

Pre-Analysis Plan

Effect of Women’s Groups on Social Cohesion in Nigeria

1 Introduction

This pre-analysis plan (PAP) builds on the original Metaketa V Nigeria study, “Women’s Action Committees and Local Services in Nigeria.” It pre-registers the analysis for a new paper on the spillovers of the study interventions on social cohesion, made possible by additional data collection not planned at the time of the original PAP registration. The majority of the data to be analyzed comprise questions to be fielded only during a second endline survey (approximately 2.25 years after the first endline); these are outcomes that may not have changed in the short run but may be impacted over a longer time horizon. Other data we analyze were instead collected during *both* the first and second endlines. These are data pertaining to outcomes that could plausibly have changed in both the short-run and the long-run, and for which we plan to analyze both short-run and long-run effects.

The study evaluates whether training interventions designed to strengthen women’s individual and collective efficacy and political participation influence social cohesion, and whether or not additionally training women’s husbands in male allyship further amplifies those effects. In a randomized controlled trial spanning 450 communities in southwestern Nigeria, we recruited 5,850 politically unaffiliated women into newly formed women’s groups (WGs) trained by ActionAid Nigeria. All communities first participated in an identical burn-in process in which women were recruited into WGs and received a basic civic education session during May-June 2023. Following this, communities were randomly assigned to one of three conditions: a control group that received no further programming; a treatment group (T1) that received additional intensive training in advocacy, leadership, and organizing designed to build women’s individual and collective efficacy (an approximately 3 hour session once monthly over the following five months); and a second treatment group (T2) that received the same training as T1, with parallel allyship training (also an approximately 3 hour session once monthly over the following five months) offered to the male partners of participating women focused on men’s role in supporting women’s empowerment and gender equality.

We carried out a first endline survey approximately 2-3 months after all trainings had concluded, and a second endline survey is planned for approximately 2.25 years after trainings conclude. The

second endline will attempt to resurvey the exact same women interviewed at baseline. Some attrition is expected, and we will assess whether attrition is differential by treatment status following the same approach outlined in the original PAP and its first amendment for the original Metaketa V study. No new data sources will be introduced.

Analysis of the first endline survey yielded encouraging evidence that intensive women’s trainings designed to foster collective action produce positive effects across both political and economic domains. On the political side, both T1 and T2 significantly improved the level and quality of women’s political participation and enhanced their perceptions of the responsiveness of community leaders, with effects appearing to operate through increased self and collective efficacy and a stronger sense of shared purpose among women. On the economic side, the trainings produced meaningful early-stage gains: trained women reported higher income and savings aspirations, as well as a higher likelihood of talking with family members about income-generating ideas, inquiring about loans, starting to save toward a business investment, seeking information about non-farm enterprises, opening bank accounts, borrowing to invest, trying new crops, and applying for government or NGO support. Notably, the marginal contribution of the male allyship training on top of the female collective efficacy training was limited for both political and economic outcomes in the short run.

This study proposes to assess whether the trainings have positive spillovers on social cohesion. There are at least three pathways through which the trainings could promote broader social cohesion. The first pathway operates through group identification and collective efficacy, which were shown to increase due to the training, at least in the short term. By bringing women together in groups, the trainings may build durable social ties and a shared sense of collective strength with other women. The stronger group connection that was forged can function as an informal insurance mechanism, enabling women to pool resources, share information, and provide mutual support in times of need, while also creating a foundation for collective advocacy around service delivery. These same group connections and the collective efficacy that T1 and T2 engendered may yield further dividends in our context, where communities are experiencing farmer-herder conflict and climate-related stress. Specifically, stronger collective efficacy and a heightened sense of shared identity among women may generalize outward over time, improving relations between the trained women and other community members and building the generalized trust that undergirds broader social cohesion. Women who feel more socially connected and collectively capable may also be better positioned to engage constructively

with community institutions and leaders as well as with out-group members.

The second pathway operates through self-efficacy, which was also shown to increase in the short-term in both T1 and T2 communities. The trainings were also designed to raise women’s sense of self-efficacy—their belief that they are capable of taking actions that improve their own and their families’ circumstances. This heightened sense of efficacy may improve social cohesion through two channels. First, women who feel more capable and in control of their circumstances may be less likely to perceive out-group members as economic threats, improving the horizontal dimension of social cohesion over time. In contexts of resource scarcity and recurrent farmer-herder conflict, zero-sum thinking can drive inter-group hostility: when people believe that their well-being depends on fixed communal resources, they may be more likely to view out-group members as competitors. Self-efficacy can disrupt this dynamic by increasing women’s sense of their own capacity to generate economic and social opportunities, thereby reducing the threat perceptions that fuel inter-group tension. Second, women with higher self-efficacy may be more likely to believe that their voices matter and that engaging with community institutions is worthwhile, strengthening the vertical dimension of social cohesion over time (assuming these interactions with leaders prove fruitful).

The third pathway operates through advocacy skills and engagement with community leaders. The T1 and T2 trainings explicitly built women’s capacity to identify community needs, communicate effectively with decision-makers, and take coordinated action. Women who developed these skills may be better equipped to navigate the institutions and relationships that govern access to services and resources. In contexts like rural Nigeria, where state capacity is limited and institutions are often perceived as unresponsive, women may disengage from formal decision-making not only because they are excluded but because they lack the civic knowledge to navigate it effectively. Advocacy training can address both barriers: by demystifying how decisions are made and who holds relevant authority, it may reduce alienation from institutions; and by building collective, coordinated voice, it may increase women’s legibility and salience as constituents, making it more likely that local leaders and service providers respond. Indeed, in the short-run, we observed that T1 and T2 increased women’s perception that political elites and institutions were responsive to them. Over time, greater institutional responsiveness may alter how women relate to political authority, promoting more productive and trusting engagement with the local state and strengthening vertical cohesion and trust (i.e., the relationship between citizens and the state).

We note that men and women may experience the effects of the trainings on social cohesion through different pathways and to different degrees. Women in both T1 and T2 received training that directly built collective efficacy and social ties among women. Men in T2 (but not in T1) received parallel group-based training that may have similarly built social ties among men. However, the content of men’s trainings was distinct. While women’s trainings focused on building their collective efficacy and civic voice, men’s trainings focused on gender norm change and equipping men to support women’s empowerment. The pathways through which this training could affect social cohesion among men may therefore differ from those operating for women—for example, norm change among men (in T1 or T2) over the long run could improve horizontal social cohesion by shifting and improving how men relate to women as community members, while network formation among like-minded men (in T2) could build trust and shared purpose among men. At the same time, T2 men may have encountered resistance or ridicule from other men in their communities who did not share the norms promoted by the training, potentially undermining their social ties with those men, thereby reducing rather than improving social cohesion overall. We therefore treat outcomes from the women’s survey and the men’s survey as co-equal and distinct primary outcomes throughout this PAP and do not assume that effects will necessarily be symmetric across genders or that the mechanisms will be the same.

Along all three pathways, we acknowledge that negative effects are possible. If the trainings triggered backlash at the household or community level—for example, if husbands or community leaders responded to T1 or T2 by restricting women’s mobility or reducing household resource sharing—this could offset or reverse any gains in social cohesion. We pre-register this possibility and will examine evidence of backlash dynamics as a candidate mechanism for negative spillover effects, drawing on data from the second endline surveys of women, men, and community elites (i.e. Baales).

2 Empirical Strategy

This PAP pre-registers the analysis we will conduct using the first and second endline data. We will analyze whether or not improvements materialized at the time of each follow-up survey round (2-3 months and 2.25 years after trainings concluded for the first and second endline surveys, respectively) for outcomes measured in both endlines, and whether there were longer-run impacts on outcomes measured only in the second endline.

For each outcome present at both endlines, we construct a pooled dataset stacking the first and

second endline surveys and estimate a set of interactions between treatment assignment and survey wave. Specifically, we estimate the following specification:

$$Y_{ict} = \alpha + \beta_1 T1_c + \beta_2 T2_c + \beta_3 \text{Second}_t + \beta_4 (T1_c \times \text{Second}_t) + \beta_5 (T2_c \times \text{Second}_t) + \mathbf{X}'_{ict} \gamma + \delta_b + \varepsilon_{ict} \quad (1)$$

where Y_{ict} is the outcome for individual i in community c at wave t ; Second_t is an indicator equal to 1 for the second endline and 0 for the first endline; $\mathbf{X}'_{ict}$ is a vector of baseline covariates when present (this applies to outcomes measured in the women’s surveys only, as there was no baseline survey conducted of men); and δ_b denotes block fixed effects. Standard errors are clustered at the community level. To estimate the short-run effect of T1, we will look at β_1 . To estimate the short-run effect of T2, we will look at β_2 . To estimate the long-run effect of T1, we will look at $\beta_1 + \beta_4$. To estimate the long-run effect of T2, we will look at $\beta_2 + \beta_5$. To assess the change between the short-run and long-run effects for T1 and T2 relative to the placebo group, we will look at β_4 and β_5 , respectively. To test whether the effect of T2 differs from the effect of T1 in the long-run, we will test whether $\beta_1 + \beta_4 = \beta_2 + \beta_5$. To test whether the effect of T2 differs from the effect of T1 in the short-run, we will test whether $\beta_1 = \beta_2$.

For the many outcomes present only at the second endline, we estimate a simpler form of equation (1), featuring only a single round of data, as follows:

$$Y_{ic} = \alpha + \beta_1 T1_c + \beta_2 T2_c + \mathbf{X}'_{ic} \gamma + \delta_b + \varepsilon_{ic} \quad (2)$$

where Y_{ic} is the outcome for individual i in community c ; $\mathbf{X}'_{ic}$ is a vector of baseline covariates when present (again, this applies to outcomes measured in the women’s surveys only, as there was no baseline survey of men); and δ_b denotes block fixed effects.¹ Standard errors are clustered at the community level. The coefficient β_1 captures the long-run effect of T1 relative to the control group, while β_2 captures the long-run effect of T2 relative to the control group. To test whether the effect of T2 differs from the effect of T1 in the long run, we will test the null hypothesis that $\beta_1 = \beta_2$.

This PAP leverages outcomes from both the women’s survey and the men’s survey. The community elite survey is not used for primary hypothesis testing, but we will explore responses in the elite survey to examine whether there are any potential backlash concerns.

¹To improve balance on multiple pre-treatment covariates, we applied Moore’s (2012) “optimal-greedy blocking” algorithm in the randomization procedure. We used several factors to group similar communities into blocks and then assigned treatment at the community level using complete random assignment within the block. The variables for determining the blocks included the number of women attending the burn-in meeting and the presence or quality of the following community-level characteristics: availability of an electricity grid, a sewage system, mobile phone service, a borehole, a post office, a school, a police station, a health clinic, market stalls, banking services, transportation services, roadblocks, and customs checkpoints, as well as the condition of the main road.

3 Primary Outcomes and Hypotheses

This section pre-registers the primary outcomes and hypotheses for our study of the effects of T1 and T2 on social cohesion. For each social cohesion related outcome, we estimate ITT effects of T1 versus control, T2 versus control, and T2 versus T1. Social cohesion comprises the vertical and horizontal relations among members of society and the state, characterized by trust, inclusive identity, and cooperation for the common good (Chan et al. 2006). Within this framework, horizontal social cohesion refers to the quality of citizen-citizen relationships, or the sense of trust, solidarity, and shared purpose that holds communities together across group boundaries. It encompasses both bonding social capital (trust and attachment within groups) and bridging social capital (relationships and networks that extend across groups and societal divisions) (Putnam 2000).

We expect that the trainings will improve horizontal and vertical social cohesion over time among women and men through the three pathways identified in Section 1: (1) group identification and collective efficacy, (2) self-efficacy, and (3) advocacy skills and engagement with community leaders. We also expect men who participated in allyship trainings (T2) to report higher levels of horizontal and vertical trust, both within their communities and across group boundaries (though anticipate it is possible that this could go in the other direction if there are growing tensions between trained and untrained men). We acknowledge that negative effects are possible in the event of household or community-level backlash against women’s empowerment, and that effects may differ across genders.

3.1 Horizontal Social Cohesion

We measure horizontal social cohesion along two dimensions—generalized communal trust and out-group social cohesion, capturing bonding trust within the community and bridging trust across group boundaries, respectively (Putnam 2000). In communities characterized by farmer-herder conflict, the primary axis of inter-group division is between farming communities and Fulani herders, who are nomadic cattle grazers who migrate seasonally in search of pastures. Competition over land and water resources between these two groups has intensified over recent decades due to climate-related stresses, desertification, and population growth, making farmer-herder conflict a central fault line of inter-group tension in this region of Nigeria (George et al. 2022). Fulani herders therefore represent one of the primary out-groups in this context, and improvements in social cohesion with Fulani community members—in the form of greater comfort with everyday social and economic interactions—

are a meaningful indicator of bridging social capital. As baseline surveys indicate that many of our communities of interest are affected by this conflict (Amare et al. 2024)—around half of our sample has experienced violent herder-related conflict events within 15km in the past 2 years (48.6 percent)—Fulani herders are the primary outgroup of interest.

3.1.1 Generalized Communal Trust

First considering generalized measures of horizontal (citizen-citizen) trust and social cohesion, we will construct indicator variables for trusting several groups: other women in the community (for women only), most people (for both men and women), and neighbors (for both men and women, specifically, trusting neighbors to look after their house if they are away). Exact question wordings are as follows:

- “Other women in my community can be trusted” (Response Options: Strongly agree = 1, Somewhat agree = 2, Neither agree nor disagree = 3, Somewhat disagree = 4, Strongly disagree = 5).
- “Most people can be trusted” (Response Options: Strongly agree = 1, Somewhat agree = 2, Somewhat disagree = 3, Strongly disagree = 4).
- “I trust my neighbors to look after my house if I am away” (Response Options: Strongly agree = 1, Somewhat agree = 2, Somewhat disagree = 3, Strongly disagree = 4).
- “Now I’m going to ask you a question about your community - both farmers and herders. To what extent can people in this area be trusted?” (Response Options: Not at all = 1, A little = 2, Somewhat = 3, A lot = 4, A great deal = 5).

Whenever a respondent has as the answer “refused to answer” or “don’t know,” we will treat this as missing data. We will create an equally weighted index of these measures, called our *Generalized Communal Trust Index* (for men, this will exclude the first question in the above list). All raw index components will be standardized by dividing by the standard deviation in the control group. In other words, we will take equally-weighted averages of the standardized index components. We then standardize the average by dividing by the standard deviation in the control group. When there are missing component measures, we will use imputation strategies to minimize the loss of observations.

3.1.2 Out-group Trust

Next considering bridging trust and social cohesion with out-group members, we will consider an additional set of variables. We will test whether treated women and men report stronger social cohesion with Fulani community members, using the following measures:

- “With regards to some from the Fulani ethnic group, would you feel comfortable if they worked in your field?”
- “With regards to some from the Fulani ethnic group, would you feel comfortable paying them to watch their animals?”
- “With regards to some from the Fulani ethnic group, would you feel comfortable trading goods with them?”
- “With regards to some from the Fulani ethnic group, would you feel comfortable sharing a meal with them?”
- “With regards to some from the Fulani ethnic group, would you feel comfortable with a close relative marrying a person of Fulani descent?”
- “How much do you trust people from the Fulani ethnic group?”

Response options for each of the questions are as follows: Not at all = 1, A little = 2, Somewhat = 3, A lot = 4, A great deal = 5. We will construct an equally-weighted index of these measures, called our *Out-Group Trust Index*, as a primary summary outcome, with higher values indicating greater out-group social cohesion. Whenever a respondent has as the answer “refused to answer” or “don’t know,” we will treat this as missing data. Again, all raw index components will be standardized by dividing by the standard deviation in the control group. In other words, we will take equally-weighted averages of the standardized index components. We then standardize the average by dividing by the standard deviation in the control group. When there are missing component measures, we will use imputation strategies to minimize the loss of observations.

3.2 Vertical Social Cohesion

Considering vertical (citizen-state) trust, for both men and women, we will construct indicator variables for trusting the state and authority figures. The specific variables to be used, and the way they will be coded, are outlined below.

- “How much do you trust: National assembly?”
- “How much do you trust: Your local government councilors?”
- “How much do you trust: Your traditional leader?”
- “To what extent do you respect the political institutions of Nigeria?”
- “To what extent do you think that citizens’ basic rights are well protected by the political system of Nigeria?”

- “To what extent do you think that women’s basic rights are well protected by the political system of Nigeria?”
- “To what extent do you feel your leaders are doing the best job possible for Nigerians?” (Response Options: Not at all = 1, A little = 2, Somewhat = 3, A lot = 4, A great deal = 5).

Response options for each of these questions are as follows: Not at all = 1, A little = 2, Somewhat = 3, A lot = 4, A great deal = 5. We will construct an equally-weighted index of these measures, called our *Political Trust Index*, as a primary summary outcome, with higher values indicating greater vertical social cohesion. Whenever a respondent has as the answer “refused to answer” or “don’t know,” we will treat this as missing data. Again, all raw index components will be standardized by dividing by the standard deviation in the control group. In other words, we will take equally-weighted averages of the standardized index components. We then standardize the average by dividing by the standard deviation in the control group. When there are missing component measures, we will use imputation strategies to minimize the loss of observations. We will also look at each component measure separately in case the trainings impact trust on specific actors and institutions and not others.

3.3 Mix of Both Horizontal and Vertical Cohesion

The three remaining outcomes incorporate aspects of horizontal and vertical cohesion that are especially relevant in conflict-affected settings. In each case, the measures draw on elements of both horizontal and vertical social cohesion:

1. Perceived conflict resilience capacity—women’s and men’s sense that they have allies and resources to draw on if farmer-herder conflict affects them—is a direct expression of horizontal and vertical social cohesion. A female community member who feels that her neighbors and fellow women (horizontal ties), and local leaders (vertical ties) would support her in times of conflict has stronger horizontal and vertical ties and a greater sense of solidarity than one who feels isolated and without resources.
2. Similarly, the perceived effectiveness of civil society and local authorities reflects the degree to which respondents see their social environment as resourced and capable of collective action in the face of conflict.
3. Finally, freedom of movement and activity in the face of insecurity captures the degree to which the strength of horizontal and vertical social relations translates into practical behavioral freedom: women and men who feel more socially connected and supported are better able to

go about their daily lives—going to market, fetching water, attending meetings—even in an insecure environment, because they can rely on community members and elites for protection, information, and mutual support. This freedom of movement is also important because it further increases social contact with both one’s own group and out-groups, bolstering social connectivity (i.e., a positive feedback loop).

3.3.1 Conflict Resilience Capacity

We expect the trainings to strengthen both women’s and men’s perceived capacity to navigate farmer-herder conflict, reflecting greater personal agency and social support. We treat this as a dimension of social cohesion: given that participants live in conflict-affected areas, conflict resilience captures the extent to which respondents feel they have allies and community resources to draw on should conflict affect them. The items of interest for this measure are as follows:

- “If I am affected by a farmer-herder conflict, I feel confident that I would know what to do”
- “If I am affected by a farmer-herder conflict, I feel women in my community would be allies that would help me”
- “If I am affected by a farmer-herder conflict, I feel men in my community would be allies that would help me”
- “If I am affected by a farmer-herder conflict, I feel my baale would help me”

Response options for each of these questions are as follows: Strongly agree = 1, Somewhat agree = 2, Somewhat disagree = 3, Strongly disagree = 4. We will construct an equally-weighted index of these measures, called our *Conflict Resilience Capacity Index*, with higher values indicating greater conflict resilience capacity. Whenever a respondent has as the answer “refused to answer” or “don’t know,” we will treat this as missing data. Again, all raw index components will be standardized by dividing by the standard deviation in the control group. In other words, we will take equally-weighted averages of the standardized index components. We then standardize the average by dividing by the standard deviation in the control group. When there are missing component measures, we will use imputation strategies to minimize the loss of observations. We will also look at each component measure separately in case the trainings impact trust on specific actors and institutions and not others.

3.3.2 Perceived Effectiveness Index

We expect that the trainings will increase women’s and men’s perceived effectiveness of power-holding institutions, authority figures, and civil society organizations in alleviating farmer-herder conflict,

reflecting a stronger sense of horizontal and vertical trust.

- The federal government
- The state government
- Male traditional and religious leaders
- Female traditional and religious leaders
- The media
- Male-led civil society organizations
- Female-led civil society organizations

Response options for each of these questions are as follows: Extremely effective = 1, Somewhat effective = 2, Not very effective = 3, Not at all effective = 4. We will reverse-code all outcomes such that a higher overall index score indicates greater perceived effectiveness. We will then construct an equally-weighted index of these measures, called our *Perceived Effectiveness Index*, with higher values indicating greater perceived effectiveness of a range of institutions. Whenever a respondent has as the answer “refused to answer” or “don’t know,” we will treat this as missing data. Again, all raw index components will be standardized by dividing by the standard deviation in the control group. In other words, we will take equally-weighted averages of the standardized index components. We then standardize the average by dividing by the standard deviation in the control group. When there are missing component measures, we will use imputation strategies to minimize the loss of observations.

We will also use principal component analysis to assess whether these measures load clearly onto two factors corresponding to the horizontal and vertical dimensions of social cohesion, which would argue for constructing two separate indices. We do not pre-assign survey measures to one dimension or the other at this stage, as the mapping of each survey item is not straightforward in the Nigerian context—for instance, it is unclear whether media is best understood as a power-wielding institution more characteristic of vertical social cohesion, or as a check on authority more characteristic of the horizontal dimension.

3.3.3 Freedom of Movement and Activity

Finally, a consequence of having stronger horizontal and vertical social relations is that a person is able to go about their daily lives despite a challenging security environment. As noted earlier, our study takes place in a conflict-affected area, and we expect that the trainings will increase women’s and men’s effective freedom of movement and activity in the face of insecurity. We will construct an

equally-weighted index of the following indicator variables (response options are Yes = 1, and No = 0), for both women and men, called our *Freedom of Movement and Activity Index*. Each index component will be reverse-coded so that higher values reflect greater freedom of movement and activity (coded as 0 for “don’t know” responses, as uncertainty about whether insecurity affected one’s activities most plausibly reflects an absence of meaningful constraint; and coded as missing for refusals):

- In the last two years, were there any areas that you avoided going to or through because of insecurity during the NIGHT?
- In the last two years, were there any areas that you avoided going to or through because of insecurity during the DAY?
- In the last two years, did insecurity ever prevent you from: Working when you wanted to work?
- In the last two years, did insecurity ever prevent you from: Going to the market?
- In the last two years, did insecurity ever prevent you from: Getting water to the household?
- In the last two years, did insecurity ever prevent you from: Going to your field/farm?
- In the last two years, did insecurity ever prevent you from: Moving your animals to grazing areas?
- In the last two years, did insecurity ever prevent you from: Moving your animals to water?
- In the last two years, did insecurity ever prevent you from: Earning money or going to work?
- In the last two years, did insecurity ever prevent you from: Going to school?

All reverse-coded index components will be standardized by dividing by the standard deviation in the control group. In other words, we will take equally-weighted averages of the standardized index components. We then standardize the average by dividing by the standard deviation in the control group. When there are missing component measures, we will use imputation strategies to minimize the loss of observations.

4 Other Analytical Considerations

4.1 Non-random attrition

Even after this procedure, we anticipate that some women interviewed at baseline may be missing from the second endline survey. We will first test for differential attrition by estimating regressions of outcome-specific missingness indicators on treatment assignment and baseline covariates, using the same specification as for our primary outcome analyses. If we find no evidence of differential attrition,

we will treat attrition as plausibly random and proceed without adjustment. If differential attrition is detected, we will estimate bounds on the treatment effects using the approach of Lee (2009) or a related bounding method. We will not rely on approaches such as multiple imputation or inverse probability weighting, as these require a conditional ignorability assumption—that, conditional on observed characteristics, missingness is effectively random—which may not be credible in our setting.

4.2 Sense-Making

A number of additional measures will help us interpret any treatment effects on social cohesion — specifically, whether such effects reflect changes in the objective shock and security environment communities face, or changes in how women and men navigate that environment. To distinguish between genuine improvements in participants’ capacity to engage with their communities and institutions versus a simple reduction in exposure to shocks such as conflict and those due to climate change, we will analyze measures capturing shock exposure, severity, and economic well-being. The relevant measures available for this sense-making exercise are detailed in a companion pre-analysis plan, available here: <https://www.socialscienceregistry.org/trials/11695/history/205983>.

References

- Amare, M., Carrillo, L., Kosec, K., and Kyle, J. (2024). Conflict, aspirations, and women’s empowerment: Household survey evidence from farmer-herder conflicts in nigeria.
- Chan, J., To, H.-P., and Chan, E. (2006). Reconsidering social cohesion: Developing a definition and analytical framework for empirical research. *Social Indicators Research*, 75(2):273–302.
- George, J., Adelaja, A., Vaughan, O., and Awokuse, T. (2022). Explaining transhumance-related violence: Fulani ethnic militia in rural nigeria. *Journal of Rural Studies*, 89:275–286.
- Lee, D. S. (2009). Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *The Review of Economic Studies*, 76(3):1071–1102.
- Moore, R. T. (2012). Multivariate continuous blocking to improve political science experiments. *Political Analysis*, 20(4):460–479.
- Putnam, R. D. (2000). *Bowling Alone: The Collapse and Revival of American Community*. Simon & Schuster, New York.