WISE Pre-Analysis Plan

June 30, 2025

1 Motivation

Research has shown that increasing women's access to, and control over, potential income increases her financial autonomy and participation in economic activities. In settings with conservative gender norms, this activity can also reshape restrictive attitudes and increase markers of female empowerment (Field et al., 2021). While the individual- and household-level impacts of such an initiative are clear, the consequences of this policy when implemented at-scale – across entire markets and communities – are less clear: Encouraging women's economic control and access to economic opportunities could increase local economic activity and liberalize restrictive norms, and direct impacts could also spillover onto untreated women. In addition, an at-scale approach could also change market wages and job availability and impact men's economic opportunities. It is possible this could cause increased resistance by community members, especially men, towards greater economic activity by women.

In this project, we examine the consequences of such a "big push" approach to increasing women's financial control and access to local work opportunities, varying the intensity of treatment by randomly assigning some locations to have a higher saturation of treatment than others. We assess the impact of this approach in rural Madhya Pradesh, India through supporting women to access MGNREGS, a guaranteed work access program that provides low-wage work close to home. Support is provided through a government-driven push to "link" women's bank accounts to the workfare payment system, as well as through a standardized training utilizing a video and facilitated discussion, delivered in women's communities through either government representatives or NGO-sponsored trainers.

2 Research Questions

Our primary research questions center on understanding impacts of linking a woman's bank account to India's workfare payment system for direct payments, alongside a training that helps women understand how to access workfare and ensure their wages are paid directly to that personal bank account along multiple dimensions. We assess the impact of this training+linking treatment both relative to a pure control and in comparison to treatment at a different level of saturation.

2.1 Primary Research Questions

1. How does linking+training (training to encourage and facilitate women's participation in workfare) impact women's labor supply in both the public and private sector?

2. How does linking+training impact gender norms and women's (and men's) attitudes related to women and work?

3. How does linking+training impact intrahousehold relationships, including women's agency within the household and their experience of controlling behavior and intimate partner violence?

4. How do the treatment effects in questions 1-3 (on labor supply, norms/attitudes, and intrahousehold relationships) vary by the share of women treated in a community or labor market?

5. What are the spillover effects on households that are not directly treated with training? (For example, in terms of their labor market outcomes or gender-related attitudes?)

2.2 Secondary Research Questions

1. What are the general equilibrium consequences of increasing women's labor supply in terms of wages, prices, and local economic activity? Are gender gaps in wages and local economic outcomes impacted?

2. How does the efficacy of training depend on whether the training is delivered by experienced NGO trainers versus directly by government representatives (mates)? What is the cost-effectiveness of each mode of training?

3. How does linking+training affect household consumption, debt levels, and savings?

4. How does linking+training affect women's (and men's) well-being and mental health?

5. How does increasing access to workfare through linking alone, rather than a linking+training approach, affect primary work-related outcomes of women and their households?

6. Does linking+training affect family planning, and if so, how?

3 Experimental Design

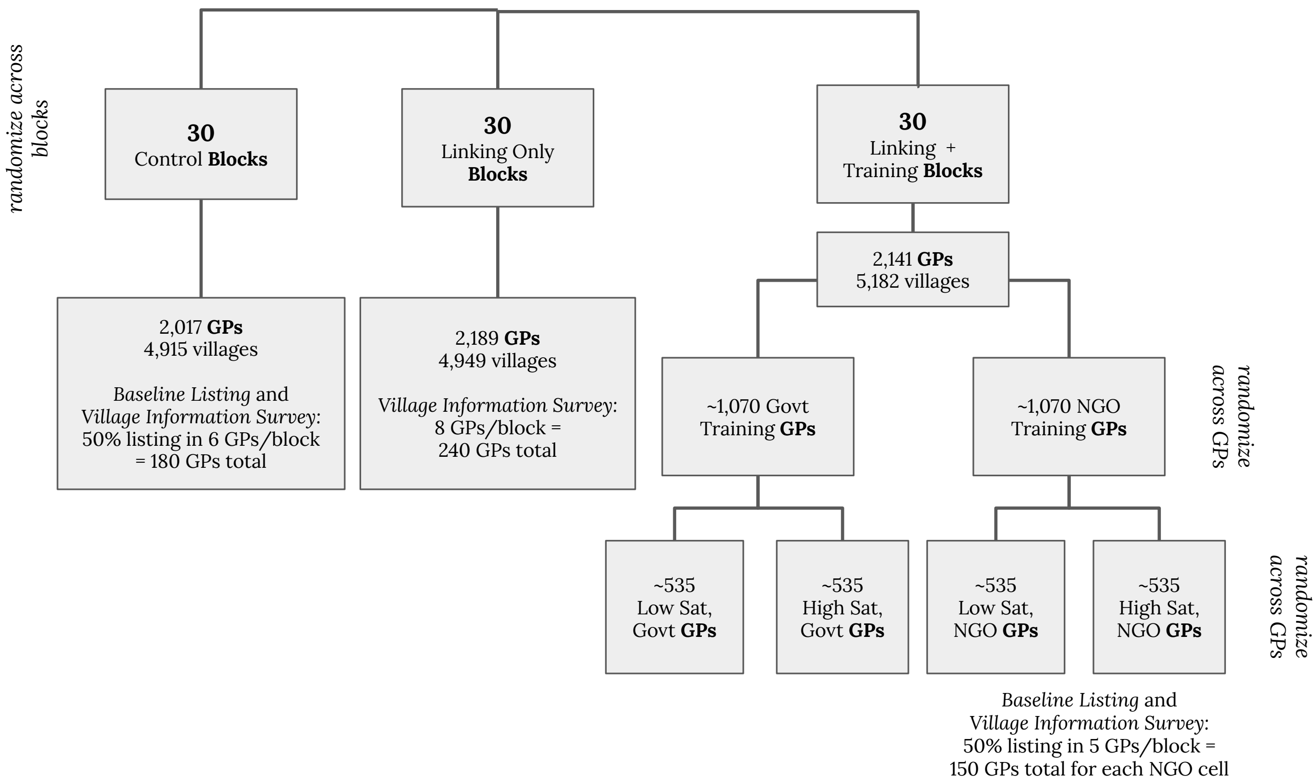

3.1 Intervention Description

There are two main interventions in the study: (i) training and (ii) linking. In the 90 blocks included in our study, 30 blocks are assigned to receive no intervention (pure control), 30 blocks are assigned to receive only the linking intervention, and 30 blocks are assigned to receive both the linking and training interventions.

Linking. We are coordinating with the state government to link women's bank accounts to their biometric identity (Aadhaar) cards, and to link their Aadhaar cards to the MGNREGS system, which allow them to avail of the MGNREGS work guarantee program. Given prior efforts to link recent MGNREGS workers, this initiative will target "inactive" female workers who have not participated in MGNREGS in the last 2-3 years.

Training. We are collaborating with our NGO partner, PRADAN, to design and deliver trainings that build women's capacity to request MGNREGS work and navigate direct deposit payments. PRADAN is a well-established Indian NGO dedicated to empowering rural communities through grassroots livelihood interventions. Within the training intervention, we test variations in (a) who delivers the training and (b) what share of women per village are trained. In each GP, we randomly select one village that receives training.

- **NGO- vs government-led trainings.** Within the training and linking blocks, half of the GPs are randomly assigned to be "NGO-led" GPs, while the other half are assigned to be "government-led" GPs. Although both the NGO- and government-led GPs will receive trainings designed by PRADAN and will have trainers who are trained by PRADAN, the identity and experience level of the the last-mile trainers will vary. In the NGO-led GPs, the trainings are delivered by PRADAN appointed “Community Resource Persons” (CRPs), locally based employees, typically women, with experience supporting local self-help groups and providing training on a variety of topics. The government-led GPs will instead have trainings that are delivered by government-selected workers nominated by a local MGNREGS official, often women known as “mahila mates.” Mahila mates are women who are informally appointed to support community members in providing access to MGNREGS workfare, and typically they have less experience organising trainings. The goal of these two variations is to test whether mahila mates can deliver trainings at the last mile as effectively as more experienced PRADAN trainers, which would make this intervention more scalable.

- **High- vs low-saturation training.** Within the 30 training and linking blocks, half of the GPs are randomly assigned to be “high saturation GPs,” while the other half of the GPs are “low saturation GPs." In low-saturation GPs, we target training approximately 25% of working-age women in the randomly selected village of the GP. In high-saturation GPs, we plan to train approximately 75% (recognizing that we may not quite achieve this target). Population numbers for saturation targets will be created using lists of women who are receiving payments through *Ladli Behena Yojana* (LBY), a recently rolled-out government program targeting low-income married women between the ages of 21 and 60, which is a good proxy for the target population of our own training intervention.

3.2 Selection and randomization of study areas

Our study state, Madhya Pradesh, has a population of approximately 90 million, spread across (from largest to smallest administrative unit) 55 districts, 313 blocks, ≈23,000 GPs, and ≈55,000 villages. Within these 55 districts, we selected the 15 districts where our NGO partner PRADAN was already active and would therefore be able to deliver the intervention, but was not already delivering a similar MGNREGS training. The subsections below describe selection and randomization at the block, gram panchayat (GP), and village level.

Block selection and randomization. The highest unit of randomization in our study is the block, an administrative unit below the district. In order to ensure that our NGO training partner, PRADAN, is able to deliver the trainings, we selected blocks from the 15 districts in which PRADAN was already operating as of August 2024, but not conducting related trainings. This resulted in 95 potential blocks. Then, 90 out of 95 blocks were randomly selected for study inclusion.

Two strata of three levels each were used to randomly assign the 90 blocks to three treatment arms: (i) control, (ii) linking only, and (iii) linking and training. First, we stratify on district clusters, which group study districts into central, southeast, and northeastern regions of the state. Second, the other stratum divides the sample based on whether PRADAN works there and provides MGNREGS support as of August 2024, works there but does not provide MGNREGS support, or doesn’t work there at all.

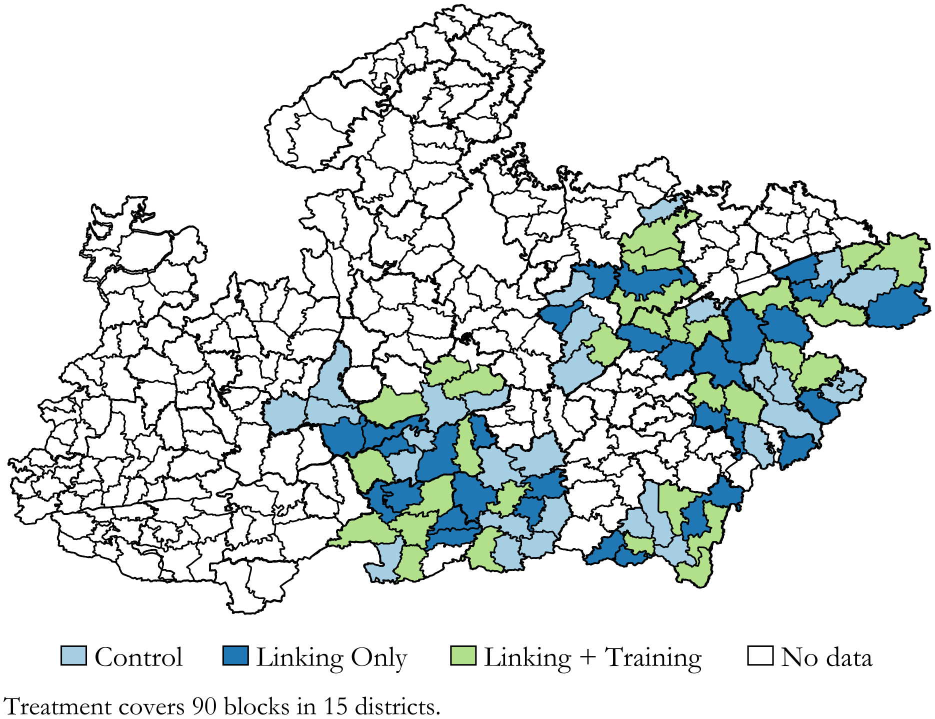

Figure 1: Block-Level Selection and randomization

Gram Panchayat (GP) randomization. The lowest level of randomization is the GP, a cluster of villages and India's lowest-level administrative unit. Assignment to NGO- vs government-led training was cross-randomized with assignment to high- and low-saturation training.

- *randomization to NGO- vs government-led training.* All GPs in the 30 training+linking blocks were randomly assigned to receive either NGO-or government-led training. randomization to these two training modes is stratified by block, eligibility for the survey sample (described in section 4.1), and population.[1]
- *randomization to high- vs low-saturation training.* All GPs within the 30 linking+training blocks were randomly assigned to either a high-saturation arm or low-saturation arm. randomization was stratified by block, NGO-led vs government-led, and, within the NGO-led arm, by baseline listing status.

Village selection. Within each GP, one village was randomly selected to receive training. Villages that had population zero according to (or were missing from) the 2011 census were excluded.

[1] There are 61 strata in total. The survey-sample-eligible GPs are grouped into two strata per block (one low population, one high population) which x30 = 60 strata. The population strata are relative to the median population of baseline-eligible GPs in the block. We include one additional stratum (the 61st) that includes all GPs across blocks which were not eligible for the survey sample.

4 Data Collection

4.1 Baseline Listing and Village Information Survey

We are conducting two forms of primary data collection at baseline in randomly selected villages: (i) a village information survey with a village leader, and (ii) a baseline listing with 40-50% of households. In total, 720 GPs are enrolled in some primary baseline data collection; in 480 GPs, both the village information survey and baseline listing are conducted, while in the remaining 240 GPs, only the village information survey is being conducted.

Primary data collection was decided by training arm. In the control arm, 6 GPs per block were randomly selected for both the household listing and village information survey (180 GPs total). In the linking only arm, 8 GPs per block (240 GPs in total) were randomly selected for the village information survey only. We did not conduct a household listing in the linking only arm because we do not anticipate doing a household endline survey in the linking only blocks given current project funds. In the training+linking arm, 10 GPs per block were randomly selected for both the household listing and village information survey (300 GPs total). A greater number of GPs per block were selected for the training+linking arm as compared to the control arm in order to increase our power to detect differences in treatment effects between high- and low-saturation training as well as government- versus NGO-led training modalities.

Village information survey. The village information survey is completed with a local MGNREGS official (Gram Rozgar Sahayaks, or GRS) or, if unavailable, the locally-elected sarpanch. The GRS is a field-level official responsible for implementing the MGNREGS at the village level, including maintaining records of work demand, job cards, and wage payments. The sarpanch is the elected head of the village panchayat (local self-government institution) and is responsible for overseeing village-level development programs and administration. The village information survey includes questions about the village leaders, labor markets, agricultural production, population, and engagement of village residents with MGNREGS.

Baseline household listing. The household listing covers 40-50% of households and is completed in 6 villages per block in the control group and 10 villages per block in the training+linking treatment arm. To determine which households are included, a team of enumerators approaches every other household following the left-hand rule and asks to speak with a knowledgeable adult man or woman within the household.

Village selection process for baseline. The baseline villages were selected through a three-step process. First, gram panchayats that had reported population of zero or that did not appear in administrative data were excluded. Second, GPs were randomly selected from each block. Third, a village from each GP was randomly selected.

1. *Eligible GPs*. GPs were eligible for the survey sample if (i) they were in the MGNREGS administrative data received from the Ministry of Rural Development (MoRD) via a custom API, (ii) they were included in the 2011 census with a non-zero population, and (iii) our team was able to match the GP to the GP-village crosswalks in administrative data. Out of the initial 6,347 GPs in the 90 study blocks blocks, 6,320 (99.6%) appeared in the data provided by MoRD. Of these GPs, 6,283 had a non-zero population according to the 2011 census (99.4%). Then there were 41 additional GPs (<0.01%) that our team was unable to match to the GPs in administrative

data mapping villages to GPs. Excluding these gave us a sample of 6,264 GPs that were eligible for baseline selection, corresponding to 98.7% of all GPs in study blocks.

2. *GP selection.* Eligible GPs were randomly selected within block while stratifying on within-block median GP population. In the control arm, 6 GPs per block were selected. In the training and linking arm, 10 GPs per block were selected. No listing was completed in the linking only arm, but the village information survey was conducted in 8 GPs per block in this arm. We also selected a "buffer" list of GPs within each stratum, which was used when the originally selected GPs were unable to be surveyed, for example in the unlikely event that the enumerator team was turned away by village officials who were opposed to data collection. [2]

Village selection. Within the eligible GPs, there were 2.4 villages per GP on average. To randomly select one of the villages for the baseline surveys, we first excluded villages with population zero according to the 2011 census (3.8% of villages in the sample). Then, one remaining village per GP is randomly selected.

4.2 Follow-Up Survey

We plan to conduct a follow-up survey with approximately 20 households in the same villages where we conducted a baseline listing. However, this follow-up survey si not currently fully funded, and so the number of villages and number of households per village may change moving forward. Depending on funding, we may also conduct follow-up surveys with additional households or in additional villages/panchayats.

5 Analysis Plan

We plan to estimate both the effects of linking only and the effects of the combined training and linking intervention, although we will concentrate resources on estimating the effects of the combined training and linking intervention, as we expect linking alone to have a relatively muted effect. Below we outline our plan for analyzing the effects of the combined training and linking intervention (henceforth referred to as "training").

The specifications and tests outlined below are our best approach given all the information we have on hand pre-intervention; we may adjust our approach to deepen understanding of impacts and mechanisms, and so do not limit ourselves solely to what is described here.

Our experiment employs a randomized saturation design to understand how the effects of training vary by the share of women treated. Within the training treatment arm, we randomize GPs to either low saturation (targeting approximately 25% of women for training in the randomly selected village) or high saturation (targeting approximately 75%). Assignment to control, low saturation, and high saturation creates village-level random variation in the intensity of treatment.

[2] As of 31 May 2025, this had occurred three times.

Our experiment also aims to estimate spillover effects of training onto women who do not directly attend the training. To generate exogenous variation in the participation of individual women and households, we randomly select some women within treatment villages to be named on an invitation list, with the number of invitees calibrated to achieve the randomly-assigned village-level saturation target. These lists will consist of women randomly selected from the LBY list as well as from our baseline household listing. Training coordinators will make special efforts to ensure that invited women attend the training sessions. We expect that actual attendance will be subject to imperfect compliance—some invited women will not attend, while some non-invited women will attend.

This two-level randomization allows us to separately identify at each level of saturation: (1) direct effects of attending training, (2) spillover effects on women who attended training, and (3) spillover effects on women who did not attend training. We also consider the combined direct and indirect effects together in a total causal effect and assess how the total effect depends on the level of training saturation.

5.0.1 GP-level Analysis with Administrative Data

We plan to use government administrative data to assess the effects of training+linking as well as the effects of linking only at the GP level using regressions such as the following:

$$Y_v = \beta_0 + \beta_1 T_v + \beta_2 L_v + X_v' \gamma + \mu_s + \varepsilon_v \tag{1}$$

where:

- Y_v is the outcome for gram panchayat v
- T_v is an indicator for GP v being assigned to a training+linking block
- L_v is an indicator for GP v being assigned to a linking only block
- X_v is a vector of baseline GP-level covariates
- μ_s are strata fixed effects
- ε_v is an error term clustered at the block level

If we do not see treatment effects of linking only, then in some specifications we may pool together the linking only arm with the control arm.

To estimate treatment effects by **saturation level**, we will also run the following regression that separates GPs assigned to low- versus high-saturation training:

$$Y_v = \beta_0 + \beta_1 (T_v \times HighSat_v) + \beta_2 (T_v \times LowSat_v) + \mu_s + X_v' \gamma + \varepsilon_v \tag{2}$$

To estimate treatment effects by **training modality**, we will also run the following regression that separates GPs assigned to NGO- versus government-led training:

$$Y_v = \beta_0 + \beta_1(T_v \times NGO_v) + \beta_2(T_v \times Gov_v) + \mu_s + X_v'\gamma + \varepsilon_v \tag{3}$$

We will also estimate regression equations that examine effects separately by both training modality and saturation level:

$$\begin{aligned} Y_v = &\beta_0 + \beta_1(T_v \times NGO_v \times HighSat_v) + \beta_2(T_v \times NGO_v \times LowSat_v) + \beta_3(T_v \times Gov_v \times HighSat_v) \\ &+ \beta_4(T_v \times Gov_v \times LowSat_v) + \mu_s + X_v'\gamma + \varepsilon_v \end{aligned} \tag{4}$$

where β_1 represents the ITT effect of NGO-led, high-saturation training, β_2 the effect of NGO-led, low-saturation training, β_3 the effect of government-led, high-saturation training, and β_4 the effect of government-led, low-saturation training. The omitted group is the control group.

5.0.2 Intent-to-Treat Effects from Survey Data

In this analysis we will primarily compare (a) control blocks to NGO-led GPs within linking+training blocks, and (b) high-saturation vs low saturation NGO-led GPs within linking+training blocks. To assess impacts, we will use the following specifications in **individual- or household-level** regressions:

$$Y_{iv} = \beta_0 + \beta_1 T_{iv} + X_{iv}'\gamma + \mu_s + \varepsilon_{iv} \tag{5}$$

$$Y_{iv} = \beta_0 + \beta_1 T_{iv} \times LowSat_{iv} + \beta_2 T_{iv} \times HighSat_{iv} + X_{iv}'\gamma + \mu_s + \varepsilon_{iv} \tag{6}$$

where:

- Y_{iv} is the outcome for individual or household i in village v
- T_{iv} is an indicator for being in a village located in a training block
- X_{iv} is a vector of baseline covariates
- μ_s are strata fixed effects
- ε_{iv} is an error term clustered at the block (treatment vs control comparisons) or GP level (high saturation vs low saturation comparisons)

In equation 5, β_1 estimates the intent-to-treat (ITT) effect of training, while the omitted group is households or individuals assigned to the control arm. As with the GP-level regressions, we also plan to estimate treatment effects separately by saturation level (equation 6). We plan to estimate these equations both without control variables as well as using control variables selected by PDS LASSO (with lagged outcomes in the selection set when available).

We may also collect **village-level** outcomes in our follow-up surveys (e.g. number of active MGNREGS worksites in the village) and run similar regressions at the village level with survey data.

5.0.3 Separating Treatment-on-Treated vs Spillover Effects

We plan to estimate IV models to separately identify treatment-on-treated (TOT) versus spillover effects of training within treatment villages. To address endogeneity in actual training attendance, we use our individual-level randomized invitations as instruments. There are two sources of individual-level randomization in attendance: (i) inclusion in a randomly-selected group of women from the baseline listing that is added to the top of the training mobilization list, and (ii) a randomly-selected group of women from the baseline listing whom we attempt to exclude from the LBY mobilization list being used in that GP. We anticipate that the instruments will have more bite in low-saturation villages, where it is easier for trainers to adhere to the invitation lists only. Our structural equation is as follows:

$$Y_{iv} = \beta_0 + \beta_1 Attend_{iv} + \beta_2 Spillover_{iv} + \mu_s + X'_{iv}\gamma + \varepsilon_{iv} \tag{7}$$

where the endogenous variables *Attend* and *Spillover* are predicted using block-level random assignment to linking+training, GP-level random assignment to training modality and saturation, and individual-level random assignment to placement on the mobilization list.

Then β_1 is the effect of training attendance (combining direct effects of training as well as spillover effects on the trained), and β_2 is the spillover effect of training on non-attendees who live in training villages. The β_1, β_2 values may be different for low vs high saturation training villages. We anticipate estimating the IV models separately by saturation levels, but we may also estimate these equations while pooling together the two different levels of saturation. We plan to test whether β_1, β_2 are the same at low- vs high-saturation levels of treatment.

We are also exploring the possibility of separately identifying and estimating spillover effects on the treated, following Baird et al. (2018). We may also employ the frameworks outlined in DiTraglia et al. (2023). We will update this PAP if we decide to pursue these strategies.

5.0.4 General equilibrium effects.

If the treatment is effective at increasing women's involvement in MGNREGS or other paid work, then we also estimate GE effects of increasing female labor force participation. We plan to release specifications for this analysis in an addendum to this initial PAP. Increasing women's labor force participation through our linking+training program could generate several types of general equilibrium effects in local labor markets:

1. **Labor supply effects**: The program directly increases the supply of female labor, which could affect:

    - Wage levels for women and for men, in and out of agriculture (potentially decreasing wages if labor demand is not perfectly elastic)

- Employment opportunities for untrained women and community men (potentially creating displacement effects)

2. **Labor demand effects**: Higher female employment could generate:

   - Increased consumption, stimulating business activity
   - Productivity gains if the increased female labor supply leads to better worker-job matches
   - New business formation or business expansion responding to the availability of female workers

3. **Equilibrium adjustments in related markets**:

   - The prices of capital inputs to production (e.g. land, animals)
   - Changes in household formation and marriage markets (e.g. if women in households of trained women or women in general delay marriage in response to broader economic effects)
   - Credit market prices (i.e. interest rates)

6 Outcomes

Our planned **first stage outcome** families are as follows:

1. Training attendance (e.g. binary indicator for training attendance, number/share of women who attend the training)
2. Aadhaar-based payment system (ABPS) enablement (e.g. number of new workers who are ABPS-linked for direct deposit of MGNREGS wages)
3. Training takeaways (e.g. understanding of or interest in MGNREGS)

These first-stage outcomes will help us to assess the quality of the intervention. In the case where we do not find effects on the outcomes below, some of these first stage outcomes may become our primary outcomes.

Our planned **primary outcome** families are as follows:

1. Labor supply of household members

   (a) Related to MGNREGS (e.g. number/share of women (men) making MGNREGS work demands, number of days worked by women (men)), payments made to women (men)

   (b) Paid work in the private labor market (e.g. any work, number of days worked, earnings)

2. Gender-related attitudes and norms (e.g. both own and beliefs about others' attitudes about women and work, the appropriate role of men vs women inside and outside the household)
3. Women's agency (e.g. related to physical mobility, financial decision making)

4. Male backlash behavior: e.g. intimate partner violence and controlling behaviour

Our planned **secondary outcome** families are as follows:

1. Household consumption patterns (e.g. shares of spending on "female" vs "male" goods, on children), as well as levels of debt, savings, and asset ownership

2. Physical and mental health (of women, men), family planning, and measures of subjective well-being.

3. General equilibrium outcomes including the price of labor (wages) for men and women, as well as prices of other inputs to production, the price of consumption goods, and local economic activity

We anticipate estimating **heterogeneous treatment effects** along several dimensions. For outcomes that are available at the GP level, we may examine how effects differ by characteristics including but not limited to previous MGNREGS activity in the GP, share of the population that belongs to a scheduled tribe (ST) group, and previous economic activity like levels of female labor force participation or average household incomes (most likely from the 2011 census). For household- and individual-level outcomes, we may examine how effects differ by characteristics such as baseline gender attitudes, previous work experience, caste, and income. Based on insights from intervention implementation and process monitoring, there may be additional dimensions of heterogeneity that we decide to add to the analysis.
We plan to aggregate outcomes into several indices (e.g. GLS-weighted within-family indices). These indices ensure we do not not overindexing on one or two results.

6.1 Power Calculations

All of the minimum detectable effects (MDEs) reported below assume 80% power to detect effects at the 95% significance level, following the method described in Glennerster and Takavarasha (2013). The MDE depends on the intra-cluster correlation (ICC) that we assume, and so we report the range of MDEs for ICCs used in each calculation.

For the power calculations below, we assume that we will survey a male and female respondent in 20 randomly selected households per GP in each of the 480 GPs where we have conducted the baseline listing during the follow-up survey. In order to be well-powered to detect differences in effects between both treatment and control, as well as within treatment (NGO-led vs government-led training, as well as low-saturation vs high-saturation training), we allocate survey GPs so that there are 10 GPs per block in the training and linking arm (6,000 households) and 6 GPs per block in the control arm (3,600 households). We may adjust the number of households we survey during the follow-up as well as the number of participants per household, at a later stage based on funding constraints.

Effects of linking. To estimate the effects of linking alone, we plan to use the MGNREGS administrative data to which we have been given API access. We are powered to detect effects of 0.09-0.2 SDs, which for an outcome such as the share of women in paid work, corresponds to a 4-9 percentage point (pp) increase on a base of 30% for rural Madhya Pradesh (Periodic Labor Force Survey 2024). This assumes an ICC in the range of 0-0.1 at the block level for GP-level observations. Because we have the same number of blocks assigned to linking only and to linking+training together, these are also the same MDEs for the difference in effects for linking only versus linking and training together.

Effects of linking+training. To assess the effect of linking and training together in comparison with the control group, we plan to use a combination of the MGNREGS administrative data as well as our own surveys. For the administrative data, we are similarly powered to detect the same MDEs for linking and training together, i.e. 0.09-0.2 SDs or 4-9 pp increase in share of women in paid work. For our survey outcomes, we are powered to detect effects of 0.059-0.11 SDs, which corresponds to a 2.7-5.0 pp increase in women's participation in paid work.

NGO-led versus government training. For government-led training, we are powered to detect effects of 0.066-0.11 SDs in comparison with the control group. This corresponds to a 3.0-5.0 pp increase in women's participation in paid work. For NGO-led training, we are powered to detect effects of 0.061-0.10 SDs in comparisons with the control group or in comparison with government-led training, which again corresponds to a 2.8-4.6 pp increase in women's participation in paid work.

Low saturation versus high saturation. We are powered to detect effects of 0.069-0.12 SDs in comparisons between GPs assigned to low saturation NGO-led versus control, or comparisons between high saturation NGO-led versus control. The same holds for comparisons of low-saturation government-led and high-saturation government-led versus control. This corresponds to a 3.2 - 5.5 pp change in the share of women participating in paid work. For comparisons between low- and high-saturation training, we are powered to detect effects of 0.072-0.12 SDs (3.3 - 5.5 pp change in female participation in paid work).

References

Baird, S., Bohren, J. A., McIntosh, C., and Özler, B. (2018). Optimal design of experiments in the presence of interference. *Review of Economics and Statistics*, 100(5):844–860.

DiTraglia, F. J., García-Jimeno, C., O'Keeffe-O'Donovan, R., and Sánchez-Becerra, A. (2023). Identifying causal effects in experiments with spillovers and non-compliance. *Journal of Econometrics*, 235(2):1589–1624.

Field, E., Pande, R., Rigol, N., Schaner, S., and Troyer Moore, C. (2021). On her own account: How strengthening women's financial control impacts labor supply and gender norms. *American Economic Review*, 111(7):2342–2375.