Evaluating the Distributive Effects of a Development Intervention

Pushkar Maitra, Sandip Mitra,
Dilip Mookherjee and Sujata Visaria

HKUST CEP Working Paper No. 2021-06

April, 2021

HKUST CEP working papers are distributed for discussion and comment purposes. The views expressed in these papers are those of the authors and do not necessarily represent the views of HKUST CEP.

More HKUST CEP working papers are available at:
http://cep.hkust.edu.hk/WP

HKUST Center for Economic Policy
The Hong Kong University of Science and Technology
Clear Water Bay, Kowloon, Hong Kong
T: 3469 3007 E:cep@ust.hk W: cep.hkust.edu.hk

Evaluating the Distributive Effects of a Development Intervention

Pushkar Maitra, Sandip Mitra, Dilip Mookherjee and Sujata Visaria

HKUST CEP Working Paper No. 2021-06
April, 2021

Abstract

Most analyses of randomized controlled trials of development interventions estimate an average treatment effect. However, the aggregate impact on welfare also depends on distributional effects. We propose a simple approach to evaluate efficiency-equity trade-offs, that follow the utilitarian tradition of Atkinson (1970). The method does not impose additional assumptions or data requirements beyond those needed to estimate the average treatment effect. We illustrate the approach using data from a credit delivery experiment we implemented in West Bengal, India.

Authors' contact information

Pushkar Maitra

Department of Economics, Monash University, Clayton Campus, VIC 3800, Australia
Email: pushkar.maitra@monash.edu.au

Sandip Mitra

Sampling and Official Statistics Unit, Indian Statistical Institute, 203 B.T. Road, Kolkata 700108, India
Email: Sandip@isical.ac.in

Dilip Mookherjee

Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215, USA
Email: dilipm@bu.edu

Sujata Visaria

Department of Economics, Lee Shau Kee Business Building, Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong
Email: svisaria@ust.hk

Evaluating the Distributive Effects of a Development Intervention*

Pushkar Maitra† Sandip Mitra‡ Dilip Mookherjee§ Sujata Visaria¶

April 2021

Abstract

Most analyses of randomized controlled trials of development interventions estimate an average treatment effect. However, the aggregate impact on welfare also depends on distributional effects. We propose a simple approach to evaluate efficiency-equity trade-offs, that follow the utilitarian tradition of Atkinson (1970). The method does not impose additional assumptions or data requirements beyond those needed to estimate the average treatment effect. We illustrate the approach using data from a credit delivery experiment we implemented in West Bengal, India.

Key words: Distributive Impacts, Program Evaluation, Agricultural Finance

JEL Codes: D82, O16, C93, H21

*Funding was provided by the Australian Agency for International Development (CF09/650), the International Growth Centre (1-VRA-VINC-VXXXX-89120), United States Agency for International Development (AID-OAA-F-13-00007) and the Hong Kong Research Grants Council (GRF16503014). We are grateful to Shree Sanchari for collaborating on the project. Jingyan Gao, Arpita Khanna, Clarence Lee, Daijing Lv, Foez Mojumder, Moumita Poddar and Nina Yeung provided excellent research assistance at different stages of the project. Elizabeth Kwok provided exceptional administrative support. We thank Gaurav Datt, Lakshmi Iyer, Christina Jenq, Xun Lu, Farshid Vahid, Diego Vera-Cossio and seminar participants at the Institute for Emerging Market Studies at HKUST, UNSW, the Workshop on The Role of the Private Sector in Development at the University of Sydney, the IEMS-CAG Workshop on Financial Inclusion in Asia and the Italian Summer School in Development Economics held in Prato, Italy for helpful feedback and comments. Internal review board clearance was received from Monash University, Boston University and The Hong Kong University of Science and Technology. The authors are responsible for all errors.

†Pushkar Maitra, Department of Economics, Monash University, Clayton Campus, VIC 3800, Australia. pushkar.maitra@monash.edu.au

‡Sandip Mitra, Sampling and Official Statistics Unit, Indian Statistical Institute, 203 B.T. Road, Kolkata 700108, India. Sandip@isical.ac.in

§Dilip Mookherjee, Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215, USA. dilipm@bu.edu

¶Sujata Visaria, Department of Economics, Lee Shau Kee Business Building, Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong. svisaria@ust.hk

1 Introduction

It is now common to conduct randomized controlled trials to evaluate the impact of development policy interventions. The focus of most such evaluations is the *average treatment effect* (ATE), or the expected change in the beneficiary's outcome due to the intervention. However, the average treatment effect masks the fact that impacts may differ for program participants, with resulting implications for inequality. A more complete evaluation of the program's effects on welfare would not only account for average effects, but also for how equitably the benefits are distributed.

When the current literature examines heterogeneity in treatment effects, it usually follows one of two methods. In the first, beneficiaries are classified into sub-groups according to certain fixed characteristics (e.g., gender) or baseline levels of well-being (e.g., baseline income or wealth), and treatment effects are estimated for each sub-group separately. However this does not provide a quantitative summary measure of the impact on distributive equity. The second method is to estimate the treatment effects on different quantiles of the outcome distribution. For example, one may estimate the effect of the intervention on the median of the outcome variable, or its 25th or 75th percentile. However, quantile treatment effects do not allow straightforward inferences about the distributional impacts of the intervention (see Bedoya et al., 2017, for an overview of the arguments). Crucially, in general the effect on a particular quantile of the outcome distribution is not equivalent to the treatment effect on that quantile of the baseline outcome distribution. For these two to be equivalent, beneficiaries must maintain the same rank in the outcome distribution in both treatment and counterfactual conditions. This is a non-trivial requirement, and there is little to suggest that it is generally satisfied in the data. For example, a treatment effect that decreases in the baseline value of the outcome variable might violate this assumption.

In this paper we propose an alternative approach rooted in the utilitarian tradition of public economics going back to Atkinson (1970), which addresses both problems described above. A well established, although sparse, literature in public economics has similarly used Atkinson welfare functions to evaluate the distributional impacts of taxes, government transfers and price changes, especially when they have general equilibrium effects (see, for example Newbery and Stern, 1987, Hughes, 1987, Newbery, 1995, Coady and Harris, 2004). Our method shows how this approach can be applied to evaluate efficiency and equity impacts even for small-scale randomized interventions.

In the Atkinson approach, social welfare is represented as the sum of the welfare of the individuals in a population, as evaluated by an impartial observer, aid donor or social planner. This incorporates both efficiency and distributive implications. The welfare of an individual is an increasing, concave, iso-elastic function of the individual's wellbeing $U_i = U(y_i) \equiv \frac{y_i^{1-\theta}}{1-\theta}$, where wellbeing y_i is proxied by income or consumption. The welfare function $U(\cdot)$ reflects ethical judgments of the external observer, in the "extended sympathy" approach to social choice theory.[1] Roberts (1980) provides axioms of cardinality and comparability of wellbeing that characterize this class of welfare functions.

For semantic convenience, in what follows we shall refer to the measure of wellbeing y_i as "income". Our specific empirical application also uses income, but the methodology can be applied to consumption as well. The parameter $\theta > 0, \neq 1$ represents the degree of *inequality aversion* incorporated into the welfare function. When $\theta = 0$, the measure reduces to the sum of incomes, thus ignoring income distribution entirely. When $\theta = 1$, welfare $U(y) = \log y$ and marginal welfare weights are inversely proportional to income; so income changes of poorer households receive greater weight. As θ increases, the social welfare function becomes more responsive to equity, by placing greater weight on the wellbeing of worse-off individuals. As θ approaches $+\infty$, the resulting expression for social welfare approaches the Rawlsian maximin criterion $\min_i\{y_i\}$, thereby placing all weight on the welfare of the worst-off individual. Hence by varying the value of θ, the external evaluator can assess how distributive considerations affect the assessment. This requires that we estimate the average impact on a given monotone function of well-being, rather than well-being itself. This is a straightforward exercise that does not require any additional assumptions beyond those used to estimate the standard average treatment effect. Note, in particular, that we do not need to assume rank-invariance. Moreover, by examining the sensitivity of assessed impacts to the selection of θ, we obtain a summary indication of the importance of distributive considerations in the welfare assessment.

Section 2 illustrates the methodology by showing how it can be applied for a general setting with a randomized policy intervention, with differing underlying assumptions about the specific context. We allow a first stage where (a subset of) individuals in treated villages are selected as beneficiaries. This may be the result of a screening procedure or explicit criteria. It may depend on household-specific observable as well as unobservable characteristics, and therefore is not necessarily random. After this, the intervention is offered to a random subset of the selected group. If the research design includes both a treatment arm of villages where the intervention is conducted and a control arm where it is not introduced, then the evaluation is straightforward. We also show how the methodology can be applied in a more parsimonious research design where there is no control arm. In that case, the intervention can be evaluated relative to a hypothetical counterfactual, provided that only a random subset of the selected individuals receive the intervention, and there are no spillovers to untreated individuals.

We apply this methodology to evaluate the distributive impacts of three different micro-credit interventions that we implemented in West Bengal, India, during 2010-13. Two of these inter-

[1] This is in contrast to the notion of a utility function, which determines the household's actual behavior, or represents the household's own subjective sense of wellbeing, incorporating considerations of status or relative income.

ventions involved appointing a commission agent and asking him to select eligible borrowers for individual liability loans. In the Trader-Agent Intermediated Lending (TRAIL) arm, the agent was selected from among private traders operating in the village. In the Gram Panchayat Agent Intermediated Lending (GRAIL) arm he was appointed by the local government. In a third intervention (Group Based Lending or GBL), borrower groups could self-form and apply for joint liability loans. In previous work (Maitra et al., 2017, 2021), we estimated the average treatment effects of these interventions on farm incomes. In this paper we apply the methodology described above to study distributive impacts.

We find that the TRAIL intervention had statistically significant positive welfare impacts across the entire range of θ values that we consider. In other words, any increase in inequality appears to be small enough that even at high levels of inequality-aversion, it does not outweigh the positive efficiency effects. In contrast, the welfare impacts of the GRAIL and GBL schemes are non-significant throughout the range of θ values. The difference in the welfare effects of the TRAIL and GRAIL schemes, and the TRAIL and GBL schemes are both statistically significant at all values of θ. To better understand the underlying mechanisms, we examine if the impacts are heterogeneous across four different landholding classes, and conduct decomposition of ATEs by land category.This shows that the larger welfare impact of the TRAIL scheme is driven by the larger conditional treatment effect on the welfare of the poorest group in the sample, namely the landless households.

The rest of the paper is organized as follows. Section 2 explains the empirical methodology. The rest of the paper is devoted to the particular application. Section 3 describes the interventions in more detail. More details on the data and descriptive statistics are presented in Section 4, while Section 5 presents the welfare ATE estimates for the three schemes and for different values of θ. In Section 6 we show the welfare decomposition of ATEs by land class. Section 7 concludes.

2 Methodology

Consider a population of villages, which are randomized into a treatment group denoted T and a control (or counterfactual) group C. In each village, individuals belong to one of two types $\sigma = s, n$ of individuals in the population, where s is the type that is selected to receive the intervention, and type n is not selected. Since the intervention is randomly assigned across villages, $\Pr(s)$ describes the expected fraction of s types in both T and C villages.

By construction, the intervention is available to none of the individuals in the C villages. It is also unavailable to n type households in T villages. Let $e \in \{0, 1\}$ denote whether a specific individual receives the intervention, and $p \equiv P(t = 1|s, T)$ denote the fraction of s types in a T village that receive the intervention. Hence we have

$$P(e=1|n,T) = P(e=1|\sigma,C) = 0 < p \equiv P(e=1|s,T); \text{ for } \sigma = s,n \tag{1}$$

Let the endline outcome (income or consumption) for an individual be represented by random variables $y(\sigma, e, T)$ and $y(\sigma, e, C)$ in T and C villages respectively. Then the social welfare in treated villages can be written as

$$W(T) = \underbrace{p\,\Pr(s)E[U(y(s,1,T))]}_{\text{selected and treated}} + \underbrace{(1-p)\,\Pr(s)E[U(y(s,0,T))]}_{\text{selected but untreated}} + \underbrace{[1-\Pr(s)]E[U(y(n,0,T))]}_{\text{not selected}} \tag{2}$$

while the social welfare in control (or counterfactual) villages can be written as

$$W(C) = \Pr(s)E[U(y(s,0,C))] + [1-\Pr(s)]E[U(y(n,0,C))] \tag{3}$$

$$= E[U(y(\sigma,0,C))] \tag{4}$$

If the research design includes control villages and data are collected from a random sample of households, then $W(C)$ can be directly estimated.

In treatment villages, we assume income is measured for random samples within each of the three relevant groups: "selected and treated" $(s,1)$, "selected but untreated" $(s,0)$ and "not selected" $(n,0)$.[2] Then $W(T)$ can be estimated; in turn this allows us to directly estimate the welfare impact of the intervention $[W(T) - W(C)]$.

The welfare impact can also be estimated if the research design is more parsimonious, in that there is no control arm, provided that two conditions hold. First, we need $p < 1$, i.e., some selected subjects do not receive the intervention. This implies there is a non-null group of selected but untreated, $(s,0)$. The second condition is that there are no spillovers from treated to untreated subjects, or that the treatment does not affect untreated subjects of either type:

$$y(n,0,T) = y(n,0,C), y(s,0,T) = y(s,0,C) \tag{5}$$

Then observe that if C denotes the counterfactual that would have occurred if the T villages had not received the intervention, then equation (5) implies that the welfare impact of the intervention

[2] This applies even when $p = 1$, or in other words, all selected subjects are treated, in which case there are no "selected but untreated".

equals

$$
\begin{aligned}
W(T) - W(C) &= p\,\Pr(s)\{E[U(y(s,1,T))] - E[U(y(s,0,C))]\} \\
&\quad + (1-p)\,\Pr(s)\{E[U(y(s,0,T))] - E[U(y(s,0,C))]\} \\
&= p\,\Pr(s)\{E[U(y(s,1,T))] - E[U(y(s,0,C))]\} \\
&= p\,\Pr(s)\{E[U(y(s,1,T))] - E[U(y(s,0,T))]\} \qquad (6)
\end{aligned}
$$

The first equality uses the absence of spillovers among the non-selected, while the second and third equalities use it for the selected but untreated. Intuitively, when $p \in (0,1)$, we are able to derive unbiased estimates of the income of both the treated and untreated eligibles in the treated villages. The "no spillover" assumption implies that the income of the selected but untreated equals the income that the selected would have had, if the intervention had not been conducted. Hence the difference between incomes of the treated and selected but untreated within treated villages $E[U(y(s,1,T))] - E[U(y(s,0,T)))]$ is an unbiased estimate of the impact of the intervention on the treated $E[U(y(s,1,T))] - E[U(y(s,0,C)))]$. By scaling this by the proportion of individuals treated $p\Pr(s)$, we obtain an unbiased estimate of the welfare impact relative to no intervention.

We apply this methodology to evaluate the welfare effect of our credit interventions. These interventions were implemented through a parsimonious design randomized controlled trial, involving three different treatment arms and no control arm, and with $p < 1$ for each of the three interventions. At most 5% of the relevant population within each village received the intervention, making it unlikely that there were spillover effects on the population that did not receive the program credit. We apply the methodology described above to assess the welfare impact of each intervention relative to a no-intervention counterfactual. Since the assignment of treatments across villages was randomized, this provides an unbiased estimate of the welfare impact.

3 The Interventions

A non-profit microfinance institution conducted the three agricultural credit interventions in the districts of Hugli and West Medinipur in the state of West Bengal, India. Here we provide a brief summary of the experimental details; these are discussed in greater detail in Maitra et al. (2017, 2021). The schemes were primarily designed to facilitate the cultivation of potatoes, the most profitable cash crop in this region. In each of the three interventions, loan size, duration, interest rate and dynamic repayment incentives were identical. In October 2010, borrowers were offered the first loans of 4-month duration at an annual interest rate of 18 percent. This rate was considerably lower than the 25% per annum average that prevailed in the informal loan market. Repayment was due in a single lumpsum at the end of 4 months. Borrowers who repaid successfully were eligible for a larger loan in the subsequent cycle; those who did not were not allowed to borrow again.

The experiment was designed to compare different approaches to identifying beneficiaries for subsidized agricultural credit. The agent-intermediated lending (AIL) approach taps into the knowledge and information about local residents that exists within a community but might be unobservable to researchers. Borrower selection is delegated to a local intermediary appointed as the MFI's agent. He/she is incentivized through commissions that depend on the interest payments by the borrowers they recommend.

We implemented two different versions of the AIL approach. Twentyfour randomly selected villages were assigned to the TRAIL intervention, where a local trader was appointed as the agent. The field research team randomly selected the agent from a list of local traders who had at least 50 clients, or had been operating in the village for longer than 3 years. Another 24 villages were assigned to the GRAIL intervention, where the agent was a political appointee. The field research team randomly selected one individual from a list recommended by the local government (gram panchayat or GP) of persons who had lived in the village for at least 3 years, were personally familiar with farmers in the village and had a good local reputation.[3] Each agent was asked to recommend 30 borrowers who owned no more than 1.5 acres of land. Ten of these 30 were randomly selected to receive the program loans. At the end of each loan cycle, the agents received a commission equal to 75 percent of the interest paid by borrowers they had recommended.

A third group of 24 villages was assigned to the Group Based Lending (or GBL) intervention. Village residents who owned no more than 1.5 acres of land could form 5-member groups, attend regular group meetings and make savings deposits. At the end of 6 months, all the members of two randomly selected groups were offered the program loans. Group members were jointly liable for each others' loans: if any member defaulted, all other group members were cut off from all future loans. The MFI that organized the group meetings received commissions equaling 75 percent of the interest paid by GBL borrowers.

The three schemes differed along various dimensions. In the two AIL schemes, borrower selection was delegated to local community members. However TRAIL and GRAIL agents had different occupations and therefore might have had different motivations for participating and different criteria for borrower selection and subsequent monitoring or assistance. In the GBL scheme borrowers self-selected into groups. Liability rules also differed: in the AIL schemes borrowers were individually liable for their own loans, whereas GBL borrowers were jointly liable for their

[3] In all TRAIL villages, the first randomly chosen trader approached accepted the contract. In the GRAIL villages, one individual refused to participate for religious reasons; he was replaced by a second randomly chosen individual from the list.

group members' loans.

To enable us to examine the mechanisms behind the outcomes of each scheme, the experiment was designed to separately identify how selected borrowers differed from those not selected (selection effects), and to estimate the effect of the intervention conditional on selection (conditional treatment effects). This required that in each intervention arm, only a subset of selected households received the program loans. Specifically, in the TRAIL and GRAIL arms, in each village loans were offered to only 10 households randomly chosen from the 30 whom the agent had recommended. We refer to these as the TRAIL Treatment and GRAIL Treatment households respectively. In the GBL arm, in each village only 2 of the joint liability groups that had formed were randomly selected, and each member offered the loans. The households of these group members are the GBL Treatment households. Of the 20 recommended TRAIL and GRAIL households that were not randomly assigned to receive the loans, 10 were randomly drawn into the survey sample; we refer to them as Control 1 households. In GBL villages, two of the groups that did not receive the loans were randomly chosen and all of their member households surveyed, these are the GBL Control 1 households.[4] Importantly, there were no pure control villages in the research design, and therefore we use the methodology discussed in Section 2.

Following expression 6, the average conditional treatment effect of each of the schemes can then be estimated as the difference in welfare associated with the farm incomes of the Treatment and Control 1 households in the villages where each intervention was implemented.

4 Data and Descriptive Statistics

We start by describing the characteristics of our sample villages. In Table 1, we report the number of households in these villages as per the 2011 Census of India, and other village characteristics using data from a 2007 pre-intervention village census conducted for a different project (see Mitra et al., 2018).[5] As of 2011, our sample villages had 365 households on average, only 17 percent of which owned more than 1.5 acres of land. Thus 83 percent of the households were eligible to participate in the loan schemes. Slightly less than one-fifth of the households were landless. The different land categories appear in the three treatment arms in similar proportions.

Like many other rural areas in India, nearly 80 percent of the villages in our sample had primary

[4] In addition, 30 non-selected households were included in the sample. We refer to these as Control 2 households. In TRAIL and GRAIL villages, Control 2 households were randomly selected from those that were not recommended by the agent. In GBL villages they were randomly selected from households that did not participate in joint-liability groups. We do not include Control 2 households in our sample when we estimate the conditional treatment effects in the next section.

[5] This survey was conducted in 72 villages. However Maoist violence in 2010 forced us to replace four of the 72 villages from our 2007 sample. Therefore Table 1 uses a sample of only 68 villages.

schools, but a substantially smaller 22 percent had primary health centres. Only about one-third of villages had all-weather roads in 2007. Sixty percent of households reported in 2007 that their houses were electrified. Bank branches existed in only 8 percent of villages, and no MFIs had offices in any of the sample villages; this is consistent with the general unavailability of microcredit in this region at the time that our study began. As the table shows, we can reject the null hypothesis that these village-level characteristics can jointly explain assignment to treatment arm. Consequently, we infer that the villages were balanced on these observables.

From 2010 to 2013 we conducted eight waves of surveys at four-monthly intervals from a sample of 50 households in each of the 72 villages where the loan schemes ran. In each village, the sample includes the 10 Treatment and 10 Control 1 households as defined above.

In Table 2 we present summary statistics from the first wave of household survey data. These statistics pertain to the characteristics of households that owned no more than 1.5 acres of land. As one might expect, in this low-wealth sample, only 42 percent of household heads had studied beyond primary school. The main occupations were agricultural cultivation or casual labour. About 17 percent of households were non-Hindu, and 39 percent belonged to underprivileged social groups.[6] Although nearly three-quarters of sample households reported that their houses were electrified, only one-half owned televisions or other electrical audio-visual appliances. Telephones were more prevalent: nearly 60% of household heads reported that they had either a landline or a mobile telephone. However, housing infrastructure tended to be basic: less than a third of houses were made of brick-and-mortar, and only half had in-house toilets. About one-half of households had a savings bank account.

As columns 2–4 show, the characteristics of households in the three treatment arms are very similar. The pair-wise differences are almost always statistically non-significant (results available upon request). Using a multinomial logit regression, we cannot reject the null hypothesis that on average, observable household characteristics do not explain assignment to treatment arm (p-value=0.998).[7]

Table 2 also shows how the summary statistics vary by land class. Looking across the columns of the table, it is clear that landholding is a good proxy for a household's economic status. Education levels, occupation, landholding and asset ownership all vary substantially across land category. In households that owned more land, heads were more likely to have completed primary

[6]This includes groups that the government lists as scheduled castes, scheduled tribes or "other backward castes".

[7] Since we drew a purposive sample of Treatment, Control 1 and Control 2 households, we do not expect our sample means to be representative of the village populations. To ensure that we estimate representative means, we re-weight the sample to inflate each household in inverse proportion to the probability that they would be selected into the sample. Thus, Treatment and Control 1 households each receive a weight of $\frac{30}{N}$ and Control 2 households receive a weight of $\frac{N-30}{N}$, where N denotes the total number of households in the village, as reported in the 2011 Census. Thus we can scale up the sample proportions in each land category to arrive at the population proportions.

school, and were more (less) likely to report their main occupation as cultivation (casual labour).[8] Households with more land lived in larger houses that were more likely to be brick-and-mortar (*pucca*), have electrical connections, and an in-house toilet. They were also more likely to own televisions or other audio-visual electrical appliances and telephones, and to have bank savings accounts. We find very few land transactions across the three survey years, so that households' land class remains largely fixed over this short time horizon. Thus, it is informative to conduct sub-group analysis across households in different land categories.

Within each land category household characteristics were balanced across treatment arms. In each land class, we cannot reject the null hypothesis that these household characteristics do not predict assignment to treatment arm.

5 Empirical Estimates

5.1 Computing the CTEs

We start by estimating the conditional treatment effects. To do this, we restrict the sample to Treated and Control 1 households in the TRAIL, GRAIL and GBL villages and run the following regression:

$$\begin{aligned} U(y_{ivt}) = \beta_0 + \beta_1 \mathrm{TRAIL}_v + \beta_2 \mathrm{GRAIL}_v + \beta_3 \mathrm{Treatment}_{iv} + \beta_4(\mathrm{TRAIL}_v \times \mathrm{Treatment}_{iv}) \\ + \beta_5(\mathrm{GRAIL}_v \times \mathrm{Treatment}_{iv}) + \xi \mathbf{X}_{ivt} + \epsilon_{ivt} \end{aligned} \tag{7}$$

Here $\mathbf{X}_{ivt}$ is a set of variables measuring household characteristics consisting of landholding, household caste and religion, the age, education and occupation of the oldest male in the household, year dummies and a dummy for the village information treatment.[9]

The dependent variable in the regression is $U(y_{ivt}) = \frac{y_{ivt}^{1-\theta}}{1-\theta}$ corresponding to a specific value of θ. Here y_{ivt} is aggregate farm income for household i in village v in year t, calculated as the sum of value-added earned from the four major crop categories: potatoes, paddy, sesame and vegetables. The explanatory variable TRAIL_v indicates whether village v was assigned to the TRAIL intervention, GRAIL_v indicates whether it was assigned to the GRAIL intervention, and $\mathrm{Treatment}_{iv}$ indicates if the household was assigned to receive the loan. This allows us to estimate

[8] Note, however, that there is an active land rental market, so that even landless households do engage in agriculture.

[9] The information intervention was undertaken for a separate project aimed at examining the effect of providing information about potato prices to farmers and is similar to the public information treatment described in Mitra et al. (2018). Villages were assigned to the information treatment randomly and orthogonally to the credit intervention that is the focus of this paper. The results are unchanged if we do not include this dummy variable in the regression specification.

the conditional treatment effects on household welfare in the TRAIL villages as $\beta_3 + \beta_4$, in the GRAIL villages as $\beta_3 + \beta_5$, and in the GBL villages as β_3.[10]

In Panel A of Table 3, each column presents results from running this regression using a different inequality-aversion parameter θ ranging from 0 to 5. When θ takes value 0, the welfare impact represents the change in average farm income. In line with the results in Maitra et al. (2017), we find in column 1 that the TRAIL scheme increased the average farm income of recommended households by ₹2546. This is significant at the 10% level. The point estimates for both the GRAIL (₹125) and the GBL (₹34) schemes are smaller in magnitude and not statistically different from zero.

At higher values of θ, the farm incomes of low income households receive greater weight in the welfare calculation. When $\theta = 1$ the welfare function is logarithmic, so that the conditional welfare impacts correspond to proportional changes in farm income. Therefore, the same increase in farm income would have a larger impact on welfare if it accrued to a lower-income household than a higher-income household. As we see in Column 2, the TRAIL scheme generated a significant increase in welfare for selected households, even according to this metric. This pattern continues to be repeated as we increase θ to higher values. This suggests that the TRAIL scheme benefited poorer households in particular.

The results in columns 2—6 also suggest that neither the GRAIL nor the GBL schemes changed welfare significantly, even when the welfare function is highly sensitive to inequality: while the point estimates continue to be positive they are never statistically significant. Thus not only is the average effect of these schemes small, there is no evidence to suggest that lower income households benefited either.

5.2 Computing the Treatment Proportion

To estimate the change in aggregate village welfare, in Panel B we multiply the estimated conditional treatment effect on welfare with the proportion of households treated, or $p\Pr(s)$.

Across the 24 TRAIL villages, 6714 households owned no more than 1.5 acres of land and thus were eligible to participate in the TRAIL scheme. As we described above, in each village, the agent was asked to recommend up to 30 households. Of these 30, our sample consists of up to 20 (10 Treatment households and 10 Control 1 households), or a maximum of 480 selected households

[10] In order to estimate the conditional treatment effects (for values of $\theta = 1$ and above), we add 10,00,000 to the farm imputed profits. This helps to ensure that we can take the log (for $\theta = 1$), and for the other values of θ we are not working with very small numbers. The regressions give us point estimates and standard errors for this transformed variable. We transform them back to arrive at estimated treatment effects on welfare. The re-transformation introduces a stochastic element, making it difficult to analytically calculate the standard errors for the point estimates on welfare. Therefore, we present bootstrapped standard errors.

across the 24 TRAIL villages. In some villages, the agent recommended fewer than 30 farmers. Also, a small number of households were subsequently disqualified when our field research team discovered their landholdings were larger than the 1.5 acre threshold. This leaves us with 461 households in TRAIL villages who had been selected by agents. Of these, 227 were randomly selected to receive the program loans. Thus the proportion actually treated in the TRAIL scheme was 3.3 percent. A similar calculation shows that the corresponding proportion in the GRAIL scheme was 3.4 percent. Selection in GBL villages took place differently. First, note that the GBL scheme placed no restriction on the number of groups that could form in a village. However a maximum of 20 households per village, or four 5-member groups per village, for a maximum total of 480 households. In reality, 449 households formed groups in GBL villages, since fewer than 20 households formed groups in some villages. Of these, a random subset of groups were assigned to treatment, so that 229 households received GBL loans. Thus in the GBL scheme the treatment proportion was 5.1 percent, higher than in the TRAIL or GRAIL schemes.

5.3 Change in Aggregate Welfare as a Result of the Interventions

Panel B of Table 3 shows the implied change in aggregate welfare, calculated as the product of the conditional treatment effects as presented in Panel A, and the treatment proportion described above. This measures the change in aggregate welfare that would be expected if the intervention were introduced in a representative village. It is clear once again that the TRAIL scheme increased aggregate welfare at all inequality-aversion levels. However neither the GRAIL nor the GBL schemes had a significant effect on welfare.

5.4 Comparing the Welfare impacts of the Interventions

In Panel C we compare the welfare impacts of the three interventions. We conduct pair-wise Mann-Whitney rank-sum tests on 2000 bootstrap replications of the aggregate welfare effects. Bootstrap samples were drawn using a stratified (by treatment arm) clustered (by village) random procedure, to ensure that each sample contained an equal number of randomly drawn TRAIL, GRAIL and GBL villages. Once a village was drawn into the sample, all original sample households from that village were included. In each bootstrap sample we estimate the treatment proportion for each scheme, and the conditional treatment effects of each scheme. The product is the treatment effect on average welfare.

At $\theta = 0$, the null hypothesis that the GRAIL and TRAIL schemes generate identical aggregate welfare effects is rejected with a p-value < 0.00. Similarly, we can reject the null that the GRAIL scheme had a larger welfare effect at the 5% level when $\theta = 1$, and at the 10% level when $\theta = 2$ and $\theta = 3$. Similarly, we reject the null that the GBL scheme generated a larger aggregate welfare effect than the TRAIL scheme for $\theta = 0, 1, 2$, and the p-values range from 0.14 to 0.20 for $\theta = 3, 4, 5$. We are never able to reject the null hypothesis that the GBL scheme generated a larger aggregate welfare effect than the GRAIL.

The cumulative distribution functions of these estimated changes in welfare for TRAIL and GBL are presented in Figure 1. These corroborate our findings from the rank-sum tests. For every value of θ, the aggregate welfare effects of the TRAIL scheme stochastically dominate those of the GRAIL and GBL schemes. However at low values of θ, there is no clear pattern of stochastic dominance between the GRAIL and GBL schemes. As the value of θ increases, the GBL scheme stochastically dominates the GRAIL scheme, i.e., the GBL scheme is relatively better for low-income households than the GRAIL scheme.

6 Welfare Decomposition

To investigate how these results come about, we decompose them into the welfare effects on different landholding sub-groups. For illustrative purposes, we conduct this decomposition for three values of the inequality aversion parameter, $\theta = 0, 1, 2$. Letting g denote the landholding sub-group, the average village welfare impact of the intervention can be written as a weighted average of the conditional treatment effects on welfare of the different groups:

$$W(T) - W(C) = \sum_g \alpha_g s_g \text{CTE}_g \tag{8}$$

where α_g denotes the demographic weight of group g, and s_g denotes the fraction of group g that was treated. The conditional treatment effect CTE_g equals $\{E[U(y)|s, 1, T, g] - E[U(y_i)|s, 0, T, g]\}$. $E[U(y)|s, 1, T, g]$ and $E[U(y_i)|s, 0, T, g]$ denote average utility among Treatment and Control 1 subjects respectively, within group g in treatment villages.

Observe that variations in treatment proportions across the different groups mainly reflect corresponding variations in selection proportions, since selected members of all groups were equally likely to be chosen for treatment.

Suppose $\theta > 0$. Then, if a household in a low-landholding group g experienced the same (absolute) conditional treatment effect on income as a household in a higher-landholding group g', it would experience a larger conditional treatment effect on welfare. All else remaining the same, an intervention where group g households had a larger chance of being selected than group g' households would increase welfare by more. This is the effect of reallocating selected beneficiaries from group g' to group g. Hence the overall distributive impact of each intervention depends both on variations in CTE_g and $\alpha_g \cdot s_g$. The decomposition represented in equation (8) shows that the overall welfare impact can be expressed as the (population share) weighted average of the product of treatment proportion and CTEs across different groups.

6.1 Population Proportions (α_g) and Treatment Proportions (s_g)

We classify eligible households into four land classes: the landless, those with 0–0.5 acres, 0.5–1 acre and 1–1.5 acres of land. The first row in each panel of Table 4 presents the population proportions for each land category: this is the number of households in that land group among all the households who owned less than 1.5 acres of land. These shares are based on our 2007 pre-intervention houselisting exercises in each sample village.

Figure 2 and the second row of Table 4 present the treatment proportions for each group g. Across the 24 TRAIL villages, agents recommended 5.9 percent of landless households for the TRAIL loans, and since a random subset were assigned to treatment, 3.7 percent of landless households were assigned to the treatment group. By contrast, in GRAIL villages, agents selected only 3.9 percent of landless households, and 1.6 percent were assigned to receive GRAIL loans. In GBL villages, 7.3 percent of landless households self-selected into joint liability groups, and then 3.8 percent received GBL loans. The likelihood of treatment declines steadily with landholding in GBL villages: 2.9 percent of those with 0–0.5 acres, 2 percent of those with 0.5–1 acres and 1.5 percent of those with 1–1.5 acres received GBL loans. In the TRAIL villages the proportions were nearly constant across all four landholding groups between 3.1 and 3.7 percent. In GRAIL villages the corresponding treatment proportions increased with landholding: 1.6 percent of the landless but 5.4 percent of those with 1–1.5 acres of land received the loans. Hence with respect to likelihood of selection and through that the likelihood of treatment, the GBL scheme was the most egalitarian, and the GRAIL scheme the least egalitarian.

6.2 Conditional Treatment Effects on Land Sub-groups

To estimate the conditional treatment effects (CTE_g) on welfare for each land group g we run the following regression:

$$
\begin{aligned}
U(y_{igvt}) = & \sum_{g=1}^{G} \gamma_g (Z_{igv}) + \sum_{g=1}^{G} \delta_g (Z_{igv} \times \text{TRAIL}_v) + \sum_{g=1}^{G} \zeta_g (Z_{igv} \times \text{GRAIL}_v) + \sum_{g=1}^{G} \eta_g (Z_{igv} \times \text{Treatment}_{igv}) \\
& + \sum_{g=1}^{G} \theta_g (Z_{igv} \times \text{TRAIL}_v \times \text{Treatment}_{igv}) + \sum_{g=1}^{G} \kappa_g (Z_{igv} \times \text{GRAIL}_v \times \text{Treatment}_{igv}) \\
& + \lambda \mathbf{X}_{ivt} + \epsilon_{igvt} \qquad (9)
\end{aligned}
$$

where Z_{igv} is an indicator for whether household i in village v belongs to land category g. The sample is restricted to Treatment and Control 1 households. The TRAIL, GRAIL and GBL conditional treatment effects for a household in category g are given by $\delta_g + \theta_g$, $\zeta_g + \kappa_g$ and η_g respectively; $\mathbf{X}_{ivt}$ is as defined earlier.

The estimates are presented in columns 1–4 of Table 4. As before, when $\theta = 0$, we are effectively estimating conditional treatment effects on farm income, measured in rupees. Starting with Panel A, we see that the conditional treatment effects of TRAIL loans vary substantially by land category. Although imprecise, the point estimate in column 1 suggests that among the selected landless households those that were offered the TRAIL loans earned ₹576 more than those who were not. The conditional treatment effects are larger for households with more land, although they are only statistically significant in land class 3. The point estimates are ₹1507 for households with 0–0.5 acres, ₹5118 for those with 0.5–1 acres, and ₹4530 for households with 1–1.5 acres.

However, as we see in the third row of Panel A, baseline farm incomes vary considerably across the four land classes; as a result these absolute changes in farm incomes correspond to very different proportional changes. When we impose $\theta = 1$, the regression in equation (9) is run on the logarithm of farm income, and the proportional increases are larger for households with less land. In particular, since Control 1 landless households earned only ₹361 in farm income, the increase of ₹519 represents a 142 percent increase. In contrast, the ₹1805 increase in farm income for selected 0—0.5 acre households represents a 32 percent increase, for 0.5–1 acre households the proportional increase in farm income is 38 percent, and for the highest land category of 1—1.5 acres, it is 20 percent. While these differences are not estimated precisely enough to be statistically significant, the numbers suggest that the TRAIL scheme benefited lower-income households by more than they benefited higher-income households. This helps explain our previous finding that the TRAIL welfare effects were positive even at high degrees of inequality aversion, which places large weights on the lowest-income households.

Panel B presents the conditional treatment effects by land class for the GRAIL scheme. Although imprecise, the point estimates for landless households and households with 0-0.5 acres of land are positive and in the range of ₹276–857. Households with 0.5 to 1 acre appear to have lost farm income as a result of the GRAIL loan; however the effect is not statistically significant. The point estimate for 1—1.5 acre households is also not significant but is a positive ₹3821.

Finally, the estimates for the GBL scheme are presented in Panel C. The absolute effect for landless households is estimated at ₹679, although it is not statistically significant. It is remarkable that the point estimates for all other land categories are smaller than this, and for the two highest land categories they are imprecisely estimated but negative. Since, once again, the baseline farm income increases monotonically with landholding, it is no surprise that landless households achieved the largest proportional increase in farm income of 49 percent. For households with 0–0.5 acres, the point estimate is 10 percent; for 0.5–1 acres it is smaller than 1 percent, and for the 1–1.5 acres households it is negative 22 percent.

We make a few observations. First, in all three schemes, there is suggestive evidence that the landless benefited disproportionately more than the landed. In the TRAIL scheme, the proportional benefits decreased gradually as landholding increased. In the GRAIL scheme, households with intermediate landholding actually lost income, and none of the other land classes experienced any significant increases. In the GBL scheme the benefits were large only for the landless, and these coefficients are either negligible or negative for households that owned land. Thus there is little evidence to suggest that any of the schemes benefited the richer households more than the poor.

Second, it is apparent that even among landless households, there is considerable difference in mean baseline income in the three schemes. Landless households in the TRAIL villages who were selected but were not assigned the loan (Control 1 households) had average farm income of only ₹362, whereas the corresponding landless households in GRAIL and GBL villages earned more than four times as much. Thus even within the landless group, the TRAIL agent appears to have targeted poorer households than the GRAIL agent did. In addition, even though there was no explicit gatekeeper in GBL preventing particular landless households from participating, it appears that landless households that earned very low incomes were either unable to or unwilling to form joint-liability groups.

This difference in baseline incomes among landless households also helps explain why a relatively small ₹519 increase in farm income translated into a 142 percent increase for landless households in the TRAIL scheme, while a similar ₹679 increase translated into a much smaller 26 percent increase for landless households in the GBL scheme. This is also conveyed in Figure 3, where we see that despite a smaller absolute increase (as seen in the left panel), the TRAIL scheme created the largest proportional conditional increase in welfare for landless households (as see in the right panel). Thus more progressive targeting within the landless class allowed the TRAIL scheme to generate larger increases in welfare than either of the two other schemes, and that this result holds at all degrees of inequality aversion we consider.

6.3 Average (Unconditional) Treatment Effects

The product of a scheme's conditional treatment effect on group g and its selection frequency for group g is the ATE on welfare of group g. As we see in Panel A of Table 4, a representative landless household could expect to earn ₹19.20 (or 5 percent) more per year if the TRAIL scheme was introduced in its village. The expected treatment effects for more landed households were larger in absolute but smaller in proportional terms: Rs 56 (1 percent), Rs 177 (1.3 percent) and 157.70 (0.7 percent) for households with 0—0.5, 0.5—1 and 1—1.5 acres of land respectively. (This is depicted in Figure 4.) In the GRAIL and GBL villages as well, the expected proportional treatment effects on welfare were substantially larger for the landless than for the other land groups.

The analysis is also repeated with the inequality-aversion parameter $\theta = 2$. The welfare of poorer households now receives even larger weight in the welfare calculation. We continue to find that the TRAIL scheme dominates the other two schemes. In the GBL scheme the negative point estimates for expected treatment effects for households with 1-1.5 acres are so large that they cause the weighted average in column 5 to be negative as well, indicating a negative effect on aggregate welfare.

7 Conclusion

Researchers and policy-makers are increasingly using randomized controlled trials to assess the impacts of development interventions, and select those that they will implement and scale-up. However most of these evaluations have focused on the average treatment effects on well-being measured by income, consumption or assets, with limited attention to attendant changes in the distribution of these indicators. This paper has presented a simple approach to a policy-relevant summary welfare assessment of an intervention, which examines both average and distributional impacts. We illustrate our approach using data from three credit delivery interventions implemented in randomly-chosen villages in West Bengal, India.

The approach has several advantages. One, its conceptual underpinnings provide a clear rationale for the welfare evaluations that it generates. Two, unlike other approaches it does not require strong assumptions for a clear-cut interpretation: the welfare effects are identified under the same assumptions by which a randomized controlled trial delivers consistent estimates of the average treatment effect, namely that a random subset of the sample is assigned to treatment. Three, the approach provides a single summary quantitative measure of the welfare impact, allowing unambiguous evaluations. Four, the approach is simple to implement and does not impose any additional data requirements beyond those used to estimate the standard average treatment effects. The method allows for decomposition analysis across population sub-groups, allowing an closer examination of the patterns that drive the aggregate effects. The more complete welfare evaluations that this method generates could be used in the cost-benefit calculations for development interventions.

One qualification to our findings for the specific application considered in this paper is that we ignored financing costs. However, the methodology can be extended in a straightforward manner to include in the analysis additional groups of citizens who might bear the financial cost of the interventions.

References

Atkinson, A. B. (1970). On the Measurement of Inequality. *Journal of Economic Theory*, 2:244 – 263.

Bedoya, G., Bittarello, L., Davis, J., and Mittag, N. (2017). Distributional Impact Analysis: Toolkit and Illustrations of Impacts Beyond the Average Treatment Effect. Policy research working paper 8139, World Bank.

Coady, D. and Harris, R. (2004). Evaluating Transfer Programs Within a General Equilibrium Framework. *Economic Journal*, 114:778 – 799.

Hughes, G. (1987). The Incidence of Fuel Taxes: A Comparative Study of Three Countries. In Newbery, D. and Stern, N., editors, *The Theory of Taxation for Developing Countries*. Oxford University Press.

Maitra, P., Mitra, S., Mookherjee, D., Motta, A., and Visaria, S. (2017). Financing Smallholder Agriculture: An Experiment with Agent-Intermediated Microloans in India. *Journal of Development Economics*, 127:306 – 337.

Maitra, P., Mitra, S., Mookherjee, D., and Visaria, S. (2021). Decentralized Targeting of Agricultural Credit Programs: Private versus Political Intermediaries. Working paper, http://people.bu.edu/dilipm/wkpap/.

Mitra, S., Mookherjee, D., Torero, M., and Visaria, S. (2018). Asymmetric Information and Middleman Margins: An Experiment with West Bengal Potato Farmers. *Review of Economics and Statistics*, 100(1).

Newbery, D. and Stern, N., editors (1987). *The Theory of Taxation for Developing Countries*. World Bank: Oxford University Press.

Newbery, D. M. (1995). The Distributional Impact of Price Changes in Hungary and the United Kingdom. *Economic Journal*, 105(431):847 – 863.

Roberts, K. (1980). Interpersonal Comparability and Social Choice Theory. *Review of Economic Studies*, 47(2):421–439.

Figure 1: Cumulative Distribution Functions of Estimated Changes in Aggregate Welfare for Different Inequality-Aversion Parameters

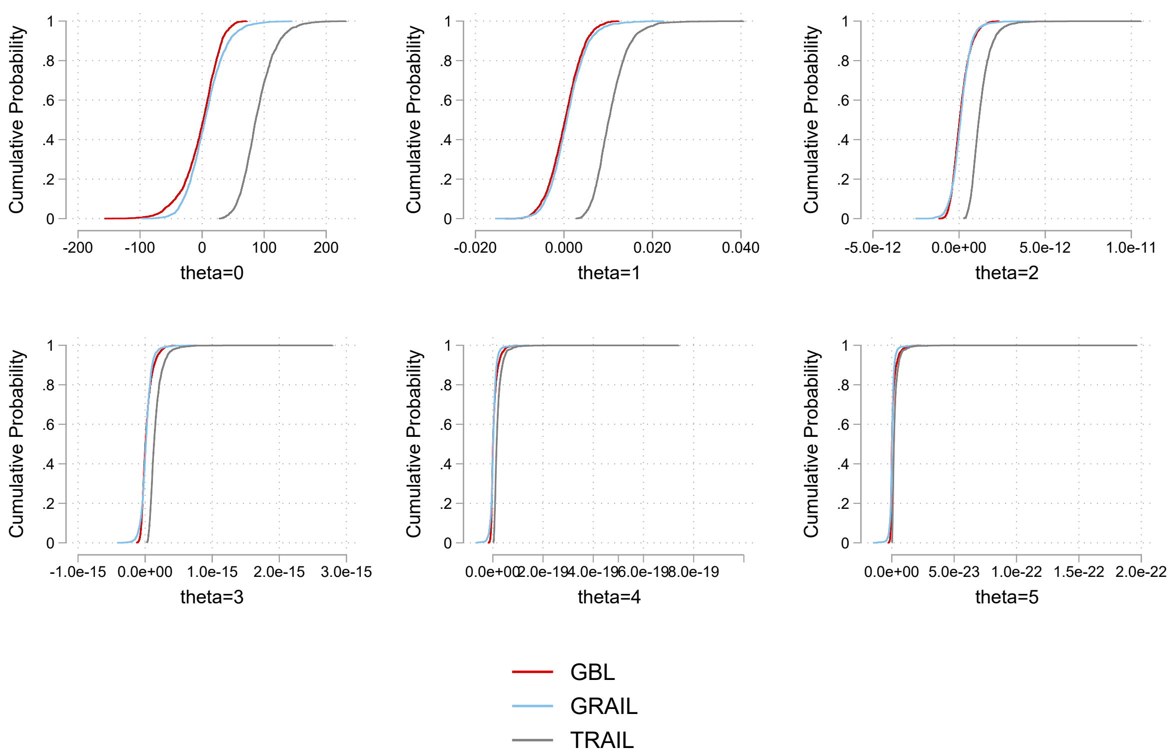

Notes: Cumulative distribution functions are drawn from 2000 bootstrap estimates of aggregate welfare impacts of the TRAIL, GRAIL and GBL schemes.

Figure 2: Treatment Proportions by Land class

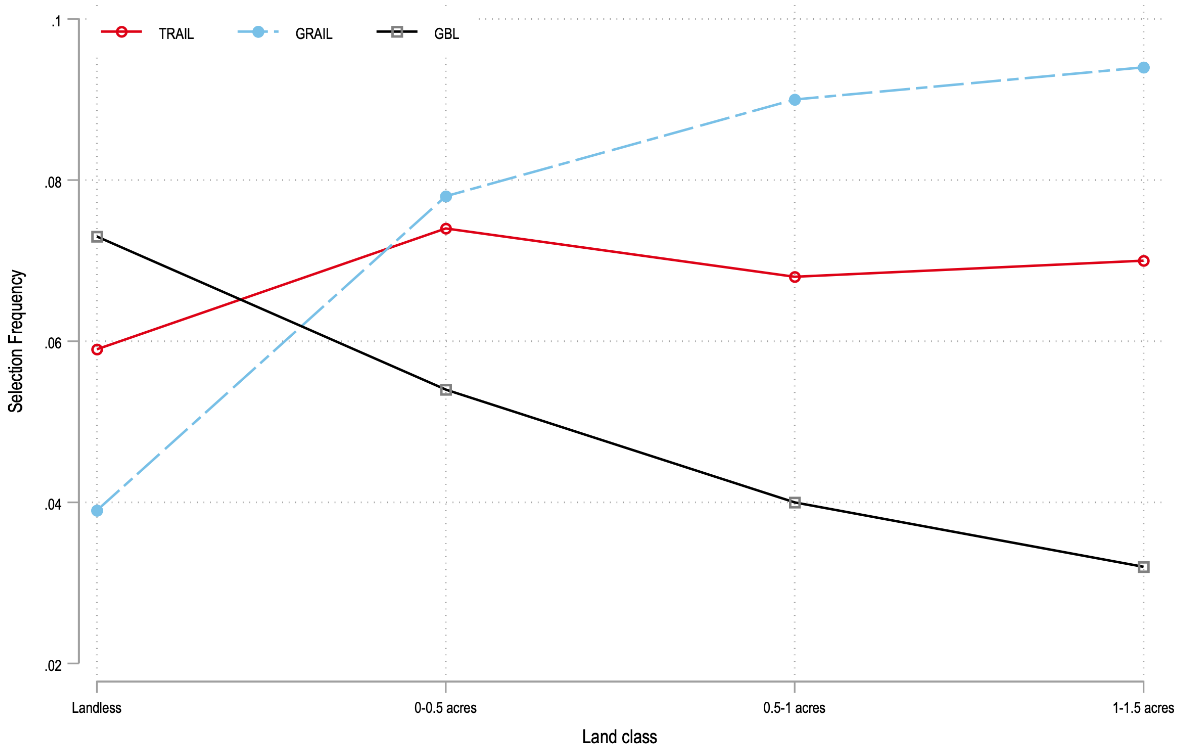

Notes: Average Selection frequencies across all TRAIL, GRAIL and GBL villages by land class presented.

Figure 3: Conditional Treatment Effects by Land Category, Intervention and θ

$\theta = 0$ $\theta = 1$

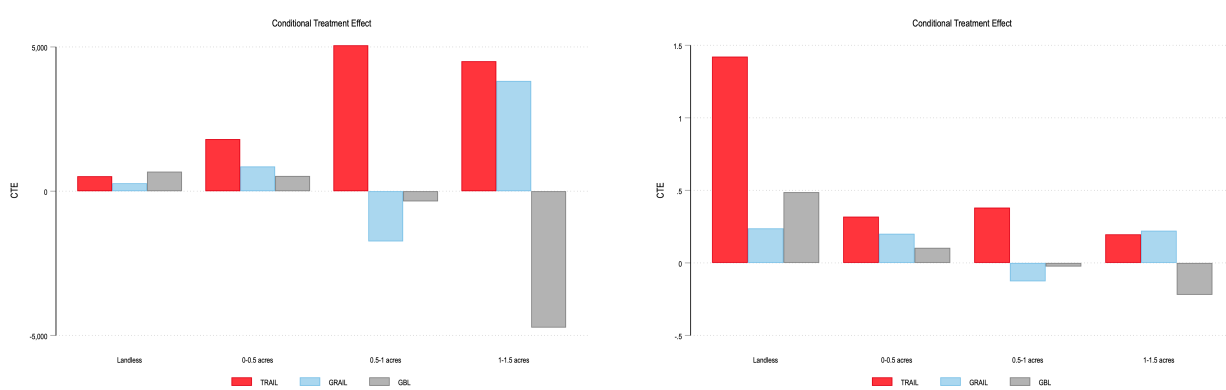

Figure 4: Treatment Effects by Land Category, Intervention and θ

$\theta = 0$ $\theta = 1$

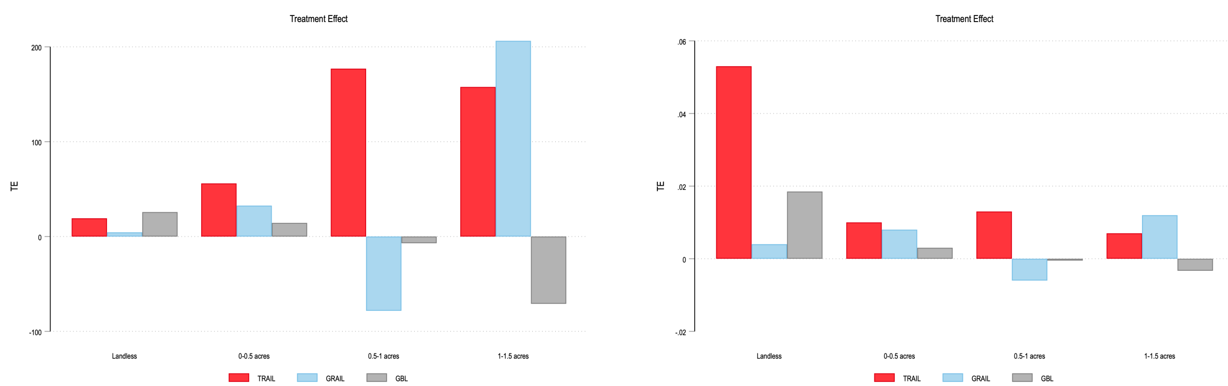

Table 1: Balance of Village-level Characteristics, by Village Treatment Arm

	All	TRAIL	GRAIL	GBL	Differences: Two-way comparisons		
	(1)	(2)	(3)	(4)	(2)-(3)	(2)-(4)	(3)-(4)
Number of households	365.32	327.63	310.71	457.58	16.92	-129.96	-146.88
	(40.66)	(52.28)	(64.87)	(88.35)	(83.32)	(102.66)	(109.61)
Proportions by landholding class							
Landless	0.18	0.18	0.18	0.17	0.00	0.02	0.02
	(0.01)	(0.03)	(0.03)	(0.02)	(0.04)	(0.03)	(0.03)
0–0.5 acres	0.34	0.32	0.36	0.34	-0.04	-0.02	0.02
	(0.02)	(0.03)	(0.03)	(0.03)	(0.04)	(0.04)	(0.04)
0.5–1 acre	0.22	0.23	0.20	0.22	-0.03	0.01	0.02
	(0.01)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
1–1.5 acres	0.10	0.10	0.11	0.11	-0.01	-0.01	0.00
	(0.01)	(0.01)	(0.02)	(0.01)	(0.02)	(0.02)	(0.02)
$>$ 1.5 acres	0.17	0.17	0.15	0.17	0.02	-0.00	0.02
	(0.01)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Percent households electrified	0.615	0.603	0.652	0.591	-0.049	0.01	0.061
	(0.03)	(0.06)	(0.05)	(0.05)	(0.08)	(0.08)	(0.08)
Has primary school	0.779	0.773	0.773	0.792	0.00	-0.02	-0.02
	(0.05)	(0.09)	(0.09)	(0.08)	(0.129)	(0.12)	(0.12)
Has primary health centre	0.221	0.273	0.182	0.208	0.09	0.06	-0.03
	(0.05)	(0.10)	(0.08)	(0.08)	(0.13)	(0.13)	(0.12)
Has bank branch	0.074	0.00	0.045	0.167	-0.05	-0.17	-0.12
	(0.03)	(0.00)	(0.05)	(0.08)	(0.05)	(0.08)	(0.09)
Has MFI	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Has *pucca* road	0.353	0.27	0.36	0.42	-0.09	-0.14	-0.05
	(0.06)	(0.10)	(0.11)	(0.10)	(0.14)	(0.14)	(0.15)
F-test of joint significance					0.45	1.11	0.51
$p - value$					*0.81*	*0.37*	*0.77*

Notes: The number of households in the 72 sample villages is taken from the 2011 Census of India village directory. Proportions of households in each landholding class are calculated from the 2007 house-listing exercise we conducted in 68 of these 72 villages for a previous studies reported in (Maitra et al., 2017, 2021). Other village-level characteristics are sample means from 68 village surveys conducted in 2007. Four villages from the (Maitra et al., 2021) study were replaced in 2010 because of Maoist conflict, and we do not have pre-intervention village census or village survey data for the replacements.

Table 2: Household Characteristics, by Village Treatment Arm and Land Category

	All (1)	TRAIL (2)	GRAIL (3)	GBL (4)	Landless (5)	0–0.5 acres (6)	0.5–1 acre (7)	1–1.5 acres (8)
Head's education: primary or more	0.420	0.407	0.420	0.433	0.234	0.356	0.564	0.650
	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)	(0.03)
Head's occupation: cultivation	0.431	0.441	0.415	0.437	0.056	0.381	0.689	0.696
	(0.01)	(0.02)	(0.02)	(0.02)	(0.01)	(0.01)	(0.02)	(0.02)
Head's occupation: Labor	0.335	0.340	0.343	0.323	0.677	0.404	0.089	0.041
	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.01)	(0.01)
Non-Hindu	0.172	0.21	0.151	0.155	0.188	0.146	0.195	0.181
	(0.01)	(0.02)	(0.01)	(0.01)	(0.02)	(0.01)	(0.04)	(0.02)
Low caste	0.387	0.383	0.355	0.423	0.565	0.417	0.286	0.199
	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)	0.023
Area of house and homestead (acres)	0.053	0.052	0.052	0.054	0.037	0.048	0.063	0.074
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Pucca House	0.294	0.288	0.334	0.259	0.207	0.280	0.344	0.379
	(0.01)	(0.01)	(0.02)	(0.01)	(0.02)	(0.01)	(0.02)	(0.03)
Electrified house	0.751	0.740	0.752	0.760	0.666	0.729	0.811	0.841
	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)	(0.02)
Separate toilet in house	0.575	0.564	0.608	0.552	0.434	0.541	0.664	0.741
	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)
Owns radio/ TV/ VCR/ DVD	0.464	0.450	0.486	0.456	0.350	0.420	0.541	0.639
	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)	(0.03)
Owns telephone (mobile or landline)	0.590	0.573	0.590	0.607	0.446	0.528	0.706	0.796
	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)
Has savings bank account	0.456	0.447	0.475	0.446	0.268	0.410	0.576	0.680
	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)
Test of joint significance for assignment to treatment								
Chi-squared statistic			10.32		32.64	28.30	24.09	32.27
p-value			*0.993*		*0.112*	*0.248*	*0.456*	*0.121*

Notes: Household characteristics data are from the first wave of household surveys conducted in the 72 sample villages in 2010. Only eligible households are included in the sample. Since we drew a purposive sample of Treatment, Control 1 and Control 2 households, we do not expect our sample means to be representative of the village populations. To correct for the non-representativeness of our sample, we assign each household a weight that is in inverse proportion to the probability that they would be selected into the sample. Thus, Treatment and Control 1 households each receive a weight of $\frac{30}{N}$ and Control 2 households receive a weight of $\frac{N-30}{N}$, where N denotes the total number of households in the village, as reported in the 2011 Census.

Table 3: Estimated Welfare Impacts of the Loan Interventions

Value of θ	$\theta = 0$	$\theta = 1$	$\theta = 2$	$\theta = 3$	$\theta = 4$	$\theta = 5$
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Conditional treatment effects on household welfare						
TRAIL	2456.36	0.283	3.205E-11	3.699E-15	4.271E-19	4.934E-23
Confidence interval	*(1493.55, 3399.04)*	*(0.134, 0.401)*	*(4.52E-12, 4.75E-11)*	*(-1.02E-15, 5.81E-15)*	*(-3.78E-19, 7.01E-19)*	*(-8.55E-13, 8.39E-23)*
GRAIL	125.83	0.017	2.205E-12	2.942E-16	3.921E-20	5.221E-24
Confidence interval	*(-1099.77, 1298.32)*	*(-0.146, 0.180)*	*(-1.97E-11, 2.44E-11)*	*(-2.81E-15, 3.45E-15)*	*(-4.17E-19, 5.01E-19)*	*(-6.65E-23, 7.23E-23)*
GBL	33.5	0.0048	7.146E-13	1.012E-16	1.428E-20	2.010E-24
Confidence interval	*(-1488.72, 1988.19)*	*(-0.251, 0.206)*	*(-4.35E-11, 2.14E-11)*	*(-7.67E-15, 2.44E-15)*	*(-1.34E-18, 2.89E-19)*	*(-2.31E-22, 3.61E-23)*
Panel B: Change in aggregate welfare (using treatment proportion)						
TRAIL	83.05	0.010	1.084E-12	1.251E-16	1.444E-20	1.668E-24
	(27.29, 114.79)	*(0.002, 0.014)*	*(-1.74E-13, 1.63E-12)*	*(-8.54E-17, 1.97E-16)*	*(-2.04E-20, 2.37E-20)*	*(-3.81E-24, 2.84E-24)*
GRAIL	4.310	0.001	7.558E-14	1.009E-17	1.344E-21	1.790E-25
	(-49.08, 47.15)	*(-0.007, 0.006)*	*(-8.75E-13, 9.00E-13)*	*(-1.25E-16, 1.29E-16)*	*(-1.87E-20, 1.83E-20)*	*(-2.85E-24, 2.68E-24)*
GBL	0.870	1.00E-04	1.858E-14	2.630E-18	3.713E-22	5.228E-26
	(-38.23, 62.32)	*(-0.006, 0.006)*	*(-1.05E-12, 6.56E-13)*	*(-1.82E-16, 7.36E-17)*	*(-3.19E-20, 8.89E-21)*	*(-5.69E-24, 1.10E-24)*
Panel C: Comparison of welfare effects						
Pr(GRAIL = TRAIL) : $p - value$	0.000	0.000	0.000	0.000	0.000	0.000
Pr(GRAIL > TRAIL) : $p - value$	*0.021*	*0.035*	*0.057*	*0.089*	*0.123*	*0.155*
Pr(GBL = TRAIL) : $p - value$	0.000	0.000	0.000	0.000	0.000	0.000
Pr(GBL > TRAIL) : $p - value$	*0.003*	*0.019*	*0.061*	*0.144*	*0.163*	*0.204*
Pr(GBL = GRAIL) : $p - value$	0.000	0.000	0.124	0.695	0.049	0.001
Pr(GBL > GRAIL) : $p - value$	*0.445*	*0.465*	*0.486*	*0.504*	*0.518*	*0.529*

Notes: Welfare impacts are estimated following the procedure outlined in Section 2, where θ indicates the value of the inequality-aversion parameter in the individual utility function that is the building block of the social welfare function. The terms in parentheses denote the bootstrapped 90% confidence interval (with 2000 replications). Panel C presents the results of Mann-Whitney rank-sum tests using 2000 bootstrap replications, comparing the aggregate welfare effects of the schemes. The p-value is for a test that the two unconditional treatment effects come from the same population. $Pr(X > Y)$ indicates an estimate of the probability that a random bootstrapped estimate of the unconditional treatment effect of scheme X is larger than a randomly drawn bootstrapped estimate of the unconditional treatment effect of scheme Y, where $X; Y \in$ {TRAIL, GRAIL, GBL}

Table 4: Decomposition of Welfare Impacts of the Loan Interventions

Land class:	Landless	0—0.5 acres	0.5—1 acres	1—1.5 acres	Weighted average
	(1)	(2)	(3)	(4)	(5)
Panel A: TRAIL					
Population proportion (α_g)	0.226	0.392	0.279	0.103	
Treatment proportion (s_g)	0.037	0.031	0.035	0.035	
Mean Aggregate Farm Income, Control 1 households	361.64	5701.46	13293.15	22137.23	8579.63
Conditional treatment effects:					
$\theta = 0$	518.81	1805.34*	5057.49***	4505.63	
	(469.86)	(934.54)	(1724.33)	(4153.25)	
$\theta = 1$	1.422	0.319*	0.381***	0.197	
	(1.290)	(0.161)	(0.129)	(0.184)	
$\theta = 2$	0.004	6.00E-05*	3.00E-05***	8.56E-06	
	(0.004)	(3.00E-05)	(9.623E-06)	(8.12E-06)	
Treatment effects:					
$\theta = 0$	19.20	55.97	177.01	157.70	91.91
$\theta = 1$	0.053	0.010	0.013	0.007	0.02
$\theta = 2$	1.48E-04	1.86E-06	1.05E-06	3.00E-07	3.45E-05
Panel B: GRAIL					
Population proportion (α_g)	0.302	0.387	0.222	0.090	
Treatment proportion (s_g)	0.016	0.038	0.045	0.054	
Mean Aggregate Farm Income, Control 1 households	1415.21	4237.07	13367.33	16750.96	
Conditional treatment effects:					
$\theta = 0$	276.33	857.23	-1740.29	3820.87	
	(937.82)	(601.43)	(1424.58)	(3944.03)	
$\theta = 1$	0.238	0.201	-0.127	0.222	
	(0.620)	(0.141)	(0.105)	(0.234)	
$\theta = 2$	1.00E-04	5.00E-05	-9.34E-06	1.00E-05	
	(0.0005)	(3.00E-05)	(7.790E-06)	(1.00E-05)	
Treatment effects:					
$\theta = 0$	4.42	32.57	-78.31	206.33	15.13
$\theta = 1$	0.004	0.008	-0.006	0.012	0.004
$\theta = 2$	1.60E-06	1.90E-06	-4.20E-07	5.40E-07	1.17E-06
Panel C: GBL					
Population proportion (α_g)	0.198	0.377	0.285	0.140	
Treatment proportion (s_g)	0.038	0.029	0.020	0.015	
Mean Aggregate Farm Income, Control 1 households	1236.29	4961.05	12580.71	21392.85	
Conditional treatment effects:					
$\theta = 0$	678.57	528.97	-352.24	-4728.52	

Continued . . .

Decomposition of Welfare Impacts of the Loan Interventions (Continued)

Land class:	Landless	0-0.5 acres	0.5-1 acres	1-1.5 acres	Weighted average
	(1)	(2)	(3)	(4)	(5)
	(708.22)	(988.53)	(2658.86)	(5767.83)	
$\theta = 1$	0.488	0.104	-0.025	-0.220	
	(0.635)	(0.197)	(0.210)	(0.269)	
$\theta = 2$	4.00E-04	2.00E-05	-1.71E-06	-1.00E-05	
	(0.0004)	(4.00E-05)	(2.00E-05)	(1.00E-05)	
Treatment effects:					
$\theta = 0$	25.79	15.34	-7.04	-70.93	-1.05
$\theta = 1$	1.85E-02	3.02E-03	-5.00E-04	-3.30E-03	-1.05
$\theta = 2$	1.52E-05	5.80E-07	-3.42E-08	-1.50E-07	3.20E-06

Notes: The conditional treatment effects are estimated for each land sub-group following equation (9) in the text. Population proportions are estimated from 2007 pre-intervention village census data from 68 of the 72 sample villages. The treatment proportions are proportions of the households in land sub-group g who were actually treated. θ indicates the value of the inequality-aversion parameter in the individual utility function that is the building block of the social welfare function.

Data Availability Statement

The data underlying this article will be made available in online supplementary material.