

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

Susan Athey, Bruno Fava, Dean Karlan,
Adam Osman, Jonathan Zinman*

December 2025

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 exacerbate this tradeoff: algorithms can maximize profits via prediction, whereas targeting impact requires the more difficult task of estimating causal effects. We estimate these trade-offs using data from randomized microcredit approvals in South Africa, the Philippines, and Bosnia. Two of the three lenders could have increased profits by 10–14% through machine learning-based targeting, but doing so would have reduced access for groups microcredit seeks to help, including female and lower-income borrowers, suggesting a tradeoff between lender profit and distributional objectives. However, we find no evidence that maximizing profits would lead to lower treatment effects on borrower income; if anything, lender profits are increasing in those effects, suggesting complementarity between lender profit and economic efficiency objectives. Simulating effects of policy constraints on lender targeting, we find that holding borrower average income fixed at baseline levels would reduce potential lender profit by two-thirds. These findings highlight the importance of quantifying tradeoffs and complementarities when deciding what to maximize.

*Contact: Susan Athey, athey@stanford.edu, Stanford University and NBER; Bruno Fava, bruno-fava@u.northwestern.edu, Northwestern University; Dean Karlan, karlan@northwestern.edu, Northwestern University, IPA, J-PAL, and NBER; Adam Osman, aosman@illinois.edu, University of Illinois at Urbana Champaign, J-PAL; Jonathan Zinman, jzinman@dartmouth.edu, Dartmouth College, IPA, J-PAL, and NBER. All errors are our own.

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 maximize shareholder welfare, which may encompass objectives beyond pure profit (Hart and Zingales, 2017, 2022). This debate has become even more salient with the rise of Environmental, Social and Governance (ESG) investing and policy initiatives (Gillan et al., 2021; Pástor et al., 2021). This raises a critical empirical question: what are the tradeoffs, if any, between these two objectives? While some suggest that “win-win” scenarios are possible, others worry that pursuing broader welfare goals inevitably compromises firm efficiency and profitability.

These tensions have intensified as machine learning and algorithmic targeting have advanced —potentially providing firms with greater ability to target their activities toward specific objectives. Importantly, targeting capabilities have advanced unevenly, as the data necessary to improve profits is easier for firms to obtain than what they need for the more difficult task of assessing social impacts. Does the ability to improve targeting of activities to more profitable opportunities affect customer welfare? Can firms who want to target on welfare maintain or exceed current levels of profitability? In other words- do these tradeoffs exist and if so how steep are they?

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 is primarily a prediction problem: given observable borrower characteristics and loan terms, lenders must predict repayment and profits. In contrast, maximizing borrower welfare is a treatment-effect problem: the lender must estimate how access to credit causally affects borrower outcomes and how those effects vary across borrowers. Identifying these heterogeneous treatment effects requires information on outcomes relative to an unobserved counterfactual, typically generated by randomized or otherwise exogenous variation in credit access. While firms often collect the administrative data necessary to forecast profits, they rarely have data on borrower or community outcomes combined with the exogenous variation required to estimate these causal effects. Financial inclusion, the largest sector within the impact investing space (Hand et al., 2023), provides a perfect context to overcome

these challenges and make headway on these questions.¹

Lenders’ allocation of scarce capital critically shapes individual and aggregate economic outcomes (Bryan et al., 2024; Hsieh and Klenow, 2009; Moll, 2014). We combine data from loan applications, repayment performance, and endline surveys with experimental variation from three randomized controlled trials of microcredit in South Africa (Karlan and Zinman, 2010), the Philippines (Karlan and Zinman, 2011), and Bosnia (Augsburg et al., 2015). These allow us to first train machine learning algorithms to predict the profitability of each loan applicant and then use the random assignment of loans to causally estimate the (counterfactual) welfare impacts of targeting credit based on these predictions. In two of the three cases, the lender could have increased profits by changing its allocation policies in line with the predictions made by the algorithms. Profits could have increased by 14% in South Africa and 10% in the Philippines. These techniques are not silver-bullets - in Bosnia current data and algorithms are unable to reliably predict which borrowers 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 female, less-educated, and lower-income borrowers. These are typically groups of society that already struggle to gain access to credit, potentially exacerbating existing gaps.

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 those offered a loan with similar people who were not. We find suggestive evidence that the borrowers that are most profitable to the bank are also those that benefit from the loans. But while the estimated effects of targeting are positive, the standard errors are large, and these results are not statistically significant.

Policymakers may be worried that profit-maximizing loan targeting policies could lead to imbalanced lending choices that would lead to inequality over some policy relevant characteristics such as gender and baseline income levels. To that end we consider what would happen if regulations insisted that targeting policies leave the average income of borrowers unchanged. We analyze a balanced lending policy where the lender splits borrowers up into 5 quintiles based on their baseline income and then

¹The report splits Financial inclusion into “Financial Services (excluding microfinance)” and “Microfinance”. When combined it is the largest sector.

lends to the most profitable 33% of borrowers within each quintile. With this policy, lenders can still 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 increase in profits 65% smaller than the unrestricted case. This may provide a way for lenders to still increase profits while minimizing the potential bias that a fully data-driven lending policy may lead to, but it comes at a nontrivial cost to lenders.

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.

A growing literature uses machine learning to design and evaluate personalized policies, often considering trade-offs between firm performance (such as profits or engagement) and social objectives. Closest to our setting is Dubé and Misra (2023), who study 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), Haushofer et al. (2025), Knaus et al. (2022), Yoganarasimhan et al. (2023), Athey et al. (2023), Athey et al. (2024), Athey et al. (2025)). Most of this literature does not analyze how optimizing supplier profits affects user welfare,² 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 service. We contribute to this literature by explicitly estimating potential tradeoffs between supplier profits and experimentally identified effects on borrowers’ incomes, and then using these estimates to examine how policy targeting affects realized

²For exceptions, see Athey et al. (2024), who use experimental data to estimate how targeting affects the tradeoffs faced by fundraising charities between user experience, dollars raised, and the size of the donor base, and Athey et al. (2025), who use a calibrated model to analyze how the (counterfactual) introduction of algorithmic recommendation systems in the Kiva lending platform changes the distribution of outcomes across borrower groups.

downstream outcomes.

We also 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 (Hart and Zingales, 2017; Kitzmüller and Shimshack, 2012; Servaes and Tamayo, 2013). As firms adopt technology that enables them to use targeting to screen customers or present customized offers, the trade-offs between profit maximization and achieving more community-focused outcomes could become more stark. In this context, we find that changing lending practices would change who is served by lenders. On the other hand, we find no evidence that maximizing profits would lead to negative impacts on average borrower outcomes, but this result is relatively noisy. It is possible that as these methods become more powerful, and more prevalent, the impacts on borrowers could become more pronounced and the costs of sustaining a double bottom line may become too large to bear. We show that regulation could force lenders to change the way they target loans, but this could come at a 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, 2024). Much of the earlier literature focused 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. 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, policy-makers will often use proxies like measures of poverty to allocate support, but this 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 aggregate 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 needed baseline information about the borrower along with repayment data to assess how different borrowers performed with the loan. Second, to assess the 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).

Appendix 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 credit worthiness and a lower probability for the group with lower credit scores.³

³72% of the treatment group took up the loan.

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.⁴

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 in the relevant time period as well as any additional fees or penalties paid by borrowers. We subtract out the present value of the capital that was lent by the bank as well as an estimate of the other origination costs of the loan (this attempts to include things like the cost of loan officer time, overhead, etc.). To approximate this additional cost we use data from the Microfinance Information Exchange Market dataset (World Bank, 2025) to estimate a fixed and variable cost for each loan.⁵ Together this provides us a value for the profitability of each loan made by the lender.

For the experiments in South Africa and Bosnia, we measure borrower income using reported income. For the experiment in the Philippines, the lender’s pre-treatment survey does not collect individual income, and so we use the borrowers profits reported in our follow up survey 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.

⁴53% of the treatment group took up the loan.

⁵For variable cost, we divide reported financial expenses by the size of the gross loan portfolio. For fixed cost, we divide operating expenses by total number of loans. In South Africa we use a measure that is more relevant for the lender we have data for (Coetzee et al., 2003).

3 Methodology

For each experiment, our methodology proceeds in three steps. First, we estimate the predicted profitability of a loan, conditional on borrower characteristics, using a flexible model estimated using a machine learning algorithm, and we use the model to estimate the (counterfactual) lender profits that would be obtained if the lender only offered loans to those classified by the algorithm to be in the top tercile of predicted profits. 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, to evaluate the distributional impact of targeting profitable borrowers, we compare outcomes from prioritizing purely by profits to outcomes from a lending policy that requires targeting to be balanced across quintiles of baseline household income. Third, we estimate conditional average treatment effects of receiving a loan on borrower outcomes using machine learning, and we use these estimates to evaluate the effects of counterfactual targeted lending policies on borrowers.

To formalize our approach, we introduce some notation. We describe our approach 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 iid sample where B is the borrower outcome (such as endline household income), L is the lender outcome (lender profits), X is a set of borrower covariates collected before randomization, A is the treatment assignment indicator, and D is the indicator for whether one actually took the loan. Denote the sample by $S = \{1, \dots, n\}$, and the subsample of borrowers by S_B , that is, $S_B = \{1 \in S : D_i = 1\}$. 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 credit worthiness.

3.1 Targeting Lender Profits

We follow the framework of Fava (2025) to train machine learning (ML) algorithms to predict lender profits and evaluate the effects of a targeting policy based on these predictions on 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 improve reproducibility of our results, we repeat the process described below 100 times and report the average estimates and standard errors, as in Fava (2025). 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 Gradient Boosting, Random Forest, Single-layer Neural Network, Elastic Net) to estimate the expected value of lender profits conditional on covariates,

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

For each observation $i \in S$, we predict lender profits using the four predictors trained on the folds that do not contain that observation. 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 B.

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. 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 population of approved applicants. We calculate robust standard errors as usual, following Fava (2025).

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 the fact that in the Philippines and South Africa, treatment assignment probabilities vary with credit score. 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, $\beta_3^L - \beta_1^L$, 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.⁶ 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 $G_{i,j}^R$.

⁶As 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.

3.3 Targeting Borrower Impact

We train ML models to identify borrowers with the largest treatment effects on household income, and evaluate their effectiveness for detecting such borrowers and the impact on lender profits of a lending policy based on these predictions. Our approach is similar to that of Subsection 3.1, where we used ML to target lender profits instead of impact in borrower income. While we can train a model to predict lender profits using only a representative sample of borrowers, predicting borrower impact leverages the randomization design to estimate individual-level treatment effects by comparing predicted outcomes under treatment and control. Specifically, we modify step 2 by running each of the four ML algorithms twice, to predict household income under treatment and control. Then, we define the predicted individual treatment effect by taking the difference between the predicted outcome under treatment and control. Again, we combine the predictions of the four algorithms 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 C.

However, while the lender profit problem is fundamentally a prediction problem—requiring only a representative sample to train accurate models—the borrower income problem is a treatment effect estimation problem. Here, randomization plays a different role: we must estimate heterogeneous treatment effects, which requires not only random treatment assignment but also a method to estimate outcomes under the counterfactual treatment status. Instead, in step 2, we run each of the four ML algorithms twice, to predict household income under treatment and control. Then, we define the predicted individual treatment effect by taking the difference between the predicted outcome under treatment and control. Again, we combine the predictions of the four algorithms following Fava (2025), and define dummies

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 C.

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 1 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 while Panel B maximizes total profits from each loan reported in absolute terms (USD).

Columns 1 and 2 in Panel A show how the total profit 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).

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 their profits on the average loan by 14% 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 profits by 9.5% 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 7% 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. We do the same for those predicted to be in the top 33% of lender profits, and report the difference in treatment effects in column 5. In Column 6 we

report how the average treatment effect on borrower income would change if the bank implemented a lending policy that only lent to the most profitable third of borrowers relative to the status quo.

These estimates test for if the profit gains that lenders can make from improved algorithmic targeting would come at the expense of borrowers. But across all specifications we report point estimates that 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. While none of these estimates reach conventional levels of statistical significance, the coefficients 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, instead of the percent relative to loan size. In this case we continue to find strong evidence of the benefits of targeting in South Africa, but no longer see differences in the Philippines.

Table 2 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 34.6 years old, which is 0.9 years 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 2% borrowers in the less profitable terciles are.

Overall we see that in South Africa targeted borrowers are older, richer, higher educated, and more likely to be male. 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. 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.

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 difference 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.

All together this suggests that existing data and machine learning strategies are sufficient to allow lenders to develop targeting strategies that would demonstrably increase profits. There is 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. 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 doesn’t 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 3. Column 1 reproduces the results from Table 1 (column 2 in Table 1), showing the difference in lender profits between those predicted to be in the top 33% of profitability and those predicted to be in the bottom 33%. Columns 4-6 consider a “balanced” targeting strategy that would split the sample into 5 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 profits by 14%, while the balanced approach would only increase profits by 2%. in the Philippines it drops from 9.5 to 5.6 percent. On average we see a drop from

12% to 4%, decreasing the benefits by nearly two thirds.

Columns 5 and 6 show how borrower impacts would change in this balanced targeting strategy. While these 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 4 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, the standard errors are large.

If we were to take the point estimates seriously we see no evidence of a tradeoff in lender profits from targeting their loans to those that benefit the most from them. As shown in Columns 6 and 9, targeting the group of borrowers who benefit the most leads to estimates of positive increases in lender profits. However, the general difficulty of this exercise suggests that current data and algorithms limit the ability to greatly improve borrower outcomes through targeting alone.

Table 5 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's 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 trade-offs lenders 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. While the example we provided 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 that even after combining three different studies, our total sample size is under ten thousand observations. 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.

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, vol. 4536641, SSRN.
- ATHEY, S., D. KARLAN, E. PALIKOT, AND Y. YUAN (2025): “Smiles in Profiles: Improving Efficiency While Reducing Disparities in Online Marketplaces,” .
- 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.
- BEAMAN, L., D. KARLAN, B. THUYSBAERT, AND C. UDRY (2023): “Selection into credit markets: Evidence from agriculture in Mali,” Econometrica, 91, 1595–1627.
- BRYAN, G. T., D. KARLAN, AND A. OSMAN (2024): “Big loans to small businesses: Predicting winners and losers in an entrepreneurial lending experiment,” American Economic Review.
- 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.
- 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.

- DUBÉ, J.-P. AND S. MISRA (2023): “Personalized pricing and consumer welfare,” Journal of Political Economy, 131, 131–189.
- FAVA, B. (2024): “Predicting the Distribution of Treatment Effects via Covariate-Adjustment, with an Application to Microcredit,” .
- (2025): “Training and Testing with Multiple Splits: A Central Limit Theorem for Split-Sample Estimators,” .
- FRIEDMAN, M. (1970): “of Business is to Increase its Profits,” New York Times Magazine, September, 13, 122–126.
- 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.
- HSIEH, C.-T. AND P. J. KLENOW (2009): “Misallocation and manufacturing TFP in China and India,” The Quarterly journal of economics, 124, 1403–1448.
- 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.
- MOLL, B. (2014): “Productivity losses from financial frictions: Can self-financing undo capital misallocation?” American Economic Review, 104, 3186–3221.
- MULLAINATHAN, S. AND J. SPIESS (2017): “Machine learning: an applied econometric approach,” Journal of Economic Perspectives, 31, 87–106.
- PÁSTOR, L., R. F. STAMBAUGH, AND L. A. TAYLOR (2021): “Sustainable investing in equilibrium,” Journal of financial economics, 142, 550–571.
- SERVAES, H. AND A. TAMAYO (2013): “The impact of corporate social responsibility on firm value: The role of customer awareness,” Management science, 59, 1045–1061.
- WORLD BANK (2025): “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.

Tables

Table 1: Targeting on Lender Profits

Panel A: Predicting Impacts on Lender Profits (%)									
Study	Lender Profits (%)			ATE on Borrower Income Change (%)			ATE on Borrower Income Change (%)		
	Bottom Tercile	Top Tercile - Bottom Tercile	Lending to Top Only vs. Current Policy	Bottom Tercile	Top Tercile - Bottom Tercile	Lending to Top Only vs. Current Policy	Bottom Tercile	Top Tercile - Bottom Tercile	Lending to Top Only vs. Current Policy
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Joint Test	-6.3	13.2*** (4.6)	7** (2.9)	-5.3	21.1 (26.5)	21.5 (15.5)	-418	2860 (3324)	1872 (2059)
South Africa	-13.0	26.9** (13)	14.3* (8.1)	-3.7	27.9 (50.5)	20.8 (29.1)	-2181	3539 (8638)	1881 (5351)
Philippines	-8.7	18.1*** (1.5)	9.5*** (0.8)	-7.2	31.6 (58)	42.4 (34.3)	150	5315 (4294)	3855 (2720)
Bosnia	2.8	-5.6 (4.8)	-2.8 (2.9)	-4.9	3.8 (20.4)	1.2 (11.9)	778	-275 (2535)	-118 (1452)
Panel B: Predicting Impacts on Lender Profits (\$)									
Study	Lender Profits (\$)			ATE on Borrower Income Change (%)			ATE on Borrower Income Change (\$)		
	Bottom Tercile	Top Tercile - Bottom Tercile	Lending to Top Only vs. Current Policy	Bottom Tercile	Top Tercile - Bottom Tercile	Lending to Top Only vs. Current Policy	Bottom Tercile	Top Tercile - Bottom Tercile	Lending to Top Only vs. Current Policy
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Joint Test	9.3	-20.4 (48.5)	-9.9 (29)	-11.1	19.1 (27)	13.1 (15.6)	-16	1708 (3312)	1088 (2067)
South Africa	-38.0	74.1** (35.9)	39.7* (23.2)	-8.8	37.9 (51.3)	25.7 (29.4)	-1086	2014 (8663)	1433 (5455)
Philippines	0.7	-0.7 (5.1)	-0.1 (2.8)	-19.5	16.9 (59.3)	13.4 (34.2)	404	3106 (4142)	1814 (2563)
Bosnia	65.2	-134.4 (141)	-69.2 (83.8)	-4.8	2.6 (20.7)	0.1 (11.9)	633	6 (2549)	16 (1454)

Notes: This table reports the results from using machine learning methods to identify borrowers based on the tercile of the predicted profits lenders would get from them. Panel A focuses on the percent change in profits relative to loan size, while Panel A replicates the analysis in levels. Column 1 reports 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 bottom tercile and those predicted to be in the top 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*.

Table 2: Who Would Be Targeted

Panel A: Targeting Top 33% for Lender Profits (%)						
	South Africa		Philippines		Bosnia	
	Targeted Group	Target Group - Everyone Else	Targeted Group	Target Group - Everyone Else	Targeted Group	Target Group - Everyone Else
	(1)	(2)	(3)	(4)	(5)	(6)
Age	34.6	0.9 (0.9)	40.3	-2.7*** (0.6)	40.3	-0.2 (0.8)
College Graduate	0.08	0.06*** (0.02)	0.76	0.25*** (0.03)	0.04	0 (0.01)
Female	0.49	-0.04 (0.05)	0.78	-0.1*** (0.02)	0.41	0 (0.03)
Income	12239	5001.3*** (800.7)	30195	18289.3*** (1392.4)	23782	426.9 (1225.8)
Loan Size	347	110.6*** (35.7)	896	349.9*** (25.1)	2110	-33.1 (134.5)
Share of Profits Coming from Penalties	0.09	0.01 (0.04)	0.22	-0.08 (0.06)	0.03	0 (0.02)
Panel B: Targeting Top 33% for Lender Profits (\$)						
	South Africa		Philippines		Bosnia	
	Targeted Group	Target Group - Everyone Else	Targeted Group	Target Group - Everyone Else	Targeted Group	Target Group - Everyone Else
	(1)	(2)	(3)	(4)	(5)	(6)
Age	35.1	1.7* (0.9)	41.3	-1.1* (0.6)	41.3	0.1 (0.8)
College Graduate	0.08	0.07*** (0.02)	0.68	0.13*** (0.03)	0.68	0 (0.01)
Female	0.50	-0.03 (0.05)	0.81	-0.1** (0.02)	0.81	0 (0.03)
Income	12824	5880.8*** (810.2)	23324	7969.1*** (1381.1)	23324	544.7 (1212.9)
Loan Size	353	123.3*** (35.6)	773	150.4*** (27.4)	773	38.1 (134.8)
Share of Profits Coming from Penalties	0.09	0 (0.04)	0.23	-0.07 (0.05)	0.23	0 (0.02)

Notes: This table 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*.

Table 3: Restricting Targeting Strategies

Panel A: Predicting Impacts on Lender Profits (%)

Study	Unrestricted Algorithmic Targeting			Restricting Targeting Strategy to Choose Top Third of Each Income Quintile			p-value for (1) - (4)
	Lender Profits (%)	Borrower Income Change (%)	Borrower Income Change (\$)	Lender Profits (%)	Borrower Income Change (%)	Borrower Income Change (\$)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Joint Test	7** (2.9)	21.5 (15.5)	1872 (16)	1.8 (2.6)	0 (15.4)	1414 (15)	0.07
South Africa	14.3* (8.1)	20.8 (29.1)	1881 (5351)	2.7 (7.3)	8 (28.8)	4029 (5008)	0.18
Philippines	9.5*** (0.8)	42.4 (34.3)	3855 (2720)	5.6*** (0.9)	-9.6 (34.3)	351 (2650)	0.00
Bosnia	-2.8 (2.9)	1.2 (11.9)	-118 (1452)	-3.1 (2.9)	1.7 (12)	-138 (1453)	0.87

Panel B: Predicting Impacts on Lender Profits (\$)

Study	Unrestricted Algorithmic Targeting			Restricting Targeting Strategy to Choose Top Third of Each Income Quintile			p-value for (1) - (4)
	Lender Profits (\$)	Borrower Income Change (%)	Borrower Income Change (\$)	Lender Profits (\$)	Borrower Income Change (%)	Borrower Income Change (\$)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Joint Test	-9.9 (29)	13.1 (15.6)	1088 (16)	-26.5 (28.9)	4.1 (15.5)	1789 (15)	0.39
South Africa	39.7* (23.2)	25.7 (29.4)	1433 (5455)	0.2 (20.1)	13 (29)	4045 (4850)	0.09
Philippines	-0.1 (2.8)	13.4 (34.2)	1814 (2563)	-1.2 (2.8)	-1.7 (34.1)	1375 (2361)	0.61
Bosnia	-69.2 (83.8)	0.1 (11.9)	16 (1454)	-78.6 (84.2)	1 (11.9)	-54 (1445)	0.86

Notes: This table compares the impacts of targeting strategies across two scenarios. In the first scenario the targeting algorithm is unrestricted - the 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 (\$). Standard errors in parentheses. Statistical Significance <0.01***, <0.05**, <0.10*

Table 4: Targeting on Borrower Income Change

Panel A: Predicting Impacts on Borrower Income Change (%)

	ATE on Borrower Income Change (%)			Lender Profits (%)			Lender Profits (\$)		
	Bottom Tercile	Top Tercile - Bottom Tercile	Lending to Top Only vs. Current Policy	Bottom Tercile	Top Tercile - Bottom Tercile	Lending to Top Only vs. Current Policy	Bottom Tercile	Top Tercile - Bottom Tercile	Lending to Top Only vs. Current Policy
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Joint Test	-7.4	12.3 (21.5)	5.9 (12.2)	-1.6	4.2 (3.6)	2.6 (2.2)	-8.3	6.1 (41.8)	5.3 (24)
South Africa	-16.2	37.6 (51.5)	17.5 (28.7)	-5.5	13.8 (13.5)	8.7 (8.1)	-12.7	27.4 (37.1)	18.3 (22.2)
Philippines	-18.2	4 (59.2)	1.5 (34)	-0.9	2.3 (1.6)	1.4 (1)	-0.6	1.7 (5)	1.1 (2.9)
Bosnia	-5.7	6.2 (20.2)	2.8 (11.7)	-0.6	0.9 (4.9)	0.3 (2.8)	-11.5	24.6 (140)	13.1 (80.7)

Panel B: Predicting Impacts on Borrower Income Change (\$)

	ATE on Borrower Income Change (\$)			Lender Profits (%)			Lender Profits (\$)		
	Bottom Tercile	Top Tercile - Bottom Tercile	Lending to Top Only vs. Current Policy	Bottom Tercile	Top Tercile - Bottom Tercile	Lending to Top Only vs. Current Policy	Bottom Tercile	Top Tercile - Bottom Tercile	Lending to Top Only vs. Current Policy
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Joint Test	647	58 (2868.4)	20.4 (1566.3)	-0.2	1.5 (3.6)	1.3 (2.1)	-13.6	11.6 (41.8)	5.5 (24.4)
South Africa	-66	-2358 (9202.5)	-1954.6 (4813.3)	-0.4	3.3 (13.5)	3.3 (7.8)	-5.0	11.9 (37.5)	10.5 (21.9)
Philippines	1984	-80.3 (4071.5)	227.9 (2482.8)	-0.4	2.7 (1.6)	2.3** (0.9)	0.2	0.9 (5.1)	1.1 (2.8)
Bosnia	984	-250.6 (2653.5)	109.9 (1476.6)	-0.4	-0.2 (4.9)	-0.6 (2.8)	-8.7	1.3 (140.1)	-7.5 (81)

Notes: 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, as well as the bottom 33% of borrowers. Panel A focuses on the change in borrower income in percent terms relative to control, while Panel B analyzes changes in levels. Column 1 reports 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 bottom tercile and those predicted to be in the top tercile with standard errors in parentheses. Column 3 reports the difference in average impacts 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 lender profits. 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 5: Characteristics of Borrowers when Targeting on Borrower Income Change

Panel A: Targeting top 33% for Borrower Income Change (%)						
	South Africa		Philippines		Bosnia	
	Targeted Group	Target Group - Everyone Else	Targeted Group	Target Group - Everyone Else	Targeted Group	Target Group - Everyone Else
	(1)	(2)	(3)	(4)	(3)	(4)
Age	33.5	-0.8 (0.9)	41.6	-0.6 (0.6)	38.4	0.9 (0.8)
College Graduate	0.04	0.01 (0.02)	0.62	0.04 (0.03)	0.04	0 (0.01)
Female	0.51	-0.01 (0.05)	0.83	-0.03 (0.02)	0.44	0.05 (0.03)
Income	9718	1216.4 (806.1)	19951	2902.5** (1256.3)	23777	419.7 (1256.9)
Loan Size	294	20.4 (36.5)	718	65.6** (27.4)	2103	-43.9 (134.7)
Share of Profits Coming from Penalties	0.11	0.05 (0.05)	0.26	-0.01 (0.06)	0.03	0 (0.02)
Panel B: Targeting top 33% for Borrower Income Change (\$)						
	South Africa		Philippines		Bosnia	
	Targeted Group	Target Group - Everyone Else	Targeted Group	Target Group - Everyone Else	Targeted Group	Target Group - Everyone Else
	(1)	(2)	(3)	(4)	(3)	(4)
Age	33.5	-0.8 (0.9)	41.7	-0.5 (0.6)	37.7	-0.1 (0.8)
College Graduate	0.04	0.01 (0.02)	0.62	0.05 (0.03)	0.04	0 (0.01)
Female	0.51	-0.01 (0.05)	0.83	-0.03 (0.02)	0.42	0.01 (0.03)
Income	9718	1216.4 (806.1)	21059	4566.4*** (1378.5)	23336	-242.3 (1220)
Loan Size	294	20.4 (36.5)	739	96.6*** (27.5)	2096	-54.9 (134)
Share of Profits Coming from Penalties	0.11	0.05 (0.05)	0.25	-0.02 (0.06)	0.03	0 (0.02)

Notes: 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, as well as the bottom 33% of borrowers. Columns 1 & 2 predict changes in percent terms, while Columns 3 & 4 predict changes in levels. Columns 1 & 3 show average characteristics for those that were targeted by the algorithm, while Columns 2 & 4 report the difference in the average characteristic between those targetted any everyone else in the sample. Standard errors in parentheses, Statistical Significance <0.01***, <0.05**, <0.10*.

Appendix A Tables

Appendix Table 1: Study Details

Study	South Africa	Philippines	Bosnia
Year of Implementation	2004	2006-2007	2008-2009
Treated/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)	\$405-\$4,054	\$312-1,562	\$52-\$1,067
Timing of Endline	6-12 months	11-22 months	14 months

Notes: South Africa: Karlan and Zinman (2010), Philippines: Karlan and Zinman (2011), Bosnia: Augsburg et al. (2015).

Appendix B Aggregating Predictions from Multiple Algorithms

We follow the approach of Fava (2025) for aggregating predictions from multiple ML algorithms. We take the following steps for aggregating lender profits 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 $\ell = 1, \dots, L$:

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 C), 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 $\ell = 1, \dots, L$:

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

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

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 lender profits for all observations in fold $\mathbf{s}_\ell$:

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

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 C Targeting Borrower Impact

We take the following steps:

1. Randomly split all **treated** observations into 3 *folds*, with roughly the same number of observations.
2. Using two folds at a time (**treated** observations only), train four ML algorithms (Extreme Gradient Boosting, Random Forest, Single-layer Neural Network, Elastic Net) for predicting borrower income. Predict borrower income for observation i using the four models trained without the fold that contains that observation. Denote the prediction of individual i with ML j ($j = 1, \dots, 4$) by $\hat{B}_{i,j}(1)$.
3. Randomly split all **control** observations into 3 *folds*, with roughly the same number of observations.
4. Using two folds at a time (**control** observations only), train the same four ML algorithms for predicting borrower income. Again, predict for observation i with models trained without the fold that contains that observation. Denote the prediction of individual i with ML j ($j = 1, \dots, 4$) by $\hat{B}_{i,j}(0)$.
5. For each of the 4 ML algorithms, calculate the predicted individual treatment effect $\hat{\tau}_{i,j} = \hat{B}_{i,j}(1) - \hat{B}_{i,j}(0)$. Combine the predictions from the four algorithms into $\hat{\tau}_i$, using the *ensemble* strategy of Fava (2025), described in Appendix B.
6. Split the borrowers into terciles of $\hat{\tau}_i$, separately by fold. Define dummies $G_{i,j}$ taking the value 1 if observation i is in the j -th tercile of their fold.

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