

# Distributional Aspects of Microcredit Expansions

Tobias G. Christiansen      Melvyn Weeks

## Abstract

Various poverty reduction strategies are being implemented in the pursuit of eliminating extreme poverty. One such strategy is increased access to microcredit in poor areas around the world. Microcredit, typically defined as the supply of small loans to underserved entrepreneurs that originally aimed at displacing expensive local money-lenders, has been both praised and criticized as a development tool Banerjee et al. (2015c).

This paper presents an analysis of heterogeneous impacts from increased access to microcredit using data from three randomised trials. In the spirit of recognising that in general the impact of a policy intervention varies conditional on an unknown set of factors, particular, we investigate whether heterogeneity presents itself as groups of winners and losers, and whether such subgroups share characteristics across RCTs. We find no evidence of impacts, neither average nor distributional, from increased access to microcredit on consumption levels. In contrast, the lack of average effects on profits seems to mask heterogeneous impacts.

The findings are, however, not robust to the specific machine learning algorithm applied. Switching from the better performing Elastic Net to the worse performing Random Forest leads to a sharp increase in the variance of the estimates. In this context, methods to evaluate the relative performing machine learning algorithm developed by Chernozhukov et al. (2019) provide a disciplined way for the analyst to counter the uncertainty as to which algorithm to deploy.

## Reference Details

CWPE                    20100  
Published            3 November 2020

Key Words            Machine learning methods, microcredit, development policy, treatment effects, random forest, elastic net

JEL Codes            D14, G21, I38, O12, O16, P36

Website                [www.econ.cam.ac.uk/cwpe](http://www.econ.cam.ac.uk/cwpe)

# Distributional Aspects of Microcredit Expansions

Tobias G. Christiansen  
Faculty of Economics  
University of Cambridge

Melvyn Weeks\*  
Faculty of Economics and Clare College,  
University of Cambridge

November 3, 2020

---

**Abstract.** Various poverty reduction strategies are being implemented in the pursuit of eliminating extreme poverty. One such strategy is increased access to microcredit in poor areas around the world. Microcredit, typically defined as the supply of small loans to underserved entrepreneurs that originally aimed at displacing expensive local money-lenders, has been both praised and criticized as a development tool Banerjee et al. (2015c).

This paper presents an analysis of heterogeneous impacts from increased access to microcredit using data from three randomised trials. In the spirit of recognising that in general the impact of a policy intervention varies conditional on an unknown set of factors, particular, we investigate whether heterogeneity presents itself as groups of winners and losers, and whether such subgroups share characteristics across RCTs. We find no evidence of impacts, neither average nor distributional, from increased access to microcredit on consumption levels. In contrast, the lack of average effects on profits seems to mask heterogeneous impacts.

The findings are, however, not robust to the specific machine learning algorithm applied. Switching from the better performing Elastic Net to the worse performing Random Forest leads to a sharp increase in the variance of the estimates. In this context, methods to evaluate the relative performing machine learning algorithm developed by Chernozhukov et al. (2019) provide a disciplined way for the analyst to counter the uncertainty as to which algorithm to deploy.

JEL Classification Codes: D14, G21, I38, O12, O16, P36

Keywords: Machine learning methods, microcredit, development policy, treatment effects, random forest, elastic net

---

---

\*Contact Author: Dr. M. Weeks, Faculty of Economics, University of Cambridge, Cambridge CB3 9DD, UK. Email: mw217@econ.cam.ac.uk

# 1 Introduction

According to the World Bank there has been marked progress in reducing poverty over the past decades, and the number of people living at or below the poverty line in 2015 was close to one third of the number in 1990. Yet, the number of people currently living in extreme poverty still amounts to 10 percent of the world’s population World Bank Group (2018). Various poverty reduction strategies have, and are, being implemented in the pursuit of eliminating extreme poverty. One such strategy is increased access to microcredit in poor areas around the world. Microcredit, typically defined as the supply of small loans to underserved entrepreneurs that originally aimed at displacing expensive local money-lenders, has been both praised and criticized as a development tool Banerjee et al. (2015c). In 2006, M. Yunus received the Nobel Prize for pioneering the concepts of microcredit as a measure to combat poverty by allowing poor entrepreneurs to expand their businesses and thereby generate economic prosperity Yunus (2006). More recently, A. Banerjee, E. Duflo, and M. Kremer jointly received the Nobel Prize in 2019 for their experimental approach to alleviating global poverty, which includes experiments concerning microcredit Duflo (2019). Others blame microcredit for harming the poor by creating over-indebtedness, and displacing existing business owners by increasing competition and eating local demand Bateman (2019).

Both the positive and negative claims may carry some truth given that initiatives, such as increased access to microcredit, rarely have homogenous effects leaving everyone better or worse off. Distributional treatment effects are therefore important as the impact of microcredit is likely to be heterogenous across individuals (see, for example, Meager. (2017a)). Understanding the heterogeneity is of great interest to policy makers as it illuminates how treatment affects sub-groups of the population differently. Solely focusing on average effects may be of little use if some groups benefit, while other lose, thereby creating a zero net effect on average when considering the entire population.

A large number of randomised evaluations of microfinance institutions (MFIs) have been deployed around the world and yet the consensus about the overall effect of microcredit has been questioned due to concerns about the external validity of the findings across RCTs Banerjee et al. (2015c); Pritchett and Sandefur (2015). Learning the characteristics of potential winners and losers across RCTs may be of use in predicting the effect of microcredit expansions in future settings that consider different populations with different characteristics Duflo (2018).

This paper presents an analysis of heterogenous impacts of increased access to microcredit, using data from three RCTs. In particular, we study heterogeneous impacts on monthly business profits and monthly consumption using data from three studies that all considered increased access to microcredit as the main intervention, employed random assignment of the access to microcredit, and collected data on large rich number of baseline covariates, thereby enabling the exploration of heterogeneity in the treatment effects. The studies we consider are Crépon et al. (2015), Attanasio et al. (2015), and Augsburg et al. (2015) which were originally deployed in Morocco, Mongolia and Bosnia & Herzegovina, respectively. We focus on monthly business profits to evaluate the claim that microcredit allow poor entrepreneurs to expand their businesses and increase profits.<sup>1</sup> Moreover, as households may benefit from microcredit in other ways, we also study heterogeneous effects on total monthly consumption which can be viewed as a crude proxy for welfare benefits.

A key contribution of this paper to the microcredit literature is the use of machine learning

---

<sup>1</sup>Another claim is that microcredit allow the poor to open businesses Meager. (2017a). The effect of entry is implicitly captured by profits as new entrants will affect the estimates.

as a tool to uncover differences in outcomes. Knowledge of how microcredit affects profits and consumption differently across households is valuable in deciding whether microcredit is worth allocating resources towards, and in particular, inform policy makers about the presence of potential winners and losers. Specifically, the use of machine learning helps in overcoming traditional issues related to multiple hypothesis testing, by providing an agnostic approach to defining subgroups with different treatment effects, in conjunction with locating the characteristics of these subgroups that are the most important in predicting the heterogeneity. The method improves on alternative ways of dealing with concerns of multiple hypothesis, such as pre-analysis plans, by allowing researchers to take full advantage of the rich baseline datasets, that are often collected in microcredit RCTs.

In addition, the method is generic with respect to the machine learning algorithms used to estimate the heterogeneity in treatment effects, as it works under biased and even inconsistent machine learning estimators. To take advantage of this, two machine learning algorithms with different strengths and weaknesses are applied. Using an Elastic Net regression and a Random Forest we analyse whether heterogeneity in treatment effects exist across studies and whether it presents itself as subgroups of winners and losers.

Across two of the three RCTs we find evidence of heterogeneous effects on monthly profits, which presents itself as groups of winners and losers, thereby supporting both the negative and positive claims of microcredit. Moreover, the heterogeneity in treatment effects can be predicted from household level covariates, and the most affected households are less indebted, engage in self-employment activities, and consume less at baseline. The findings are, however, not robust to a change in the machine learning algorithm used in conjunction with the method developed by Chernozhukov et al. (2019). Turning to monthly consumption, there is no evidence of average nor distributional effects from increased access to microcredit.

The paper proceeds as follows. Section 2 gives a brief review of the microcredit literature and the studies analysed in this paper. Section 3 provides a thorough outline of the method applied. Section 4 presents the results and section 5 concludes.

## 2 Microcredit Expansions: Experiments and Data

In the past decade numerous RCTs have been implemented in a variety of settings in the pursuit of understanding the effects of microcredit (see Crépon et al. (2015); Attanasio et al. (2015); Augsburg et al. (2015); Banerjee et al. (2015a); Karlan and Zinman (2012); Tarozzi et al. (2015); Angelucci et al. (2015)). The main focus has been on average effects on a myriad of different outcome variables, spanning from profits and consumption to educational attainment and women's empowerment. Even though most studies find some evidence towards increased take-up rates of microloans from treatment, there is a lack of evidence of significant effects on the average household. Most studies therefore find little or no support for neither the positive nor the negative claims of the effects of microcredit. Some of the studies also present a limited analysis of heterogeneous effects, primarily by estimating quantile treatment effects. Other methods such as sub-sample analysis or multiplicative models based on interactions between covariates and the treatment dummy are also provided. In addition, Pritchett and Sandefur (2015) conduct a meta-analysis of six experiments and concludes that attention should be pointed towards heterogeneous treatment effects as little can be learned from analyses that focus exclusively on, for the most part insignificant, estimates of average effects.

Meager (2019); Meager. (2017a) also conduct a meta-analysis in which she aggregates the data from the seven RCTs above using a Bayesian Hierarchical model. In particular, Meager. (2017a)

analyses heterogeneous impacts of microcredit using quantile treatment effects and finds that microcredit has negligible impacts on profits, revenues, expenses and consumption for households below the 75th percentile across settings. Meager (2019) focuses upon average treatment effects and finds that average effects on profits, revenues, expenses and different measures of consumption from increased access to microcredit may be limited. Turning to heterogeneity, she finds that approximately 60 percent of the observed variation in treatment effects across studies is due to sampling variation as opposed to genuine heterogeneity in effects. In an attempt to explain the variation in treatment effects across studies she interacts the treatment indicator with previous business experience and finds that existing business owners benefit the most from increased access to microcredit. That previous business experience is important in predicting heterogeneity in treatment effects is also found by Banerjee et al. (2015b).

## 2.1 The Randomised Experiments

Table 1 summarizes the three studies analysed in this paper. The studies all meet the following inclusion criteria: the main intervention studied is an expansion of microcredit access, the assignment of access is randomised, and data includes a large number of baseline covariates which can be used to uncover heterogeneity in treatment effects.<sup>2</sup> Apart from these requirements the studies differ across a range of factors. In particular, the studies were deployed across three different countries (Morocco, Mongolia and Bosnia & Herzegovina). In addition, and as detailed in Table 1, they differ in their treatment methods, sampling frames, eligibility criteria and microloan liabilities and therefore target different populations. Comparing treatment effects from studies that differ to such an extent, while still maintaining the same overall goal, has the potential to reveal commonalities in the factors driving the effects from increased access to microcredit. To the extent that heterogeneous treatment effects are manifest as subgroups of winners and losers that share similar characteristics, is useful in identifying mechanisms that underpin these differences.

We study the effect of randomly assigned increase in *access* to credit from MFIs on monthly business profits and total monthly consumption, both measured in USD (PPP) indexed to 2009.<sup>3</sup> The increase in access to microfinance occurred through MFIs randomly opening branches at the community level in the RCTs in Morocco and Mongolia, and from random offers made at the individual level in Bosnia & Herzegovina. Importantly, the main parameter of interest is not actual loan take up but increased access to microcredit which may be interpreted as the intent-to-treat effect. The reason for this is that the stable unit treatment value assumptions (SUTVA) are likely to fail due to informal financial links between households and general equilibrium effects in the RCTs where randomisation were done at the community level. Indeed, Banerjee et al. (2015a) argue that the possibility of general equilibrium effects or spillovers, such as effects on prices and wages, are likely. Moreover, as pointed out by Attanasio et al. (2015), focusing on the effects from access to credit allows a policy maker to learn about the effects on the population initially targeted and not just those who take up loans.

An important feature of the three studies considered in this paper is the collection of pre-treatment characteristics which can be used to uncover heterogeneity in treatment effects, by defining subgroups of the population under consideration. Importantly, characteristics used to create subgroups must be immutable, such as age or gender, or recorded at baseline (before randomisation) to rule out the possibility that they are affected by treatment, and thereby potential mediators distorting treatment effect estimates (see Montgomery et al. (2018)). Tables 5-7 show the covariates used in our analysis along with comparisons of the mean values for treated

---

<sup>2</sup>The studies were published in a special edition of AEJ: Applied Economics. The original articles and datasets are available at <https://www.aeaweb.org/issues/360>.

<sup>3</sup>Data on the PPP conversion factors can be found at the World Bank Open Data (ID: PA.NUS.PRVT.PP).

and untreated. Notably, for the study in Morocco an additional 1433 observations were randomly sampled at the endline, and no baseline covariates are available for these individuals. As these observations are missing completely at random we include them by setting the missing values equal to a constant (zero) and adding a dummy variable indicating whether the observation is missing or not.<sup>4</sup>

In contrast to Attanasio et al. (2015) who solely focus on people in group treatment (i.e. joint liability lending) we include people from both group and individual treatment schemes to enrich the sample. We then add a dummy indicating whether an individual was assigned to either type of treatment. In addition, and in contrast to Crépon et al. (2015), we abstain from trimming the sample to only include people with a high probability of borrowing as our focus is on *access* to microcredit rather than actual loan take up. Apart from these changes we conform to the decisions made by the original authors regarding the analysis and definitions of the outcome variables throughout the paper.

---

<sup>4</sup>In addition to the random missingness, there is a small number of likely non-random missing observations across a large number of covariates for the study in Morocco. We drop these observations, which amounts to approximately 3 percent of the sample, reducing the overall sample size to 5329 observations. Table 5 suggests that the covariates remain balanced across treated and untreated after doing so.

**TABLE 1**  
Summary of the Randomised Control Trials

Country	Morocco	Mongolia	Bosnia & Herzegovina
Study citation	Crépon et al. (2015)	Attanasio et al. (2015)	Augsburg et al. (2015)
Treatment	Open branches	Open branches target likely borrowers	Lend to marginally rejected borrowers
Randomisation level	Community (162 villages)	Community (40 villages)	Individual
Urban or Rural	Rural	Rural	Both
Target women?	No	Yes	No
Target microentrepreneurs?	Yes	Yes	Yes
MFI already operating	No	No	Yes
Microloan liability type	Group (3-4 people)	Both group and individual (7-15 people in groups)	Individual
Collateralized	No	Yes	Yes
Loan take up in treatment	17 pct.	54 pct.	100 pct.
Average loan size per borrower	\$1374	\$588	\$1848
Interest rate <sup>a)</sup>	14.5 percent APR	26.8 percent APR	22 percent APR
Sampling frame	Random sample plus likely borrowers	Women who registered interest in loans and met eligibility criteria.	Marginal applicants considered too risky to be offered credit as regular borrowers.
Loan eligibility	Men and women ages 18-70 who holds a national ID card, a residency certificate and have had an economic activity other than non-livestock agriculture for at least 12 months.	Women who own less than MNT 1 million (\$1948) in assets and earn less than MNT 0.2 million (\$389) in monthly profits from a business.	Men and women with sufficient collateral, repayment capacity, credit worthiness. business capacity etc.
Response rate at endline	92 pct.	84 pct.	83 pct.
Sample size <sup>b)</sup>	5524	960	995
Baseline covariates exist	Yes <sup>c)</sup>	Yes	Yes
Study starting date	April 2006	March 2008	January 2009
Study duration	24 months	19 months	14 months

Notes: The table presents characteristics of the original studies. All dollar (\$) values are USD (PPP) indexed to 2009.

<sup>a)</sup>Annual Percentage Rate (APR) given by the upper bound of the interest rate ranges reported for each study.

<sup>b)</sup>Sample sizes are defined as the endline sample sizes i.e. after attrition has been accounted for.

<sup>c)</sup>For the RCT in Morocco an additional 1433 individuals were randomly sampled at the endline to increase the sample size. For these additional individuals no baseline covariates exist.

### 3 Estimating Heterogenous Treatment Effects

The strategy employed in this paper is a method recently developed by Chernozhukov et al. (2019) for estimating heterogenous treatment effects in RCTs using machine learning. The method builds upon the potential outcomes framework introduced by Neyman (1923), and extended by Rubin (1974), and defines the main causal functions as the Baseline Conditional Average (BCA),

$$b_0(Z) \equiv \mathbb{E}[Y(0)|Z] \tag{3.1}$$

and the Conditional Average Treatment Effect (CATE),

$$s_0(Z) \equiv \mathbb{E}[Y(1)|Z] - \mathbb{E}[Y(0)|Z] \tag{3.2}$$

where  $Y(1)$  and  $Y(0)$  are the potential outcomes in the treatment state ( $D = 1$ ) and non-treatment state ( $D = 0$ ), defined by the binary treatment indicator  $D \in \{1, 0\}$ ,  $Z$  is a vector of covariates that characterise the observational unit, and  $\mathbb{E}[\cdot]$  is an expectation operator.<sup>5</sup> The data observed is  $(Y_i, Z_i, D_i)_{i=1}^N$ , consisting of  $N$  i.i.d. draws of the random vector  $(Y, Z, D)$ . In this framework  $s_0(Z)$  can be used to discover heterogeneity in treatment effects at the individual level or for groups by comparing outcomes between subgroups defined by  $Z$ . However, as is well known in the causal inference literature, and was labeled the fundamental problem of causal inference by Holland (1986), the key issue to estimating  $s_0(Z)$  is the fact that either  $Y(1)$  or  $Y(0)$  is observed for each observational unit, but never both. In order to identify  $s_0(Z)$  the following assumptions of unconfoundedness and overlap are made.

#### *Unconfoundedness*

Unconfoundedness requires that all confounding information for the relation between the treatment and the potential outcomes is captured by the observed covariates  $Z$ . In other words,  $D$  has to be as good as random conditional on  $Z$ . Mathematically, it can be expressed as,

$$Y(1), Y(0) \perp D|Z \tag{3.3}$$

where  $\perp$  denotes independence.

#### *Overlap*

This requires any observation to have a positive probability of being assigned to either the treatment or control group. Alternatively put, there has to be individuals who are treated and untreated for all possible values of  $Z$ ,

$$0 < \mathbb{P}(D = 1|Z) = \mathbb{E}(D = 1|Z) = p(Z) < 1 \quad \forall z \in Z \tag{3.4}$$

where  $p(Z)$  denotes the propensity score which measures the probability that an individual receives treatment conditional on covariates. These assumptions together are referred to as

---

<sup>5</sup> $D$  indicates whether or not an individual experiences increased *access* to microcredit.



the assumption of strong ignorability in the paper by Rosenbaum and Rubin (1983). Both assumptions are naturally satisfied in a randomised experiment given proper randomisation as this ensures that treated and untreated individuals are comparable (in expectation) across both observed and unobserved covariates, with the only difference being the treatment. This makes it possible to interpret differences in outcomes  $Y$  across treated and untreated, conditional on  $Z$ , as the causal effect of treatment.<sup>6</sup>

Under the assumption of strong ignorability, the causal functions in (3.1) and (3.2) are identified.

### 3.1 Linear Framework

An estimate of the CATE can be obtained by including interactions of the treatment indicator and the covariates of interest in a linear interaction model. Letting  $X_i$  be a single observed covariate then one model of the CATE is given by,

$$Y_i = \rho_0 + \rho_1 D_i + \rho_2 X_i + \rho_3 (D_i \times X_i) + \nu_i \quad (3.5)$$

In (3.5)  $\rho_1$  is the ATE and  $\rho_1 + \rho_3 \times X_i$  is a linear approximation of the CATE which measures how the marginal effect of treatment depends on the value of  $X_i$  Gelman and Hill (2006). The estimate of the CATE obtained from (3.5) is a linear function of  $X_i$  by construction. Moreover, the interaction involves only a single covariate. This is often the case in practice where  $X_i$  equals things such as gender, age, or previous business experience, thereby limiting the analysis of heterogeneous effects to a single dimension of the covariate space at a time. Examples from the microcredit literature include Karlan and Zinman (2012) who analyse heterogeneous treatment effects of multiple outcomes, including business size, by interacting a loan assignment indicator and a gender dummy in an OLS setting, and Tarozzi et al. (2015) who analyse heterogeneous impacts on labor supply by restricting the sample to different age groups before applying OLS.<sup>7</sup>

An issue with this approach is that researchers often have little knowledge, a priori, about which dimensions of the covariate space  $Z$  that are the most relevant in detecting heterogeneity in treatment effects. Moreover, the true functional form of  $s_0(Z)$  is rarely known, and hence the standard linear model may not be the most suitable for accurately depicting the heterogeneity in treatment effects Imbens and Wooldridge (2009). A concern of multiple hypothesis testing and invalid inference then arises if researchers search over the covariate space and alternative model specifications until they find heterogeneity of interest and reports these findings as confirmatory results, without giving an external party the ability to check the validity. The issue is that the size of the tests in such scenarios are invalid, and hence the probability of falsely rejecting a null hypothesis is larger than what it appears to be Wallach et al. (2017). Methods that adjust the sizes of the test, such as the Bonferroni and the Hochbergs corrections Bonferroni (1935); Benjamini and Hochberg (1995), only correct the sizes of the test statistics presented. If

---

<sup>6</sup>Given proper randomisation it is not necessary to condition on  $Z$  for (3.3) to hold. However, in practice, and as discussed by Deaton and Cartwright (2018) and Deaton (2019), some variables may be imbalanced in *finite* RCTs. Controlling for  $Z$  can therefore help with avoiding bias, as well as improve on the precision of the estimated treatment effect.

<sup>7</sup>Most studies in the microcredit literature also include an analysis of heterogeneous impacts by estimating quantile treatment effects (Banerjee et al. (2015a), Crépon et al. (2015), Tarozzi et al. (2015) and others)

more tests were conducted but not presented, the adjustment fails and inference becomes invalid.

One way to overcome such concerns is to introduce pre-analysis plans which consists of a report prior to the analysis stating the covariates to be used in the analysis of heterogeneity in treatment effects. This makes it clear what additional tests were run beyond those originally planned, and thus it makes multiple testing adjustment more credible Meager (2017b). Some, including Athey (2016), Chernozhukov et al. (2019) and Deaton (2019), have however pointed out that pre-analysis plans amount to throwing away lots of valuable information by restricting heterogeneity analyses to pre-registered subgroups. This has given rise for other techniques to gain ground, some of which apply machine learning.

### 3.2 The Machine Learning Approach

Machine learning is a broad term that covers a variety of different statistical learning methods, including Random Forests, Elastic Net Regularization, Neural Networks and others. Machine learning is centered around prediction and distinguishes itself from the traditional frequentist approach by using sample splitting for model selection, estimation and testing the model of interest. This is inherently different from traditional econometrics which uses all the available data to estimate the parameters of a pre-specified model and aims at drawing causal inference Athey (2016). Machine learning has proven to be excellent at handling high-dimensional data as it is effective at exploring various functional forms and discovering the most relevant covariates in order to yield accurate predictions.

Redirecting the goal of machine learning from prediction to causal inference allows researchers to reap these benefits. In particular, it allows researchers to take an agnostic and hypothesis-free approach in discovering heterogeneity in treatment effects across subgroups of the population in high-dimensional settings, by letting the data drive the model and covariate selection (Athey and Imbens (2017), Chernozhukov et al. (2019)). In particular, it is beneficial in relation to the microcredit literature as it allows researchers to take full advantage of the rich baseline characteristics that are often collected during randomised experiments.

Ideally, machine learning would therefore be used to estimate general functional forms of the BCA. Under the assumption of strong ignorability the general functional forms of the BCA and CATE can be identified as,

$$b_0(Z) = \mathbb{E}[Y|D = 0, Z] = f_b(Y, Z, D) \tag{3.6}$$

and

$$s_0(Z) = \mathbb{E}[Y|D = 1, Z] - \mathbb{E}[Y|D = 0, Z] = f_s(Y, Z, D) \tag{3.7}$$

Multiple different machine learning algorithms could be used to provide estimators for  $b_0(Z)$  and  $s_0(Z)$ . The estimators of  $b_0(Z)$  and  $s_0(Z)$  from naively applying machine learning are, however, likely to be both biased and inconsistent.

### 3.2.1 Naive Machine Learning Estimator of the CATE

A naive way to obtain estimators of  $b_0(Z)$  and  $s_0(Z)$  (i.e.  $B(Z)$  and  $S(Z)$ ) using machine learning, involves splitting the sample into two, with one part containing all the treated and the other all the untreated individuals.  $B(Z)$  is then given by the estimator obtained from fitting a given machine learning algorithm (e.g. Elastic Net, Random Forest) to the subsample with the untreated individuals, and  $S(Z)$  is given by the difference between the two estimators obtained from fitting the machine learning method to each subsample. The estimators obtained in this way are naive in the sense that they apply methods developed for prediction in an attempt to draw inferences about model parameters (treatment effects) without accounting for the model selection (regularization) that comprises the machine learning algorithms Belloni et al. (2014). In particular, most machine learning methods fail to provide valid inference as they fail to yield consistent estimators of  $s_0(Z)$  in high-dimensional settings without strict, and often not testable assumptions, such as assumptions about the sparsity of the relation between the outcome and the covariates Chernozhukov et al. (2019). The issue with machine learning estimators of conditional treatment effects stems from biases that result from model selection (i.e. regularization).<sup>8</sup>

#### *Regularization Bias*

Most machine learning methods have prediction at their core and use the Mean Squared Error (MSE) as the objective to minimize. This minimization often relies on regularization of the machine learning methods to avoid overfitting the data used to estimate the model and thereby keeping the variance of the resulting estimator from exploding. It does, however, also induce a bias in  $S(Z)$ . The issue with the naive machine learning estimators described above arises as the treatment indicator is implicitly forced to remain in the models (as it defines the subsamples) and hence  $D$  becomes immune to regularization. Any variable that is highly correlated with the treatment variable therefore tends to be excluded or neglected from the model since the variable does not add much predictive power for the outcome given the treatment indicator is already controlled for. If such variables also correlate with the outcome then regularization is effectively excluding confounding variables from the model leading to a biased estimator of  $s_0(Z)$  Belloni et al. (2014). Counter to what is commonly presumed in RCTs, the presence of confounding variables in RCTs may indeed be an issue as the likelihood that some variables are imbalanced across treated and untreated, and thereby correlated with  $D$ , is high in *finite* RCTs Deaton and Cartwright (2018); Deaton (2019).<sup>9</sup> The regularization bias from neglecting confounders will in general converge to zero at very slow rates, thereby invalidating inference that relies on assumptions of asymptotic normality Chernozhukov et al. (2017).

---

<sup>8</sup>An exception is the Causal Forest algorithm developed by Athey and Imbens (2018) which is an adaptation of the Random Forest that aims at providing valid pointwise inference for the CATE. However, as Chernozhukov et al. (2019) point out the consistency of the Causal Forest hinges on the dimension of  $Z$  being less than  $\log(n)$ . Causal Forests are therefore only guaranteed to provide valid results in low-dimensional settings which contrasts the settings considered in this paper.

<sup>9</sup>By inspection of Tables 5-7 it is evident that a few variables in each RCT are imbalanced and correlates with treatment.

### 3.3 Machine Learning Proxies

The method developed by Chernozhukov et al. (2019) allows for estimating heterogeneous treatment effects in randomised experiments using machine learning which is valid in high-dimensional settings. In particular, the method avoids imposing strong conditions on the naive machine learning estimators as it works under biased and even inconsistent learning. Instead, Chernozhukov et al. (2019) simply treat the naive machine learning estimators described in the previous subsection as proxies that serve in place of  $b_0(Z)$  and  $s_0(Z)$ . The machine learning proxy of  $s_0(Z)$  is then used to create an orthogonalized variable which together with an orthogonalized treatment indicator identify features or objects of the CATE function in separate regression models. The method is therefore a two-step procedure in which machine learning is used as a part of a less ambitious goal that aims at describing a few properties of the CATE function (3.7) rather than the entire CATE function.

Chernozhukov et al. (2019) propose the following steps to obtain the machine learning proxies,

1. Randomly split one of the RCT samples into a main sample ( $Data_M$ ) and a auxiliary sample ( $Data_A$ ) of equal size.
2. Use  $Data_A$  to estimate a machine learning algorithm (e.g. Elastic Net, Random Forest) to predict  $Y$  using  $Z$ . In particular,  $Data_A$  is split into treated and untreated individuals and the machine learning method is fit to each subsample.
3. For each  $i \in Data_M$  predict  $Y_i$  for each observation in  $Data_M$  using both prediction models from the previous step to obtain two predicted outcomes,  $\hat{Y}_i^{D=1}$  and  $\hat{Y}_i^{D=0}$ .
4. Calculate the proxy of  $b_0(Z)$  using the model trained on the control group  $B(Z) = \hat{Y}_i^{D=0}$  and the proxy of  $s_0(Z)$  as the difference between the predictions  $S(Z) = \hat{Y}_i^{D=1} - \hat{Y}_i^{D=0}$ .

The procedure is done twice with  $Y$  equal to monthly profits and monthly consumption for each RCT, and two pair of the proxies  $B(Z)$  and  $S(Z)$  are obtained. In both cases,  $Z$  is given by the covariates presented in Tables 5-7.<sup>10</sup> The proxies obtained from the above procedure are similar to the naive machine learning estimators described, and  $S(Z)$  is likely to be a biased and inconsistent estimator of  $s_0(Z)$  due to regularization. Nevertheless, Chernozhukov et al. (2019) show that  $S(Z)$  can still be used to extract important *features* of  $s_0(Z)$ .

#### 3.3.1 Key Features of the CATE

##### *Best Linear Predictor*

The first feature is the Best Linear Predictor (BLP) of the CATE function based on the proxy  $S(Z)$ ,

$$\beta_1 + \beta_2(S(Z) - \mathbb{E}[S(Z)]) \tag{3.8}$$

---

<sup>10</sup>For the RCTs in Morocco and Mongolia, we follow the advice by Chernozhukov et al. (2019) and use stratified sample splitting.

Here,  $\beta_1$  is the average treatment effect (ATE),

$$\beta_1 = \mathbb{E}[s_0(Z)] \quad (3.9)$$

and  $\beta_2$  captures any additional heterogeneity in treatment effects. The BLP of the CATE therefore provides individualized treatment effects. Chernozhukov et al. (2019) show that,

$$\beta_2 = \frac{\text{Cov}(s_0(Z), S(Z))}{\text{Var}(S(Z))} \quad (3.10)$$

From (3.10) it is evident that  $\beta_2$  captures how well the proxy  $S(Z)$  approximates  $s_0(Z)$ . If  $S(Z)$  is a perfect proxy then  $\beta_2 = 1$  and if  $S(Z)$  is complete noise then  $\beta_2 = 0$ . In addition, if there is no heterogeneity in the treatment effects, such that  $s_0(Z) = s_0 \in \mathbb{R}$ , then  $\beta_2 = 0$ . Rejection of the null hypothesis of  $\beta_2 = 0$  therefore implies that heterogeneous treatment effects exist and  $S(Z)$  is an effective predictor of it.

Empirically, the coefficients of the BLP can be estimated from a weighted least squares (WLS) regression using the main sample and the proxies  $B(Z)$  and  $S(Z)$ ,

$$Y_{ij} = \alpha_0 + \omega_j + \alpha_1 B(Z_i) + \alpha_2 S(Z_i) + \beta_1 (D_i - p(Z_i)) + \beta_2 (D_i - p(Z_i))(S(Z_i) - \overline{S(Z)}) + \epsilon_{ij} \quad (3.11)$$

Here, the subscript  $i$  refers to the individual and  $j$  to the village/province. The weights used are the inverse of the variance of  $D_i$  i.e.  $w(Z_i) = (p(Z_i)(1 - p(Z_i)))^{-1}$  where  $p(Z_i)$  denotes the propensity score. The estimator of  $p(Z_i)$  is set to the proportion of treated individuals in each RCT because treatment was randomly assigned.  $\overline{S(Z)} = |M|^{-1} \sum_{i \in M} S(Z_i)$  is the empirical expectation of  $S(Z_i)$  with respect to the main sample. Equation (3.11) is estimated for both monthly profits and monthly consumption, with clustered standard errors at the village level for the RCTs in Morocco and Mongolia to account for the possibility of correlated shocks within villages. For the RCT in Bosnia & Herzegovina standard errors robust to heteroskedasticity are used.  $\omega_j$  are strata dummies at the village level for the RCT in Morocco, and at the province level for the RCT in Mongolia. No strata dummies are included for the RCT in Bosnia & Herzegovina as randomisation was not stratified on regions.

Ignoring the term  $\alpha_1$ ,  $\alpha_2$  and the strata dummies  $\omega_j$  which are included to improve precision, (3.11) resembles the standard linear regression approach in estimating conditional average treatment effects (3.5). In particular, the coefficient on the (orthogonalized) treatment indicator ( $\beta_1$ ) is the average treatment effect and the coefficient on the interaction term ( $\beta_2$ ) captures heterogeneity in treatment effects. There are, however, three important things that distinguish (3.11) from the standard framework.

First, the variables of interest are orthogonalized to combat regularization bias in the machine learning proxies. Subtracting the empirical counterpart of the conditional expectations of  $S(Z_i)$  and  $D_i$  given  $Z_i$  (i.e.  $\overline{S(Z)}$  and  $p(Z_i)$ ) from the variables themselves produces orthogonalized variables. The orthogonalized variables are, by construction, orthogonal to all the covariates in  $Z_i$ , and any function of these, under any  $Z_i$ -dependent weight. Any confounding effects induced by regularization are therefore partialled out under the weight  $w(Z_i)$ , as the terms  $(D_i - p(Z_i))(S(Z_i) - \overline{S(Z)})$  and  $(D_i - p(Z_i))$  are orthogonal to each other and all other functions of  $Z_i$ .<sup>11</sup>

---

<sup>11</sup>The orthogonalization technique is closely related to Double/Debiased Machine Learning Chernozhukov et al. (2017), and builds on ideas that go all the way back to Frisch and Waugh (1933) and Lovell (1963). It is clear

Next, the WLS regression is estimated using the main sample ( $Data_M$ ) while the proxies  $B(Z_i)$  and  $S(Z_i)$  are estimated using the auxiliary sample ( $Data_A$ ). If instead the machine learning algorithms used to create the proxies were fit using the entire sample it would be likely to induce a relation between these and the model errors in (3.11), where they are used as plug-ins, as data for observation  $i$  is used in forming both. This may then lead to a bias in, and slower convergence of, the estimator of  $\beta_2$  Chernozhukov et al. (2019). Using different observations to estimate the machine learning proxies and the regression model breaks any link between the proxies ( $B(Z_i)$ ,  $S(Z_i)$ ) and the model errors ( $\epsilon_i$ ).

Finally, the interaction term is not limited to a single covariate, and thereby focusing on heterogeneity in treatment effects in a single dimension at a time, which tend to be the case in a standard linear regression framework. Instead, the (orthogonalized) treatment indicator is interacted with the (orthogonalized) proxy of the CATE which is a function of the entire covariate space and thereby captures heterogeneity in treatment effects across all dimensions of  $Z_i$ . This latter fact is important as  $\beta_2$  can be used to detect heterogeneity across all dimensions of  $Z_i$  at once, rather than searching over the covariate space leading to concerns of multiple hypothesis.

#### *Sorted Group Average Treatment Effects*

The second feature of  $s_0(Z)$  is the Group Average Treatment Effects (GATES) which measures the average treatment effects across groups which are defined using  $S(Z)$ .<sup>12</sup> In particular, we estimate the average treatment effects for four non-overlapping groups  $G_1$  to  $G_4$  where  $G_4$  is the group that consists of the top 25 pct. with the highest predicted treatment effects (most affected) and  $G_1$  which consists of individuals with the bottom 25 pct. lowest predicted treatment effects (least affected) according to  $S(Z)$ . The parameters of interest are then,

$$\mathbb{E}[s_0(Z)|G_k], \quad k = \{1, 2, 3, 4\} \quad (3.12)$$

where  $G_k$  denotes the  $k^{th}$  group. These parameters can be estimated by using WLS on the following regression equation,

$$Y_{ij} = \alpha_0 + \omega_j + \alpha_1 B(Z_i) + \alpha_2 S(Z_i) + \sum_{k=1}^4 \gamma_k (D_i - p(Z_i)) 1\{i \in G_k\} + \nu_{ij} \quad (3.13)$$

The weights are defined as for (3.11) and  $1\{i \in G_k\}$  indicates whether individual  $i$  belongs to group  $k$ . Equation (3.13) is estimated for both monthly profits and monthly consumption for each RCT and includes strata dummies and uses clustered/robust standard errors as described for equation (3.11). The parameters of interest are  $\{\gamma_k\}_{k=1}^4$ . Importantly, and as described, the treatment indicator is orthogonalized to combat regularization bias, and the regression is estimated using only observations from the main sample ( $Data_M$ ) to avoid overfitting.

#### *Classification Analysis*

If the BLP and GATES detect significant heterogeneity in treatment effects then the Classification Analysis (CLAN) estimator can be used to describe the average characteristics of the most ( $G_4$ ) and least ( $G_1$ ) affected individuals. Information on how effects vary with covariates is useful

---

that orthogonalization is only as good as the covariates available as the orthogonalized variables need not be orthogonal to *unobserved* confounders.

<sup>12</sup>A drawback with using  $S(Z)$  to define the groups is that  $S(Z)$  is likely to be a biased estimator of the true treatment effects, so the groups defined may not capture the true groupings. The challenge with estimating treatment parameters of this nature is that they depend on the *joint* distribution of the potential outcomes of treated and untreated, which is not identified by randomisation Abadie et al. (2018).

if the goal is to explore mechanisms or predict what the result of an intervention may be in a specific population. Let  $g(Y, Z)$  be a vector of characteristics of an individual. Then the parameters of interest are,

$$\delta_1 = \mathbb{E}[g(Y, Z)|G_1] \quad \text{and} \quad \delta_4 = \mathbb{E}[g(Y, Z)|G_4] \quad (3.14)$$

which measure the average characteristics of the most and least affected individuals. The CLAN estimates are obtained by regressing each characteristic onto a pair of dummies indicating whether an individual belongs to the most or least affected as defined by  $S(Z)$ ,

$$z_{ij} = \delta_1 1\{i \in G_1\} + \delta_4 1\{i \in G_4\} + \eta_{ij} \quad (3.15)$$

where  $z_{ij}$  denotes the  $j^{\text{th}}$  variable in  $Z$  that belongs to individual  $i \in \text{Data}_M$ . The CLAN is conducted for both profits and consumption for each RCT using the covariates in Table 5-7.

### 3.3.2 Inference

To make valid inference on the features of the CATE function, it is important to note that the estimators of these, such as  $\hat{\beta}_2$ , inherits randomness from both the standard *estimation uncertainty* and from the *random splitting* of the sample into a main and auxiliary sample. In particular, the parameters depend implicitly on the auxiliary sample as it is used to estimate the machine learning proxies. Different partitions of the data therefore yield different estimators. Chernozhukov et al. (2019) show that the estimators are normally distributed conditional on the sample split and under mild regularity conditions. For instance,

$$\hat{\beta}_{2,s} | \text{Data}_A \stackrel{a}{\sim} N(\beta_{2,s}, \sigma_{2,s}^2)$$

for  $n \in \text{Data}_M \rightarrow \infty$ . Here the subscript  $s$  indicates the split used to estimate  $\hat{\beta}_2$ . The split-dependent estimator  $\hat{\beta}_{2,s}$ , along with related confidence intervals, are therefore treated as random *conditional* on the data due to the random partitioning. To address this randomness Chernozhukov et al. (2019) propose to repeat steps 1-4 and re-estimate equations (3.11), (3.13) and (3.15) multiple ( $S$ ) times. The final point estimators and confidence intervals then consists of taking the median of the  $S$  split-dependent quantities.<sup>13</sup> In particular, the final point estimator of  $\hat{\beta}_2$  is the median of the split-dependent estimators  $\hat{\beta}_2 = \text{Med}(\{\hat{\beta}_{2,s}\}_{s=1}^S)$ . Likewise, medians of split-dependent upper and lower confidence bounds define the final confidence interval. Importantly, the nominal confidence bounds have to be adjusted in order to account for both estimation uncertainty and splitting uncertainty induced by random partitioning of the data. Therefore, to guarantee uniform validity, the final confidence level is  $(1 - 2\alpha)\%$ . Similarly, the final sample-splitting adjusted  $p$ -value is twice the median of the split-dependent  $p$ -values.

The benefit of using multiple sample splits is two-fold. Using many splits is a way to obtain more versions of the estimator of a particular key feature. Letting the final estimator be a function of these estimators (such as the median) helps in utilizing all the data available to estimate the parameters of interest, compared to an estimator obtained from using a single split which utilizes only half of the sample Chernozhukov et al. (2017); Athey (2016). Another reason for using many splits of the data into main and auxiliary samples is to produce robust estimators. Only using a single split could lead to results that do not hold in general (i.e. for other splits).

<sup>13</sup>The median is used, in contrast to using the mean, as the former is more robust to outliers.

### 3.4 Elastic Net Regression and the Random Forest Algorithm

An appealing feature of the method described is that it is generic with respect to the machine learning algorithm used to estimate the proxies  $S(Z)$  and  $B(Z)$  as it does not rely on strong and specific assumptions on the method employed. To take advantage of this generality we apply two techniques that differ substantially in their approach to predict  $s_0(Z)$  and  $b_0(Z)$ . These are the Elastic Net regularization and the Random Forest algorithm.

Elastic Net Regularization is a parametric machine learning algorithm that was developed by Zou and Hastie (2005) and builds on top of the linear regression framework by regularizing the parameters. In particular, it extends the standard least squares estimator by adding two additional terms to the minimization criterion,

$$\hat{\theta}_{EN} = (1 + \lambda_2) \underset{\theta}{\operatorname{argmin}} \{L(\lambda_1, \lambda_2, \theta)\} \quad (3.16)$$

where

$$L(\lambda_1, \lambda_2, \theta) = \|Y - Z\theta\|_2^2 + \lambda_2 \|\theta\|_2^2 + \lambda_1 \|\theta\|_1, \quad \lambda_1, \lambda_2 \geq 0 \quad (3.17)$$

Here,  $\theta$  is a  $p$ -dimensional vector of parameters and  $\lambda_1, \lambda_2$  are positive constants associated to the L1-norm ( $\|\theta\|_1 = \sum_{j=1}^p |\theta_j|$ ) and the L2-norm ( $\|\theta\|_2^2 = \sum_{j=1}^p \theta_j^2$ ). Elastic Net regularization therefore corresponds to a penalized least squares method where the number and the absolute size of the parameters are reduced by the regularization terms  $\lambda_1$  and  $\lambda_2$ . In other words, the method adds automatic variable selection, in which it can select groups of correlated variables, and continuous shrinkage to the standard OLS case. Whereas OLS will estimate non-zero coefficients  $\theta_j$  for all covariates and thereby increasing the chance of overfitting (i.e. fitting noise in the response variables), the Elastic Net will set some coefficients to zero and shrink others to avoid this. The benefit of doing so is an increase in the accuracy of prediction as measured by the mean squared error (MSE) as well as a more parsimonious model that increases interpretability Zou and Hastie (2005). Notably, the Elastic Net works well in high-dimensional settings as it yields a sparse solution vector with only the parameters of the most relevant variables or groups of variables being nonzero, thereby providing an effective data-driven way to detect the variables that appear to be the most important in predicting outcomes linearly. The Elastic Net is better suited to provide an estimator for  $s_0(Z)$  if it is well-approximated by a linear combination of pre-specified functions of  $Z$ . If  $s_0(Z)$  is believed to be non-linear other machine learning algorithms are better.

A Random Forest is a non-parametric tree-based machine learning algorithm developed by Breiman (2001) that builds upon his earlier work on the Classification and Regression Tree (CART) algorithm Breiman et al. (1984). Random Forests are known to handle both non-linear relations and high-dimensional settings well Hastie et al. (2009). The Random Forest algorithm is an ensemble method that average over a large number of individual regression trees. In particular, each tree in the Random Forest is estimated on a randomly drawn bootstrap sample, and is composed of partitions of the data based on binary indications of the covariates. These partitions are found using Recursive Binary Splitting (RBS). Each tree then fits a simple model, such as the average outcome, in each of its partitions, and the Random Forest averages over all the trees. Importantly, each tree in the Random Forest uses a random subset of the available covariates  $Z$  as potential splitting candidates at each split-point. Growing individual trees on random bootstrap samples with a random subset of covariates available at each splitting point helps in creating decorrelated trees, which leads to lower variance of the resulting Random Forest



estimator. In particular, for an individual with a vector of covariates  $Z_i$  the estimator is,

$$\hat{Y}_i = \frac{1}{B} \sum_{b=1}^B T_b(Z_i) \quad (3.18)$$

where  $B$  is the number of trees and  $T_b(Z_i)$  is the estimated outcome from the  $b^{\text{th}}$  tree. The Random Forest improves on individual trees as it is less prone to overfitting and yield more stable estimators. In contrast to the Elastic Net the Random Forest is bad at modeling linearly additive structures Hastie et al. (2009).

### 3.4.1 Evaluation of the Machine Learning Algorithms

Chernozhukov et al. (2019) provide two measures for deciding on the better performing machine learning algorithm in a disciplined way. The best machine learning method for fitting the BLP of  $s_0(Z)$  is the one that maximizes,

$$\Lambda \equiv |\beta_2|^2 \text{Var}(S(Z)) \quad (3.19)$$

using the main sample. Maximizing  $\Lambda$  is equivalent to maximizing the correlation between the machine learning proxy predictor  $S(Z)$  and the true CATE  $s_0(Z)$ . For GATES, the best machine learning algorithm is the one that, given four groups of equal size, maximizes,

$$\bar{\Lambda} \equiv \frac{1}{4} \sum_{k=1}^K \gamma_k^2, \quad k \in \{1, 2, 3, 4\} \quad (3.20)$$

which captures the part of the variation of  $s_0(Z)$  explained by  $\widetilde{S(Z)} = \sum_{k=1}^4 \gamma_k 1(S(Z) \in G_k)$ . This corresponds to maximizing the  $R^2$  in a regression of  $s_0(Z)$  on  $\widetilde{S(Z)}$  without a constant.

The measures in (2.20) and (2.21) can only be used for ordinal comparisons as the measures do not tell how well a given algorithm fits the BLP or GATES, nor how much better a given algorithm is compared to another. This issue arises as machine learning methods are commonly being evaluated using a MSE criteria. As the true treatment effect is never observed such a criterion is infeasible.<sup>14</sup>

## 4 Results

The identification of heterogeneity in treatment effects from increased access to microcredit, and in particular describing the characteristics of subgroups of winners and losers, is an important part of understanding the impacts of microcredit and its usefulness as a development tool. This knowledge can be of use in predicting the impacts of access to microcredit in future settings, and thereby help inform policy makers whether or not to allocate resources towards microcredit.

By applying the methods described in the previous section we start out by presenting an analysis of the existence of heterogeneous treatment effects on profits and consumption from increased access to microcredit. We then investigate the existence of groups of winners and losers. Finally, given the existence of winners and losers, we depict key characteristics of these subgroups in an attempt to uncover the most important characteristics in predicting heterogeneous effects from

<sup>14</sup>As an exception, (n.d.) have developed a tree based estimator that uses an adjusted MSE which is an unbiased estimator of the MSE.

increased access to microcredit.

Before turning to the actual analysis, Table 8 summarises the performance of the Elastic Net and Random Forest algorithms. From Table 8 it is evident that the Elastic net outperforms the Random Forest across profits and consumption for both the metrics discussed in section 3.4.1. One possible explanation is that the CATE function is linear, or close to linear, and therefore better approximated by the Elastic Net than the Random Forest. An alternative explanation is the fact that the Elastic Net is a parametric method opposed to the non-parametric Random Forest. In general, parametric methods require less data to produce precise estimates of the mapping functions from  $Z$  to  $B(Z)$  and  $S(Z)$  compared to non-parametric methods (see Brownlee (2016)). Given the sample sizes considered in this paper this may be the deciding factor. Based on the results of Table 8, we choose to focus on the estimates provided by the Elastic Net in the coming subsections. The estimates where the proxies  $S(Z)$  and  $B(Z)$  are estimated using a Random Forest will be considered as a measure of robustness at the end of the section.<sup>15</sup>

#### 4.1 Heterogeneous Effects from Microcredit

Table 2 presents estimates of  $\hat{\beta}_1$  and  $\hat{\beta}_2$ . The former measures the average treatment effect (ATE) of increased access to microcredit on monthly profits and total monthly consumption both measured in USD (PPP), and the latter measures the relevance of the proxy  $S(Z)$  and whether there is heterogeneity in the treatment effects (HET). The estimates are obtained from the BLP of the CATE using equation (3.11). Notably, and in line with the original studies, the average treatment effects on profits and consumption across all RCTs are not significantly different from zero. Solely focusing on the average effects of increased access to microcredit therefore indicates that it has no effect on the abilities of microentrepreneurs to generate profits. Nor does it impact the well-being of individuals by increasing consumption.

Turning to the estimates of  $\beta_2$  we observe a different story. In particular for profits,  $\hat{\beta}_2$  is positive and significantly different from zero with  $p$ -values of 0.002 and 0.070 for the RCTs in Morocco and Mongolia while it is insignificant for Bosnia & Herzegovina. This implies that heterogeneity in treatment effects are present in two out of the three RCTs considered. The absence of average treatment effects for profits may be due to the presence of groups of winners with high treatment effects and groups of losers with low treatment effects that cancel out when averaging over the entire population.

Interestingly,  $\hat{\beta}_2$  for consumption remains insignificant across all three RCTs, leaving no sign of heterogeneous impacts from increased access to microcredit on total monthly consumption.

---

<sup>15</sup>We implement the Elastic Net (EN) and Random Forest (RF) using the CARET package in R Kuhn et al. (2020). All reported results are medians over 50 sample splits. The tuning parameters for the EN are found using twice repeated 2-fold cross-validation and a random grid search. For the RF the number of splitting variables is set to  $p/3$  where  $p$  denotes the dimension of  $Z$ , and the number of trees is set to 1000. The outcomes and covariates are rescaled to be between 0 and 1 before estimating the models. The machine learning configurations and the number of sample splits are chosen to accommodate computational time. It takes approx. 18 hours to estimate all quantities at the current settings.

**TABLE 2**  
Coefficients of the Best Linear Predictor (BLP) of the CATE across RCTs

	Morocco		Mongolia		Bosnia & Herzegovina	
	ATE ( $\hat{\beta}_1$ )	HET ( $\hat{\beta}_2$ )	ATE ( $\hat{\beta}_1$ )	HET ( $\hat{\beta}_2$ )	ATE ( $\hat{\beta}_1$ )	HET ( $\hat{\beta}_2$ )
Profits	26.959 (-25.386,77.802) [0.607]	0.260 (0.104,0.423) [0.002]	-1.893 (-4.017,0.277) [0.180]	0.186 (0.012,0.357) [0.070]	74.259 (-62.844,213.014) [0.497]	0.088 (-0.317,0.518) [0.930]
Consumption	-14.078 (-36.867,8.542) [0.424]	0.146 (-0.066,0.356) [0.383]	109.152 (-46.484,255.036) [0.343]	0.279 (-0.107,0.699) [0.271]	-38.654 (-183.837,113.039) [1.000]	0.069 (-0.222,0.409) [0.946]

Notes: The table presents estimates of the average treatment effects (ATE) and heterogeneity in treatment effects (HET) for profits and consumption. The estimates are given by the median of the estimates of  $\beta_1$  and  $\beta_2$  obtained by applying WLS to equation (3.11) over 50 splits. Similarly, the sample splitting adjusted confidence intervals and  $p$ -values are obtained by the median over 50 splits with the  $p$ -values being twice this value to correct for sample splitting uncertainty. Clustered standard errors at the village level for Morocco and Mongolia, and robust standard errors for Bosnia & Herzegovina are used. Strata dummies at the village level for Morocco and at the province level for Mongolia are included. The proxies  $S(Z)$  and  $B(Z)$  that enters equation (3.11) are estimated using an Elastic Net on the covariates presented in Tables 5, 6, and 7, for Morocco, Mongolia and Bosnia & Herzegovina respectively. Profits and consumption are monthly profits from business activities and total monthly consumption, both measured in USD PPP indexed to 2009.

- 90% confidence intervals in parentheses.
- $p$ -values for the hypothesis that the parameter is equal to zero in brackets.

## 4.2 Groups of Winners and Losers

To analyse the nature of the heterogeneity (or lack thereof) Figure 1 shows the point estimates and related 90 pct. confidence intervals for the Sorted Group Average Treatment Effects (GATES) across four groups G1 to G4: G4 consists of top 25 pct. of the individuals with the highest treatment effect and G1 the bottom 25 pct. all defined by the proxy  $S(Z)$ .<sup>16</sup> In particular, it is evident that there are no heterogeneity in treatment effects on consumption as there is substantial overlap between the confidence intervals of the treatments effects for all groups across all RCTs. The idea that microcredit may benefit poor people by allowing them to raise their consumption and thereby their well-being therefore seems to be unobtainable, at least in the short run, when considering both distributional and average impacts given the three RCTs examined in this paper.<sup>17</sup>

For profits the GATES estimates in Figure 1 do indicate that heterogeneity in treatments exist. For the RCT in Morocco it comes about as a group of winners whereas there are significant losers for the RCT in Mongolia, thereby supporting both the negative and positive claims on the effects of microcredit. Finally, there is no evidence of heterogenous impacts for the RCT in Bosnia & Herzegovina. It is difficult to explain this finding given that the RCTs differ across a variety of parameters (see Table 1). One obvious difference, however, is that the RCT in Bosnia & Herzegovina is randomised at the individual level by focusing on marginal clients all over Bosnia & Herzegovina. This contrasts the RCTs in Morocco and Mongolia that randomised at the community level and thereby created strata of treated and untreated people. Any spill-overs, general equilibrium effects, or social network effects inside villages are therefore captured by only the latter two. Banerjee et al. (2015a) argue that spill-overs are important, while Crépon et al. (2015) argue for it having limited effects. Nevertheless, if spill-overs are important drivers

<sup>16</sup>As mentioned in section 3.3.1, G1-G4 may not reflect the true groupings when defined using  $S(Z)$ . The positive estimates of  $\hat{\beta}_2$  and the monotonically increase in group treatment effects depicted in Figure 1 indicate that  $S(Z)$  and  $s_0(Z)$  are positively correlated. Moreover, there is evidence that  $S(Z)$  do capture groups that differ substantially in their treatment effects, which is the main object of interest.


<sup>17</sup>Given that our focus is upon total monthly consumption we cannot rule out that individuals shift consumption spending across goods .

of treatment effects it won't be captured by the sampling and randomisation schemes used in Bosnia & Herzegovina.

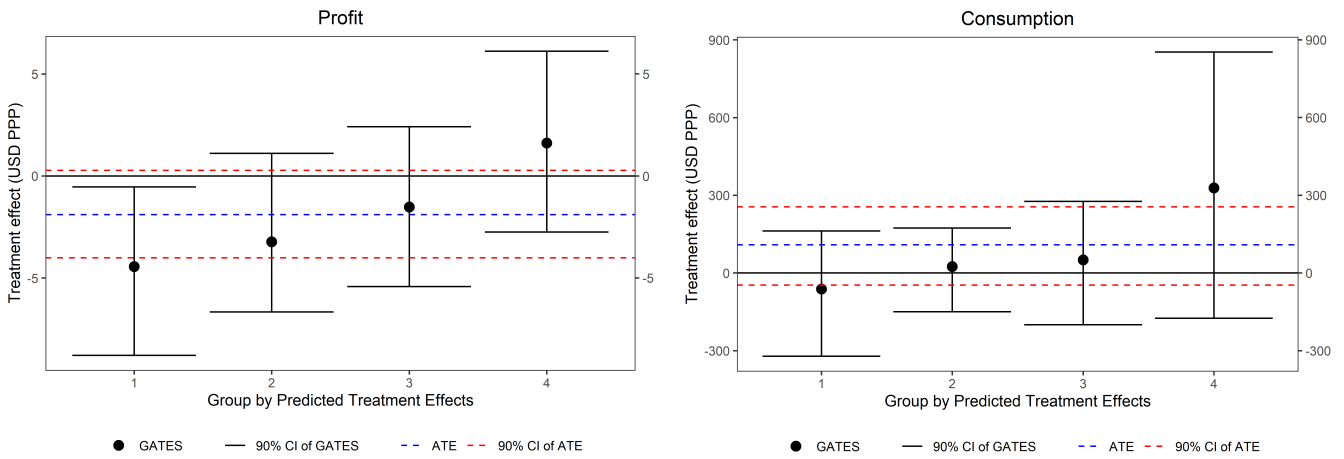
**Figure 1**

Sorted Group Average Treatment Effects (GATES) for profits and consumption across RCTs


(a) Morocco



(b) Mongolia



(c) Bosnia & Herzegovina



Notes: Figures (a), (b) and (c) presents the Sorted Group Average Treatment Effects (GATES) point estimates for monthly profits and monthly consumption measured in USD (PPP) indexed to 2009 along with 90 pct. confidence intervals for the RCTs in Morocco, Mongolia and Bosnia & Herzegovina, respectively. The final estimates and confidence intervals are found by estimating (3.13) using WLS over 50 splits and then taking the median of the split-dependent estimates. Clustered standard errors at the village level for Morocco and Mongolia, and robust standard errors for Bosnia & Herzegovina are used. Strata dummies at the village level for Morocco and at the province level for Mongolia are included. The four non-overlapping groups G1 to G4 are defined by the proxy  $S(Z)$  which is estimated using an Elastic Net. G4 consists of the top 25 pct. most affected and G1 of the bottom 25 pct. least affected. The figures also show the Average Treatment Effect (ATE) along with 90 pct. confidence intervals.

Table 3 shows the estimated treatment effects on profits and consumption for the most and least affected as well as the difference between the groups. Notably, the most affected in Morocco experience a significant increase in profits of 171.7 USD (PPP) or 103 percent of the average value in the control group. Comparing this to the least affected with a point estimate of -57.7 USD (PPP) gives a significant difference of 227 USD (PPP). In other words, increased access to microcredit seem to generate groups of winners and losers in Morocco.

Turning to Mongolia shows that increased access to microcredit can have harmful effects for a part of the population as the least affected individuals experiences a statistically significant drop in monthly profits of 4.4 USD (PPP), corresponding to a decrease of 102 percent of the average value in the control group. Here, it is worth remembering that the RCT in Mongolia focused on the poorest of the poor women (see Table 1). Additionally, when comparing the effects of the most and least affected individuals in Mongolia we observe a significant difference just shy of 7 USD (PPP), indicating that increased access to microcredit produces groups of winners and losers. Finally, for Bosnia & Herzegovina there is no significant difference.

The heterogeneity in treatment effects on monthly profits indicated by the estimates of the BLP of the CATE represents the presence of winners and losers across the RCTs in Morocco and Mongolia. Heterogeneous treatment effects may be driven by gains of entrants at the expense of incumbent business owners. Alternatively, existing business owners may reap most of the benefits from increased access to microcredit due to expertise or skill, whereas unexperienced entrants see most of their profits evaporate in the presence of high interest rates on microloans (see Table 1).

To gain a better understanding of the heterogeneity in treatment effects we turn to the Classification Analysis (CLAN) which describes the average characteristics of the most and least affected individuals. If there exist characteristics that are pronounced in either group across the RCTs, then knowledge of such characteristics could be of relevance in predicting whether access to microcredit is favourable or not, and in particular who it may help and who it may harm in future studies.

**TABLE 3**  
GATES estimates for profits and consumption across RCTs

<b>Panel A: Morocco</b>			
	25 % most ( $\hat{\gamma}_4$ )	25 % least ( $\hat{\gamma}_1$ )	Difference ( $\hat{\gamma}_4 - \hat{\gamma}_1$ )
Profits	171.739	-57.724	226.847
	(31.203,302.328)	(-155.692,49.026)	(49.537,409.779)
	[0.030]	[0.591]	[0.028]
Consumption	10.934	-61.296	76.580
	(-30.591,56.462)	(-124.001,7.433)	(-8.028,158.341)
	[0.853]	[0.162]	[0.156]
<b>Panel B: Mongolia</b>			
	25 % most ( $\hat{\gamma}_4$ )	25 % least ( $\hat{\gamma}_1$ )	Difference ( $\hat{\gamma}_4 - \hat{\gamma}_1$ )
Profits	1.609	-4.442	6.561
	(-2.746,6.115)	(-8.798,-0.539)	(1.387,12.048)
	[0.840]	[0.049]	[0.039]
Consumption	328.517	-61.879	392.856
	(-174.687,853.155)	(-321.696,162.486)	(-281.516,1051.871)
	[0.371]	[1.000]	[0.554]
<b>Panel C: Bosnia</b>			
	25 % most ( $\hat{\gamma}_4$ )	25 % least ( $\hat{\gamma}_1$ )	Difference ( $\hat{\gamma}_4 - \hat{\gamma}_1$ )
Profits	107.882	47.919	50.830
	(-184.485,372.267)	(-240.218,328.258)	(-365.110,431.934)
	[0.824]	[1.000]	[1.000]
Consumption	5.851	-60.888	66.265
	(-244.900,282.417)	(-351.425,232.547)	(-358.288,459.754)
	[1.000]	[1.000]	[1.000]

Notes: The table presents estimates of the Sorted Group Average Treatment Effects (GATES) for monthly profits and monthly consumption measured in USD (PPP) indexed to 2009. Panel A, B and C shows the estimates for the RCTs in Morocco, Mongolia and Bosnia & Herzegovina respectively. The final estimates, confidence intervals and  $p$ -values are found by estimating (3.13) by WLS over 50 splits and then taking the median of the split-dependent estimates. Clustered standard errors at the village level for Morocco and Mongolia, and robust standard errors for Bosnia & Herzegovina are used. Strata dummies at the village level for Morocco and at the province level for Mongolia are included. The two non-overlapping groups G1 and G4 are defined by the proxy  $S(Z)$  which is estimated using an Elastic Net. G4 consists of the 25 pct. most affected and G1 of the 25 pct. least affected. The difference is found by testing whether the group estimates are equal. The final estimate, confidence interval and  $p$ -value for this quantity are found in the same way as the estimates of  $\gamma_1$  and  $\gamma_4$ .

- 90% confidence intervals in parentheses.
- $p$ -values for the hypothesis that the parameter is equal to zero in brackets.

## Household Characteristics or Aggregate Factors

Before turning to the CLAN estimates it is of interest to analyse the relative role of sociodemographic factors in explaining the heterogeneity discovered so far. To do so, we analyse whether the heterogeneity is primarily explained by household level characteristics, or by aggregate or village level factors. The villages considered in both RCTs differ on a number of parameters including location, population density, infrastructure, access to education, and many more. For the RCT in Mongolia a few village characteristics are available (see Table 6). However, for the RCT in Morocco no village characteristics are available. To overcome this, we use village/province-level strata dummies as proxy variables for aggregate-level characteristics. These dummies also capture other factors such as the efficacy of the MFIs branch managers.

To quantify the relative importance of the household and aggregate level covariates in predicting heterogeneity in treatment effects we report the adjusted  $R^2$  from regressing a dummy that equals 1 if an individual belongs to the most affected and 0 if an individual belongs to the least affected, defined from the quartiles of the proxy  $S(Z)$ , on aggregate level covariates (including strata dummies), household level covariates, and all covariates, respectively. From Table 9 it is evident that household covariates explain much of the variation in the heterogeneity in both RCTs as household level covariates explain 81 (49) percent of the variation in Mongolia (Morocco). In comparison, aggregate level factors such as manager quality, spillover effects and general equilibrium effects explain 67 and 87 percent of the variation in Mongolia and Morocco, respectively. Even though it is not clear-cut which factors that are most important, it seems as if household covariates are informative in explaining the heterogeneity in treatment effects in both RCTs, and hence the forthcoming CLAN analysis is indeed of relevance.

### 4.3 Group Characteristics

As there are no sign of heterogeneity in treatment effects on profits in Bosnia & Herzegovina, nor for consumption across all three RCTs we limit the CLAN to monthly profits in Morocco and Mongolia, and Table 4 presents the estimates. In particular, we report the CLAN estimates of the five baseline covariates that have the highest correlation with the proxy  $S(Z)$  as these seem to be the most important in *predicting* heterogeneity.<sup>18</sup> This is done to avoid any concerns about the choice of covariates presented in the CLAN being arbitrary and thereby the possibility of searching over variables. We choose to limit the CLAN to five variables for each RCT to reduce the number of hypothesis tests.

In Panel A and Panel B of Table 4, we observe that the most important predictors of heterogeneity relates to self-employment activities, consumption levels, and indebtedness across both RCTs. In particular, the most affected across the RCTs seem to be less likely to be engaged in loan activities prior to the trials compared to the least affected. In Morocco 1 out of 6 of the most affected had outstanding loans compared to just shy of 1 out of 3 of the least affected, and 1 out of 5 of the people in the control

---

<sup>18</sup>Due to regularization in the machine learning algorithms used it is not possible to infer causal relations from the covariates. Regularization focus on some covariates and not others based on correlation and it is hard to infer whether the actual choice cause or correlates with the outcome Athey (2016).



group. Similarly, in Mongolia, the most affected had a current debt that was 857.5 USD (PPP) lower than the least affected at baseline. This difference may result from the most affected being less likely to borrow (as in Morocco) or because they borrow smaller amounts. In any case, the magnitude of the loan activities of the most affected are significantly lower compared to the least affected in Mongolia. It is important, however, to point out that the most affected are still slightly more indebted than the control group. As no MFI were operating in any of the villages in either country prior to the trials the difference in borrowings between the most and least affected indicates that individuals belonging to the former group were less likely to engage in loan activities other than microcredit. Such other activities are primarily made up of informal loans from local money-lenders or relatives but other sources are also recorded such as outstanding loans from utility companies.

Across the RCTs a number of self-employment activity related variables are important predictors of the heterogeneity in treatment effects on profits. The fact that previous self-employment activities are related to the treatment effects is consistent with the findings of Meager (2019), previous business owners seem to gain the most from access to microcredit across RCTs, and Banerjee et al. (2015b) who finds that long-term effects on business scale and performance also depend on previous business ownership. The CLAN, however, allows for a more detailed description of the factors related to self-employment that seem to be important in predicting heterogeneity. In particular, the most affected in Morocco are business owners that are more likely to be women (8.3 pp.), less likely to be engaged in animal husbandry (18.4 pp.) and more likely to be running a business that is not related to agriculture (4.9 pp.) compared to the least affected.<sup>19</sup> In the case of Mongolia although the most affected are more likely to be women with prior business experience (16.9 pp.), they are also more likely to receive income from agricultural activities (12.3 pp.).<sup>20</sup> Moreover, the most affected had significantly smaller businesses as measured by monthly business revenues at baseline as the revenues are 634 USD (PPP) which is almost 730 USD (PPP) lower than those of the least affected, and 143 USD (PPP) lower than the control group. This latter finding could indicate that it is small business owners that are able to expand their businesses when given access to microcredit that gain the most in Mongolia.

In summary, it seems as if previous self-employment activities are important in predicting winners and losers across RCTs. Apart from women with previous business experience, the exact nature of the self-employment activities are, on the other hand, harder to generalize across the RCTs. If previous self-employment activities work as moderators for the causal effect of increased access to microcredit then it is not surprising to find differences in the importance of the specific self-employment activities as these may operate with different strengths and effectiveness across contexts Deaton and Cartwright (2018). It may also be the case that some common characteristics are missed when considering only the five most effective predictors of heterogeneity.

Finally, it seems as if the most affected across both RCTs have lower monthly consumption compared

---

<sup>19</sup>These findings also hold, although to a lesser extent, when compared to the control group.

<sup>20</sup>Recall that the study in Mongolia focus on women. Households with prior business experience are therefore necessarily female. It is, however, uncertain with the data at hand whether women would outperform men.

to the least affected. In Morocco, the most affected consume 371 USD (PPP) on average each month which is 123 USD (PPP) lower than the least affected, and slightly lower than the control group (see Table 5). In Mongolia, the point estimate of the difference between most and least affected is -44 USD (PPP) but this is insignificant at conventional significance levels. Total monthly consumption therefore seems to be an important predictor of heterogeneity, and there are weak indications towards the most affected having lower consumption levels across settings. Lower consumption, if driven by lower consumption of temptation goods, could imply that households are more patient or disciplined in their spending. Indeed, the lower consumption of the most affected in Morocco is driven by significantly lower consumption of non-durables (not shown in Table 3.15). However, this finding does not apply for the most affected households in Mongolia.

**TABLE 4**  
Classification Analysis (CLAN) for profits across RCTs

	<b>Panel A: Morocco</b>		
	25 % most ( $\hat{\delta}_4$ )	25 % least ( $\hat{\delta}_1$ )	Difference ( $\hat{\delta}_4 - \hat{\delta}_1$ )
Total monthly consumption of household (USD)	371.250 (343.631,398.706)	490.974 (463.194,518.755)	-123.132 (-162.309,-84.504)
Self-emp. activity run by women	0.174 (0.148,0.200)	0.089 (0.063,0.114)	0.083 (0.047,0.120)
Household has animal husbandry as self-emp. activity	0.361 (0.323,0.398)	0.561 (0.524,0.598)	-0.184 (-0.237,-0.131)
Household has non-agriculture as self-emp. activity	0.172 (0.145,0.198)	0.129 (0.101,0.158)	0.049 (0.010,0.088)
Household has any outstanding loans	0.167 (0.135,0.200)	0.294 (0.262,0.326)	-0.130 (-0.174,-0.087)
	<b>Panel B: Mongolia</b>		
	25 % most ( $\hat{\delta}_4$ )	25 % least ( $\hat{\delta}_1$ )	Difference ( $\hat{\delta}_4 - \hat{\delta}_1$ )
Current debt of household (USD)	519.697 (152.503,873.556)	1393.255 (1039.424,1738.899)	-857.551 (-1338.041,-360.266)
Respondent receives income from agricultural activities	0.178 (0.123,0.233)	0.059 (0.006,0.112)	0.123 (0.037,0.200)
Household has a business activity	0.644 (0.556,0.732)	0.492 (0.404,0.580)	0.169 (0.044,0.294)
Total monthly revenues of respondents business (USD)	634.100 (161.356,1094.912)	1307.063 (870.046,1748.833)	-729.946 (-1393.561,-77.445)
Total monthly consumption of household (USD)	564.639 (462.490,751.718)	614.079 (501.044,765.852)	-44.380 (-182.971,87.231)

Notes: The RCT estimates are found by estimating equation (3.15) over 50 splits and taking the median of the estimates. In particular, the estimates for the most and least affected are found by regressing the given (baseline) covariate on two indicators using the main samples. The first equals 1 if the individual belongs to the top 25 pct. who experienced the largest treatment effect and the second equals 1 if an individual belong to the bottom 25 pct who experienced the lowest treatment effect as measured by the Elastic Net Proxy  $S(Z)$ . The estimate of the difference in treatment effects is found by testing whether the coefficients of the before mentioned indicators are equal to each other and taking the median. All the covariates presented in the table were recorded at baseline. Panel A and B present the estimates for the RCTs Morocco and Mongolia, respectively. The covariates presented for each panel are the five variables that have the highest correlation with the Elastic Net proxy predictor  $S(Z)$ . All monetary measures are in USD (PPP) indexed to 2009.

- 90% confidence intervals in parentheses.

#### 4.4 Robustness: Alternative Machine Learning Algorithm

As a measure of robustness we consider the estimates obtained from using a Random Forest to estimate the proxies  $B(Z)$  and  $S(Z)$ . Given that in some cases the Random Forest algorithm outperformed Elastic Net (Table 8), it is worth analysing whether the estimates are robust to a change of method. Tables 10 and 11 present the BLP and GATES estimates obtained from using proxies  $B(Z)$  and  $S(Z)$  from a Random Forest. Crucially, any sign of heterogeneity discovered using the Elastic Net disappears when the proxies are estimated using a Random Forest. The  $\hat{\beta}_2$  estimates for consumption and profits from (3.11) are insignificant across all RCTs. Moreover, and in line with this finding, there are no significant difference between the most and least affected for profits and consumption across all RCTs. In addition, cases where Random Forest appear to perform better than the Elastic Net give no indication of heterogeneity in treatment effects.

In addition, cases where Random Forest appear to perform better than the Elastic Net give no indication of heterogeneity in treatment effects.

It is important to underline that in this study our results depend crucially on the machine learning algorithm used for estimating the proxies  $B(Z)$  and  $S(Z)$ . In this context, the measures to determine the better performing machine learning algorithm developed by Chernozhukov et al. (2019), and discussed in Section 3.4.1, provide a disciplined way for the analyst to counter the uncertainty as to which algorithm to deploy. It does, however, not rule out that alternative and better performing methods were implemented but disregarded due to null-findings.

One possibility is to introduce pre-analysis plans that specify the machine learning algorithms to be applied. Therefore, rather than substituting for pre-analysis plans, the method developed by Chernozhukov et al. (2019) may in fact be a complement to them, as it limits the restrictions of the pre-analysis plan to more general factors, such as a list of the machine learning algorithms to be used. These machine learning algorithms can in turn be chosen to cover each others strengths and weaknesses which is often known in advance to some degree.

## 5 Concluding Remarks

Microcredit has been both praised and vilified as a development tool. Understanding heterogeneity in the effects of microcredit is therefore key in informing policy makers whether microcredit is worth allocating resources towards as a means to alleviate poverty. This paper has sought to uncover heterogeneous impacts by presenting evidence from three very different RCTs using modern machine learning methods.

Across the RCTs we find no evidence of distributional impacts on consumption from increased access to microcredit. Coupling this with zero average effects across all settings therefore points towards microcredit having no effect on household well-being through increases or decreases in consumption. On the contrary the absence of average effect on profits from increased access to microcredit do seem to mask the presence of heterogeneous effects. Across two of the three RCTs we find evidence of heterogeneous effects on profits which present itself as groups of winners and losers thereby supporting both the negative and positive claims of microcredit.

We find that the heterogeneity in treatment effects can be predicted from household level covariates and that the most affected households seem to be less indebted, engage in self-employment activities and consume less as baseline. These findings could be of use in analysing what the effects of microcredit may be in future settings in which the targeted population has different characteristics.

We observe that the findings of heterogeneous effects are not robust to the choice of machine learning method. This raises the possibility of researchers searching over algorithms to validate specific preconceived notions. One way to avoid such concerns is to pre-register the analysis, including the machine learning algorithms to be applied in conjunction with the method developed by Chernozhukov et al. (2019). This would limit concerns of model dependence while still offering a disciplined way for researchers to leverage analyses of heterogeneous treatment effects through the use of multiple machine learning methods.

## References

n.d..

- Abadie, A., et al.** 2018. “Endogenous stratification in randomized experiments.” The Review of Economics and Statistics, 100(4): 567–580.
- Angelucci, M., et al.** 2015. “Microcredit Impacts: Evidence from a Randomized Microcredit Program Placement Experiment by Compartamos Banco.” American Economic Review: Applied Economics, 7(1).
- Athey, S.** 2016. “Susan Athey on Machine Learning and Econometrics.” EconTalk Episode no. 542. <https://www.econlib.org/econtalk-archives-by-date/>.
- Athey, S., and G. Imbens.** 2018. “Estimation and Inference of Heterogenous Treatment Effects Using Random Forests.” Journal of the American Statistical Association, 113(523): 1228–1242.
- Athey, S., and Guido Imbens.** 2017. “Chapter 3 - The Econometrics of Randomized Experiments.” Handbook of Economic Field Experiments, 1: 73–140.
- Attanasio, O., Augsburg Britta, Ralph De Haas, Emla Fitzsimons, and Heike Harmgart.** 2015. “The Impacts of Microfinance: Evidence from Joint-Liability Lending in Mongolia.” American Economic Journal: Applied Economics, 7(1): 90–122.
- Augsburg, B., Ralph de Haas, Heike Harmgart, and Costas Meghir.** 2015. “The Impacts of Microcredit: Evidence from Bosnia and Herzegovina.” American Economic Journal: Applied Economics, 7(1): 183–203.
- Banerjee, A., E. Duflo, R. Glennerster, and C. Kinnan.** 2015a. “The Miracle of Microfinance? Evidence from a Randomized Evaluation.” American Economic Journal: Applied Economics, 7(1): 22–53.
- Banerjee, A., et al.** 2015b. “Do credit constraints limit entrepreneurship? Heterogeneity in returns to microfinance.”
- Banerjee, A., et al.** 2015c. “Six Randomized Evaluations of Microcredit: Introduction and Further Steps.” American Economic Journal: Applied Economics, 7(1): 1–21.
- Bateman, M.** 2019. “Impacts of the Microcredit Model. Does reflect actual practice.” Ch. 3 in The Rise and Fall of Global Microcredit: Development, Debt and Disillusion.
- Bejamini, Y., and Y. Hochberg.** 1995. “Controlling for False Discovery Rate: a Practical and Powerful Approach to Multiple Testing.” J.R. Statist. Soc. B, 57(1).
- Belloni, A., et al.** 2014. “High-Dimensional Methods and Inference on Structural and Treatment Effects.” Journal of Economics Perspectives, 28(2): 29–50.

- Bonferroni, C. E.** 1935. “Il calcolo delle assicurazioni su gruppi di teste.” Studi in Onore del Professore Salvatore Ortu Carboni.
- Breiman, L.** 2001. “Random Forests.” Machine Learning, 45(1): 5–32.
- Breiman, L., J. Friedman, C.J. Stone, and R.A. Olshen.** 1984. “Classification and Regression Trees.” Chapman and Hall/CRC.
- Brownlee, J.** 2016. “Master Machine Learning Algorithms: Discover How They Work and Implement Them From Scratch.”
- Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins.** 2017. “Double/debiased machine learning for treatment and structural parameters.” The Econometrics Journal.
- Chernozhukov, V., et al.** 2019. “Generic Machine Learning Inference on Heterogenous Treatment Effects in Randomized Experiments.” arXiv:1712.04802v4.
- Crépon, B., Florencia Devoto, Esther Duflo, and William Parienté.** 2015. “Estimating the Impact of Microcredit on Those Who Take It Up: Evidence from a Randomized Experiment in Morocco.” American Economic Journal: Applied Economics, 7(1): 123–150.
- Deaton, A.** 2019. “Randomization in the tropics revisited: a theme and eleven variations.” Randomized controlled trials in the field of development: a critical perspective, Oxford University Press.
- Deaton, A., and N. Cartwright.** 2018. “Understanding and misunderstanding randomized controlled trials.” Social Sciences & Medicine, 210: 2–21.
- Duflo, E.** 2018. “Machinistas meet randomistas: usefule ML tools for emperical researchers.” NBER summer institute. Slides: <http://users.nber.org/programs/dev/slides/18Duflo.pdf>.
- Duflo, E.** 2019. “Nobel Prize Lecture.” <https://www.nobelprize.org/prizes/economic-sciences/2019/duflo/lecture/>.
- Frisch, R., and F.V. Waugh.** 1933. “Partial Time Regressions as Compared with Individual Trends.” Econometrica, 1(4): 387–401.
- Gelman, A., and J. Hill.** 2006. “Data analysis using regression and multilevel/hierachical models.” Cambridge University Press.
- Hastie, T., et al.** 2009. “The elements of statistical learning: data mining, inference and prediction.” Springer Series in Statistics, 2nd Edition.
- Holland, P. W.** 1986. “Statistics and Causal Inference.” Journal of the American Statistical Association, 81: 945–960.

- Imbens, G., and J. Wooldridge.** 2009. “Recent Developments in Econometrics of Program Evaluation.” Journal of Economic Literature, 47: 5–86.
- Karlan, D., and J. Zinman.** 2012. “Microcredit Theory and Practice: Using Randomized Credit Scoring for Impact Evaluation.” Science, 332.
- Kuhn, M., et al.** 2020. “Package ‘caret’.”
- Lovell, M.C.** 1963. “Seasonal Adjustment of Economic Time Series and Multiple Regression Analysis.” JASA, 58(304): 993–1010.
- Meager., R.** 2017a. “Aggregating Distributional Treatment Effects: A Bayesian Hierarchical Analysis of the Microcredit Literature.”
- Meager, R.** 2017b. “Improved Specification: Pre-registration and Pre-Analysis Plans.” BITSS. Rachael Meagers talk at Berkeley Initiative for Transparency in the Social Sciences (BITSS), <https://www.youtube.com/watch?v=gQ9---37WGU>.
- Meager, R.** 2019. “Understanding the Average Impact of Microcredit Expansions: A Bayesian Hierarchical Analysis of Seven Randomized Experiments.” American Economic Journal: Applied Economics, 11(1): 57–91.
- Montgomery, J.M., et al.** 2018. “How Conditioning on Posttreatment Variables Can Ruin Your Experiment and What to Do about It.” American Journal of Political Science (AJPS), 62(3): 760–775.
- Neyman, J. S.** 1923. “On the Application of Probability Theory to Agriculture Experiments. Essay on Principles. Section 9.” Annals of Agricultural Sciences. Translated and edited by D.M. Dabrowska and T.P. Speed from the Polish original, which appeared in Roczniki Nauk Rolniczych Tom X (1923) 1-51 (Annals of Agricultural Sciences).
- Pritchett, L., and J. Sandefur.** 2015. “Learning from Experiments when Context Matters.” American Economic Review: Papers & Proceedings, 5(105).
- Rosenbaum, P.R., and D.B. Rubin.** 1983. “The Central Role of the Propensity Score in Observational Studies for Causal Effects.” Biometrika, 70: 41–55.
- Rubin, D. B.** 1974. “Estimating causal effects of treatments in randomized and nonrandomized studies.” Journal of Educational Psychology, 66(5): 688–701.
- Tarozzi, A., et al.** 2015. “The Impacts of Microcredit: Evidence from Ethiopia.” American Economic Journal: Applied Economics, 7(27): 54–89.
- Wallach, J.D., et al.** 2017. “Evaluation of Evidence of Statistical Support and Corroboration of Subgroup Claims in Randomized Clinical Trials.” JAMA Intern Med, 7(1): 1–21.
- World Bank Group.** 2018. “Piecing Together the Poverty Puzzle.” Poverty and Shared Prosperity.



**Yunus, M.** 2006. “Muhammed Yunus - Nobel Lecture Speech.” Nobel Peace Prize, Oslo. <https://www.nobelprize.org/prizes/peace/2006/yunus/lecture/>.

**Zou, H., and T. Hastie.** 2005. “Regularization and variable selection via the elastic net.” Royal Statistical Society, 67(2): 301–320.

TABLE 5

Baseline covariates used for detecting heterogeneity for the RCT in Morocco

	Total obs.	Control group		Treatment - Control		
		Obs.	Mean	SD	Coefficient	<i>p</i> -value
<i>Household Composition</i>						
Members resident	5329	2687	3.837	3.216	-0.049	0.822
Number of adults (age 16 or older)	5329	2687	2.560	2.261	-0.018	0.900
Head age	5329	2687	35.470	25.008	-0.029	0.987
Male head	5329	2687	0.697	0.460	-0.011	0.713
Head with no education	5329	2687	0.463	0.499	-0.019	0.518
<i>Access to Credit</i>						
Has borrowed from any source	5329	2687	0.193	0.395	0.025	0.508
Total amount of outstanding loans	5329	2687	\$195.055	\$678.525	\$10.776	0.768
<i>Self-employment activities</i>						
Has self-employment activity	5329	2687	0.578	0.494	-0.027	0.454
Animal husbandry self-emp. activity	5329	2687	0.390	0.488	0.025	0.458
Agricultural self-emp. activity	5329	2687	0.438	0.496	0.009	0.793
Non-Agricultural self-emp. activity	5329	2687	0.156	0.363	-0.035	0.073
Number of activities managed by women	5329	2687	0.116	0.321	0.002	0.936
<i>All businesses</i>						
Total profits from self-emp. (past 12 months)	5329	2687	\$1189.550	\$8282.596	-\$207.538	0.560
Total expenses from self-emp. (past 12 months)	5329	2687	\$2390.893	\$7288.378	-\$293.282	0.281
Total output from self-emp. (past 12 months)	5329	2687	\$3580.443	\$10491.722	-\$500.820	0.292
Total current stock of assets related to self-emp.	5329	2687	\$2064.711	\$4121.195	\$98.681	0.669
<i>Farms</i>						
Investment (past 12 months)	5329	2687	\$1.908	\$12.715	\$0.146	0.795
Sales (past 12 months)	5329	2687	\$1363.546	\$6468.694	\$79.829	0.803
Expenses (past 12 months)	5329	2687	\$518.910	\$1499.275	\$84.112	0.331
Savings	5329	2687	\$202.106	\$651.827	\$4.702	0.923
Days worked (past 12 months)	5329	2687	44.339	90.157	4.157	0.535
<i>Animal Husbandry</i>						
Investment (past 12 months)	5329	2687	\$61.119	\$337.831	\$1.421	0.902
Sales (past 12 months)	5329	2687	\$545.986	\$1647.877	\$33.509	0.692
Expenses (past 12 months)	5329	2687	\$636.300	\$1980.628	\$78.348	0.424
Savings	5329	2687	\$1700.391	\$3370.306	\$141.433	0.479
Days worked (past 12 months)	5329	2687	81.153	142.850	4.595	0.633
<i>Consumption</i>						
Total monthly consumption	5329	2687	\$383.800	\$342.228	-\$1.810	0.935
Monthly expenditure on durables	5329	2687	\$7.242	\$43.819	\$0.603	0.733
Monthly expenditure on non-durables	5329	2687	\$376.558	\$331.803	-\$2.413	0.910
Household per capita cons. is among lowest 25 pct.	5329	2687	0.193	0.395	-0.007	0.774
<i>Other</i>						
Total income from any source (past 12 months)	5329	2687	\$3598.951	\$9623.042	\$248.820	0.562
Household owns land	5329	2687	0.405	0.491	0.010	0.771
Household rents land	5329	2687	0.063	0.243	-0.003	0.783
Distance in km from home to marketplace	5329	2687	8.569	9.457	-0.780	0.447

Note: All variables are recorded at baseline (pre-treatment) and all monetary variables are measured in USD (\$) PP indexed to 2009. In addition to the household characteristics presented in the table 81 village-pair dummies are included to capture effects at the aggregate level. Following the authors of the original study we derive the coefficient and *p*-value by regressing the respective baseline covariate on the treatment dummy and using standard errors clustered at the village level.

**TABLE 6**

Baseline covariates used for detecting heterogeneity for the RCT in Mongolia

	Total obs.	Control group		Treatment - Control		
		Obs.	Mean	SD	Coefficient	<i>p</i> -value
<i>Household Composition</i>						
Members resident	960	259	4.908	1.739	0.001	0.995
Number of adults (age 16 or older)	960	259	3.162	1.532	-0.007	0.970
Number of female adults (age 16 or older)	960	259	1.838	1.025	-0.075	0.438
Number of children (age 15 or younger)	960	259	1.757	1.257	0.006	0.961
Age of respondent	960	259	40.919	9.358	-1.304	0.183
Respondent with no education	960	259	0.849	0.358	-0.028	0.439
Respondent opted for vocational training	960	259	0.216	0.412	-0.021	0.554
Respondent married or living with partner	960	259	0.591	0.493	0.090	0.229
Respondent is Buddhist	960	259	0.757	0.430	-0.028	0.570
Respondent has Hahl ethnicity	960	259	0.656	0.476	0.101	0.519
<i>Access to Credit</i>						
Has borrowed from any source	960	259	0.575	0.495	0.089	0.105
Total amount of outstanding loans	960	259	\$473.475	\$891.398	\$428.532	0.002
Total expenditure on instalments last month	960	259	\$61.813	\$125.061	\$24.738	0.051
<i>Self-employment activities</i>						
Any type of enterprise	960	259	0.602	0.490	-0.007	0.920
Respondent has own enterprise	960	259	0.378	0.486	-0.009	0.896
Expenses of any enterprise (past 12 months)	960	259	\$660.035	\$1973.116	\$48.614	0.793
Revenues of respondent's enterprise (past 12 months)	960	259	\$777.497	\$2065.149	\$229.689	0.316
Profit of respondent's enterprise (past 12 months)	960	259	\$243.740	\$1753.716	-\$3.919	0.960
Total yearly hours worked (self and wage employment)	960	259	84.124	79.564	-0.608	0.951
<i>Other employment activities</i>						
Number of income sources	960	259	0.622	0.851	-0.061	0.562
Agriculture as income source	960	259	0.120	0.325	-0.023	0.561
Private business as income source	960	259	0.077	0.267	0.028	0.190
Teaching as income source	960	259	0.093	0.291	0.009	0.759
Mining as income source	960	259	0.012	0.107	0.013	0.276
Any other income source	960	259	0.243	0.430	-0.083	0.094
<i>Consumption and savings</i>						
Total monthly expenditures on consumption	960	259	\$592.442	\$922.301	-\$35.242	0.606
Total monthly expenditures on temptation goods	960	259	\$14.558	\$28.237	\$3.363	0.268
Educational expenses last month	960	259	\$37.831	\$96.361	-\$4.942	0.482
Total household savings	960	259	\$98.767	\$314.582	\$19.314	0.502
<i>Housing</i>						
Household owns a house	960	259	0.390	0.489	-0.078	0.319
Household owns separate land	960	259	0.367	0.483	0.100	0.157
Type of dwelling <sup>a)</sup>	960	259	1.363	0.542	0.066	0.385
Household owns or rents other dwelling	960	259	0.506	0.501	-0.046	0.541
Household owns large household appliances	960	259	0.529	0.500	-0.005	0.929
Household owns small household appliances	960	259	0.776	0.418	0.010	0.810
Distance to province center (km)	960	259	113.321	52.360	-0.985	0.959
<i>Other</i>						
Number of cattle, sheep etc. owned by household	960	259	39.614	52.705	8.582	0.235
Number of other animals owned by household	960	259	0.243	0.703	0.216	0.077
Number of livestock in village	960	259	131180.800	49372.390	35388.340	0.075
Number of people in village	960	259	3609.351	3609.351	283.450	0.464
Number of people in village center	960	259	1008.529	418.624	104.661	0.562
Number of families in village	960	259	1002.915	312.139	101.908	0.418
Number of doctors in village	960	259	5.942	3.578	-1.310	0.311

Note: All variables are recorded at baseline (pre-treatment) and all monetary variables are measured in USD (\$) PPP indexed to 2009. In addition 5 province dummies are included to capture effects at the aggregate level. A dummy indicating whether individuals were assigned group or individual treatment as well as dummies indicating follow-up dates are also included. Following the authors of the original study we derive the coefficient and *p*-value by regressing the respective baseline covariate on the treatment dummy and using standard errors clustered at the village level.

<sup>a)</sup>Categorical variable: 1 = *Ger*, 2 = *House*, 3 = *Apartment* and 4 = *Other*.

**TABLE 7**

Baseline covariates used for detecting heterogeneity for the RCT in Bosnia & Herzegovina

	Total obs.	Control group			Treatment - Control	
		Obs.	Mean	SD	Coefficient	<i>p</i> -value
<i>Respondent Characteristics</i>						
Age of respondent	994	443	37.097	11.954	1.290	0.092
Female respondent	995	444	0.405	0.492	0.008	0.789
Respondent has never married	995	444	0.236	0.425	-0.008	0.772
Respondent is married	995	444	0.619	0.486	-0.008	0.803
Respondent is living with partner	995	444	0.009	0.095	-0.005	0.298
Respondent is separated	995	444	0.059	0.235	0.014	0.372
Respondent is widowed	995	444	0.077	0.266	0.005	0.767
Respondent is employed	995	444	0.568	0.496	-0.007	0.831
Respondent is unemployed	995	444	0.261	0.440	0.006	0.844
Respondent is studying	995	444	0.023	0.149	-0.013	0.098
Respondent is retired	995	444	0.092	0.290	0.004	0.837
Respondent has no education	995	444	0.315	0.465	0.031	0.296
<i>Household Characteristics</i>						
Number of female household members	995	444	1.712	1.007	0.009	0.890
Number of male household members	995	444	1.736	0.962	0.169	0.009
Number of adults (age 16 or older)	995	444	2.543	1.045	0.147	0.032
Number of children (age 15 or younger)	995	444	0.842	1.024	0.045	0.494
Number of household members attending school	995	444	0.723	0.937	0.146	0.017
Number of household members employed	995	444	1.097	0.915	0.072	0.221
Number of household members unemployed	995	444	0.685	0.884	0.021	0.710
Number of household members retired	995	444	0.313	0.515	-0.001	0.979
Number of household female members employed	995	444	0.336	0.518	0.036	0.286
<i>Access to Credit</i>						
Any type of loan	995	444	0.583	0.494	-0.019	0.549
Number of loans	995	444	0.802	0.864	-0.030	0.582
Total amount of outstanding loans	995	444	\$4611.661	\$9625.364	-\$1015.933	0.072
<i>Self-employment activities</i>						
Any income from self-employment activity	995	444	0.770	0.421	0.001	0.969
Respondent owns a business	995	444	0.613	0.488	0.019	0.540
Respondent owns a secondary business	995	444	0.081	0.273	0.022	0.223
Income from self-employment activity	995	444	\$7935.749	\$11952.370	\$269.835	0.745
Income from agricultural activity	995	444	\$413.517	\$1610.300	-\$104.851	0.279
<i>Consumption and savings</i>						
Total consumption of food last week	995	444	\$118.397	\$93.179	-\$0.045	0.994
Respondent eats higher nutritional diet than peers <sup>a)</sup>	995	444	2.374	0.708	0.005	0.906
Savings	995	444	\$1243.773	\$3244.363	\$14.101	0.944
<i>Housing</i>						
Respondent owns primary dwelling	995	444	0.858	0.349	0.035	0.101
<i>Other</i>						
Respondent thinks last year was financially successful <sup>a)</sup>	995	444	2.689	0.614	0.035	0.361
Respondent thinks next year will be financially successful <sup>a)</sup>	995	444	2.872	0.367	0.000	0.983

Note: All variables are recorded at baseline (pre-treatment) and all monetary variables are measured in USD (\$) PPP indexed to 2009. Following the authors of the original study we derive the coefficient and *p*-value by regressing the respective baseline covariate on the treatment dummy and using standard errors robust to heteroskedasticity.

<sup>a)</sup>Categorical variable: 1 = *Disagree*, 2 = *Neutral*, and 3 = *Strongly Agree*

**TABLE 8**  
Performance comparison of the Elastic Net and Random Forest

	Morocco		Mongolia		Bosnia & Herzegovina	
	Best BLP	Best GATES	Best BLP	Best GATES	Best BLP	Best GATES
	$\Lambda$	$\bar{\Lambda}$	$\Lambda$	$\bar{\Lambda}$	$\Lambda$	$\bar{\Lambda}$
<b>Profit</b>						
Elastic Net	105.035*	8680.990*	2.568*	28.927*	58.394	42957.210*
Random Forest	65.109	4793.797	0.936	10.935	68.511*	36946.750
<b>Consumption</b>						
Elastic Net	22.810	1228.565*	225.227*	15365.680	64.229*	9996.229
Random Forest	29.820*	1063.052	194.757	30605.060*	51.640	13882.880*

Notes: The table presents estimates of  $\Lambda$  and  $\bar{\Lambda}$  which measures how well the Elastic Net and Random Forest algorithms fit the Best Linear Predictor (BLP) and GATES functions for profits and consumption, respectively, in a given RCT. The estimates are derived from equation (3.19) and (3.20) and given by the medians over 50 splits. Profits and consumption are monthly profits from business activities and total monthly consumption, both measured in USD PPP indexed to 2009. The asterisk indicates the better performing machine learning method for each category (e.g. for a particular measure (BLP or GATES) for a particular outcome (profits or consumption) in a particular RCT (Morocco, Mongolia or Bosnia & Herzegovina)).

**TABLE 9**  
Predictive Power of Covariates for Treatment Effect Heterogeneity in Profits

	Morocco	Mongolia
Aggregate level covariates	0.87	0.67
Household level covariates	0.49	0.81
All covariates	0.94	0.91

Notes: The table presents the medians over 50 splits of the adjusted  $R^2$  values from an regression of an indicator dummy that equals one if an individual belongs to the top 25 pct. of people who experienced the highest treatment ( $G_4$ ) effects on profits and zero if an individual belongs to the bottom 25 pct. ( $G_1$ ) defined by the Elastic Net proxy predictor  $S(Z)$ . The first row show the adjusted  $R^2$  from the regression that controls for aggregate (village/province) level covariates and strata dummies only. The second row shows the adjusted  $R^2$  from the regression that controls for household level covariates only. Finally, the last row shows the adjust  $R^2$  from the regression controlling for all covariates (i.e. both aggregate and household level). The exact household covariates and aggregate covariates / strata dummies used in the RCTs are available in Table 5 and 6.

**TABLE 10**

ATE and HET estimates for profits and consumption across RCTs (Random Forest)

	Morocco		Mongolia		Bosnia & Herzegovina	
	ATE ( $\hat{\beta}_1$ )	HET ( $\hat{\beta}_2$ )	ATE ( $\hat{\beta}_1$ )	HET ( $\hat{\beta}_2$ )	ATE ( $\hat{\beta}_1$ )	HET ( $\hat{\beta}_2$ )
Profits	27.159	0.135	-1.567	0.204	76.812	0.107
	(-23.107,80.741)	(-0.022,0.305)	(-3.780,0.676)	(-0.213,0.669)	(-70.168,208.003)	(-0.420,0.639)
	[0.581]	[0.186]	[0.305]	[0.632]	[0.615]	[0.926]
Consumption	-10.036	0.062	129.373	0.517	-29.249	0.069
	(-32.451,14.143)	(-0.110,0.405)	(-29.102,278.685)	(-0.056,0.921)	(-180.348,114.931)	(-0.323,0.500)
	[0.694]	[0.451]	[0.224]	[0.153]	[1.000]	[1.000]

Notes: The table presents estimates of the average treatment effects (ATE) and a measure of heterogeneity in treatment effects (HET) for profits and consumption. The estimates are given by the median of the estimates of  $\beta_1$  and  $\beta_2$  obtained by applying WLS to equation (3.11) over 50 splits. Similarly, the sample splitting adjusted confidence intervals and  $p$ -values are obtained by the median over 50 splits with the  $p$ -values being twice this value to correct for sample splitting. Clustered standard errors at the village level for Morocco and Mongolia, and robust standard errors for Bosnia & Herzegovina are used. Strata dummies at the village level for Morocco and at the province level for Mongolia are included. The proxies  $S(Z)$  and  $B(Z)$  that enters equation (3.11) are estimated using an Random Forest on the covariates presented in Tables 5, 6, and 7, for Morocco, Mongolia and Bosnia & Herzegovina respectively. Profits and consumption are monthly profits from business activities and total monthly consumption, both measured in USD PPP indexed to 2009.

- 90% confidence intervals in parentheses.
- P-values for the hypothesis that the parameter is equal to zero in brackets.

**TABLE 11**  
GATES estimates for profits and consumption across RCTs (Random Forest)

<b>Panel A: Morocco</b>			
	25 % most ( $\hat{\gamma}_4$ )	25 % least ( $\hat{\gamma}_1$ )	Difference ( $\hat{\gamma}_4 - \hat{\gamma}_1$ )
Profits	144.754 (11.225,295.095) [0.070]	-6.117 (-113.139,109.946) [1.000]	149.716 (-46.900,335.198) [0.305]
Consumption	0.167 (-54.352,57.458) [0.866]	-47.423 (-114.765,19.014) [0.308]	45.517 (-51.346,146.552) [0.729]
<b>Panel C: Mongolia</b>			
	25 % most ( $\hat{\gamma}_4$ )	25 % least ( $\hat{\gamma}_1$ )	Difference ( $\hat{\gamma}_4 - \hat{\gamma}_1$ )
Profits	-0.200 (-4.486,3.808) [1.000]	-3.284 (-7.839,1.281) [0.323]	3.079 (-2.708,9.138) [0.569]
Consumption	377.988 (-92.411,850.317) [0.239]	75.083 (-101.806,253.215) [0.817]	298.610 (-204.444,811.183) [0.439]
<b>Panel D: Bosnia</b>			
	25 % most ( $\hat{\gamma}_4$ )	25 % least ( $\hat{\gamma}_1$ )	Difference ( $\hat{\gamma}_4 - \hat{\gamma}_1$ )
Profits	133.482 (-158.902,432.404) [0.701]	20.853 (-296.607,298.607) [0.885]	126.804 (-340.747,567.488) [0.887]
Consumption	5.265 (-265.275,291.638) [0.986]	-31.944 (-343.649,298.058) [1.000]	-1.524 (-425.770,447.089) [1.000]

Notes: The table presents estimates of the Sorted Group Average Treatment Effects (GATES) for monthly profits and monthly consumption measured in USD (PPP) indexed to 2009. Panel A, B and C shows the estimates for the RCTs in Morocco, Mongolia and Bosnia & Herzegovina respectively. The final estimates, confidence intervals and  $p$ -values are found by estimating (3.13) by WLS over 50 splits and then taking the median of the split-dependent estimates. Clustered standard errors at the village level for Morocco and Mongolia, and robust standard errors for Bosnia & Herzegovina are used. Strata dummies at the village level for Morocco and at the province level for Mongolia are included. The two non-overlapping groups G1 and G4 are defined by the proxy  $S(Z)$  which is estimated using a Random Forest. G4 consists of the 25 pct. most affected and G1 of the 25 pct. least affected. The difference is found by testing whether the group estimates are equal. The final estimate, confidence interval and  $p$ -value for this quantity are found in the same way as the estimates of  $\gamma_1$  and  $\gamma_4$ .

- 90% confidence intervals in parentheses.
- $p$ -values for the hypothesis that the parameter is equal to zero in brackets.