine for clients to escape poverty and, more generally, regarding the many hopes pinned on microenterprises as a way for large numbers of people to improve their lives. There are other indicators that point in this same direction: the absence of microfinance impacts on other household outcomes, the failure of microfinance businesses to grow, and the

frequent closure of these businesses. Our results reinforce a sense in the microfinance literature that there has been misplaced enthusiasm for entrepreneurship as a way to substantially improve the lives of the poor.

ACKNOWLEDGMENTS

This project received funding from the National Institutes of Health (grant PO1 HD061315-04), and received IRB approval from MIT (#0609001924) and Harvard University (#F18180-104). We thank Aparna Krishnan and Prathap Kasina for outstanding management of this difficult project, and Mahvish Shaikat and Madeline Duhon for tireless research assistance. For comments and suggestions, we thank David Cutler, Jishnu Das, Pascaline Dupas, Andrew Foster, Rachel Glennerster, Jon Gruber, Dilip Mookherjee and Ben Olken.

NOTES

1. Banerjee *et al.* (2015c) summarize the key findings, and Meager (2016a,b) provide meta-analysis across microfinance studies.
2. This region abuts Andhra Pradesh, the location of one of the previous randomized controlled trial studies and the location of the Breza and Kinnan (2016) study.
3. We use the traditional nomenclature of 'treatment' and 'control' to indicate treatment with the new health insurance requirement, even though in a sense our intervention was to keep a group of villages untouched by the new health insurance requirement.
4. Something similar should also be true if clients left microfinance and then sought out other more expensive sources of credit.
5. The two districts are Bidar and Gulbarga, which are a few hours' drive from Hyderabad, the capital of Andhra Pradesh and the location of SKS's headquarters.
6. This number reflects the average loan size on renewal in control villages following the roll-out in treatment villages.
7. SKS operation across villages is grouped within branches, of which there are seven in our sample. Within each branch, we also stratified by whether a village had more clients or fewer clients than the branch median.
8. The surveys can be downloaded at <http://dx.doi.org/10.7910/DVN/25890> (accessed 31 March 2018).
9. Surveyors visited households multiple times to interview each adult (over the age of 14), though in some cases they did not find all adults reported to be in the household.
10. Following Kling *et al.* (2007), we create each index in the paper by calculating an equally weighted average across the component characteristics' z -scores. The z -score itself is calculated by subtracting that characteristic's mean in the control group and dividing by the standard deviation in the control group, orienting the sign of each z -score to be in the same conceptual direction (e.g. a larger business). Differences in the index then reflect an average difference in the standard deviation across each component characteristic.
11. We attribute this low attrition rate to the relative stability of these households, and our ability to find households with the help of prominent village members.
12. We experimented with several definitions, but found this one to be most successful at identifying the major health incidents in which we were most interested and that might be underrepresented in the baseline and endline surveys.
13. Though at the beginning we asked about all major health incidents since January 2008, in July we switched to asking about all major health incidents in the last 30 days, in order to improve recall ability of clients and to allow us to visit villages more frequently.
14. Unfortunately, we do not have data on labour supplied to the household business, though in the interpretation section we draw on estimates from Banerjee *et al.* (2015b).
15. This correction generally multiplies the business outcome index p -value by a factor of 3, given that its p -value is generally the lowest among the three outcome indices.
16. Clients' annual loans are repaid over 50 weeks, so our clients would have been eligible to renew between 7 June 2007 and June 2008, since they all had a loan as of 21 June 2007. Since renewal can take place within a short grace period, we have included a six-week period for clients to renew their loan.
17. Based on clients' previous loan expiration dates and the dates of pilot roll-out, we calculate the fraction of clients who would have faced the health insurance requirement when their previous loan expired. If clients' renewal decisions are affected only when the health insurance requirement is binding at the time of their first opportunity for renewal, then the implicit first-stage impact of the treatment is 0.73. We do not observe roll-out dates for 20 villages, but make the conservative assumption that roll-out was immediate in these

- villages. Clients whose previous loan expired prior to June 2007 are assumed not to face the health insurance requirement.
18. In particular, there is no treatment effect on whether clients report missing any payments or the number of missed payments. Further, there is no detectable difference in effect on clients with 'older' businesses (started before SKS began operations) and clients with 'newer' businesses (started after SKS began operations).
 19. Across all control villages, only one client is reported to be enrolled in the insurance programme (and is not reported to receive any insurance benefit).
 20. We ask each adult about their difficulty in performing 15 daily activities, rated on a 5-point scale. We create an index for each adult, averaging across the responses by activity (each normalized to have mean 0 and standard deviation 1), and assign an index for each household by averaging across the adult member indices.
 21. Dow *et al.* (1997) find a similar effect, in reverse, in Indonesia: an increase in health facility fees led to an increase in self-reported health status, as people were less likely to visit the hospital. Zwane *et al.* (2011) find that asking people a long series of baseline survey questions on health tended to make them more likely to buy health insurance, perhaps because it made them aware of the risks. The information campaign could have had the same effect.
 22. Note that we asked business owners about profits directly, rather than calculating the difference between reported revenues and reported costs, so the outcome in column (5) contains additional information compared to the previous columns.
 23. If the impacts on business outcomes were driven solely by impacts on clients that owned a business at endline, then the estimates might be scaled up by factors of 5.6 (for panel A of Table 7) and 3.2 (for panel B). In fact, this is what we find when we restrict the sample to businesses that are still in existence at endline (a potentially endogenous outcome), in which case the estimated impacts on business outcomes are also all statistically significant.
 24. Older businesses are in activities that are broadly similar to those in newer businesses, though they are a few percentage points more likely to be in transportation or a food retail shop, and a few percentage points less likely to be in construction/crafts or a hotel/tea shop.
 25. Table 10 reports corresponding estimates that are expressed in per capita terms that reflect the number of adult equivalent household members based on conversion factors used by Townsend (1994) for rural Andhra Pradesh and Maharashtra.
 26. We do not estimate greater impacts on consumption after a longer period of time, comparing households surveyed further after being exposed to the treatment (after 27 months) to those households surveyed more recently after the treatment roll-out (within 27 months).
 27. We define 'progressive attitudes' by creating an index across girls' responses to five questions concerning whether men should be more educated than women, whether men should eat before women, the ideal age of marriage for women, whether women should have children immediately after marriage, and the ideal number of children.

REFERENCES

- ANGELUCCI, M., KARLAN, D. and ZINMAN, J. (2015). Microcredit impacts: evidence from a randomized microcredit program placement experiment by Compartamos Banco. *American Economic Journal: Applied Economics*, 7(1), 151–82.
- ATTANASIO, O., AUGSBURG, B., DE HAAS, R., FITZSIMONS, E. and HARMGART, H. (2015). The impacts of microfinance: evidence from joint-liability lending in Mongolia. *American Economic Journal: Applied Economics*, 7(1), 90–122.
- AUGSBURG, B., DE HAAS, R., HARMGART, H. and MEGHIR, C. (2015). The impacts of microcredit: evidence from Bosnia and Herzegovina. *American Economic Journal: Applied Economics*, 7(1), 183–203.
- BANERJEE, A., BREZA, E., DUFLO, E. and KINMAN, C. (2015a). Do credit constraints limit entrepreneurship? Heterogeneity in the returns to microfinance. Mimeo.
- , DUFLO, E. and HORNBECK, R. (2014). Bundling health insurance and microfinance in India: there cannot be adverse selection if there is no demand. *American Economic Review: Papers and Proceedings*, 104(5), 291–7.
- , ———, GLENNERSTER, R. and KINMAN, C. (2015b). The miracle of microfinance? Evidence from a randomized evaluation. *American Economic Journal: Applied Economics*, 7(1): 22–53.
- , KARLAN, D. and ZINMAN, J. (2015c). Six randomized evaluations of microcredit: introduction and further steps. *American Economic Journal: Applied Economics*, 7(1), 1–21.
- BREZA, E. and KINMAN, C. (2016). Measuring the equilibrium impacts of credit: evidence from the Indian microfinance crisis. Mimeo.
- BUERA, F. J., KABOSKI, J. P. and SHIN, Y. (2014). The macroeconomics of microfinance. Mimeo.

- CREPON, B., DEVOTO, F., DUFLO, E. and PARIENTE, W. (2015). Estimating the impact of microcredit on those who take it up: evidence from a randomized experiment in Morocco. *American Economic Journal: Applied Economics*, **7**(1), 123–50.
- DOW, W. H., GERTLER, P., SCHOENI, R. F., STRAUSS, J. and THOMAS, D. (1997). Health care prices, health and labor outcomes: experimental evidence. RAND Working Paper no. 97-01.
- HOCHBERG, Y. (1988). A sharper Bonferroni procedure for multiple tests of significance. *Biometrika*, **75**(4), 800–2.
- KARLAN, D. and ZINMAN, J. (2016). Long-run price elasticities of demand for credit: evidence from a countrywide field experiment in Mexico. Mimeo.
- KLING, J. R., LIEBMAN, J. B. and KATZ, L. F. (2007). Experimental analysis of neighborhood effects. *Econometrica*, **75**(1), 83–119.
- MEAGER, R. (2016a). Aggregating distributional treatment effects: a Bayesian Hierarchical analysis of the microcredit literature. Mimeo.
- (2016b). Understanding the average impact of microcredit expansions: a Bayesian hierarchical analysis of 7 randomized experiments. Mimeo.
- TAROZZI, A., DESAI, J. and JOHNSON, K. (2015). The impacts of microcredit: evidence from Ethiopia. *American Economic Journal: Applied Economics*, **7**(1), 54–89.
- THORNTON, R. L., HATT, L. E., FIELD, E. M., ISLAM, M., SOLIS DIAZ, F. and GONZALEZ, M. A. (2010). Social security health insurance for the informal sector in Nicaragua: a randomized evaluation. *Health Economics*, **19**, 181–206.
- TOWNSEND, R. M. (1994). Risk and insurance in Village India. *Econometrica*, **62**(3), 539–91.
- WYDICK, B. (2016). Microfinance on the margin: why recent impact studies may understate average treatment effects. *Journal of Development Effectiveness*, **8**(2), 257–65.
- ZWANE, A. P., ZINMAN, J., VAN DUSEN, E., PARIENTE, W., NULL, C., MIGUEL, E., KREMER, M., KARLAN, D. S., HORNBECK, R., GINE, X., DUFLO, E., DEVOTO, F., CREPON, B. and BANERJEE, A. (2011). Being surveyed can change later behavior and related parameter estimates. *Proceedings of the National Academy of Arts and Sciences*, **108**(5), 1821–6.