

Journal of Development Economics

Registered Report Stage 1: Proposal (REVISION 2)

Candidate Entry into Local Government

Authors and institutional affiliation: Katherine Casey (Stanford University, corresponding author), Abou Bakarr Kamara (IGC Sierra Leone), Niccolò Meriggi (IGC Sierra Leone), and Andrés Rodríguez (Stanford University)

Date of latest draft: 2/13/2024

Abstract

This proposal explores the governance challenge of how to get high human capital, high integrity, representative citizens to put themselves forward for consideration as political candidates. It evaluates an intervention designed to tackle this challenge in partnership with government and civil society during the June 2023 Local Council Elections in Sierra Leone. The intervention: i) identified, screened, and encouraged high quality potential candidates to enter politics; and ii) shared information about these aspirants with political parties. Field teams identified potential candidates via a combination of structured community nominations and screening on technocratic merit. The initiative was randomly assigned across two levels: first, across 250 rural local government wards (the most local administrative unit); and second, with varied saturation across 92 host constituencies that contain these wards to account for potential spillover effects. Key outcomes of interest concern the number and quality of aspirants, selected candidates, and elected officials.

Keywords: Political economy, economic development, local government, elections, candidates

JEL codes: H07, O17

Study pre-registration: This study was pre-registered in the American Economic Association's registry for randomized controlled trials, under the name "Candidate Entry into Local Government," RCT identifier AEARCTR-0010006, which can be accessed here: <https://www.socialscienceregistry.org/trials/10006>.

Proposed timeline: Implementation of the intervention was completed earlier this year. Data collection for key outcome variables is ongoing and will be gated from the co-author team by our data collection partner, Innovations for Poverty Action (IPA), until February 21, 2024 (please see letter in Appendix A).

1. Introduction

Many low-income countries struggle to improve the technical competence and representativeness of elected government. This is true for our empirical context of Sierra Leone, where government performance is weak (in the 10th percentile of the World Banks' government effectiveness), corruption is endemic (in the 22nd percentile for rule of law), and political candidates are drawn from the socioeconomic, predominantly male elite (e.g. Parliamentary candidates are much wealthier than voters and 89% of them are men, see Casey, Kamara and Meriggi 2021).

The roots of this problem run deep: it involves both the willingness of high-quality potential candidates to step forward; and the mechanisms through which they are successful (or not) in becoming registered candidates and subsequently winning elected seats. Political parties are critical gatekeepers in this process: as in most partisan elections, it is near impossible to win when running as an independent candidate, making both the pool of potential candidates that parties consider and their selection processes key determinants of the ultimate characteristics of candidates and elected officials.

This project focuses on ways to expand and deepen the pool of potential candidates, referred to as "aspirants." We ask how high-quality potential aspirants can be identified, screened, encouraged to run, and brought under consideration by political parties. We focus consideration on two dimensions of aspirant qualifications: technical merit and representation. The former encompasses human capital and integrity; while the latter reflects the will of local voters. Towards these ends, the intervention evaluated, called the "Local Champions Initiative," followed a three-pronged approach. It first deployed field teams who identified locally popular potential aspirants via a structured community nomination process. Second, teams privately screened the community nominees on their technical merits. And third, the initiative generated curated profiles of the highest performing nominees and shared the profiles with political party leaders.

The initiative targeted candidates for local-level elected government, where there are 456 ward-level seats across 22 distinct Local Councils nationwide. Local elections are a useful focus for several reasons. First, the barriers to entry are lower than national office: i) fewer resources are required to register as a candidate and fund a campaign; ii) the job is part-time (so no need to abandon one's professional career to serve); and iii) there are minimal eligibility criteria (e.g. no minimum education). Second, the responsibilities of Councillors involve overseeing local development projects and line ministry activities. Human capital and administrative competence are beneficial for this work, yet it does not require highly specialized expertise; and representation is also useful, to align the provision of basic goods and services with local demand. Third, there

appears to be room for improvement on both dimensions of interest: earlier data indicate that more than one third of Local Councillors have not completed high school, and only 17 percent are women. Lastly, turnover is high (only 12 percent of Councillors elected in 2012 ran for re-election in 2018), suggesting ample room for new entrants.

Our interest in the role of parties as gatekeepers stems from a few factors. Parties face non-trivial information, logistical and resource challenges in identifying and recruiting competent, popular candidates in all 456 wards. Moreover, party leaders speak of misalignment between central party officials (as principals) and the local-level party executives (as agents) whom they rely upon to screen aspirants. A sitting MP went so far as to describe these local delegates as “the most wicked members of our party,” in reference to their responsiveness to money and status in a way that favors incumbents and established “big men,” at the expense of younger, more competent aspirants. This potential misalignment in turn likely deters some high-quality potential aspirants from entering politics. The combination suggests scope to alleviate search and information frictions, while simultaneously widening the gatekeeping of parties to dimensions of competence and representation, in a way that creates space for new entrants into local politics.

2. Related Literature

Improving the quality of personnel working for the state in developing countries is a relatively new and exciting literature (Finan, Olken and Pande 2015). Previous studies evaluate whether better salaries or career prospects attract higher quality applicants for government jobs, and if this also induces less prosocial candidates to show up (Dal Bó, Finan, and Rossi 2013; Ashraf et.al. 2020, Leaver et.al. 2021). For political candidates, work in Pakistan looks at how financial constraints, information on benefits, and social motivations affect candidate entry to village councils (Gulzar and Khan 2018, 2021). Studies in the Philippines show how candidates for local office are more likely to come from well-connected families (Cruz, Labonne, and Querubin 2017); and Banerjee et.al. (2013) show how an intervention that weakens incumbents through voter information can lead to a better and more diverse pool of local candidates.

Most of the work on elected office concerns nonpartisan elections, and we propose to extend consideration to the role of political parties as gatekeepers in this process. Prior work shows the information constraints parties face in identifying popular aspirants for national office (Casey, Kamara, and Meriggi 2021), which may be even more salient in local elections. This relates to research on patronage relationships within political parties

(Jia, Kadamatsu, and Seim 2015; Larreguy, Marshall, and Querubin 2016; Colonnelli, Prem, and Teso 2020).

Our focus on screening potential aspirants on their technical merits aims to enhance the human capital base of the elected government. The nominations and testing process we propose builds on previous work in identifying competent leaders of local development projects (Casey, Glennerster, Miguel and Voors 2023), and extends it to the political arena. This approach complements other attempts to improve the human capital of local governments, like the policy studied by He and Wang (2017), which uses a central assignment of college graduates to village governments.

The candidate nomination and screening processes studied pay special attention to identifying promising female aspirants. Women's underrepresentation in politics may stem from bias at several stages of the selection process. Voters may statistically discriminate against women in places with low historical female representation in politics or have preferences for male leaders that may arise out of gender norms (Beaman et al. 2009). Party elites, who are predominantly male, may prefer to select candidates who are more like them (Niven, 1998; Cheng and Tavits 2011, Karpowitz et al. 2017; Phillips 2021) or only nominate women in areas where the party is likely to lose (Fujiwara, Hilbig and Raffler 2021). The supply of female candidates may also matter if there are fewer women inside traditional pipelines that feed politics (Thomsen and King 2020); household constraints discourage women from running for political office (Bernhard, Shames and Teele 2021); or females have higher risk or competition aversion, which discourages them from running in elections (Kanthak and Woon 2015).

3. Intervention

The “Local Champions Initiative” generated curated lists of high-quality potential aspirants to share with party leaders with encouragement to include them in the pool of individuals under their consideration. This eases the information and logistical constraints on identifying potential candidates while placing no direct obligation on parties. It proceeded in the following stages:

- **Community nominations:** A key aspect of representation is the support of local voters. Field teams elicited the names of popular individuals via private household visits. Enumerators first explained the skills and competencies that are needed to be an effective Local Councillor and then asked respondents to list up to 5 individuals from their ward that they thought would make good candidates. Respondents were then asked to think of up to 5 additional names of female

potential candidates in particular. Teams implemented these private elicitations in around 80 households from up to three of the largest communities in each ward.¹

- ***Technocratic quality screening:*** The research team selected the most popular individuals (top 10 by total nominations across households) generated by the community elicitations and screened these individuals privately using a comprehensive set of questions designed to gauge the nominees' skills and preparation for local public office. The screening instrument includes questions that have been empirically verified in related work to positively correlate with the public spending performance of Members of Parliament (MPs), the quality of local infrastructure grant proposals drafted by community members, and anonymous peer reviews conducted among currently serving members of Local Council. This multi-faceted quality screening narrows the list of potential candidates to those of the highest quality and capability. Nominees were informed about the details of the Local Champions Initiative and given the option to participate (or not). If interested, nominees chose which party (or parties) they were interested in having their profile shared with.
- ***Information provision:*** The two-track process above generated lists of high quality, representative potential candidates to share with political parties for their consideration as they screened aspirants and awarded party symbols. The research team compiled short profiles of the top two community nominees per ward who scored the highest on the technical screening, met eligibility requirements, and were interested in putting themselves forward for consideration as candidates. These profiles were then shared with party officials from the relevant party in each ward. The profiles present information under free disposal, with no obligation on parties to select any of the nominees.

Within the research sample (more details below), some wards received the full treatment, which comprises the nominations and screening procedures, followed by information provision to parties. Other wards received partial treatment, which is the nominations and screening only, with no information provision to parties. All remaining wards form a pure control group, where similar data was collected but no interventions were implemented.

The Political Parties Regulation Commission (PPRC), the government agency with the authority to regulate the conduct of political parties with respect to their members and

¹ More specifically, teams visited the headquarter community of each of the 3 largest chiefdom sections per ward. In wards that contain fewer than 3 sections, teams visited fewer communities but conducted more surveys per community visited, as these tend to be in larger town locations.

the broader electorate, invited all registered parties to join the initiative and associated research. Party leaders decided whether or not to opt into the initiative. Both major parties, the All People's Congress (APC) and the Sierra Leone People's Party (SLPP), opted into the initiative and were assigned equal numbers of treatment and control wards.

4. Institutional Context


The Government of Sierra Leone unexpectedly changed the electoral system shortly before the 2023 election. This policy change required us to make some adjustments to our research design midstream. It is thus useful to describe how the electoral system was originally structured and the key policy changes that were enacted, before delving into the research design and subsequent adjustments.

When we initially designed this experiment, elections for district councils were organized as first-past-the-post, single member jurisdictions at the ward level (the most local administrative unit). This is how district council elections have been run since the end of the civil war (in 2002) and reintroduction of decentralization (in 2004). Elected Local Councillors represent their ward in the relevant district council. The Local Champions Initiative operated in 14 of the 15 district councils nationwide.² The initial random assignment was done across wards inside these councils, treating them as races that were fully independent of one another.

Local Council wards nest neatly inside Parliamentary constituencies, which is the next higher unit of formal state administration. Constituencies in turn nest neatly inside district boundaries (see Figure 1 for an example). Nationwide, there are on average 3.4 wards per constituency. Up until 2023, Parliamentary elections were also organized as first-past-the-post, single member jurisdictions, at the constituency level. Elected MPs represent their constituency in the national Parliament. Thus voters in a given ward elected one person to represent their ward in the district council, and also elected one MP (along with voters in the other wards housed within their constituency) to represent their constituency in Parliament.

² The study sample excludes the 7 urban city councils as well as the Western Area Rural District Council, which is adjacent to the capital and relatively more urban than the rural district councils in the sample.

Figure 1: Map of Nested Wards and Constituencies for Moyamba District



Shortly before the 2023 election, the government changed the electoral system to district-block, proportional representation (PR). For Local Council elections, each party was asked to compile a list of candidates that was twice as long as the number of Council seats (where number of seats is equivalent to number of wards) in a given district. For Parliamentary elections, the parties were asked to compile lists at the district level that were twice as long as the number of constituencies inside the district's boundaries (hence the term "district-block," which for our purposes is equivalent to district).

This research focuses on the Local Council list generation process. For this, there were no clear guidelines about how the electoral policy changes would affect geographic representation: e.g., while parties could simply nominate two candidates per ward for the list, which would align closely with the old system, there was no requirement to do so. Parties were free, for example, to abandon the use of wards as the basic unit of representation, moving instead to a higher electoral administrative unit, like the Parliamentary constituency, or perhaps to an alternative unit, like the traditional chieftaincies. This raises the question of whether the ward-level processes would remain independent of one another, and thus whether the intervention studied could introduce partial interference across units (e.g., the intervention might increase geographic representation of treated wards at the potential expense of control wards). We thus adjusted our random assignment midstream in anticipation of potential interference (explained further below).

In addition, there was little guidance issued about how parties should rank candidates within their district-level lists: e.g., should they rank 1 candidate per ward to cover all wards, and then start again with the second candidate per ward, or cluster all candidates from a particular ward at the top of the list? And, at the same time, the government stipulated a women's empowerment quota, calling for 1 of every 3 candidates to be female. It was not clear how this gender-based allocation relates to the underlying ward structure, either.

In terms of timeline, the incumbent administration first publicly announced their intention to change the electoral system in October 2022. The opposition party immediately condemned the change and challenged the validity of the reform via a series of court cases over the following months. This combination means that all potential aspirants and party officials were both aware of the risk of electoral reform, but also uncertain about the likelihood that it would actually be enacted, before the intervention studied began. The community nominations launched on 20 February 2023, while the challenges were ongoing, and ran until 10 April. In March, it became clear that the opposition was running out of options to block the electoral system change. This is when we, as a research team, decided to implement the saturation design for profile sharing. Candidate applications were due to the parties by 4 April. Nominee profiles were distributed to parties between 11 and 17 April. The deadline for parties to file candidate applications to the Electoral Commission for Sierra Leone (ECSL) was 17 April. ECSL then conducted background checks on the list of candidates submitted by parties. The period for parties to confirm candidate nominations to district offices was 19 to 28 April. The election was held on 24 June 2023. Note that the actions captured by outcome variables (namely, decisions to file candidate applications and party selections of candidates) occurred in April, after it was clear to all involved—community nominees, SQ aspirants and party officials, in both treatment and control areas—that the electoral reform would indeed be implemented.


5. Research Design

5.1 Sample

In order to study the Local Champions Initiative, we first selected an experimental sample of wards within the district councils. We focused the sample on wards whose boundaries did not cut across chiefdom boundaries (chiefdoms are traditional authority divisions that are more salient to citizens than ward boundaries) and excluded wards where the intervention had been piloted. This generated a study sample of 250 wards, located in 92 distinct Parliamentary constituencies.

Since the Local Champions Initiative was done in collaboration with the two major political parties in Sierra Leone, each unit in our sample was available for both parties. Thus, the full sample of treatment assignments covers 500 party-wards and 184 party-constituencies. Nonetheless, as there are clear regional patterns of support for each major party, which map directly to which party community nominees tend to be willing to have their profile shared with, we delineate a stronghold sample (250 wards, 92 constituencies) that assigns wards to the party that historically dominates politics in that locality (defined at the district level). See Figure 2. We will also analyze impacts in weakhold areas, using the national party assignments, as it is possible that the intervention will be particularly impactful where parties are less popular and must work harder to find high quality, interested candidates.

Figure 2: National Party Samples and the Stronghold Sample



5.2 Treatment Assignment


At a high level, there were two rounds of treatment assignment. The first, at the ward-level, randomly allocated the nominations plus screening process (treatment) versus status quo (control group). This was completed under the presumption that the preexisting majoritarian electoral system would remain intact. The second, at the constituency-level, randomly assigned the profile information sharing component among the treated units, which created three groups (full treatment, partial treatment, and pure control). This was completed after it became clear that the switch to a PR system would go through, and in response, uses a saturation design to measure potential local spillovers across wards within a given district-level candidate list.

More specifically, the first randomization assigned approximately two thirds of the 250 sample wards to the nominations/screening procedure and the remaining third to a pure control group, stratified by district and partisan competitiveness (as measured by 2018 electoral returns). In treatment wards, field teams conducted the nominations and screening procedures. Community nominees indicated which party they were potentially interested in having their profile shared with, thereby making each treated unit available for both parties (conditional on finding enough qualified and willing nominees per party).

The second assignment refocuses allocation of the profile information sharing component upwards, to the constituency level, to be able to estimate potential spillover effects across wards (following Baird et al 2018). As both parties indicated that they intended to compile candidate lists that were broadly representative of the full geographic breadth of the district, the constituency seemed like a reasonable level of clustering to contain local spillovers.³ This second stage of randomization grouped wards into their 92 distinct host constituencies and assigned constituencies to three levels of treatment saturation, stratified by district: (a) zero saturation constituencies (no information sharing in any wards), (b) low saturation constituencies (treat $\frac{1}{3}$ of wards with information sharing), and (c) high saturation constituencies (treat $\frac{2}{3}$ of wards). The assignment of profile sharing was done independently for the two parties. Figure 3 details the full research design for the stronghold sample.

³ While nationwide, there are on average 3.4 wards per constituency, in our sample the average is 2.7 wards per constituency as we drop wards that cross chiefdom boundaries.

Figure 3 – Experimental Design for the Stronghold Sample




Notes: This figure displays the random assignments undergirding the research design. Brackets at the top of the figure delineate the stronghold sample, which assigns wards and constituencies to the locally dominant party based on 2018 electoral data measured at the district level. The lottery on the left-hand side indicates ward-level random assignments, which first allocated wards to pure control versus the nominations and screening treatment, and second allocated treated wards to profile information sharing versus no sharing. The lottery at the top of the figure allocates constituencies to three levels of saturation for the information-sharing component. The grid in the middle displays the number of wards in each of the resulting assignment cells.

Note that the constituency-level assignment was blind to the results of the nomination process, so was not conditional on the field team finding any minimum number of high-quality nominees who were willing to participate. Expecting that some wards would have an insufficient number of willing nominees, we assigned more constituencies to the high saturation arm than in the low saturation arm, and more to the low saturation

arm than to pure controls. This resulted in a division of 27% of constituencies in pure control, 33% in low saturation, and 40% in high saturation.

Figure 4 displays the distribution of *effective* saturation levels, after accounting for how many wards in a given constituency were available for the information treatment (e.g., if a constituency contains only a single ward where nominations were held, high levels of effective saturation are not possible).⁴ Table 1 shows that the constituency-level randomization achieved reasonable balance on measures of population, population density, electoral outcomes in the 2018 elections, and number of wards per constituency.

Figure 4: Effective Saturation Levels Across Constituencies



Notes: This figure plots the distribution of effective saturation levels of full treatment wards across constituencies. The X-axis shows the share of wards located within a given constituency that were assigned to full treatment, i.e., assigned to both the nomination and screening process in the first round of ward-level randomization and then assigned to the information-sharing treatment in the second round of ward-level randomization. The group of zero saturation constituencies are omitted from the figure as they contain no full treatment wards (saturation share = 0.0).

⁴ Another restriction on high saturation is that effective saturation is measured for all wards within a constituency, even those that cross chiefdom boundaries and are thus out of the research sample.

Table 1: Constituency-level Measures of Balance Across Saturation Assignment

VARIABLES	(1) Population	(2) Pop. Density	(3) Turnout	(4) Win Margin	(5) Num. Wards
Low Saturation	-1,261 (1,417)	17.01 (18.48)	-968.0 (768.9)	-0.0140 (0.0324)	-0.0267 (0.104)
High Saturation	-2,314* (1,379)	4.337 (17.99)	-509.7 (748.5)	-0.00277 (0.0315)	-0.0662 (0.101)
Constant	54,371*** (1,049)	96.68*** (13.68)	16,754*** (569.3)	0.630*** (0.0240)	3.568*** (0.0768)
Observations	184	184	184	184	184
District FE	YES	YES	YES	YES	YES
Ftest pvalue	0.246	0.613	0.453	0.892	0.797

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Notes: This table tests for balance on observable characteristics across constituencies assigned to the three levels of full treatment saturation. Assignments for the two major political parties were conducted independently, which are pooled in this analysis, yielding a sample of 184 party-constituency observations. Population data comes from the 2015 census, and population density is measured in people per square kilometer. Turnout and Win margin are based on the presidential race of 2018.

Once the constituency-level saturation was assigned, we randomly assigned the required number of nomination wards within each party-constituency to the information sharing condition, or full treatment. Given the effective saturation levels shown above, we end up with 40% of the party-wards in our sample being fully treated. As Figure 3 shows, for the stronghold sample this is 99 fully treated wards, 88 partially treated wards (only nominations), and 63 pure control wards. Table 2 shows that the assignments to partial and full treatment are reasonably balanced across population, population density, and 2018 electoral outcomes at the ward-level.

Table 2: Ward-level Measures of Balance Across Treatment Assignment

VARIABLES	(1) Population	(2) PopDensity	(3) Turnout	(4) Win Margin
Partial Treatment	529.9* (272.3)	-105.6 (101.1)	-45.87 (124.3)	-0.0124 (0.0136)
Full Treatment	254.1 (242.5)	-109.0 (89.98)	-135.3 (119.7)	-0.00472 (0.0131)
Constant	14,752*** (193.5)	293.5*** (71.79)	4,649*** (74.52)	0.624*** (0.00814)
Observations	500	500	650	650
Const FE	YES	YES	YES	YES
Ftest pvalue	0.143	0.458	0.518	0.655

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table tests for balance on observable characteristics across wards assigned to treatment arms. Assignments for the two major political parties were conducted independently, which are pooled in this analysis, yielding a sample of 500 party-ward observations. Population data comes from the 2015 census. Turnout and Win margin are based on the presidential race of 2018.

6. Data

We are in process of collecting data from two types of sources: (i) surveys that we designed to capture characteristics of potential candidates related to quality, representation and demographics; and (ii) administrative records from the political parties and the Electoral Commission for Sierra Leone (ECSL).

6.1 Individual Survey Data

We designed a comprehensive survey instrument to collect information on politicians and potential candidates. We use this instrument to collect information from three groups of respondents: i) sitting incumbent Local Councillors (LC); ii) community nominees (CN) identified by the ward-level elicitation process; and iii) individuals outside the Local Champions initiative who filed a candidate application with the parties, whom we refer to as status quo (SQ) aspirants. Data collection for LC and CN is complete, data collection for SQ is ongoing.

The survey aims to capture observable indicators of politician quality. As quality is an amorphous and multi-dimensional object, we selected measures that we have empirically validated in related work to positively correlate with observable outputs of performance. Specifically, we pulled measures from three separate validation exercises.

We first used data from Casey, Meriggi and Kamara (2021) to select measures of Parliamentary candidates that positively predict the subsequent public spending of elected MPs, as verified by field audits. Second, we select measures from a screening instrument used in Casey, Glennerster, Miguel and Voors (2023) to identify high skill members of rural communities in Sierra Leone. These individuals were involved in drafting proposals for local infrastructure that were submitted to a district-level grants competition. We choose measures that positively correlate with the quality of the submitted proposals and the likelihood that their proposal was selected as a winning grant. Third, during the piloting activity for this current research program, we surveyed all sitting Local Councillors. We collected both self-reported performance measures (like amount of development spending in their wards and number of visits with their constituents) and anonymous peer ratings of competence in office and lack of corruption. We select measures that positively predict either type of outcome.

We compile these verified measures into a screening instrument. We use this instrument to i) rank community nominees on their technical merits as part the treatment intervention; and ii) to compare the merits of individuals in different pools (e.g. how do the nominees compare to SQ aspirants, how do they compare to incumbent Councillors?). In order to avoid overfitting our prediction of quality, particularly with variables that are biased towards groups over-represented among elected politicians (namely older, elite males), we include additional measures that are widely perceived to be associated with quality and human capital, like years of education, professional experience, and IQ proxies. We roll the full set of indicators up into an equally weighted index, following Kling, Liebman and Katz (2007). See complete list of indicators in Section 7.1 below.

6.2 Community Elicitation Data

We measure popularity among local constituents—as one dimension of representation—directly from the frequency with which individual names were put forward by household respondents during the structured community elicitation process. This dataset by construction includes all nominees, but will also include some SQ aspirants and incumbent Councillors (both of which will need to be matched on name and home ward location).

6.3 Administrative Data

To study the effects of the intervention on entry at multiple points along the political process, we use administrative records from parties and government.

Administrative data from the political parties includes: i) lists of all candidate applications received in the 14 district councils studied (i.e. the aspirant pool); and ii) official ranked lists of selected candidates submitted to ECSL (i.e. those put forward for election). The key administrative data from ECSL is the official list of elected Local Councillors.

As most of these lists do not include ward identifiers, we plan to complement the lists with information gathered via phone calls to party officials, followed by surveys of applicants, candidates and elected Councillors. We will validate individuals' reported home ward by collecting their voter registration location as part of the individual-level quality surveys (using the same instrument as we used for community nominees).

6.4 Data Gatekeeping

In order to estimate treatment effects on key outcomes of interest (as outlined in Sections 7.2 and 7.3 below), we require data collection on two outstanding pieces of information: (i) ward identifiers of all individuals (aspirants, candidates and elected officials); and (ii) survey data on all SQ aspirants, candidates and elected officials within our study sample of wards. We have agreed to a field plan with our partner, Innovations for Poverty Action (IPA), to collect this data over the coming months. IPA has agreed to gatekeep this data away from the PI team until the review process for this registered report is completed (see letter in Appendix A).

Note that we earlier collected ward identifiers and quality data on all SQ aspirants for a small subsample of 50 party-wards (which is 10 percent of the 500 party-wards in the full research sample). We commit to showing robustness checks in the ultimate analysis that excludes these early identified party-wards for all outcomes of interest.

7. Analysis

We divide our proposed analysis into four parts: i) descriptive analysis that gauges whether the intervention did what it was intended to do; ii) causal analysis of the impacts of the intervention on candidate entry into the aspirant pool; iii) causal analysis of party responsiveness to the intervention and related impacts on selected candidates; and iv)

estimation of potential spillover effects. We refer to descriptive analysis as questions (e.g., Section 7.1, Q1), and causal estimates as hypotheses (e.g., Section 7.2, H1).

In what follows, i denotes individual, w denotes ward, and c denotes constituency. Throughout the analysis we will consider a small number of dimensions for potential heterogeneous effects, namely: i) partisan stronghold, swing and weakhold areas; ii) gender; and iii) party. For all estimates, we will include a robustness check that controls for the measures used in balance check Tables 1 and 2.

7.1: Descriptive analysis of whether the intervention did what it was designed to do

In the spirit of a “first stage” analysis, this section sets out a series of empirical tests to evaluate whether the design and implementation of the Local Champions initiative established a reasonable foundation for subsequent analysis of causal impacts.

Q1: Does the intervention identify high quality potential candidates?

Regression analysis will compare the characteristics of community nominees (CN) to Local Councillors (LC) and status quo aspirants (SQ), using the following specification:

$$Q_{iw} = \beta_0 + \beta_1 CN_{iw} + \beta_2 LC_{iw} + s_w + \varepsilon_{iw} \quad (1)$$

where Q is an index of quality measures, s is a vector of randomization strata (by district and competitiveness), and ε is an idiosyncratic error term. Construction of the index follows Kling, Liebman and Katz (2007) by: orienting each variable so that higher values indicate “better” outcomes; translating variables into standard deviation units by subtracting the mean and dividing by the standard deviation of the omitted group, which is SQ aspirants; imputing missing values at the respective respondent group mean; and giving equal weight to each variable in compiling the index. The estimand of interest is β_1 , where $\beta_1 \geq 0$ indicates that the Local Champions intervention identified nominees who were at least as high quality as SQ aspirants. This analysis will only be possible to conduct once SQ data collection is completed. We will also test whether $\beta_1 \geq \beta_2$, which would indicate that the intervention identified nominees that were at least as high quality as sitting Councillors.

The sample of wards includes the 187 wards in the partial and full treatment groups where nominations and screening were conducted. We will first estimate Equation (1) for the top picks per ward, or those who scored the highest on the screening, met eligibility criteria and were willing to participate (i.e., those who would have cleared the bar to have

their profile shared with parties). We will then compare this to estimates for the two highest scoring nominees per ward, regardless of their willingness to participate, to investigate the degree to which talent is lost from high-quality individuals self-screening out of politics.

For this and other analyses of quality, our primary specification will be for the index. We will also present, as secondary analysis, treatment effect estimates for each individual component measure of the index. Variables included in Q are as follows, where we flag indicators that predict performance outcomes in any of our three validation exercises (described above), with * denoting an indicator that predicts the performance of sitting Local Councillors, † denoting one that predicts the public spending performance of elected MPs, and + denoting one that predicts the quality of local development grant proposals developed by community members:

1. Human capital:
 - a. Years of education
 - b. Indicator for literacy and numeracy*,+
 - c. Score in Ravens Test
 - d. Score in digit memorization test
2. Work experience:
 - a. Wage in previous job
 - b. Indicator of experience in sectors like health, education, NGOs*
3. Local experience
 - a. Number of ward development projects involved in before holding elected office*
 - b. Number of ward community activities involved in before holding elected office
 - c. Number of leadership roles in ward before holding elected office*
 - d. Indicator for knowledge of NGOs working in ward*
4. Managerial capital
 - a. Score in project proposal exercise+
 - b. Indicator for experience managing projects (e.g. budgets, supervising people)
5. Conscientiousness – collected only for individuals in SQ and CN (not LC)
 - a. Returned extra transport subsidy†

Q2: Does the intervention identify popular potential candidates?

While the answer to this question is largely “yes, by design,” we will produce descriptive analysis of the average number (and percentage) of households who nominated a particular individual in *CN*, and compare these to rates for those in *SQ* and *LC*, by estimating Equation (1) for popularity outcomes.

Q3: Does the intervention identify new entrants into politics?

Here we will estimate a version of Equation (1) that uses measures of tenure of participation in politics as outcome variables. The primary measure is the likelihood that the individual has previously run for Local Council, where $\beta_1 \leq 0$, would suggest that the intervention found people (in *CN*) who were less likely to be under the consideration of parties in absence of the intervention (in *SQ* or *LC*). We will also run this for measures of connectedness (e.g., family connection to incumbent politicians or prominent party members) and measures associated with under-representation in politics (e.g., female, lower wealth).

Q4. Do the nominees enter politics?

For this we will produce summary statistics on the frequency with which nominees (i) agreed to have their information shared with parties, and (ii) filed an application with a party. We will further break these out by the quality scores of nominees. For (i), willingness to share information with parties (if selected for the initiative) was collected at the end of the screening instrument for all nominees. For (ii), while all nominees were free to file an application, the top picks were given encouragement to do so as part of the treatment. Specifically, the research team called top picks to inform them that they were indeed top picks, that their profiles would be shared with parties, and encouraged them to apply before the relevant party's application deadline (which was announced at short notice by parties). We will run and report compliance statistics for these encouragement phone calls. Heterogeneity by party and gender is especially relevant for (ii), as the parties charged different amounts and one waived fees for female candidates. We will also explore other observable characteristics that predict a positive response for these measures.

Q5. Did party leaders accept the nominee profiles?

The purpose of this measure is to evaluate party leader compliance with the information sharing component of the intervention. For the sample of full treatment wards, we will report the frequency with which a party executive signed the ledger to confirm that he/she had received the profiles. Here we will also report the "hit rate" of how many of the full treatment wards had at least one nominee who cleared all steps involved in having their details shared with parties.

7.2: Causal analysis of impacts on aspirant entry

This section estimates potential impacts of the intervention on the aspirant pool, seeking to understand whether it was effective in increasing the size and quality of people available for parties to consider in advance of their candidate selection processes. Analyses will be run at the ward-level (which will only be possible once the ward identification process has been completed). Note that potential spillovers are not relevant here, as there were no limits on how many individuals could apply to the parties, so increasing applications from one ward does not detract from applications being filed from any other ward.

H1: The intervention expands the aspirant pool

For this first hypothesis we are interested in two measures related to geographic representation: (i) a binary indicator of whether there are any aspirants available for consideration from a given ward; and (ii) its continuous counterpart that measures the size of the aspirant pool per ward. We define an aspirant as someone who is registered to vote in a given ward and has filed a candidate application with a party. Outcome data comes from the district-level lists of applicants shared by the political parties, with subsequent validation checks on ward location as described above.

For the stronghold sample (of 250 wards), we will estimate:

$$Y_w = \beta_0 + \beta_1 \text{Nominations}_w + \beta_2 \text{ShareProfiles}_w + \mathbf{s}_w + \varepsilon_w \quad (2)$$

where outcome Y captures the ward-level aspirant pool, $\mathbf{s}$ is a vector of randomization strata, and ε is an idiosyncratic error term. Coefficient estimates of interest are $\beta_1 \geq 0$, which would indicate that nominations plus screening expanded the aspirant pool, and $\beta_2 \geq 0$, which would suggest that the encouragement aspect of the profile component had an additional marginal effect on expanding the pool. For β_2 , note that applications were due before the profiles were delivered to parties, so it captures encouragement on the nominees' side (e.g., from knowing they were a top pick and that their profile would be shared), as opposed to any recruitment action taken by party leaders. For β_1 , note that all nominees were informed during the screening process of how many households in their community had nominated them and were read a quotation from one about why the respondent thought they would be a good candidate. This kind of feedback could spur prosocial motivation to enter politics (see for example, Gulzar and Khan 2023).

The counterpart regression for the full sample of 500 ward-party observations is:

$$Y_{wp} = \beta_0 + \beta_1 \text{Nominations}_{wp} + \beta_2 \text{ShareProfiles}_{wp} + \mathbf{s}_w + p_w + \varepsilon_{wp} \quad (3)$$

which indexes ward-level observations by party p and adds a party fixed effect p to Equation (2). Coefficients of interest remain as defined above.

H2: The intervention enhances the quality of the aspirant pool

The second hypothesis tests whether aspirant entry induced by the intervention improves the quality of applicants available for parties to consider in a given ward. To do so, we replace the outcome in Equation (2) with two measures of quality: the average quality of the pool and the maximum quality observed in the pool. We construct a quality index as in Equation (1), however now standardizing with respect to the mean and standard deviation of aspirant quality in the control wards. Coefficients of interest remain as in Equation (2) and we will run for both the stronghold sample of wards and full sample of ward-party observations. Analysis for impacts on quality cannot be run until both the ward identification process and the SQ survey data collection are completed.

7.3: Causal analysis of party responsiveness to the intervention

This section explores whether the intervention affected party leaders' selection of candidates. It considers both geographic representation and the quality of selected candidates. It further investigates whether any potential effects on candidate selection flow through to the set of elected Local Councillors. Running this analysis (again) requires completing the ward identification process.

H3: The intervention enhances geographic representation on the candidate lists

To test this hypothesis, we estimate Equation (2) for three distinct outcome variables: i) a binary indicator for the presence of any candidate from a given ward on the relevant district-level list of candidates submitted by parties to the ECSL; ii) the continuous counterpart or number of candidates on the list from a given ward; and iii) the (inverse) rank of candidates from a given ward on the list, as lower number ranks have a higher likelihood of getting elected. The key coefficient of interest is now $\beta_2 \geq 0$, which would suggest that sharing profiles with party leaders increased the likelihood that they selected a candidate from treated wards. We will continue to also estimate $\beta_1 \geq 0$, which would indicate that nominations plus screening increased the likelihood that the parties selected a candidate from treated wards. These regressions will be estimated for both the stronghold sample of wards and full sample of ward-parties. Note that if the parties

adhered to the old electoral system by nominating exactly two candidates per ward, then estimates for (i) and (ii) will be moot.

Note that this analysis does not depend on whether a selected candidate for a treated ward is one of the nominees identified by the intervention or not. One could imagine, for example, that the presence of a high-quality nominee in the aspirant pool induces parties to select a higher quality status quo nominee than they would have otherwise. To unpack these multiple channels, we will estimate and report the share of top nominees who became candidates, for both nominations only and profile sharing.

H4: The intervention increased the quality of selected candidates

This hypothesis test replicates the analysis outlined for H2 above but restricts attention to the sample of selected candidates. As such, it again replaces the outcome in Equation (2) with two measures of quality: the average quality of the pool and the maximum quality observed in the pool, while redefining the pool from aspirants to selected candidates representing a given ward. We standardize the quality index with respect to the mean and standard deviation of candidate quality selected to represent control wards. Coefficient estimates of interest are $\beta_2 \geq 0$, which would suggest that sharing profiles with party leaders increased the likelihood that they selected a high-quality candidate from treated wards; and $\beta_1 \geq 0$, which would indicate that nominations plus screening increased the likelihood that the parties selected a high-quality candidate even in the absence of information provision. This analysis also depends on completing the SQ survey data collection.

H5: The intervention enhanced geographic representation in the elected Local Council

Testing this hypothesis will take a similar form to that for H3 above yet restricting attention to those candidates at the top of party lists who won a seat in the elected Council. We will estimate the regression for both the presence and number of elected representatives outcomes. Coefficient estimates of interest are $\beta_2 \geq 0$, which would suggest that sharing profiles with party leaders increased the likelihood that a treated ward gained representation in Council, and $\beta_1 \geq 0$, which would indicate that nominations plus screening increased the likelihood a treated ward is represented in Council in the absence of information provision. These regressions will be estimated for both the stronghold sample of wards and full sample of ward-parties.

H6: The intervention increased the quality of elected Councillors

This hypothesis test replicates the analysis outlined for H4 above but restricts attention to the quality of elected Councillors. As such, it again replaces the outcome in Equation (2) with the average and the maximum quality of elected Councillors. We will standardize the quality index with respect to the mean and standard deviation of elected Councillor quality representing control wards. Coefficient estimates of interest are $\beta_2 \geq 0$, which would suggest that sharing profiles increased the likelihood that treated wards were represented by high quality elected Councillors, and $\beta_1 \geq 0$, which would indicate that nominations plus screening improved the quality of elected officials representing treated wards without information provision.

7.4: Estimation of potential spillover effects

A key issue with the switch from ward-level majoritarian elections to district-level PR is it creates scope for potential interference between treatment and control wards when analyzing how parties responded to the receipt of nominee profiles, as in essence all wards are competing for representation on fixed-length district-level candidate lists. The presence of interference would introduce bias in the estimation of treatment effects for H3 through H6 as outlined above.

To see how bias could arise, consider first estimates for geographic representation in H3. Suppose that in the absence of treatment, parties would select exactly 2 candidates from each ward to include on the district list. Suppose further that when they received a nominee profile from a treated ward, they added that nominee to the list and removed a candidate from a control ward to maintain the specified length of the list (which recall is two times the number of wards in the district). In this scenario, the coefficient estimate for β_2 under H3 would be biased upwards, as it combines the positive effect observed for treated wards with the negative spillover effect on control wards.

On the other hand, partial interference could lead to downward bias for estimates of treatment effects on candidate quality in H4. Extending the scenario above, suppose that in the absence of treatment, parties select one high quality ($Q = 1$) and one low quality ($Q = 0$) candidate per ward. After seeing the profile, they add a second high quality candidate (the nominee) to a treated ward and drop a low quality candidate from a control ward. As a result, average quality in treated wards becomes $2/3$ while in control wards it is equal to 1 . In this case the treatment effect estimate would be biased downwards, and would be negative in sign, even though the intervention in fact

succeeded in improving the quality of the aggregate pool (which under equal allocation of wards to treatment and control would have increased from $\frac{1}{2}$ to $\frac{3}{4}$).

To explore the relevance of potential interference, we will leverage the saturation design at the constituency level. We illustrate the method for H3 and the outcome of number of selected candidates per ward on the district list. If we find no evidence of spillovers here, there is little cause for concern about bias in estimates for any of the other three hypotheses in Section 7.2. Moreover, if the candidate lists submitted by parties to ECSL adhere closely to two candidates per ward, spillovers are unlikely to matter. Yet if we do find evidence of interference, a similar specification can be applied to all four hypotheses.

To do so, we follow the approach in Baird et al (2018) where our randomized saturation design can be characterized by saturation vector $\mathbf{I} = \{0, 0.33, 0.67\}$ and cluster share vector $f = \{0.27, 0.33, 0.40\}$. Let indicator $S_{wc} = 1$ indicate a ward that is a within-cluster control, i.e., a ward in a cluster with positive treatment saturation that did not have any nominee profile shared with party leaders. Let $L_c = 1$ indicate a constituency cluster c allocated to low saturation treatment ($p = 0.33$), and $H_c = 1$ indicate a cluster allocated to high saturation treatment ($p = 0.67$). We will estimate:

$$Y_{wcs} = \beta_0 + \beta_{1L} \text{ShareProfiles}_{wcs} L_c + \beta_{1H} \text{ShareProfiles}_{wcs} H_c + \gamma_{1L} S_{wc} L_c + \gamma_{1H} S_{wc} H_c + \alpha_s + \varepsilon_{wcs} \quad (4)$$

where $\gamma_{1L} \neq 0$ or $\gamma_{1H} \neq 0$ would indicate the presence of interference whereby the treatment of wards within a cluster has a spillover effect on control wards within the same cluster; $\beta_{1L} \geq 0$, $\beta_{1H} \geq 0$ would indicate positive intention-to-treat effects for treated wards in low, high saturation clusters respectively; and $\gamma_{1L} < \gamma_{1H}$ would indicate the spillover effects are increasing with saturation. If the spillover effects are non-zero, the key estimate of interest for policy makers is the weighted averages $0.33\beta_{1L} + (1 - 0.33)\gamma_{1L}$ and $0.67\beta_{1H} + (1 - 0.67)\gamma_{1H}$ which are the total causal effect of the profile sharing for each saturation level.

7.5 Adjustments for Multiple Inference

Concerns about multiple inference and the risk of Type I errors arise when estimating treatment effects for multiple outcomes, subgroups and/or interventions on the same dataset (Viviano, Wuthrich and Niehaus 2023). We will use false discovery rate (FDR) adjustments to mitigate this risk (following Anderson 2008 and Benjamini, Krieger and Yekutieli 2006). This section designates a core set of estimates that constitute our

primary analysis and accompanying adjustment strategy for that set. All other estimates are considered secondary, or more exploratory in nature, which we will leave unadjusted.

A note about the structure of the data is useful here. Notice that the causal analyses presented in Sections 7.2 and 7.3 estimate treatment effects for the same 4 to 5 outcome variables on a series of distinct subsamples of the same “master” dataset. Here we can think of the largest dataset, which is the pool of applicants for the full sample of party-wards, as the master dataset. This is used for hypothesis H1 and H2. The pool of candidates (for H3 and H4), is a selective subset of the aspirant pool; and the pool of elected Councillors (for H5 and H6) is in turn a selective subset of the candidate pool. Across all of these data slices, the stronghold sample of 250 wards is a strict subset of the full sample of 500 party-wards.

Our primary analysis will include estimates for: one outcome each for measures of representation and quality; in two samples, namely aspirants and selected candidates; under two treatment interventions, namely nominations/screening and information sharing; all using the full sample dataset. Thus, we will apply FDR adjustments to the p -values associated with estimates of β_1 and β_2 in Equation (3) for the number of aspirants in a ward (H1), the average quality of the aspirant pool (H2), the number of selected candidates for a ward (H3), and the average quality of selected candidates for a ward (H4). This implies adjustments over a group of 8 estimates in the full sample dataset.

We will use our preliminary data analysis to determine whether either of two refinements are empirically justified. First, if descriptive analysis outlined in Section 7.1 provides strong evidence that the intervention worked as intended, we will use one-sided tests in the direction of positive treatment effects. Second, if analysis in Section 7.4 uncovers significant spillover effects, we will redirect FDR corrections for H3 and H4 to the more policy-relevant weighted coefficient combinations that recover the total causal effect of the intervention.

7.6 Power Calculations

The table below presents minimum detectable effect (MDE) sizes for estimating impacts on outcomes for the stronghold sample of 250 study wards and on the full sample of 500 party-wards, using sources of data that pre-date project implementation. These calculations are done with a power of 80% and significance of 5%.

Table 3: Power Calculations

Row	Outcome variable	MDE (N=250)	MDE (N=500)
1	Number of aspirants per party-ward, mean = 4.875	0.89	0.63
2	Years of education; mean = 12.77	0.74	0.52
3	Managerial capital test score, mean = 73.96	5.68	4.04
4	Conscientiousness: Likelihood returns transport subsidy; mean = 0.46	0.17	0.12

Notes: data sources for these power calculations are: i) row 1 is from the 2018 Parliamentary elections analyzed in Casey, Kamara and Meriggi 2021; ii) row 2 is from a 2004 survey of Local Councillors; iii) row 3 is from a community development nomination and screening process analyzed in Casey et al 2023; and iv) row 4 is from the 2018 Parliamentary elections analyzed in Casey, Kamara and Meriggi 2021.

Following Bohren et.al (2016), we use the same outcomes above to calculate power for the saturation design. These are for a power of 80%, significance of 5%, and assuming a within cluster correlation of our outcomes of interest of 0.2. The tables below show the MDEs for all four coefficients of interest estimated in our spillover specification (Equation 4). These coefficients represent the direct effect of the profile sharing treatment for low and high saturation constituencies (T Low and T High), as well as the spillover effects on control wards in low and high saturation constituencies (S Low and L High).

Table 4: Saturation Design Power Calculations

	MDE (250 wards, 92 constituencies)					MDE (500 party-wards, 184 constituency-party units)			
Outcome	T Low	S Low	T High	S High		T Low	S Low	T High	S High
Number of Aspirants	1.57	1.28	1.25	1.38		1.18	0.96	0.93	1.03
Years of Education	1.31	1.07	1.05	1.15		0.98	0.80	0.78	0.86
Managerial Capital Test	10.11	8.22	8.06	8.85		7.58	6.16	6.00	6.64
Returns Transport Subsidy	0.32	0.26	0.25	0.28		0.24	0.19	0.19	0.21

8. Administrative information

Funding: This work was supported by the Governance Initiative at JPAL; the King Center for Global Development at Stanford University; and the International Growth Centre at LSE.

Institutional Review Board (ethics approval): This project has IRB approval from Stanford University (Protocol # 64314) and Sierra Leone Ethics and Scientific Review Committee (SLESRC Protocol # 015/09/2022).

Declaration of interest: The authors have no competing interests to declare.

Acknowledgments: We are grateful the collaboration of our partners at Innovation for Poverty Action (IPA), the Political Parties Regulation Commission of Sierra Leone, the National Electoral Commission of Sierra Leone, the All People's Congress, the Sierra Leone People's Party, and Talking Drum Studios. For excellent field and research assistance, we thank Naasey Kanko Arthur, Kevin Dunne, Pablo Guzmán Lizardo, Kiet Le, and Monica Shandal. We are grateful for detailed comments from Fred Finan.

References

- Anderson, Michael L. (2008). "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American Statistical Association*. 103(484): 1481-1495.
- Ashraf, N., Bandiera, O., Davenport, E., & Lee, S. S. (2020). "Losing prosociality in the quest for talent? Sorting, selection, and productivity in the delivery of public services." *American Economic Review*, 110(5), 1355-94.
- Baird, Sarah, J. Aislinn Bohren, Craig McIntosh, Berk Özler. 2018. "Optimal Design of Experiments in the Presence of Interference." *The Review of Economics and Statistics*, 100 (5): 844–860.
- Banerjee, A., Duflo, E., Imbert, C., & Pande, R. (2013). "Entry, Exit and Candidate Selection: Evidence from India." Working Paper.
- Beaman, Lori, Raghabendra Chattopadhyay, Esther Duflo, Rohini Pande, and Petia Topalova. "Powerful women: does exposure reduce bias?." *The Quarterly journal of economics* 124, no. 4 (2009): 1497-1540.
- Benjamini, Y., Krieger, A., and Yekutieli, D. (2006), "Adaptive Linear Step-Up Procedures That Control the False Discovery Rate," *Biometrika*, 93, 491–507.
- Bernhard, Rachel, Shauna Shames, and Dawn Langan Teele. "To Emerge? Breadwinning, Motherhood, and Women's Decisions to Run for Office." *American Political Science Review* 115, no. 2 (2021): 379–94.
- Bohren, A., P. Staples, S. Baird, C. McIntosh, and B. Özler, (2016). Power Calculation Software for Randomized Saturation Experiments, Version 1.0. Available from <http://pdel.ucsd.edu/tools/index.html>).
- Casey, Katherine, Abou Bakarr Kamara, and Niccoló F. Meriggi. 2021. "An Experiment in Candidate Selection." *American Economic Review*, 111 (5): 1575-1612.
- Casey, Katherine, Rachel Glennerster, Edward Miguel, Maarten Voors. 2023. "Skill Versus Voice in Local Development." *The Review of Economics and Statistics*, 105 (2): 311–326.
- Cheng, Christine, and Margit Tavits. "Informal influences in selecting female political candidates." *Political Research Quarterly* 64, no. 2 (2011): 460-471.
- Colonnelli, E., Prem, M., & Teso, E. (2020). Patronage and selection in public sector organizations. *American Economic Review*, 110(10), 3071-99.

- Cruz, C., Labonne, J., & Querubin, P. (2017). Politician family networks and electoral outcomes: Evidence from the Philippines. *American Economic Review*, 107(10), 3006-37.
- Dal Bó, E., Finan, F., & Rossi, M. A. (2013). Strengthening state capabilities: The role of financial incentives in the call to public service. *The Quarterly Journal of Economics*, 128(3), 1169-1218.
- Dal Bó, E., & Finan, F. (2018). Progress and perspectives in the study of political selection. *Annual Review of Economics*, 10, 541-575.
- Finan, F., B. A. Olken & R. Pande, 2015. "The Personnel Economics of the State," NBER Working Paper 21825.
- Fujiwara, Thomas, Hanno Hilbig, and Pia Raffler. "Party Nominations and Female Electoral Performance: Evidence from Germany." (2021)
- Gulzar S, & Khan MY. (2018). Barriers to political entry: experimental evidence from local government elections in Pakistan. Working Paper. Dep. Political Sci., Princeton University.
- Gulzar S, & Khan MY. (2023). "Good Politicians": Experimental Evidence on Motivations for Political Candidacy and Government Performance. Working Paper. Dep. Political Sci., Princeton.
- He, G., & Wang, S. (2017). Do college graduates serving as village officials help rural China?. *American Economic Journal: Applied Economics*, 9(4), 186-215.
- Jia, R., Kudamatsu, M., & Seim, D. (2015). Political selection in China: The complementary roles of connections and performance. *Journal of the European Economic Association*, 13(4), 631-668.
- Karpowitz, C. F., Monson, J. Q., & Preece, J. R. (2017). How to elect more women: Gender and candidate success in a field experiment. *American Journal of Political Science*, 61(4), 927-943.
- Kanthak, Kristin, and Jonathan Woon. "Women don't run? Election aversion and candidate entry." *American Journal of Political Science* 59, no. 3 (2015): 595-612.
- Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz. 2007. "Experimental Analysis of Neighborhood Effects," *Econometrica*, 75(1): 83-119.
- Larreguy, H., Marshall, J., & Querubin, P. (2016). Parties, brokers, and voter mobilization: How turnout buying depends upon the party's capacity to monitor brokers. *American Political Science Review*, 110(1), 160-179.

- Leaver, C., Ozier, O., Serneels, P., & Zeitlin, A. (2021). Recruitment, Effort, and Retention Effects of Performance Contracts for Civil Servants: Experimental Evidence from Rwandan Primary Schools. *American Economic Review*, 111(7): 2213-46.
- Niven, David. "Party elites and women candidates: The shape of bias." *Women & Politics* 19, no. 2 (1998): 57-80.
- Pande, Rohini, and Deanna Ford. "Gender Quotas and Female Leadership: A Review. World Development Report on Gender," 2011.
- Phillips, M. L. (2020). Family matters: gendered candidate selection by party gatekeepers. Working Paper. Dep. Political Sci., University of Cal. Berkeley, Berkeley, CA.
- Thomsen, Danielle M., and Aaron S. King. "Women's Representation and the Gendered Pipeline to Power." *American Political Science Review* 114, no. 4 (2020): 989–100.
- Viviano, Davide, Kaspar Wuthrich, and Paul Niehaus, "When should you adjust for mulitpl hypothesis testing?" manuscript, 2023.

Appendix A: Data Collection and Gatekeeping Letter from IPA

January 21, 2024

To Whom It May Concern,

MORE EVIDENCE, LESS POVERTY



We are partnering with the PI team Katherine Casey, Abou Bakarr Kamara, Niccolo Meriggi, and Andres Rodriguez on the project titled “Candidate Entry into Local Government.” We have been collaborating with them since the beginning of the project and remain their primary partner for all matters related to data collection in Sierra Leone.

As articulated in their pre-registered analysis, there are two remaining sources of data that are necessary to collect before the PIs can begin to estimate causal effects on their key outcomes of interest, which relate to the number and quality of aspirants and candidates at the ward level. These are:

- Ward-level identifiers for all aspirants and candidates in their research sample
- Survey data on the quality of all status quo (SQ) aspirants and candidates in their research sample

We plan to collect this data in the following steps:

1. Secure from leaders of both major political parties the ward identifier and/or phone number for all individuals listed in their applicant dataset and the official gazette of registered district council candidates
2. Use ward identifiers to narrow the population of applicants and candidates to those registered in any of the 250 wards in the research sample
3. Cross-check this list to identify SQ aspirants not in the existing community nominee or Local Councilor survey datasets and flag for fieldwork. Our best estimate is that this sample includes roughly 800 SQ aspirants.
4. Contact these individuals via field visits to verify that the ward identifier provided by parties matches the location listed on their voter registration card
5. Conduct quality screening surveys with these individuals

This process is underway, and we will not provide full datasets to the PI team for at least one month from this date (February 21, 2024). Please do not hesitate to contact me with any questions.

Sincerely,



Walker Higgins, Country Director - IPA Sierra Leone & Liberia