

Profits and Social Impacts: Complements vs. Tradeoffs for Lenders in Three Countries

[Susan Athey](#)

Stanford University

[Bruno Fava](#)

Stanford University

[Dean Karlan](#)

Northwestern University and IPR

[Adam Osman](#)

University of Illinois Urbana-Champaign

[Jonathan Zinman](#)

Dartmouth College

Version: September 25, 2026

DRAFT

Please do not quote or distribute without permission.

Abstract

The canonical approach to corporate governance posits a tradeoff between maximizing shareholder profits and maximizing other aspects of shareholder welfare such as social impact. Advances in machine learning that enhance the firm's ability to target specific customers may exacerbate or ameliorate this tradeoff. The authors estimate these tradeoffs using data from randomized microcredit approvals in South Africa, the Philippines, and Bosnia. They examine social impact on two dimensions: credit access (i.e., reaching disadvantaged groups that typically have less access to financial markets) and impact (i.e., treatment effects on household income). Two of the three lenders could have increased average loan profit margin by 7–9 percentage points through machine learning-based targeting. However, such targeting pushes against social impact goals vis-à-vis credit access; specifically, women and lower-income households would have been more excluded (the implied impact on borrowers' income, also central to the dual goal of social impact, is too imprecisely estimated to draw an inference). To gauge the magnitude of the first tradeoff, the authors examine how profits would change if the lenders altered whom they lent to within each quintile of baseline borrower income, thus holding the distribution broadly similar but fine-tuning the targeting within income bands. Such a constraint would lower the profit gains of targeting by about half. These findings highlight the importance of quantifying tradeoffs and complementarities when deciding what to maximize.

1 Introduction

A central debate in corporate governance concerns the proper objective(s) of the firm. While Friedman (1970) argues that a firm’s sole social responsibility is to maximize profits for its shareholders, a growing literature argues that firms should instead consider shareholder welfare more broadly, including objectives beyond pure profit (Hart and Zingales, 2017, 2022). This debate has become more salient with the rise of impact investing and business initiatives to adopt environmental, social, and governance objectives (Gillan et al., 2021; Pástor et al., 2021). At its heart sits a critical empirical question: do firms actually face a tradeoff between profits and the welfare of the people they serve, and if so, how steep is it?

We study this tradeoff using data from three lenders providing financial services for low-income households and enterprises. Financial services is the largest sector within the impact investing space.¹ We combine application and repayment data with experimental variation and household and enterprise surveys from three randomized controlled trials. More broadly, our machine learning approach works for any client-facing, double-bottom-line enterprise weighing profits against pro-social goals such as reaching underserved populations or improving client welfare.

Estimating the tradeoffs between a singular focus on profits and a broader focus on shareholder welfare is challenging. These objectives map into two different kinds of econometric problems. Maximizing profits, in our case for lenders, is primarily a prediction problem: given observable borrower characteristics and loan terms, lenders work to predict repayment and profits. In contrast, assessing the social consequences of profit targeting requires examining both who receives credit and how credit affects them. While administrative data can provide insight on *who* gets credit, estimating how credit access changes borrower welfare instead requires causal inference. Although firms often collect the administrative data necessary to forecast profits, they rarely have the necessary data on borrower or community outcomes combined with the exogenous variation required to estimate welfare effects, making measuring the tradeoff in the same setting difficult. But understanding this tradeoff is becoming even more important as machine learning and richer data steadily improve lenders’ ability to target on profits (Berg et al., 2020; Fuster et al., 2022), fueling concerns among pol-

¹Hand et al. (2023) reports industry data and splits financial inclusion into “Financial Services (excluding microfinance)” and “Microfinance.” When combined, it is the largest sector.

icymakers and regulators that algorithmic targeting may systematically disadvantage certain groups (Consumer Financial Protection Bureau et al., 2023). Without comparable progress on these social objectives, it is difficult to know whether sharper profit targeting comes at the expense of borrower inclusion and outcomes, or whether profits and social objectives can be advanced together.

The three microcredit studies are from South Africa (Karlan and Zinman, 2010), the Philippines (Karlan and Zinman, 2011), and Bosnia (Augsburg et al., 2015).² We first train machine learning algorithms to predict the profitability of each loan applicant and then use the random assignment of loans to estimate the welfare impacts of targeting credit based on these predictions. In each study the experimental sample consists of marginal applicants—those who would otherwise have been rejected or who fell near the lender’s approval threshold—so all of our estimates describe what targeting could achieve within this marginal pool rather than across the lender’s full applicant base. In two of the three cases, the lender could have increased profits within this marginal pool by changing its allocation policies in line with the predictions made by the algorithms. The average loan profit margin would have increased by 7 percentage points in South Africa and 9 percentage points in the Philippines. These techniques are not silver bullets—in Bosnia the data and algorithms are unable to reliably predict which marginal applicants to target in order to maximize lender profits.

We then assess how downstream outcomes would change if lenders allocated funds to maximize profits based on these algorithms. First, we find that if lenders approved only the highest profitability clients, they would shift lending away from groups that often struggle to access credit: female, less-educated, and lower-income borrowers. This would exacerbate existing gaps in financial inclusion (Demirgüç-Kunt et al., 2022).

Next, thanks to the randomization, we can estimate how lender adoption of profit-maximizing targeted policies affects the distribution of borrower outcomes across applicant groups. We group applicants by the predicted profit score and, within each group, compare the income of those offered a loan with similar people who were not. The estimates are positive but noisy. We find increases of 18% and 47% for South Africa and the Philippines, respectively, with standard errors of 31% and 37%.

We then consider what would happen if regulation required that any firm’s targeting strategy leave the average baseline income of borrowers unchanged. Such a rule targets

²We focus on these studies because we require publicly available data with individual-level randomization along with administrative data that allows for the construction of lender profits.

potential policymaker concerns that profit-maximizing lending could widen gaps by characteristics like gender or income. We analyze a balanced lending policy where the lender splits borrowers up into five quintiles based on their baseline income and then lends to the most profitable 33% of borrowers within each quintile. With this policy, lenders can still try to increase their own profits without changing the average income of their borrower pool. This balanced approach limits firms’ ability to benefit from algorithmic targeting, with an estimated increase in profits 36% smaller than the unrestricted case in the Philippines. This policy is even more restrictive in South Africa, erasing about half of the lender’s increase in profits from targeting.

Next, we ask if it is possible for lenders to implement a policy that directly targets their loans to individuals who would experience the largest positive change in their household income in response to the loan, irrespective of their profitability to the lender. We find that this goal is much more difficult to achieve with current machine learning strategies and existing data. We are unable to reliably predict who will benefit more from the loan in any of the three datasets.

We contribute to three literatures in economics and finance. First, a growing literature uses machine learning to design and evaluate personalized policies, often considering tradeoffs between firm performance (such as profits or engagement) and consumer welfare. Closest to our setting is Dubé and Misra (2023), which studies the welfare implications of profit-maximizing personalized pricing. They use data from a randomized pricing field experiment to estimate a structural demand model and then simulate profits and consumer surplus under alternative pricing policies. In parallel, a large methodological and applied literature develops and implements data-driven treatment rules and targeting policies in field settings, including labor markets, social programs, and digital platforms (e.g., Mullainathan and Spiess (2017), Kitagawa and Tetenov (2018), Knaus et al. (2022), Athey et al. (2023), Yoganarasimhan et al. (2023), Athey et al. (2024), Athey et al. (2025), Haushofer et al. (2025)). Most of this literature does not analyze how optimizing supplier profits affects user welfare,³ and it typically estimates social impacts captured either through model-based measures of surplus or through intermediate outcomes such as program engagement or equitable access to a

³An important exception is Athey et al. (2025), which uses a calibrated model to analyze how the (counterfactual) introduction of algorithmic recommendation systems in the Kiva lending platform highlights tradeoffs between speed of raising money (similar to supplier profits) and targeting objectives (similar to social impact).

service. We contribute to this literature by explicitly estimating potential tradeoffs between supplier profits and experimentally identified effects on borrowers' incomes as well as targeting success to underserved populations, and using these estimates to examine how policy targeting affects realized downstream outcomes.

Second, we contribute to the literature on impact investing and corporate social responsibility. Many microfinance organizations claim a dual mission—helping their communities while making a profit. Most Fortune 500 companies have made large investments in corporate social responsibility that introduce elements of a double bottom line (Kitzmüller and Shimshack, 2012; Barber et al., 2021). In emerging markets, perhaps the largest “impact investor” is the International Finance Corporation under the World Bank Group, and estimates of its financial returns indicate it surpassed a risk-adjusted comparison to the U.S. market index, i.e., the inclusion of “impact” goals did not lead to a reduction in financial returns (Cole et al., 2026). Brest et al. (2019) argues that no reduction in risk-adjusted financial returns implies no additionality, and thus no social impact. We build on this debate by examining the social impact directly, rather than inferring it from a combination of theory and inference from the financial returns. Furthermore, we posit that as firms adopt technology that enables them to use targeting to screen customers or present customized offers, the tradeoffs between profit maximization and achieving more community-focused outcomes could become starker. We show how a combination of machine learning and experimental data can help tackle assessing the direction and magnitude of this tradeoff. As these methods become more powerful, and more prevalent, the prescriptions to business policy could get more precise and that could of course lead to starker and more precisely estimated tradeoffs. We also show how regulations that aim to limit lenders' use of targeting could come at a real cost to the lender.

Finally, we also contribute to the literature on the heterogeneous effects of microcredit (Banerjee et al., 2019; Beaman et al., 2023; Bryan et al., 2024; Crépon et al., 2024; Fava, 2026). Much of the literature focuses on how some borrowers can benefit from loans more than others. In contrast, we show how certain borrowers can benefit the *lender* more than others. Our findings imply that the distribution of microcredit's benefits may be influenced not only by borrower characteristics and behaviors, but also by lenders' targeting strategies. This insight underscores the need for a more nuanced understanding of microcredit's impact, one that considers both sides of the

lending equation, particularly as for-profit lenders are far more common now than before. Moreover, it raises important questions about the long-term sustainability and social impact of microcredit programs in an era of data-driven decision-making. This is also true more generally for a wide variety of social programs, where policy-makers will often use proxies like measures of poverty to allocate support that may not align with which individuals are most likely to benefit from those programs (Haushofer et al., 2025). We show how allocation decisions at for-profit entities can also lead to nuanced changes in who gains access, with implications not only for equity across groups but also for the effectiveness of programs designed to expand economic opportunity.

2 Data

We have two goals that require different kinds of data. First, to assess how machine learning strategies could affect lender profits we need baseline information about the borrower along with repayment data to assess how different borrowers performed with the loan. Second, to assess the causal impacts that a change in lending policies would have on borrower outcomes we need data that includes exogenous variation in borrowing, along with information on changes in borrower income over time. Fortunately, a few randomized evaluations of credit interventions have all of these features and publicly available data. We identified three such datasets: Augsburg et al. (2015), Karlan and Zinman (2011), and Karlan and Zinman (2010).

Table 1 below provides some summary statistics about the different experiments and associated datasets. Each experiment and dataset has its own set of idiosyncrasies. They each used their own survey instruments and collected different sets of variables. Since there is limited overlap in variables collected across datasets, we chose to analyze each experiment separately. This enables us to analyze sources of heterogeneity experiment by experiment, identifying differences in empirical patterns as well as common lessons across contexts. As we will see below, different contexts indeed show different patterns of heterogeneity.

While each study generates exogenous variation in whether an individual in the experimental sample receives credit, each study does so differently. The experiment in Bosnia and Herzegovina focused on extending credit to a poorer segment of the population typically deemed too risky to lend to. Loan applicants who were marginally

rejected by a large microfinance institution were randomly assigned to a treatment group that was offered a loan or a control group that was not. 98% of the treatment group took the loan. We discuss incomplete loan take-up below.

In the Philippines researchers worked in collaboration with a for-profit bank providing individual liability microloans to entrepreneurs near Manila. The study used a credit-scoring model to identify marginally creditworthy applicants. These marginal applicants were then randomly assigned to either receive a loan offer (treatment) or be rejected (control), with a stratified design that had a higher treatment probability for the group with higher creditworthiness and a lower probability for the group with lower credit scores.⁴

In South Africa, researchers partnered with a for-profit lender offering high-interest consumer credit to employed individuals. Loan officers first labeled rejected applications as “egregious” or “marginal.” A digital randomizer then instructed officers to approve a randomly selected subset of the marginally rejected applicants, with probabilities based on credit score. Control applicants remained rejected. Loan officers had final discretion, so not all randomized approvals resulted in disbursed loans.⁵

We consider one primary outcome for lenders, Lender Profits, and one for borrowers, Borrower Income. In each experiment, we measure Lender Profits using the lender’s administrative data on repayment. We calculate the net present value of all repayments made during the loan term, including interest and any fees or penalties paid by borrowers. We subtract the present value of the disbursed capital along with an estimate of the loan’s origination costs (this attempts to include things like the cost of loan officer time, overhead, etc.). We estimate these costs using data from the Microfinance Information Exchange Market dataset (World Bank, 2024): a variable cost, computed as financial expenses divided by the gross loan portfolio and then multiplied by the individual loan amount, and a fixed cost per loan, computed as total operating expenses divided by the number of loans.⁶ We analyze profits in two forms: in levels (USD PPP) and as a percentage of loan size. For profits in levels, we subtract the mean profit across borrowers within each dataset so that the average is normalized to zero. For profits as a percentage, we first divide each borrower’s profit in levels (demeaned)

⁴72% of the treatment group took up the loan.

⁵53% of the treatment group took up the loan.

⁶In South Africa we use a fixed cost estimate more relevant for the specific lender in our data (Coetzee et al., 2003).

by their loan size and then subtract the cross-borrower mean of these ratios, again normalizing the average to zero. In both cases, positive values indicate above-average profitability and negative values indicate below-average profitability.⁷

For the experiments in South Africa and Bosnia, we measure borrower income using reported income. For the experiment in the Philippines, the lender’s survey does not collect household income, and so we use the borrowers’ reported business profits instead.

Since not everyone assigned to treatment in these studies actually borrowed, we define treatment as being *offered* credit and estimate intention-to-treat (ITT) effects throughout. In this setting, incomplete take-up is part of the object of interest rather than a nuisance. Individuals in the treatment group who do not take up the loan experience zero treatment effect from the loan offer itself. These zeros are policy-relevant: under any targeting rule that directs loan offers to similar individuals, they would continue to decline the product; therefore the social impact of such offers would be zero.

3 Methodology

For each experiment, our methodology proceeds in three steps. First, we estimate predicted loan profitability conditional on borrower characteristics using a flexible machine learning model. We then use this model to calculate lender profits under a counterfactual policy of offering loans only to borrowers predicted to be in the top tercile of profitability. Specifically, we evaluate (i) the performance of this strategy for improving lender profits; (ii) the impact on borrowers’ income; and (iii) the characteristics of borrowers who are prioritized under this targeting rule.

Second, we evaluate distributional impacts by comparing two lending policies: one that targets only the most profitable borrowers, versus one that ensures balanced lending across baseline income quintiles. Third, we estimate heterogeneous treatment effects on borrower outcomes using machine learning, then evaluate counterfactual targeting policies designed to maximize borrower impact. All results were produced using

⁷Because lender profit distributions contain extreme outliers, we winsorize profits at the 5th and 95th percentiles for our main specifications. Results using unwinsorized profits are reported in Appendix D and are qualitatively similar.

the `ensembleHTE` R package.⁸

To formalize our approach, we introduce some notation. We describe our methodology using one dataset to simplify exposition, and we implement it across all three datasets. We use n for the sample size, assume $(B_i, L_i, X_i, A_i, D_i)_{i=1}^n$ is an i.i.d. sample where B is a borrower outcome (such as endline household income), L is a lender outcome (e.g., lender profits from the experimental loan), X is a set of borrower covariates collected before randomization, A is the treatment assignment indicator, and D is an indicator equal to 1 if a treated individual took the loan, and 0 otherwise. We implement the same methodology for different outcomes. For example, B may denote borrower income either in levels or logs, and L lender profits, either in levels or per dollar lent. Denote the sample by $S = \{1, \dots, n\}$, and the subsample of borrowers by S_B , that is, $S_B = \{i \in S : D_i = 1\}$. We use $p(x)$ to denote the propensity score, the probability that an individual with characteristics x is assigned to treatment. In our datasets, $p(x)$ is constant in Bosnia, but not in the Philippines or South Africa where treatment probabilities depended on creditworthiness.

3.1 Targeting Lender Profits

We follow the framework of Fava (2025) to train machine learning (ML) algorithms to predict lender profits and evaluate how a targeting policy based on these predictions affects both borrowers and lenders. Fava (2025) establishes a central limit theorem for a large class of split-sample estimators under weak regularity conditions, and lays out a new strategy to learn heterogeneous treatment effects within the Generic ML framework of Chernozhukov et al. (2025). Compared to the latter, the strategy in Fava (2025) (i) uses more data for training the ML predictors, (ii) uses the entire sample for constructing confidence intervals, and (iii) combines predictions from multiple ML algorithms, potentially improving accuracy and avoiding issues of multiple hypothesis testing. To ensure reproducibility of our results, we repeat the following procedure 500 times and report average estimates and standard errors. We take the following steps:

1. Randomly split the sample into three *folds* of approximately equal size.
2. Using two folds at a time and only sample S_B (individuals who were treated and took the experimental loan), we train four ML algorithms (Extreme Gradi-

⁸<https://github.com/bfava/ensembleHTE>

ent Boosting, Regression Forest, Single-layer Neural Network, Elastic Net)⁹ to estimate the expected value of lender profits (for the borrowers) conditional on covariates,

$$\mu_L(x) = \mathbb{E}[L|X = x, D = 1].$$

For each observation $i \in S$, we predict lender profits using the four predictors trained on the folds that do not contain that observation.¹⁰ This gives predictions $\hat{L}_{i,j}$ for each observation i and ML algorithm $j = 1, \dots, 4$.

We combine the four predictions into a final $\hat{L}_i$ using the *ensemble* strategy of Fava (2025), described in Appendix A.

3. Split the observations into terciles of $\hat{L}_i$ within each fold. Define indicator variables $G_{i,j}$, where $G_{i,j} = 1$ if observation i falls into the j -th tercile ($j = 1, 2, 3$), and 0 otherwise. For instance, $G_{i,3} = 1$ indicates that the predicted lender profits of observation i , $\hat{L}_i$, is among the top one-third predictions within that fold.
4. Define two targeting policies, where the output of the policy function $\pi(\cdot)$ is an indicator of whether credit is offered (1) or not (0). The first is a targeted policy where the lender offers credit exclusively to the top tercile of predicted performers,

$$\hat{\pi}_L(X_i) = G_{i,3}.$$

The second is a benchmark policy where the lender offers credit to the entire population of approved applicants,

$$\bar{\pi}(X_i) = 1.$$

Next, we describe how we evaluate our ML models' predictive performance and compare the policies $\hat{\pi}_L$ and $\bar{\pi}$ in terms of their impact on both lenders and borrowers.

⁹Implemented using the R packages `xgboost` (Chen et al., 2025), `grf` (Tibshirani et al., 2025), `nnet` (Venables and Ripley, 2002), and `glmnet` (Friedman et al., 2010), respectively.

¹⁰An alternative approach trains the prediction model on all treated individuals, including those who were offered a loan but did not take it up, with their lender profits set to zero (i.e., the portfolio mean after normalization; see Appendix E). This intent-to-treat (ITT) specification captures the expected profitability of *offering* a loan rather than the profitability conditional on *borrowing*. Results using ITT predictions are reported in Appendix E and are qualitatively similar.

We start by estimating the regression

$$L_i = \beta_1^L G_{i,1} + \beta_2^L G_{i,2} + \beta_3^L G_{i,3} + \varepsilon_i \quad (i \in S_B) \quad (1)$$

on the subset of individuals who took the experimental loan, from all three folds. The coefficients from this regression recover the average profits generated for the lender by lending to the individuals in each tercile of predicted profits.

We evaluate the predictive performance of our ML models by comparing the average lender profits generated by the predicted top and bottom performers, which is captured by the difference $\beta_3^L - \beta_1^L$. To compare policies $\hat{\pi}_L$ and $\bar{\pi}$, we estimate the difference in expected lender profit,

$$\beta_3^L - \left(\frac{1}{3}\beta_1^L + \frac{1}{3}\beta_2^L + \frac{1}{3}\beta_3^L \right).$$

The first term (β_3^L) is the profit of the average loan for the lender under the targeted policy $\hat{\pi}_L$, while the second term is the profit of the average loan under the benchmark policy $\bar{\pi}$, which is equivalent to the average profit when lending to the entire experimental sample of marginal applicants.

We calculate robust standard errors.

To evaluate the impact on borrowers, we use the full sample to estimate the regression

$$B_i = \beta_0^B + \sum_{j=1}^3 \beta_j^B G_{i,j} (A_i - p(X_i)) + \varepsilon_i \quad (i \in S) \quad (2)$$

with weights $W_i = (p(X_i)[1-p(X_i)])^{-1}$, as in the ‘‘GATES’’ regression of Chernozhukov et al. (2025). These weights adjust for unequal treatment assignment probabilities, which vary with credit score in the Philippines and South Africa datasets. With this weighting, the coefficients in (2) identify the average treatment effect (ATE) on household income (B_i) of offering the experimental loan for each subgroup of predicted lender profits, that is, the intention-to-treat effect. Similarly, we estimate the difference in ATE between top and bottom terciles, $\beta_3^B - \beta_1^B$, and difference in ATE under policy $\hat{\pi}_L$ and $\bar{\pi}$,

$$\beta_3^B - \left(\frac{1}{3}\beta_1^B + \frac{1}{3}\beta_2^B + \frac{1}{3}\beta_3^B \right).$$

Finally, we assess the characteristics of individuals selected by the targeting policy $\hat{\pi}_L$. Specifically, we compare mean values between those in the top tercile versus

the bottom two terciles of predicted lender profits for the following characteristics: age, college graduate status, gender, household income, loan size, and share of profits coming from penalties. Standard errors are the sample standard deviation divided by the square root of the sample size.

3.2 Balanced Targeting Strategies

We analyze a restricted lending policy that requires balance across quintiles of baseline household income, denoted $\hat{\pi}_R$. The approach is very similar to the main specification in subsection 3.1, with one key modification in step 3 and in how we define the targeting policy. Let $Q(i) \in \{1, 2, 3, 4, 5\}$ denote which quintile of the baseline income distribution observation i belongs to, so $Q(i) = 1$ represents the lowest quintile and $Q(i) = 5$ the highest. In Subsection 3.1, the indicators $G_{i,j}$ equal 1 if observation i is in the j -th tercile of predicted lender profits $\hat{L}_i$ calculated across the entire sample. For the balanced policy $\hat{\pi}_R$, we instead calculate terciles of $\hat{L}_i$ separately within each income quintile.¹¹ Specifically, $G_{i,j}^R$ equals 1 if observation i is in the j -th tercile of predicted profits among all observations with the same baseline income quintile $Q(i)$. Hence, a lender following the policy $\hat{\pi}_R(X_i) = G_{i,3}^R$ will select one-third of borrowers from each quintile of baseline income, ensuring a balanced composition of borrowers across the baseline income distribution. After defining the dummies $G_{i,j}^R$ in this way, we run the same specifications in (1) and (2), replacing $G_{i,j}$ with $G_{i,j}^R$.

3.3 Targeting Borrower Impact

We train ML models to detect borrowers with the largest treatment effects on household income. We then evaluate how well these models perform at detecting such borrowers, and how a lending policy based on their predictions affects lender profits. Our approach is similar to that of Subsection 3.1, where we used ML to target lender profits.

While we can train a model to predict lender profits using only a representative sample of borrowers, predicting borrower impact leverages the experimental variation from random assignment to estimate individual-level treatment effects. Specifically, we

¹¹As in subsection 3.1, we also calculate terciles separately by fold. That is, we first separate all observations into quintiles of baseline income. Within quintile, we separate observations by which of the three folds they belong to. Within quintile-fold, we split observations into terciles of predicted lender profits.

modify step 2 by adopting an R-learner strategy (Nie and Wager, 2021): each of the four ML algorithms provides first-stage nuisance estimates to residualize the outcome, after which we estimate the conditional average treatment effect (CATE) function using causal forests as implemented in the R package `grf` (Athey et al., 2019), yielding a predicted individual treatment effect for each borrower.

Again, we combine the predictions of the four R-learners following Fava (2025), and define dummies $G_{i,j}$ based on terciles of predicted treatment effects. Finally, we run the same specifications as in (1) and (2). We provide a detailed explanation in Appendix B.

3.4 Aggregating Estimates

We report estimates separately by dataset, as well as a “joint test” that aggregates all datasets. The aggregated point estimate is the unweighted average of the three dataset-specific point estimates, and the resulting aggregated standard error assumes the three dataset-specific standard errors are independent. That is, the aggregated squared standard error is the average of the squared standard errors divided by 3.

4 Results

Targeting Lender Profits

Table 2 describes what could happen if lenders used algorithms specifically designed to maximize their profits. Panel A focuses on maximizing profits as a percentage of the loan size (the loan profit margin) while Panel B maximizes total profits from each loan reported in absolute terms (USD).

Columns 1 and 2 in Panel A show how the loan profit margin differs for those predicted to be in the top 33% of profitable borrowers versus those who are predicted to be in the bottom 33% of profitable borrowers. In both South Africa and the Philippines there is an identifiable group of borrowers who bring in a sizable profit, and a group who cost the lender money. The difference is reported in Column 2 with standard errors in parentheses. The differences between the top and bottom terciles are statistically and economically significant, providing evidence of heterogeneity in line with the methods

used in Chernozhukov et al. (2025).¹²

In Column 3 we consider how a policy that lends only to those in the top tercile of profitability would affect lender outcomes relative to the status quo. In South Africa the lender can increase the average loan profit margin by 7.4 percentage points by focusing their lending on the most profitable borrowers and ending their lending to the most costly borrowers. In the Philippines the lender could increase the loan profit margin by 8.6 p.p. by following this strategy. In Bosnia the data and underlying heterogeneity are not sufficient for the lender to improve their targeting. The top row provides an unweighted average of all three studies, showing an average 4.8 p.p. increase.

Columns 4–6 and 7–9 hold the terciles fixed while looking at how the loans causally impact borrower income. For example, column 7 looks at the difference in income between individuals in treatment who were predicted to be in the bottom 33% of lender profits with those in control who were also predicted to be in the bottom 33% of lender profits. In Column 6 we report how the average treatment effect on borrower income would change if the lenders implemented a lending policy that only lent to the most profitable third of borrowers relative to the status quo.

These estimates test whether the profit gains that lenders can make from improved algorithmic targeting would come at the expense of borrowers. In South Africa and the Philippines—the two contexts where the algorithm successfully identifies more profitable borrowers—the point estimates suggest that when lenders target those who would yield the greatest lender profits, the borrowers themselves are the type who benefit more from the loans. In Bosnia the algorithm fails to predict who is more profitable to begin with (Column 3), so the borrower-impact comparison in Columns 6 and 9 does not separate genuinely more from less profitable borrowers and the negative point estimates there are uninformative about the underlying tradeoff. While none of the borrower-impact estimates reach conventional levels of statistical significance, the coefficients in the contexts where targeting works all point towards a lack of a tradeoff between the lender’s bottom line and borrower outcomes. This is true when we consider targeting based on lender profits in percent or in levels, and when looking at borrower outcomes in percent and levels.

Panel B replicates the analysis but instead attempts to maximize the total amount of lender profits rather than the percent relative to loan size. In this case we find strong

¹²As mentioned above we winsorize at the 5th and 95th percentiles to decrease variance. Appendix D replicates the analysis without winsorizing and finds qualitatively equivalent results.

evidence of the benefits of targeting in South Africa and the Philippines.

Table 3 looks at how the composition of borrowers would change if lenders exclusively targeted borrowers who are predicted to be in the top 33% of profitability for the lender. In Column 1 of Panel A we report that the average South African borrower in the high profitability group is 35 years old, which is about one year older than clients who are in the other two less profitable terciles for the lender. Similarly, 8% of the targeted borrowers are college educated, while only 1% of borrowers in the less profitable terciles are.

Overall we see that in South Africa targeted borrowers are richer and more educated. The loans are also larger on average. This pattern holds for both strategies of targeting, either on profits in levels or percent. In the Philippines the story is similar, with the main difference being that high profit borrowers are younger relative to everyone else and more likely to be male. These results show that strategies that target borrowers based on profits for the lender are very likely to lead to compositional shifts in *who* is getting access to credit, although how it changes borrower composition can differ based on the context.

Finally, we also consider how the share of profits from penalties changes across individuals in the most and least profitable groups. One potential concern is that the way lenders maximize profits is by targeting borrowers who often struggle to repay on time and so are paying more in penalties relative to standard interest charges. This could be seen as a predatory lending strategy that comes directly from the algorithm. The table shows that there are no statistically significant differences in the proportion of profits that are coming from penalties for the highest profit borrowers relative to everyone else, and that the estimates are small and go in the opposite direction in the Philippines.

Altogether this suggests that existing data and machine learning strategies are sufficient to allow lenders to develop targeting strategies that would demonstrably increase profits. We find no evidence that targeting to maximize profits would interfere with a potential “double bottom line” concern that lenders need to sacrifice borrower outcomes to maximize their own profits, or vice versa. However, the borrower-income estimates are too imprecise to rule out a meaningful tradeoff, and this finding should therefore be read as inconclusive rather than as evidence that profit and borrower welfare are aligned. Nonetheless, if implemented, these strategies would change the

composition of borrowers in ways that could meaningfully change who gets access to credit and who does not.

Balanced Targeting Strategies

Policymakers could be concerned that algorithm-based profit-maximizing strategies could lead to undesired changes in credit access. For example, these strategies could favor richer borrowers, as was the case in our samples in South Africa and the Philippines. If lenders were restricted in how they could target their loans, for example, to ensure that the average income of their borrowers does not change, would they still be able to improve profits, and if so, how might it affect other related outcomes?

We address this hypothetical in Table 4. Column 1 reproduces the results from Table 2 (column 3 in Table 2), showing the difference in lender profits between lending to only those predicted to be in the top 33% of profitability relative to current policy. Columns 4–6 consider a “balanced” targeting strategy that would split the sample into five quintiles based on borrower income at baseline. The lender would then implement a strategy that targets the top 33% *in each baseline income quintile*, instead of the top 33% irrespective of their baseline income. This would ensure that the average income of those that get loans is similar to what it was before the targeting exercise.

In Panel A we see that in each country a balanced lending strategy would decrease the profitability of targeting. For example, in South Africa unrestricted targeting increased loan profit margin by 7.4 p.p., while the balanced approach would increase the profit margin half as much, by only 3.6 p.p. In the Philippines, it drops from 8.6 to 5.5 p.p., a 36% reduction. On average we see a drop from 4.8 p.p. to 2.5 p.p., decreasing the benefits by around half.

Columns 5 and 6 show how borrower impacts would change in this balanced targeting strategy. While these borrower-impact differences are not statistically different from those in Columns 2 and 3, they are suggestive of the balanced targeting strategies being *worse* for both the lender *and* the borrower, but this result is noisy.

Policymakers could suggest any number of restrictions to targeting rules that could affect which characteristics are regulated. Here we only consider the most salient one in this context—baseline income—and find notable tradeoffs.

Targeting on Loan Impacts on Borrower Income

Alternatively, a lender could try to explicitly target on improving borrower outcomes. Table 5 presents this analysis, splitting borrowers into terciles based on the predicted impact of the loan on their income. We find that predicting borrower impact is much more difficult than predicting lender profit. Across all three settings we are unable to reliably find strategies that can target borrowers who benefit more from the loans. While point estimates are typically positive in South Africa, the standard errors are large. For the Philippines we find a negative estimate in Panel A and a positive one in Panel B. We see the opposite in Bosnia. The general difficulty of this exercise suggests that current data and algorithms limit the ability to greatly improve borrower outcomes through targeting alone.

Table 6 reveals the characteristics of borrowers who are predicted to benefit most from the loans. Given the struggle for the algorithms to effectively target on this outcome it is not surprising that there are minimal differences between groups, although it once again suggests that those with high incomes are more likely to benefit from the loans.

5 Discussion and Conclusion

This work sheds light on the complex tradeoffs lenders may face as they incorporate machine learning into their loan allocation processes. Our findings reveal that while these algorithms can significantly boost lender profits, they also change the composition of borrowers, potentially in ways that could exacerbate socioeconomic inequalities. Though the algorithms appear effective in predicting profitability, they are less successful in identifying borrowers who would experience the greatest improvements in household income. This could change as algorithms improve over time and as more and different kinds of data become available to lenders.

The balanced lending strategy we propose shows that there are potential pathways for lenders to achieve profitability gains while lessening impacts on borrower diversity. But these restricted targeting strategies are not as effective, making them costly to the lenders. Although our example focuses on ensuring balance on borrower baseline income, future work could consider other strategies including multi-dimensional approaches that aim to balance several variables that are typically considered important,

including gender and education.

One limitation of our work is statistical power. Because each study uses a different survey instrument and collects a different set of covariates, we train and evaluate the ML algorithms separately within each experiment, which leaves us with effective samples of roughly 500 to 1,100 observations per study—and considerably fewer when the analysis is restricted to those who took up the loan, as is the case for the lender-profit predictors. Machine learning algorithms, especially those that assess heterogeneous treatment effects, typically require large amounts of data to effectively model the complex, high-dimensional patterns in treatment effect heterogeneity involving multiple covariates. Identifying opportunities to implement these strategies at a much larger scale (which many lenders already have access to) may provide further insights into the potential and pitfalls of these algorithmic targeting strategies.

These tradeoffs are likely to only get starker over time. Predicting lender profits is easier to do with the kinds of data lenders already collect and more in line with the traditional bottom line. Predicting borrower impact, in contrast, requires experimental variation in credit access that few lenders will ever generate on their own. Hence, the profit side of the frontier will likely advance faster than the impact side, and the gap between what is easy for lenders to target and what is worth targeting will widen. As that gap widens, the double bottom line may be harder for lenders to credibly maintain on their own and more a standard that regulators, impact investors, or mission-driven capital may have to actively enforce.

References

- ATHEY, S., U. BYAMBADALAI, M. CERSOSIMO, K. KOUTOUT, AND S. NATH (2024): “The Heterogeneous Impact of Changes in Default Gift Amounts on Fundraising,” Available at SSRN 4796598.
- ATHEY, S., S. A. COLE, S. NATH, AND S. J. ZHU (2023): “Targeting, personalization, and engagement in an agricultural advisory service,” Working paper, SSRN 4536641.
- ATHEY, S., D. KARLAN, E. PALIKOT, AND Y. YUAN (2025): “Smiles in Profiles: Improving Fairness and Efficiency Using Estimates of User Preferences in Online Marketplaces,” NBER Working Paper 30633.
- ATHEY, S., J. TIBSHIRANI, AND S. WAGER (2019): “Generalized random forests,” The Annals of Statistics, 47, 1148 – 1178.
- AUGSBURG, B., R. DE HAAS, H. HARMGART, AND C. MEGHIR (2015): “The impacts of microcredit: Evidence from Bosnia and Herzegovina,” American Economic Journal: Applied Economics, 7, 183–203.
- BANERJEE, A., E. BREZA, E. DUFLO, AND C. KINNAN (2019): “Can microfinance unlock a poverty trap for some entrepreneurs?” Tech. rep., National Bureau of Economic Research.
- BARBER, B. M., A. MORSE, AND A. YASUDA (2021): “Impact investing,” Journal of financial economics, 139, 162–185.
- BEAMAN, L., D. KARLAN, B. THUYSBAERT, AND C. UDRY (2023): “Selection into credit markets: Evidence from agriculture in Mali,” Econometrica, 91, 1595–1627.
- BERG, T., V. BURG, A. GOMBOVIĆ, AND M. PURI (2020): “On the Rise of FinTechs: Credit Scoring Using Digital Footprints,” The Review of Financial Studies, 33, 2845–2897.
- BREST, P., R. J. GILSON, AND M. A. WOLFSON (2019): “Essay: How Investors Can (and Can’t) Create Social Value,” The Journal of Corporation Law, 44, 205–228.

- BRYAN, G., D. KARLAN, AND A. OSMAN (2024): “Big loans to small businesses: Predicting winners and losers in an entrepreneurial lending experiment,” American Economic Review, 114, 2825–2860.
- CHEN, T., T. HE, M. BENESTY, V. KHOTILOVICH, Y. TANG, H. CHO, K. CHEN, R. MITCHELL, I. CANO, T. ZHOU, M. LI, J. XIE, M. LIN, Y. GENG, Y. LI, J. YUAN, AND D. CORTES (2025): xgboost: Extreme Gradient Boosting, r package version 3.1.2.1.
- CHERNOZHUKOV, V., M. DEMIRER, E. DUFLO, AND I. FERNÁNDEZ-VAL (2025): “Fisher–Schultz Lecture: Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments, With an Application to Immunization in India,” Econometrica, 93, 1121–1164.
- COETZEE, G., R. BATES, AND N. MOKOBORI (2003): “Innovative Approaches to Delivering Microfinance Services: Credit Indemnity in South Africa,” Tech. rep., MicroSave.
- COLE, S., M. MELECKY, F. MÖLDERS, AND T. REED (2026): “Long-Run Returns to Private Equity in Emerging Markets,” Management Science, 72, 2041–2063.
- CONSUMER FINANCIAL PROTECTION BUREAU, U.S. DEPARTMENT OF JUSTICE, EQUAL EMPLOYMENT OPPORTUNITY COMMISSION, AND FEDERAL TRADE COMMISSION (2023): “Joint Statement on Enforcement Efforts Against Discrimination and Bias in Automated Systems,” .
- CRÉPON, B., M. EL KOMI, AND A. OSMAN (2024): “Is it who you are or what you get? comparing the impacts of loans and grants for microenterprise development,” American Economic Journal: Applied Economics, 16, 286–313.
- DEMIRGÜÇ-KUNT, A., L. KLAPPER, D. SINGER, AND S. ANSAR (2022): The Global Findex Database 2021: Financial Inclusion, Digital Payments, and Resilience in the Age of COVID-19, Washington, DC: World Bank.
- DUBÉ, J.-P. AND S. MISRA (2023): “Personalized pricing and consumer welfare,” Journal of Political Economy, 131, 131–189.

- FAVA, B. (2025): “Training and Testing with Multiple Splits: A Central Limit Theorem for Split-Sample Estimators,” [ArXiv:2511.04957](#).
- (2026): “Predicting the Distribution of Treatment Effects: A Covariate-Adjustment Approach,” [ArXiv:2407.14635](#).
- FRIEDMAN, J., T. HASTIE, AND R. TIBSHIRANI (2010): “Regularization Paths for Generalized Linear Models via Coordinate Descent,” [Journal of Statistical Software](#), 33, 1–22.
- FRIEDMAN, M. (1970): “The Social Responsibility of Business is to Increase its Profits,” [New York Times Magazine](#), September, 13, 122–126.
- FUSTER, A., P. GOLDSMITH-PINKHAM, T. RAMADORAI, AND A. WALTHER (2022): “Predictably Unequal? The Effects of Machine Learning on Credit Markets,” [The Journal of Finance](#), 77, 5–47.
- GILLAN, S. L., A. KOCH, AND L. T. STARKS (2021): “Firms and social responsibility: A review of ESG and CSR research in corporate finance,” [Journal of Corporate Finance](#), 66, 101889.
- HAND, D., S. SUNDERJI, AND N. M. PARDO (2023): “2023 Market GIINsight: Impact Investing Allocations, Activity & Performance,” Tech. rep., Global Impact Investing Network (GIIN), New York.
- HART, O. D. AND L. ZINGALES (2017): “Companies Should Maximize Shareholder Welfare Not Market Value,” [Journal of Law, Finance, and Accounting](#), 2, 247–275.
- (2022): “The new corporate governance,” Tech. rep., National Bureau of Economic Research.
- HAUSHOFER, J., P. NIEHAUS, C. PARAMO, E. MIGUEL, AND M. WALKER (2025): “Targeting impact versus deprivation,” [American Economic Review](#), 115, 1936–1974.
- KARLAN, D. AND J. ZINMAN (2010): “Expanding credit access: Using randomized supply decisions to estimate the impacts,” [The Review of Financial Studies](#), 23, 433–464.

- (2011): “Microcredit in theory and practice: Using randomized credit scoring for impact evaluation,” Science, 332, 1278–1284.
- KITAGAWA, T. AND A. TETENOV (2018): “Who should be treated? empirical welfare maximization methods for treatment choice,” Econometrica, 86, 591–616.
- KITZMUELLER, M. AND J. SHIMSHACK (2012): “Economic perspectives on corporate social responsibility,” Journal of economic literature, 50, 51–84.
- KNAUS, M. C., M. LECHNER, AND A. STRITTMATTER (2022): “Heterogeneous employment effects of job search programs: A machine learning approach,” Journal of Human Resources, 57, 597–636.
- MULLAINATHAN, S. AND J. SPIESS (2017): “Machine learning: an applied econometric approach,” Journal of Economic Perspectives, 31, 87–106.
- NIE, X. AND S. WAGER (2021): “Quasi-oracle estimation of heterogeneous treatment effects,” Biometrika, 108, 299–319.
- PÁSTOR, L., R. F. STAMBAUGH, AND L. A. TAYLOR (2021): “Sustainable investing in equilibrium,” Journal of financial economics, 142, 550–571.
- TIBSHIRANI, J., S. ATHEY, E. SVERDRUP, AND S. WAGER (2025): grf: Generalized Random Forests, r package version 2.5.0.
- VENABLES, W. N. AND B. D. RIPLEY (2002): Modern Applied Statistics with S, New York: Springer, fourth ed., ISBN 0-387-95457-0.
- WORLD BANK (2024): “MIX Market (Microfinance Information Exchange) Dataset,” <https://databank.worldbank.org/source/mix-market>, accessed: 2024-06-17.
- YOGANARASIMHAN, H., E. BARZEGARY, AND A. PANI (2023): “Design and evaluation of optimal free trials,” Management Science, 69, 3220–3240.

Table 1: Study Details

Study	South Africa	Philippines	Bosnia
Year of Implementation	2004	2006–2007	2008–2009
Treatment/Control Observations	235/313	891/222	551/444
Female/Male Observations	282/266	948/165	408/587
Number of Covariates	33	24	48
Avg Loan Size (USD PPP/year)	\$286	\$674	\$2,133
Range of Loan Size (USD PPP)	\$52–\$1,067	\$312–\$1,562	\$405–\$4,054
Timing of Endline	6–12 months	11–22 months	14 months

South Africa: Karlan and Zinman (2010), Philippines: Karlan and Zinman (2011), Bosnia: Augsburg et al. (2015). Treatment/control and female/male counts refer to our analysis sample (observations with non-missing outcomes and covariates used in the main specifications), and may differ from the full randomized sample reported in the original studies.

Table 2: Targeting on Lender Profits

	Lender Profits (%)			ATE on Borrower Income Change (%)			ATE on Borrower Income Change (\$)		
	Bottom Tercile (1)	Lending to Top vs. Bottom (2)	Lending to Top vs. Current Policy (3)	Bottom Tercile (4)	Lending to Top vs. Bottom (5)	Lending to Top vs. Current Policy (6)	Bottom Tercile (7)	Lending to Top vs. Bottom (8)	Lending to Top vs. Current Policy (9)
Panel A: Predicting Impacts on Lender Profits (%)									
Pooled	-4.0	11.0*** (2.3)	4.8*** (1.2)	-0.5	18.7 (27.9)	20.9 (16.4)	569	604 (1300)	841 (779)
South Africa	-6.3	18.4*** (5.4)	7.4*** (2.8)	14.2	17.4 (52.2)	18.0 (30.5)	677	877 (2420)	1007 (1472)
Philippines	-9.9	18.0*** (1.4)	8.6*** (0.6)	-14.5	42.6 (61.8)	47.1 (36.5)	496	1892 (2143)	1902 (1321)
Bosnia	4.3	-3.4 (3.8)	-1.5 (2.2)	-1.3	-3.8 (21.4)	-2.4 (12.1)	533	-958 (2180)	-387 (1245)
Panel B: Predicting Impacts on Lender Profits (\$)									
Pooled	35.2	-10.2 (29.3)	-4.3 (17.3)	1.2	11.5 (28.7)	15.1 (16.2)	405	597 (1304)	661 (779)
South Africa	-26.5	44.7*** (10.3)	20.2*** (5.4)	5.6	35.4 (52.6)	27.7 (30.2)	851	705 (2444)	1014 (1494)
Philippines	-8.2	7.5*** (2.0)	3.6*** (1.1)	-0.6	4.0 (64.7)	21.0 (36.3)	-33	2040 (2132)	1488 (1289)
Bosnia	140.4	-82.9 (87.3)	-36.9 (51.7)	-1.3	-5.0 (21.4)	-3.5 (12.2)	397	-955 (2189)	-519 (1250)

This table reports the results from targeting borrowers based on the tercile of predicted lender profits, winsorized at the 5th/95th percentile. Panel A predicts lender profits as a percent of loan size, while Panel B predicts profits in dollar levels. Column 1 reports the average for those predicted to be in the bottom 33% of predicted lender profits, while Column 2 reports the difference between those predicted to be in the top tercile and those predicted to be in the bottom tercile with standard errors in parentheses. Column 3 reports the difference in average outcomes if lending only to the top tercile relative to the full sample. Columns 4–6 report the change in borrower household income as a percent change relative to the control group, and Columns 7–9 report the same change in dollar levels; in both cases income is measured from survey data and winsorized at the 5th/95th percentile. In the Philippines we do not have estimates of household income and use borrower business profits instead. When calculating lender profits we use data from MIX to estimate fixed and variable costs of each loan and then normalize remaining average profits across the portfolio to 0 to aid with comparisons. Statistical Significance <0.01***, <0.05**, <0.10*.

Table 3: Who Would Be Targeted

	South Africa		Philippines		Bosnia	
	Targeted Group (1)	Target Group - Everyone Else (2)	Targeted Group (3)	Target Group - Everyone Else (4)	Targeted Group (5)	Target Group - Everyone Else (6)
Panel A: Targeting Top 33% for Lender Profits (%)						
Age	34.6	1.0 (0.9)	40.4	-2.5*** (0.6)	38.5	1.0 (0.8)
College Graduate	0.08	0.07*** (0.02)	0.76	0.25*** (0.03)	0.05	0.01 (0.01)
Female	0.48	-0.05 (0.05)	0.79	-0.10*** (0.02)	0.42	0.02 (0.03)
Income	13174	6405*** (875)	30158	18209*** (1398)	22652	-1269 (1214)
Loan Size	346	110*** (35)	897	352*** (25)	2117	-22 (134)
Share of Profits Coming from Penalties	0.09	0.01 (0.04)	0.22	-0.08 (0.06)	0.03	-0.01 (0.02)
Panel B: Targeting Top 33% for Lender Profits (\$)						
Age	34.8	1.2 (0.9)	41.1	-1.4** (0.6)	38.4	0.8 (0.8)
College Graduate	0.09	0.08*** (0.02)	0.79	0.30*** (0.03)	0.04	0.01 (0.01)
Female	0.46	-0.08* (0.05)	0.79	-0.09*** (0.02)	0.43	0.02 (0.03)
Income	14625	8585*** (846)	27967	14922*** (1441)	22646	-1278 (1224)
Loan Size	378	169*** (34)	850	276*** (26)	2123	-15 (134)
Share of Profits Coming from Penalties	0.09	0.00 (0.04)	0.23	-0.07 (0.06)	0.03	-0.01 (0.02)

This table reports the results from using machine learning methods to identify the top 33% of borrowers based on predicted lender profits, winsorized at the 5th/95th percentile, in % (Panel A) and in \$ (Panel B). In each country we show the average characteristics for individuals who would have been targeted if the lender only provided funding to the top 33% of borrowers on that dimension. We compare their characteristics to everyone else who would have otherwise gotten funding if there was no targeting. We show the difference in the average characteristic between groups by subtracting the average value of the non-targeted group from the value of the targeted group. Standard errors in parentheses, Statistical Significance <0.01***, <0.05**, <0.10*.

Table 4: Restricting Targeting Strategies

	Unrestricted Algorithmic Targeting			Restricting Targeting Strategy to Choose Top Third of Each Income Quintile			
	Lender Profits (%) (1)	Borrower Income Change (%) (2)	Borrower Income Change (\$) (3)	Lender Profits (%) (4)	Borrower Income Change (%) (5)	Borrower Income Change (\$) (6)	p-value for (1) - (4) (7)
Panel A: Predicting Impacts on Lender Profits (%)							
Pooled	4.8*** (1.2)	20.9 (16.4)	841 (779)	2.5* (1.4)	2.3 (16.3)	-18 (752)	0.04
South Africa	7.4*** (2.8)	18.0 (30.5)	1007 (1472)	3.6 (3.4)	6.3 (30.1)	593 (1393)	0.20
Philippines	8.6*** (0.6)	47.1 (36.5)	1902 (1321)	5.5*** (0.7)	3.1 (36.7)	-223 (1261)	0.00
Bosnia	-1.5 (2.2)	-2.4 (12.1)	-387 (1245)	-1.6 (2.2)	-2.4 (12.1)	-425 (1251)	0.98
Panel B: Predicting Impacts on Lender Profits (\$)							
Pooled	-4.3 (17.3)	15.1 (16.2)	661 (779)	-10.6 (17.5)	3.7 (16.0)	278 (742)	0.53
South Africa	20.2*** (5.4)	27.7 (30.2)	1014 (1494)	5.0 (6.7)	17.8 (29.5)	906 (1392)	0.03
Philippines	3.6*** (1.1)	21.0 (36.3)	1488 (1289)	2.6** (1.1)	-3.6 (35.9)	429 (1202)	0.30
Bosnia	-36.9 (51.7)	-3.5 (12.2)	-519 (1250)	-39.5 (52.0)	-3.1 (12.1)	-499 (1255)	0.93

This table compares the impacts of targeting strategies across two scenarios, using lender profits and borrower income (\$) winsorized at the 5th/95th percentile. In the first scenario the targeting algorithm is unrestricted—the 33% of people with the highest predicted profits for the lender get the loan. In the second scenario the lender pursues a “balanced” targeting strategy, by targeting the top 33% of each baseline income quintile, ensuring that the group of targeted borrowers is balanced on baseline income. Each cell reports how the new targeting policy affects the outcome at the top of the column relative to existing policy (which is lending to everyone in our sample). Panel A targets on lender profits (%), while Panel B targets on lender profits (\$). Column 7 reports the *p*-value for the test that the lender-profit gain under unrestricted targeting (Column 1) equals the gain under the balanced strategy (Column 4). Standard errors in parentheses. Statistical Significance <0.01***, <0.05**, <0.10*.

Table 5: Targeting on Borrower Income Change

	ATE on Borrower Income Change			Lender Profits (%)			Lender Profits (\$)		
	Bottom Tercile (1)	Lending to Top vs. Bottom (2)	Lending to Top vs. Current Policy (3)	Bottom Tercile (4)	Lending to Top vs. Bottom (5)	Lending to Top vs. Current Policy (6)	Bottom Tercile (7)	Lending to Top vs. Bottom (8)	Lending to Top vs. Current Policy (9)
Panel A: Predicting Impacts on Borrower Income Change (%)									
Pooled	-8.0	9.8 (28.1)	3.8 (16.1)	1.7	1.2 (2.3)	0.7 (1.3)	29.0	1.1 (29.9)	0.8 (17.2)
South Africa	-1.7	30.6 (53.0)	15.1 (29.4)	2.6	4.7 (5.7)	2.6 (3.2)	-6.7	10.5 (11.5)	5.9 (6.6)
Philippines	-18.5	-4.0 (61.9)	-5.4 (36.2)	0.1	-0.8 (1.4)	-0.3 (0.8)	-4.1	-0.4 (2.1)	-0.1 (1.2)
Bosnia	-3.9	2.9 (20.9)	1.7 (12.0)	2.5	-0.3 (3.8)	-0.2 (2.2)	97.9	-6.9 (88.8)	-3.4 (51.2)
Panel B: Predicting Impacts on Borrower Income Change (\$)									
Pooled	253	126 (1297)	29 (759)	2.0	0.8 (2.3)	0.7 (1.3)	32.2	-7.2 (29.8)	-4.2 (17.3)
South Africa	439	139 (2423)	22 (1407)	5.2	-0.4 (5.7)	0.1 (3.2)	-1.3	-0.0 (11.5)	0.7 (6.5)
Philippines	39	696 (2098)	201 (1246)	-1.9	4.1*** (1.4)	2.6*** (0.8)	-5.0	1.8 (2.1)	1.2 (1.2)
Bosnia	281	-457 (2209)	-137 (1286)	2.8	-1.2 (3.8)	-0.7 (2.2)	103.0	-23.3 (88.6)	-14.7 (51.6)

This table reports the results from using machine learning methods to identify the top 33% of borrowers based on changes in borrower income, as measured through surveys. Panel A focuses on the change in borrower income in percent terms relative to control, while Panel B analyzes changes in levels using income winsorized at the 5th/95th percentile. Column 1 reports the average for those predicted to be in the bottom 33% of predicted impacts of treatment relative to control, while Column 2 reports the difference between those predicted to be in the top tercile and those predicted to be in the bottom tercile with standard errors in parentheses. Column 3 reports the difference in average impacts if lending only to the top tercile relative to the full sample. Columns 4–6 report the change in lender profits as a percent of loan size, and Columns 7–9 report it in dollar levels; both are winsorized at the 5th/95th percentile. In the Philippines we do not have estimates of household income and use borrower business profits instead. When calculating lender profits we use data from MIX to estimate fixed and variable costs of each loan and then normalize remaining average profits across the portfolio to 0 to aid with comparisons. Statistical Significance <0.01***, <0.05**, <0.10*.

Table 6: Characteristics of Borrowers when Targeting on Borrower Income Change

	South Africa		Philippines		Bosnia	
	Targeted Group (1)	Target Group - Everyone Else (2)	Targeted Group (3)	Target Group - Everyone Else (4)	Targeted Group (5)	Target Group - Everyone Else (6)
Panel A: Targeting top 33% for Borrower Income Change (%)						
Age	33.8	-0.3 (0.9)	41.7	-0.5 (0.6)	37.7	-0.1 (0.8)
College Graduate	0.04	0.01 (0.02)	0.61	0.02 (0.03)	0.04	0.00 (0.01)
Female	0.52	0.02 (0.05)	0.85	-0.01 (0.02)	0.43	0.03 (0.03)
Income	10167	1891** (841)	17246	-1160 (1102)	23909	618 (1229)
Loan Size	305	37 (37)	664	-16 (26)	2124	-12 (135)
Share of Profits Coming from Penalties	0.10	0.03 (0.05)	0.27	0.00 (0.06)	0.03	0.00 (0.02)
Panel B: Targeting top 33% for Borrower Income Change (\$)						
Age	33.9	-0.1 (0.9)	41.8	-0.4 (0.6)	37.4	-0.6 (0.8)
College Graduate	0.05	0.01 (0.02)	0.60	0.01 (0.03)	0.04	0.00 (0.01)
Female	0.53	0.03 (0.05)	0.83	-0.03 (0.02)	0.40	-0.02 (0.03)
Income	9718	1216 (788)	22637	6937*** (1334)	25544	3072** (1275)
Loan Size	293	20 (36)	752	116*** (28)	2125	-11 (135)
Share of Profits Coming from Penalties	0.09	0.00 (0.05)	0.25	-0.03 (0.06)	0.03	0.00 (0.02)

This table reports the results from using machine learning methods to identify the top 33% of borrowers based on predicted changes in borrower income. Panel A targets on income change in percent terms, while Panel B targets on income change in dollar levels using income winsorized at the 5th/95th percentile. In each country we show the average characteristics for individuals who would have been targeted if the lender only provided funding to the top 33% of borrowers on that dimension. We compare their characteristics to everyone else who would have otherwise gotten funding if there was no targeting. We show the difference in the average characteristic between groups by subtracting the average value of the non-targeted group from the value of the targeted group. In the Philippines we do not have estimates of household income and use borrower business profits instead. Standard errors in parentheses, Statistical Significance <0.01***, <0.05**, <0.10*.

Appendix A Aggregating Predictions from Multiple Algorithms

We use the R package `ensembleHTE` to follow the approach of Fava (2025) for aggregating predictions from multiple ML algorithms. We take the following steps for aggregating lender profit predictions in Subsection 3.1:

1. Split the sample into five folds (independent of the previous 3 folds), denoted by $(\mathbf{s}_\ell)_{\ell=1}^5$. Denote by $\tilde{\mathbf{s}}_\ell$ the complement of $\mathbf{s}_\ell$, all observations not contained in $\mathbf{s}_\ell$. Let $\hat{L}_{i,j}$ be the predicted lender profits coming from ML model j ($j = 1, \dots, 4$).
2. For $\ell = 1, \dots, 5$:

2.1 Run the regression

$$L_i = \gamma_0 + \sum_{j=1}^4 \gamma_j \hat{L}_{i,j} + \varepsilon_i$$

using only the borrowers in $\tilde{\mathbf{s}}_\ell$, that is, all folds except $\mathbf{s}_\ell$. Denote the estimates by $(\hat{\gamma}_j)_{j=0}^4$.

2.2 Calculate the final predicted lender profits for all observations in fold $\mathbf{s}_\ell$:

$$\hat{L}_i = \hat{\gamma}_0 + \sum_{j=1}^4 \hat{\gamma}_j \hat{L}_{i,j}.$$

For aggregating predicted individual treatment effects, let $k(i) \in \{1, 2, 3\}$ denote the original fold that observation i belongs to (first step in Appendix B), and denote those folds by $(\mathbf{s}'_1, \mathbf{s}'_2, \mathbf{s}'_3)$. We take steps similar to Subsection 3.3:

1. Split the sample into five folds (independent of the previous 3 folds), denoted by $(\mathbf{s}_\ell)_{\ell=1}^5$, and let $\hat{\tau}_{i,j}$ be the predicted individual treatment effect coming from ML model j ($j = 1, \dots, 4$).
2. For $\ell = 1, \dots, 5$:

2.1 Using only the observations in $\tilde{\mathbf{s}}_\ell$, run the regression

$$B_i = \gamma_0 + \sum_{j=1}^4 \gamma_j (\hat{\tau}_{i,j} - \bar{\tau}_{j,k(i)}) (A_i - p(X_i)) + \gamma_5 (A_i - p(X_i)) + \varepsilon_i$$

with weights $W_i = (p(X_i)[1 - p(X_i)])^{-1}$, where $\bar{\tau}_{j,k}$ is the average of $\hat{\tau}_{i,j}$ over all observations in fold $\mathbf{s}'_k$. B_i denotes household income of observation i . Denote the estimates by $(\hat{\gamma}_j)_{j=0}^4$.

2.2 Calculate the final predicted individual treatment effect for all observations in fold $\mathbf{s}_\ell$:

$$\hat{\tau}_i = \sum_{j=1}^4 \hat{\gamma}_j \hat{\tau}_{i,j}.$$

The regression for calculating the weights $(\hat{\gamma}_j)_{j=0}^4$ is similar to the Best Linear Predictor regression of Chernozhukov et al. (2025).

Appendix B Targeting Borrower Impact

We adopt an R-learner strategy (Nie and Wager, 2021) to estimate individual treatment effects on borrower income, implemented through the `ensembleHTE` package in R. We take the following steps:

1. Randomly split the sample into three *folds* of approximately equal size.
2. For each of the four ML algorithms j ($j = 1, \dots, 4$: Extreme Gradient Boosting, Regression Forest, Single-layer Neural Network, Elastic Net), using two folds at a time, estimate the conditional expectation of borrower income $\mathbb{E}[B \mid X]$ using the full sample (both treated and control observations). For each observation i , compute the predicted outcome $\hat{B}_{i,j}$ using the model trained on the folds that do not contain that observation.
3. Residualize the outcome and the treatment assignment:

$$\tilde{B}_{i,j} = B_i - \hat{B}_{i,j}, \quad \tilde{A}_i = A_i - p(X_i),$$

where $p(X_i)$ is the known propensity score.

4. Estimate the CATE function using a causal forest (Athey et al., 2019) from the R package `grf`, passing the first-stage estimates $\hat{B}_{i,j}$ and $p(X_i)$ directly.¹³ The causal forest then predicts, for each observation i , the CATE evaluated at its covariates X_i ; we denote this predicted individual treatment effect by $\hat{\tau}_{i,j}$.
5. Combine the predictions from the four algorithms into $\hat{\tau}_i$, using the *ensemble* strategy of Fava (2025), described in Appendix A.
6. Split the observations into terciles of $\hat{\tau}_i$ within each fold. Define dummies $G_{i,t}$ taking the value 1 if observation i is in the t -th tercile of their fold.

With the new definitions of $(G_{i,1}, G_{i,2}, G_{i,3})$, we run the same specifications as in (1) and (2).

¹³Specifically, we call `grf::causal_forest` with arguments `Y.hat =  $\hat{B}_{i,j}$` and `W.hat =  $p(X_i)$` , so that the causal forest uses our ML-estimated residuals rather than its own internal estimates.

Appendix C Calibration of Machine Learning Predictors

This appendix reports calibration diagnostics for the machine learning models used in the main analysis. For the models predicting individual treatment effects, we apply two standard procedures from the generic machine learning framework of Chernozhukov et al. (2025): the Best Linear Predictor (BLP) and Group Average Treatment Effects (GATES). For the models predicting lender profits, we use a similar BLP approach, and calculate Group Average Values (GAVS) instead of GATES. The statistical methodology, including ensemble construction and repeated cross-fitting, follows Fava (2025).

Appendix C.1 Best Linear Predictor (BLP)

The BLP test assesses whether the ML-predicted treatment effects contain systematic signal about actual heterogeneity. The BLP specification for borrower outcomes is the weighted regression:

$$Y_i = \alpha + \beta_1 (A_i - p_i) + \beta_2 (A_i - p_i) (S(X_i) - \bar{S}) + \varepsilon_i, \quad (3)$$

with weights $w_i = 1/(p_i(1-p_i))$, where Y_i is the outcome, A_i is the treatment indicator, p_i is the propensity score, $S(X_i)$ is the ML-predicted conditional average treatment effect (CATE), and $\bar{S}$ is its sample mean. The coefficient β_1 is the average treatment effect (ATE). The coefficient β_2 is a heterogeneity loading that tests whether the ML predictions capture genuine variation in treatment effects: under correct specification of the CATE, $\beta_2 = 1$. A statistically significant $\hat{\beta}_2$ indicates that the ML model is detecting real heterogeneity in treatment effects.

For lender outcomes, which are predicted for the treated subsample rather than estimated as treatment effects, the BLP test reduces to an OLS regression of the observed outcome on the ML predictions:

$$Y_i = \beta_1 + \beta_2 S(X_i) + \varepsilon_i, \quad (4)$$

where $\beta_2 = 1$ under perfect calibration. Similar to before, a statistically significant $\hat{\beta}_2$ indicates that the model has real power for predicting the outcome.

Appendix Table 1 reports the BLP results separately for each dataset. All regressions use HC1 heteroskedasticity-robust standard errors and estimates are averaged

over repeated sample splits. No ATE estimate is statistically significant, and the only statistically significant heterogeneity loading coefficients are Lender Profits (\$) for the South Africa dataset, and Lender Profits (%) for the Philippines and South Africa datasets.

Appendix Table 1: BLP calibration test.

	Log HH Income	HH Income	Business Profits	Lender Profits (\$)	Lender Profits (%)
$\hat{\beta}_1$					
BOS	-0.03 (0.09)	230 (1170)	1061 (777)	0.04 (20.83)	0.00 (0.71)
PHI	-0.17 (0.26)	1650 (1677)	—	-0.01 (1.18)	0.58 (0.36)
SA	0.14 (0.22)	279 (3353)	—	2.46 (4.57)	0.39 (1.48)
$\hat{\beta}_2$					
BOS	-0.03 (0.38)	-0.39 (0.44)	-0.11 (0.35)	-0.12 (0.16)	-0.13 (0.15)
PHI	-0.06 (0.36)	-0.33 (0.39)	—	-0.13 (0.24)	0.55*** (0.04)
SA	-0.09 (0.32)	0.10 (0.20)	—	0.18** (0.08)	0.13** (0.07)

For borrower outcomes (Log HH Income, HH Income, Business Profits), $\hat{\beta}_1$ is the ATE and $\hat{\beta}_2$ the heterogeneity loading.

For lender outcomes (Lender Profits (\$) and (%)), $\hat{\beta}_1$ is the intercept and $\hat{\beta}_2$ the predictive loading on the ML score.

Under correct specification, $\hat{\beta}_2 \approx 1$ in both cases.

In the Philippines we do not observe household income, so the household income columns report business profits. We observe business profits directly in Bosnia; it is not available in South Africa.

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

Appendix C.2 Group Average Treatment Effects (GATES)

The GATES analysis, explained in Section 3, can be used as a calibration diagnostic for the machine learning models predicting individual treatment effects. It consists

of sorting observations by their predicted individual treatment effects $S(X_i)$, dividing them into terciles, from the lowest to the highest predicted treatment effect, and calculating the average treatment effect within each tercile. If the model is well-calibrated, the GATES should increase monotonically from the 1st to the 3rd tercile.


Appendix Figure 1 displays the GATES estimates for borrower outcomes by project. Points represent the estimated GATE for each tercile, with 95% confidence intervals. The dashed horizontal line marks the overall ATE. All panels show that the ML models targeting treatment effect heterogeneity are not well-calibrated, since the GATES estimates are noisy and have a flat pattern instead of increasing monotonically.

Appendix C.3 Group Average Values (GAVS)

For lender outcomes, where the relevant predictions are fitted values rather than treatment effects, we report Group Average Values (GAVS), as described in Section 3.1. GAVS is the prediction-problem analogue of GATES: observations are sorted by their ML-predicted outcome $S(X_i)$, divided into terciles, and the average of the actual outcome is calculated within tercile. Well-calibrated predictions should yield monotonically increasing group averages.


Appendix Figure 2 displays the GAVS estimates for lender outcomes by project. The ML models predicting lender profits (\$) in the South Africa dataset, and lender profits (%) in the Philippines and South Africa datasets, are somewhat well-calibrated, showing monotonically increasing patterns.

Appendix Figure 1: GATES Estimates for Borrower Outcomes



Each panel shows group average treatment effects by tercile of predicted treatment effects (group 1 = lowest, group 3 = highest predicted effect). Error bars are 95% confidence intervals. The dashed horizontal line marks the overall average treatment effect. Rows correspond to datasets; columns to outcomes. In the Philippines we do not observe household income, so the household income columns report business profits. We observe business profits directly in Bosnia; it is not available in South Africa.

Appendix Figure 2: GAVS Estimates for Lender Outcomes



Each panel shows group average values by tercile of ML-predicted outcomes (group 1 = lowest, group 3 = highest predicted value). Error bars are 95% confidence intervals. The dashed horizontal line marks the overall sample mean. Rows correspond to datasets; columns to outcomes.

Appendix D Robustness: Unwinsorized Lender Profits

The following tables replicate Tables 2–4 using unwinsorized lender profit measures instead of profits winsorized at the 5th and 95th percentiles.

Appendix Table 2: Targeting on Lender Profits (Unwinsorized)

	Lender Profits (%)			ATE on Borrower Income Change (%)			ATE on Borrower Income Change (\$)		
	Bottom Tercile (1)	Lending to Top vs. Bottom (2)	Lending to Top vs. Current Policy (3)	Bottom Tercile (4)	Lending to Top vs. Bottom (5)	Lending to Top vs. Current Policy (6)	Bottom Tercile (7)	Lending to Top vs. Bottom (8)	Lending to Top vs. Current Policy (9)
Panel A: Predicting Impacts on Lender Profits (%)									
Pooled	-6.5	12.2*** (4.5)	5.8** (2.5)	-5.7	23.0 (27.7)	20.0 (16.2)	200	1699 (3421)	1267 (2144)
South Africa	-12.2	23.1* (12.4)	11.3 (6.9)	4.0	32.8 (52.6)	23.3 (30.0)	362	502 (8602)	591 (5448)
Philippines	-9.9	18.6*** (1.6)	8.7*** (0.8)	-19.4	40.0 (60.8)	39.3 (36.2)	-609	5816 (4783)	3810 (2999)
Bosnia	2.6	-5.2 (4.8)	-2.6 (2.9)	-1.7	-3.7 (21.4)	-2.7 (12.1)	848	-1221 (2912)	-600 (1638)
Panel B: Predicting Impacts on Lender Profits (\$)									
Pooled	17.4	-32.5 (48.0)	-13.8 (28.5)	-7.9	17.6 (28.2)	12.0 (16.3)	1053	-286 (3435)	82 (2134)
South Africa	-40.8	79.8** (33.5)	42.7** (18.6)	4.3	35.4 (53.5)	26.5 (30.3)	1674	-2159 (8808)	-776 (5528)
Philippines	1.3	-1.5 (5.3)	-0.3 (2.6)	-26.2	21.1 (61.9)	12.3 (36.5)	949	2017 (4488)	1427 (2765)
Bosnia	91.8	-175.6 (140.0)	-83.8 (83.4)	-1.8	-3.8 (21.2)	-2.8 (12.1)	536	-716 (2912)	-406 (1672)

This table replicates Table 2 using unwinsorized lender profit measures. Panel A focuses on the percent change in profits relative to loan size, while Panel B replicates the analysis in levels. Column 1 reports the average for those predicted to be in the bottom 33% of predicted lender profits, while Column 2 reports the difference between those predicted to be in the top tercile and those predicted to be in the bottom tercile with standard errors in parentheses. Column 3 reports the difference in average outcomes if lending only to the top tercile relative to current lending policies (which is lending to everyone in our sample). Columns 4–6 report the change in borrower household income as a percent change relative to the control group, and Columns 7–9 report the same change in dollar levels; in both cases income is measured from survey data. In the Philippines we do not have estimates of household income and use borrower business profits instead. When calculating lender profits we use data from MIX to estimate fixed and variable costs of each loan and then normalize remaining average profits across the portfolio to 0 to aid with comparisons. Statistical Significance <0.01***, <0.05**, <0.10*.

Appendix Table 3: Who Would Be Targeted (Unwinsorized)

	South Africa		Philippines		Bosnia	
	Targeted Group (1)	Target Group - Everyone Else (2)	Targeted Group (3)	Target Group - Everyone Else (4)	Targeted Group (5)	Target Group - Everyone Else (6)
Panel A: Targeting Top 33% for Lender Profits (%)						
Age	34.7	1.1 (0.9)	40.4	-2.6*** (0.6)	38.4	0.9 (0.8)
College Graduate	0.08	0.06*** (0.02)	0.75	0.24*** (0.03)	0.04	0.01 (0.01)
Female	0.50	-0.02 (0.05)	0.79	-0.10*** (0.02)	0.42	0.01 (0.03)
Income	12769	5797*** (807)	30309	18437*** (1394)	22629	-1303 (1219)
Loan Size	354	123*** (36)	896	350*** (25)	2115	-26 (134)
Share of Profits Coming from Penalties	0.08	0.00 (0.04)	0.22	-0.08 (0.06)	0.03	-0.01 (0.02)
Panel B: Targeting Top 33% for Lender Profits (\$)						
Age	34.7	1.1 (0.9)	41.0	-1.5*** (0.6)	37.9	0.1 (0.8)
College Graduate	0.09	0.08*** (0.02)	0.65	0.09*** (0.03)	0.04	0.00 (0.01)
Female	0.49	-0.03 (0.05)	0.83	-0.04* (0.02)	0.41	0.00 (0.03)
Income	14132	7844*** (799)	23616	8396*** (1302)	23716	329 (1224)
Loan Size	379	171*** (34)	762	136*** (27)	2142	14 (134)
Share of Profits Coming from Penalties	0.09	0.00 (0.04)	0.23	-0.07 (0.05)	0.03	0.00 (0.02)

This table replicates Table 3 using unwinsorized lender profit measures. It reports the results from using machine learning methods to identify the top 33% of borrowers based on predicted lender profits in % (Panel A) and in \$ (Panel B). In each country we show the average characteristics for individuals who would have been targeted if the lender only provided funding to the top 33% of borrowers on that dimension. We compare their characteristics to everyone else who would have otherwise gotten funding if there was no targeting. We show the difference in the average characteristic between groups by subtracting the average value of the non-targeted group from the value of the targeted group. Standard errors in parentheses, Statistical Significance <0.01***, <0.05**, <0.10*.

Appendix Table 4: Restricting Targeting Strategies (Unwinsorized)

	Unrestricted Algorithmic Targeting			Restricting Targeting Strategy to Choose Top Third of Each Income Quintile			
	Lender Profits (%) (1)	Borrower Income Change (%) (2)	Borrower Income Change (\$) (3)	Lender Profits (%) (4)	Borrower Income Change (%) (5)	Borrower Income Change (\$) (6)	p-value for (1) - (4) (7)
Panel A: Predicting Impacts on Lender Profits (%)							
Pooled	5.8** (2.5)	20.0 (16.2)	1267 (2144)	1.1 (2.6)	4.5 (16.2)	1344 (2016)	0.08
South Africa	11.3 (6.9)	23.3 (30.0)	591 (5448)	0.8 (7.2)	12.4 (29.5)	3629 (5016)	0.19
Philippines	8.7*** (0.8)	39.3 (36.2)	3810 (2999)	5.3*** (0.8)	4.1 (36.5)	1040 (2947)	0.00
Bosnia	-2.6 (2.9)	-2.7 (12.1)	-600 (1638)	-2.7 (2.9)	-3.1 (12.1)	-638 (1649)	0.94
Panel B: Predicting Impacts on Lender Profits (\$)							
Pooled	-13.8 (28.5)	12.0 (16.3)	82 (2134)	-32.0 (28.6)	4.0 (16.2)	1277 (1881)	0.29
South Africa	42.7** (18.6)	26.5 (30.3)	-776 (5528)	-4.6 (19.6)	8.0 (29.6)	2929 (4766)	0.06
Philippines	-0.3 (2.6)	12.3 (36.5)	1427 (2765)	-1.8 (2.6)	6.2 (36.6)	1299 (2521)	0.44
Bosnia	-83.8 (83.4)	-2.8 (12.1)	-406 (1672)	-89.7 (83.6)	-2.2 (12.1)	-397 (1669)	0.90

This table replicates Table 4 using unwinsorized lender profit measures. It compares the impacts of targeting strategies across two scenarios. In the first scenario the targeting algorithm is unrestricted—the 33% of people with the highest predicted profits for the lender get the loan. In the second scenario the lender pursues a “balanced” targeting strategy, by targeting the top 33% of each baseline income quintile, ensuring that the group of targeted borrowers is balanced on baseline income. Each cell reports how the new targeting policy affects the outcome at the top of the column relative to existing policy (which is lending to everyone in our sample). Panel A targets on lender profits (%), while Panel B targets on lender profits (\$). Column 7 reports the *p*-value for the test that the lender-profit gain under unrestricted targeting (Column 1) equals the gain under the balanced strategy (Column 4). Standard errors in parentheses. Statistical Significance <0.01***, <0.05**, <0.10*.

Appendix E Robustness: Intent-to-Treat Predictions

In our main specification (Subsection 3.1), we train ML models to predict lender profits using only the subsample of individuals who were offered and took the experimental loan (S_B). The resulting predictor estimates $\mu_L(x) = \mathbb{E}[L \mid X = x, D = 1]$, the expected profitability of a loan conditional on the borrower actually taking it up.

An alternative is to train the prediction model on all treated individuals, including those who were offered a loan but did not take it up. For non-takers, we set lender profits to zero. Because our profit measures have been normalized so that the mean across borrowers is zero (see Section 3), this is equivalent to imputing the average borrower profit for individuals who decline the offer. For profits in levels, this amounts to assuming that capital not lent to a non-taker can be deployed elsewhere in the portfolio, earning the average return. For profits as a percentage of loan size, this corresponds to imputing the mean percentage return across borrowers. The resulting ITT predictor estimates $\mathbb{E}[L \mid X = x, A = 1]$, the expected profitability of *offering* a loan to an individual with characteristics x . This quantity incorporates both the probability that the individual takes up the loan and the conditional profitability if they do, and is the more relevant object when the lender bears costs associated with making an offer regardless of whether the borrower accepts.

The remainder of the methodology is identical to the main specification: we form terciles of predicted profits, define the same targeting policies, and evaluate them using the same regressions. The following tables replicate Tables 2–4 using these ITT predictions. Results are qualitatively similar to the main specification.

Appendix Table 5: Targeting on Lender Profits (ITT)

	Lender Profits (%)			ATE on Borrower Income Change (%)			ATE on Borrower Income Change (\$)		
	Bottom Tercile (1)	Lending to Top vs. Bottom (2)	Lending to Top vs. Current Policy (3)	Bottom Tercile (4)	Lending to Top vs. Bottom (5)	Lending to Top vs. Current Policy (6)	Bottom Tercile (7)	Lending to Top vs. Bottom (8)	Lending to Top vs. Current Policy (9)
Panel A: Predicting Impacts on Lender Profits (%)									
Pooled	-3.7	6.6*** (2.1)	3.5*** (1.2)	-3.4	23.0 (28.0)	22.5 (16.3)	-251	2510 (3381)	1641 (2091)
South Africa	-7.1	10.4* (5.4)	5.3* (3.0)	5.3	32.0 (52.7)	23.9 (30.5)	-627	2251 (8481)	1366 (5273)
Philippines	-6.1	12.9*** (1.0)	7.0*** (0.6)	-13.5	40.9 (61.9)	46.6 (36.3)	-780	6087 (4751)	3932 (2971)
Bosnia	2.0	-3.6 (3.0)	-1.7 (1.8)	-2.1	-3.8 (21.4)	-3.2 (12.2)	656	-807 (2899)	-376 (1652)
Panel B: Predicting Impacts on Lender Profits (\$)									
Pooled	15.5	-18.9 (26.7)	-6.5 (15.8)	-3.6	12.8 (28.6)	11.5 (16.1)	86	1131 (3490)	555 (2136)
South Africa	-19.6	33.9** (14.8)	21.4** (8.7)	-0.1	39.7 (53.6)	26.3 (30.5)	344	-445 (8909)	-387 (5461)
Philippines	-3.1	4.8** (2.1)	2.7** (1.3)	-7.7	0.4 (63.6)	10.3 (35.4)	-456	4310 (4669)	2381 (2903)
Bosnia	69.1	-95.3 (78.8)	-43.6 (46.6)	-3.0	-1.8 (21.1)	-2.0 (12.1)	370	-473 (2907)	-329 (1682)

This table replicates Table 2 using the ITT (intent-to-treat) version of lender profit predictions, where the model is trained on all borrowers including controls. Panel A focuses on the percent change in profits relative to loan size, while Panel B replicates the analysis in levels. Column 1 reports the average for those predicted to be in the bottom 33% of predicted lender profits, while Column 2 reports the difference between those predicted to be in the top tercile and those predicted to be in the bottom tercile with standard errors in parentheses. Column 3 reports the difference in average outcomes if lending only to the top tercile relative to current lending policies (which is lending to everyone in our sample). Columns 4–9 hold the groups fixed while considering the change in borrower household income using survey data. In the Philippines we do not have estimates of household income and use borrower business profits instead. When calculating lender profits we use data from MIX to estimate fixed and variable costs of each loan and then normalize remaining average profits across the portfolio to 0 to aid with comparisons. Statistical Significance <0.01***, <0.05**, <0.10*.

Appendix Table 6: Who Would Be Targeted (ITT)

	South Africa		Philippines		Bosnia	
	Targeted Group (1)	Target Group - Everyone Else (2)	Targeted Group (3)	Target Group - Everyone Else (4)	Targeted Group (5)	Target Group - Everyone Else (6)
Panel A: Targeting Top 33% for Lender Profits (%)						
Age	35.5	2.3*** (0.9)	40.3	-2.6*** (0.6)	38.2	0.6 (0.8)
College Graduate	0.08	0.07*** (0.02)	0.76	0.25*** (0.03)	0.04	0.01 (0.01)
Female	0.49	-0.03 (0.05)	0.78	-0.10*** (0.02)	0.41	0.01 (0.03)
Income	13350	6669*** (833)	30103	18151*** (1398)	23143	-532 (1219)
Loan Size	361	138*** (35)	898	352*** (25)	2119	-20 (134)
Share of Profits Coming from Penalties	0.09	0.01 (0.04)	0.22	-0.08 (0.06)	0.03	-0.01 (0.02)
Panel B: Targeting Top 33% for Lender Profits (\$)						
Age	35.1	1.7* (0.9)	41.1	-1.4*** (0.6)	37.8	-0.0 (0.8)
College Graduate	0.09	0.08*** (0.02)	0.76	0.25*** (0.03)	0.04	0.00 (0.01)
Female	0.48	-0.05 (0.05)	0.80	-0.08*** (0.02)	0.41	0.00 (0.03)
Income	14467	8347*** (802)	27050	13565*** (1429)	23884	580 (1241)
Loan Size	386	183*** (34)	825	233*** (27)	2130	-4 (134)
Share of Profits Coming from Penalties	0.09	0.01 (0.04)	0.23	-0.06 (0.06)	0.03	0.00 (0.02)

This table replicates Table 3 using ITT (intent-to-treat) predictions of lender profits. It reports the results from using machine learning methods to identify the top 33% of borrowers based on predicted lender profits in % (Panel A) and in \$ (Panel B). In each country we show the average characteristics for individuals who would have been targeted if the lender only provided funding to the top 33% of borrowers on that dimension. We compare their characteristics to everyone else who would have otherwise gotten funding if there was no targeting. We show the difference in the average characteristic between groups by subtracting the average value of the non-targeted group from the value of the targeted group. Standard errors in parentheses, Statistical Significance <0.01***, <0.05**, <0.10*.

Appendix Table 7: Restricting Targeting Strategies (ITT)

	Unrestricted Algorithmic Targeting			Restricting Targeting Strategy to Choose Top Third of Each Income Quintile			
	Lender Profits (%) (1)	Borrower Income Change (%) (2)	Borrower Income Change (\$) (3)	Lender Profits (%) (4)	Borrower Income Change (%) (5)	Borrower Income Change (\$) (6)	p-value for (1) - (4) (7)
Panel A: Predicting Impacts on Lender Profits (%)							
Pooled	3.5*** (1.2)	22.5 (16.3)	1641 (2091)	1.1 (1.3)	3.5 (16.2)	1423 (1929)	0.02
South Africa	5.3* (3.0)	23.9 (30.5)	1366 (5273)	0.7 (3.3)	11.2 (29.8)	3446 (4700)	0.13
Philippines	7.0*** (0.6)	46.6 (36.3)	3932 (2971)	4.4*** (0.6)	2.7 (36.3)	1258 (2943)	0.00
Bosnia	-1.7 (1.8)	-3.2 (12.2)	-376 (1652)	-1.8 (1.8)	-3.3 (12.3)	-434 (1656)	0.89
Panel B: Predicting Impacts on Lender Profits (\$)							
Pooled	-6.5 (15.8)	11.5 (16.1)	555 (2136)	-16.2 (15.8)	-0.4 (15.9)	1154 (1858)	0.28
South Africa	21.4** (8.7)	26.3 (30.5)	-387 (5461)	-2.5 (8.6)	6.9 (29.8)	2492 (4563)	0.02
Philippines	2.7** (1.3)	10.3 (35.4)	2381 (2903)	1.6 (1.2)	-6.4 (35.0)	1285 (2731)	0.27
Bosnia	-43.6 (46.6)	-2.0 (12.1)	-329 (1682)	-47.6 (46.7)	-1.6 (12.1)	-315 (1675)	0.87

This table replicates Table 4 using ITT (intent-to-treat) predictions of lender profits. It compares the impacts of targeting strategies across two scenarios. In the first scenario the targeting algorithm is unrestricted—the 33% of people with the highest predicted profits for the lender get the loan. In the second scenario the lender pursues a “balanced” targeting strategy, by targeting the top 33% of each baseline income quintile, ensuring that the group of targeted borrowers is balanced on baseline income. Each cell reports how the new targeting policy affects the outcome at the top of the column relative to existing policy (which is lending to everyone in our sample). Panel A targets on lender profits (%), while Panel B targets on lender profits (\$). Column 7 reports the *p*-value for the test that the lender-profit gain under unrestricted targeting (Column 1) equals the gain under the balanced strategy (Column 4). Standard errors in parentheses. Statistical Significance <0.01***, <0.05**, <0.10*.