CANDIDATE ENTRY INTO LOCAL GOVERNMENT

Katherine Casey[*] Abou Bakarr Kamara[†] Niccolò F. Meriggi[‡] Andrés F. Rodriguez[§]

THIS VERSION: November 5, 2026

Abstract

The competence of public officials affects state performance and economic growth, yet building this capacity is difficult: it requires protocols to identify more capable, representative individuals; their willingness to run for office; and party interest in supporting their candidacies. We evaluate an experiment designed to parse these problems of identification, entry and gatekeeping in advance of the 2023 local elections in Sierra Leone. A combination of community nominations and technocratic screening identifies individuals who are substantially higher quality and enjoy broader popular support than both incumbents and status quo applicants. One quarter of top nominees formally enter politics, positively self-selected on quality, and their entry improves the maximum quality observed in candidate applicant pools. Party gatekeepers, however, pass over these more capable nominees to instead favor their own re-running incumbents in choosing candidates. Results suggest that the challenge for political renewal is less about supply than demand in this context: there are indeed better potential candidates out there who are willing to run, but institutional incentives work against selecting them.

Key words: Political entry and selection, economic development, local government, elections

JEL codes: H70, O17

---

[*] Stanford University and NBER. Corresponding author: kecasey@stanford.edu, 655 Knight Way, Stanford CA 94305, +1 (650) 725-2167. [†] International Growth Centre. [‡] University of Oxford. [§] Stanford Impact Labs.

We are grateful for the collaboration of our partners at Innovations for Poverty Action, the Political Parties Regulation Commission of Sierra Leone, the National Electoral Commission of Sierra Leone, the All People's Congress party, the Sierra Leone People's Party, and Talking Drum Studios. We thank Jeremy Bowles, Mike Callen, Fred Finan, Thomas Fujiwara, Maggie Jones, Yasir Khan and seminar participants for invaluable comments. We benefited from excellent field and research assistance by Naasey Kanko Arthur, Kevin Dunne, Rania Ezzane, Pablo Guzmán Lizardo, Kiet Le and Monica Shandal. This project has IRB approval from Stanford University (protocol 64314) and the Sierra Leone Ethics and Scientific Review Committee (protocol 015/09/2022). This study was pre-registered in the American Economic Association's RCT registry, accessible at https://www.socialscienceregistry.org/trials/10006. Our results-blind, pre-registered report was accepted by the *Journal of Development Economics* without restriction on where to publish the final analysis. We gratefully acknowledge funding from Stanford King Center on Global Development, J-PAL Governance Initiative and the International Growth Centre.

I. Introduction

Worldwide, citizen esteem for politicians and representative institutions is perilously low (Valgarðsson et al., 2025). Only 16 percent of Americans trust the federal government and roughly a third identify politicians as the "biggest problem with the political system" (Pew Research Center, 2023). Across dozens of wealthy nations, those who lack trust in government outnumber those who have trust (by 44 to 39 percent, OECD, 2024). And while a strong majority (75 percent) of citizens in thirty African countries believe that elections are the best way to choose leaders, only 37 percent are satisfied with how their democracy is working (Afrobarometer, 2024). Politicians are viewed as unresponsive, incompetent or self-serving, creating a disconnect between voters and the state (Mahmoud, 2014; Mettler, 2018; Schäfer & Zürn, 2023).

The contribution of inclusive political institutions to economic growth makes this fracture particularly troubling for lower-income countries (Acemoglu & Robinson, 2012). Characteristics of individual leaders are key drivers of the link between politics and economics at both the national (Besley, Montalvo, & Reynal-Querol, 2011; Jones & Olken, 2005) and local level (Besley, Pande, & Rao, 2005; Martínez-Bravo, 2017). This raises the question of whether it is possible to improve state performance by bringing "better"—higher human capital, more honest, more representative—people into elected office.

The challenge in doing so is threefold. First, societies may lack the protocols and technology to identify individuals who would make for more effective politicians. Second, even if highly capable individuals can be identified, they may be unwilling to run for office. And third, even if they can be convinced to run, political parties' incentives may not align with selecting these new entrants as their candidates. We experimentally evaluate an intervention that aims to parse these problems—of identification, entry and gatekeeping—and assess which one(s) pose binding constraints on political renewal for the 2023 Local Council elections in Sierra Leone.

Collaborating with the government and both major parties, we situate this exploration in local politics for a few reasons. The work of Council focuses on the provision of basic public goods and services, which is important for economic development, and benefits from general competence rather than specialized expertise. Barriers to entry are relatively low (requiring few resources to contest, only part-time service and minimal eligibility requirements), suggesting a viable pathway in for new talent, and yet parties retain significant control over who runs and gets elected.

The intervention studied approaches the first issue of identifying political talent from the perspectives of representation and technocratic merit. Field teams began by asking a representative sample of voters to nominate individuals in their jurisdictions who have the skills needed to perform well in local government, people they think would do a good job in office, and whom they would want to represent them. Individuals could name up to three people in a first unconstrained round and up to an additional three women in a second round. Nominations were elicited via household visits

in the three largest communities per ward, which is the administrative unit tied to a Local Council seat.[1]

Technocratic screening then winnowed the field of popular names to the two highest quality nominees per ward, whom we refer to as top nominees. Enumerators privately screened up to ten of the most frequently named individuals on skills that empirically correlate with on-the-job performance. The screening instrument collates measures from multiple studies in this context and focuses on those that predict favorable Parliamentary audit reports (from Casey, Kamara, & Meriggi, 2021), success in winning competitive infrastructure grants (Casey, Glennerster, Miguel, & Voors, 2023), and positive anonymous peer reviews among sitting incumbents. We use the same instrument to assess all incumbents and status quo applicants (those not involved in the intervention who filed a candidate application with a political party), which enables a rare profiling of the universe of potential candidates.

To assess whether this strategy identifies "better" potential candidates, community nominations and technocratic screening were randomized across 250 wards within fourteen of the fifteen rural district Councils nationwide. As each treated ward was available to both participating parties, this generates a full national sample of 500 party-ward units. We focus much of the discussion on observations in the respective regional stronghold of each party, as this is where they are most likely to win seats in the general election (all full sample counterparts are in the Appendix).

Headline estimates show that the intervention succeeded in finding individuals of substantially higher quality than both incumbents and status quo applicants. In the full national sample, top community nominees scored on average 0.28 standard deviation units (standard error of 0.02) higher on the quality index than status quo applicants and 0.10 units higher than incumbents (p-value for rejecting equality of 0.001). Their quality premium reflects differences across multiple components including measures of human capital like education, numeracy and cognitive ability; exercises that capture managerial ability; and a behavioral measure of how they handle discrepancies in the token compensation provided for survey participation, which we refer to as conscientiousness but can also be thought of as a proxy for integrity or attention to detail.

These potential candidates command broad support among local residents. Top community nominees were named by an average of 16 percent of households in their respective wards, which is 9 percentage points higher than incumbents and 13 percentage points higher than status quo applicants (both differences are significant at the 99 percent confidence level). While these comparisons suggest improvements in representation, interpretation is subject to the caveat that status quo applicants were not known until after the nominations process concluded, so any errors matching their names back to the nomination dataset may underestimate their popularity.

[1] Census data estimates suggest that this household sampling strategy is representative of roughly 29 to 37 percent of the total ward population.

Community nominees are new to electoral politics. Compared to status quo applicants, they are less likely to be registered with a party (by 22 percentage points), to have held political office (at 28 versus 41 percent), or have relatives in a party office (by 5 points). Other characteristics position them as an alternate elite within the local political landscape. Compared to voters, they are wealthier, better educated and more likely to be male. They are disproportionately drawn from so-called "ruling" families—63 percent of top nominees versus 21 percent of status quo applicants—which are traditional governance structures headed by chiefs who wield authority over land, labor and justice. While studies suggest that chiefs themselves may not be optimal for development outcomes (Acemoglu, Reed, & Robinson, 2014; Casey, Glennerster, Miguel, & Voors, 2023), this intervention identifies their younger relatives and descendants. Survey data indicates substantial support for traditional authorities in Sierra Leone (Afrobarometer, 2024), and they have been productively brought into state service elsewhere (Balán, Bergeron, Tourek, & Weigel, 2022). How open political parties are to candidates associated with these parallel institutions, however, is unclear (Baldwin, 2016; Henn, 2023).

Evidence thus far suggests that there indeed exists a talented, popular cadre of potential candidates on the margins of politics. The next question is whether they run for office. We define entry as filing a candidate application with a political party, which requires time, effort and financial dues.

There are several reasons why the intervention might encourage entry. First, household respondents may internalize that someone needs to be a good fit for politics to be worth nominating, thereby identifying individuals near the entry margin. Consistent with this, 89 percent of top nominees said they were "interested" or "very interested" in running for Local Council. Second, screening protocols provided information to nominees about how many people had put their name forward and shared testimonials from the household data about why respondents thought they would make a good Councillor. Enumerators further explained that the purpose of the initiative was to work with parties to "identify high-quality candidates" and that there was a possibility, but no guarantee, that information collected during screening would be shared with party leadership. Eighty-nine percent of top nominees gave consent to share their information if such an opportunity arose. Third, for a subset of top nominees, we randomized a follow-up nudge in the form of a phone call explaining that they had attained the highest merit score and that the research team would indeed share their profile with party leaders.

Using administrative data obtained from the parties, we calculate an entry rate of 13 percent for community nominees ranked below the top two positions on the merit-based screening. Strong positive self-selection on quality raises this rate to 23 percent for top nominees, a difference that is statistically significant at the 99 percent confidence level. The phone call nudge increased entry among top picks by an additional 5 percentage points, which is sizeable but not statistically distinct from zero (standard error 3.8).

The entry induced by the intervention had limited impacts on the overall size of ward-level candidate

applicant pools but did increase the maximum quality of individuals found in the pools. We estimate a treatment effect in the stronghold sample of 0.22 additional applicants per ward (on a control mean of 2.83), which moves in the expected direction but is not statistically distinct from zero. The maximum quality score found among all who filed candidate applications from a particular ward is 0.09 standard deviation units higher in the nominations and screening arm compared to control wards (p-value 0.08). The corresponding estimate for impacts on the mean score is also positive (0.05 units), but not statistically significant (p-value 0.23).

We use our rich data to simulate impacts that would have obtained under higher realized entry rates for the best quality nominees. These counterfactuals retain the top nominees who applied in practice, and then for wards where no top picks applied, randomly boost the entry rate of the highest scoring nominee. Boosting these top non-applicants' entry by 25 percent generates a 0.12 standard deviation unit increase in the maximum quality of the pool (p-value 0.02) and 0.06 unit increase in the mean (p-value 0.15). Simulating a 50 percent boost yields a 0.15 unit increase in the maximum (p-value 0.01) and 0.07 unit increase in the mean (p-value 0.09). These simulations reveal additional value that could be unlocked by investing more in converting the best nominees to applicants, which was the lightest touch link in the intervention (i.e., the phone call nudge).

Whether the observed entry rates are large or small is difficult to say. While entry is a rare event, its rarity depends on the reference pool, ranging from all eligible individuals to the politically ambitious (Gulzar, 2021). There are only a few related estimates in the literature. In the U.S., the civil society organization Run for Something aims to identify young progressives and encourage them to run for local office. Their conversion rate from website intake forms to candidacy in 2018 was 11.6 percent (Green, Conroy, & Hammond, 2024). In Brazil, Dal Bó, Ferraz, Finan, & Pessoa (2024) find that 16 percent of people who applied for a candidate training program ran for municipal office, and that participating in the training doubles this figure. In Pakistan, Gulzar & Khan (2025) find entry rates of 2 to 4 percent among randomly selected male household heads for new, nonpartisan village council elections. In Sierra Leone, the fact that high-quality community nominees entered politics at nontrivial rates suggests that the nominations and screening protocols developed with government partners helped solve the technology challenge in this context.

As neither finding new talent nor encouraging them to run appear to be binding constraints on political renewal, we turn to the third issue of party gatekeeping. Gatekeeping is strong: of those elected to Local Council in 2018, only 12 percent were from a minor party and 1 percent ran as Independents. Party power increased in 2023 with an abrupt shift from first-past-the-post ward-level elections to district-level proportional representation (PR), where parties submitted closed candidate lists twice as long as the number of wards in the district. The reform was controversial: the opposition decried it as a brazen power grab to wrest control of Parliament, while the incumbent regime cited previous experience with PR (in the 1996 and 2002 elections) and potential benefits for reducing one-party dominance in strongholds. Notable for our purposes is how the closed-list

feature severs the link between voters and the selection of individual candidates, placing party leaders firmly in control (Ahuja & McClendon, 2026). This isolates gatekeeping from electoral accountability and allows us to observe what party leaders prioritize when their choices over individuals are largely unconstrained by voter preferences.

The last set of estimates thus moves from investigation of political entry decisions (on the supply side) to downstream impacts on party selection of candidates (on the demand side). In the raw data, top community nominees who entered politics were more likely than status quo applicants to be selected onto candidate lists (at 69 versus 59 percent) and, once on a list, won seats at the same rate (33 percent). Almost all corresponding regression estimates, however, are null. While modest realized entry rates limit statistical power for this analysis, the main lever constraining the advancement of new talent—both nominees and status quo challengers—is the overwhelming preference parties exhibited toward incumbents.

The privileging of incumbents marks a substantial departure from previous elections. Under the old majoritarian system, only 14 percent of incumbents ran for re-election in 2018 and 37 percent of those re-won a seat. Both rates increased with the change to PR: 48 percent ran again and 64 percent of those won in 2023, thereby reducing the effective rate of political renewal from 95 to 69 percent in a single electoral cycle. One of the strongest predictors of party selection of candidates is incumbency, accompanied by mixed or null results for the quality measures collated in the screening instrument.

The fact that parties passed over top nominees in favor of re-running incumbents can be interpreted in multiple ways. It can be seen as bad for voters in that nominees are more popular and higher quality than incumbents. But perhaps parties have better information about the returns to experience, which means re-selecting incumbents is the most expeditious route to improving Council performance. This seems unlikely in our data, as years in Council do not favorably predict anonymous peer reviews among sitting Councillors, while several components of the quality index do. It also runs counter to evidence that electoral turnover incentivizes better economic policy and performance (Marx, Pons, & Rollet, 2025; Bazzi et al., 2026). Or maybe the rushed electoral reforms left parties with too little time to both adjust to new rules and vet new types of candidates, meaning the null results reflect a temporary shock. Lastly, it could be the institutional rules that matter, as research conducted under the previous majoritarian system showed that parties were open to accommodating information about voter preferences (Casey, Kamara, & Meriggi, 2021), which they do not seem to do under PR. Setting interpretation aside, what is clear is that for this election, efforts to bring new people into public office largely stalled at party gates.

One last feature of the research design that we cannot say much about due to low statistical power is how sharing profiles impacted party selections. The profiles promised in the phone call nudge were produced and given to party leaders but faced two implementation challenges. First, to mitigate potential econometric issues raised by the shift to PR that places wards in zero-sum competition

over finite slots on the list, we assigned profiles to only a subset of treated units using a saturation design (to measure spillovers, as in Baird, Bohren, McIntosh, & Özler, 2018). Second, fewer than a third of profiled nominees (72 of 245) ended up filing an application with a party, leaving a small number of profiles with a real chance of shaping candidate selection decisions. Thus, while the potential to have one's profile shared may have operated on the supply side in encouraging nominees to enter, we have little to say about demand-side effects of delivering profiles on party choices.

This study contributes to the literature on political selection (see Dal Bó & Finan, 2018 and Gulzar, 2021 for reviews). Demographic representation, often studied through the lens of quotas, shapes policy choices and economic outcomes (Chattopadhyay & Duflo, 2004; Clayton, 2021; Pande, 2003). Party gatekeeping has been shown to limit entry for women and minorities, independent of their qualifications (Cheng & Tavits, 2011; Fujiwara, Hilbig, & Raffler, 2025; Karpowitz, Monson, & Preece, 2017). Under weak institutions, patronage and kinship networks distort selection toward politically connected candidates over those stronger on merit (Colonnelli, Prem, & Teso, 2020; Jia, Kudamatsu, & Seim, 2015; Querubín, 2016; Riaño, 2023). While such frictions suggest that institutional reforms can improve both representation and competence simultaneously (Besley, Folke, Persson, & Rickne, 2017), questions remain about potential tradeoffs therein (Dal Bó et al., 2017). Our data show that it is possible to promote the entry of highly qualified people who also enjoy broad support among local voters. While their ultimate electoral success was limited by parties stacking incumbents atop candidate lists, this finding itself is interesting as it captures party behavior in response to an abrupt shift to PR, a change with potentially broader implications for policy (Chang & Golden, 2007; Kunicová & Rose-Ackerman, 2005; Persson, Tabellini, & Trebbi, 2003), political turnover and voter accountability (Bowles & Marx, 2023).

Our analysis contributes most directly to work on improving the supply of individuals who run for office. Stronger financial incentives have been found to attract higher ability candidates who perform better in political and front-line service jobs (Ashraf, Bandiera, Davenport, & Lee, 2020; Dal Bó, Finan, & Rossi, 2013; Gagliarducci & Nannicini, 2013; Leaver, Ozier, Serneels, & Zeitlin, 2021). Other successful approaches focus on talent re-allocation (He & Wang, 2017), pro-social messaging in recruiting candidates and party volunteers (Gulzar & Khan, 2025; Gulzar, Pathak, Thompson, & Tóth, 2025), and training workshops (Dal Bó, Ferraz, Finan, & Pessoa, 2024; Ravanilla, 2016). We add new evidence about the role of community participation in identifying political talent and self-selection in shaping the quality profile of entrants. Returning to the three-pronged challenge of improving political selection, our results suggest that identifying high caliber potential politicians and inducing them to enter are margins that can be constructively moved, while unlocking party gatekeeping proved more difficult in this context.

II. Institutional Context

II.A. Local Government

Sierra Leone is a low-income country with weak state capacity, ranking in the 10th percentile for government effectiveness (World Bank, 2024). Local government is organized into 22 district and town councils, with representation tied to ward population, the most local administrative unit of the formal state. Local Council wards, of which there are 446 nationwide, nest neatly inside Parliamentary constituencies, the next higher unit of state administration, which in turn nest neatly inside district boundaries (see Appendix Figure A1).

Local Councils promote economic development primarily through providing local public infrastructure and monitoring the performance of line ministries in delivering public services. Eligibility criteria to serve on Council are minimal: one must be at least 18 years of age, registered to vote, up to date in tax payments, and ordinarily resident in the ward for the past three years (Government of Sierra Leone, 2022). The work of Council is part time, so elected officials can keep their primary employment while serving.

Voters express high levels of dissatisfaction with the politicians currently serving in local government. In the last round of Afrobarometer (2024) before the 2023 elections, 70 percent of rural respondents "disapprove" or "strongly disapprove" of the performance of their Local Council, 35 percent claim that "most" or "all" Local Councillors are corrupt, and 33 percent say they do not trust their Councillor "at all." This contrasts with voter opinions of the chieftaincy system, which is an older, parallel system of traditional authority, open to descendants of what are known locally as "ruling" families. The same dataset shows that only 19 percent of rural respondents claim that "most" or "all" traditional leaders are corrupt, and only 11 percent claim to not trust them "at all."

Although the chieftaincy runs largely in parallel to elected government, it has formal points of contact with the state. Nationally, one Paramount Chief from each provincial district sits in Parliament. These members are elected not by voters but by Chiefdom Councillors: the elders, section chiefs and village heads who make up the traditional chiefdom administration (Government of Sierra Leone, 1991). Paramount Chiefs are more directly involved at the district level, where they sit as ex officio members of the district council and attend its meetings on a rotating basis (Government of Sierra Leone, 2022). There they are expected to articulate their chiefdoms' development priorities and advocate for these before the district council. Chiefly authorities are still responsible for collecting some local taxes and passing a share of them to the district council,

one of several areas in which their roles formally overlap with those of local government.[2] This framework has been in place since decentralization was reintroduced in 2004 and was not changed by the 2023 switch to PR.

One way to address voter dissatisfaction is to recruit new talent into local politics. If party incentives align with voter preferences, then improving the supply of people who file applications to be considered as candidates could be a sufficient strategy. Previous research on Parliamentary elections in Sierra Leone suggests that voters and party leaders have similar views about what makes for a good candidate, like having a strong record of participating in local development projects (Casey, Kamara, & Meriggi, 2021). This idea motivates the main treatment arm studied here (T1), which combines community nominations to identify individuals with broad popular support with technocratic screening to elevate the most competent nominees. It complements the previous study, which focuses on primary-stage selections over a fixed set of applicants, by turning to supply-side considerations about how to expand and enhance the quality of applicant pools.

Parties also face information and logistical constraints in screening over 400 applicant pools in a setting with high transport and communication costs and few, if any, firms engaged in political polling. Indeed, when the earlier study asked party leaders which applicant would win if voters were to select one directly, 90 percent said it would be their own first choice applicant, which was in fact true in voter polling data only about half the time. When the voter data was shared with party leadership, they became 23 percentage points more likely to select the applicant who ranked first among voters, for a 61 percent increase on the base rate. Party responsiveness to information motivates the inclusion of a second treatment arm here (T2) that shares information about top nominees' popularity and quality with party leaders via one-page profiles. Profile sharing was randomly assigned to a subset of T1 wards.

To engage parties in the design and implementation of these interventions, the research team worked in partnership with the Political Parties Regulation Commission (PPRC), the government agency with the constitutional authority to regulate the conduct of parties with respect to their members and the broader electorate. The PPRC invited all registered parties to join the initiative and associated research. Party leaders decided whether or not to opt in. Both major parties, the All People's Congress (APC) and the Sierra Leone People's Party (SLPP), joined the initiative and were assigned equal numbers of treatment and control units in the experimental design.

[2]Paramount Chiefs also sit on the ward development committees in their chiefdom. Under the Local Government Act, Chiefdom Councils must cooperate with district councils and may carry out functions the council delegates to them, while keeping their customary roles of making bye-laws and holding land in trust (Government of Sierra Leone, 2022). The Act nonetheless keeps the two systems distinct at the level of individual office, as Paramount Chiefs themselves cannot stand for election as Local Councillors.

II.B. Electoral Reforms and Research Adjustments

Since the end of the civil war (in 2002) and the reintroduction of decentralization (in 2004), Local Councils in Sierra Leone have been organized as single-member jurisdictions with first-past-the-post (FPTP) elections held at the ward level. In a controversial decision, the Government of Sierra Leone changed from this majoritarian system to proportional representation organized at the district level shortly before the 2023 election studied. This reform was hotly contested by the opposition party, including Supreme Court challenges, but was ultimately enacted for both Local Council and Parliamentary elections. For each district council, parties were asked to compile a closed list of candidates that was twice as long as the preexisting number of seats (equivalent to the number of wards) (see Appendix Figure A2). Voters saw only party labels and not the list of individual candidate names on the ballot. The government also enacted a long-discussed quota for thirty percent women's representation in politics. In practice the quota translated into every third slot on the candidate list being occupied by a female candidate.

While switches between electoral systems are uncommon, they have occurred recently in Sierra Leone, which used PR for the 1996 and 2002 elections, switched to FPTP for 2007, 2012 and 2018, and went back to PR in 2023. Neighboring Liberia also switched from FPTP to PR in 1997 and back again in 2005. Among countries that use PR, three quarters use closed lists and one quarter use open lists (Cruz, Keefer, & Scartascini, 2021). While systematic data on the visibility of candidate lists is harder to come by, web search identifies several African countries that join Sierra Leone in printing only party labels on the ballot (including Angola, Benin, Burkina Faso, Burundi, Namibia and Niger).

It was unclear *ex ante* how the switch to PR might impact geographic representation, which is an outcome of interest and affects the way we structured our research design. It could be minimal, as an easy response for party leaders is to adhere to the old system and simply pick two candidates per ward to place on their district list. Or, parties could switch to a somewhat higher administrative unit, like Parliamentary constituency or traditional chiefdom, still broadly reflecting the geographic variation of the district. Alternatively, parties could abandon geographic considerations altogether and pick candidates from the highest density town, ignoring outlying areas. Conversations with party leaders at the time made this last concentrated approach seem unlikely; however, their ultimate strategy was not revealed until they submitted finalized candidate lists to the National Electoral Commission, which was, by construction, after this field experiment concluded.[3]

Econometrically, movement away from the ward as the organizing unit for the parties' lists introduces the risk of partial interference between units for estimating treatment effects of profile sharing (T2). This is because PR places wards in competition with each other for representation via a finite

[3]Note that the home ward of all eventual candidates is knowable, as voter identification and registration remained tied to wards, so answering these questions was an issue of timing and not missing data.

number of slots on the district-level list. So, if parties respond to an applicant profile by adding that person to the list and making space for them by dropping someone from a control ward, this biases upwards estimates of how profiles enhance the geographic representation of T2 wards. This issue should not affect estimates of how nominations and screening (T1) impact candidate entry, as there is no cap on how many people can file candidate applications from a particular ward.

To proactively mitigate this risk, we followed Baird, Bohren, McIntosh, & Özler (2018) and adjusted the random assignment for profile sharing upwards to the constituency-level with varying degrees of saturation to enable estimation of potential spillovers. A constituency contains 3.4 wards on average, so this seemed like a reasonable level to capture the likely scope of local spillovers. *Ex post*, both parties largely adhered to the old ward-based system: the modal ward had two candidates on each party's respective district-level list and ended up with exactly one elected representative, so the risk of partial interference proved to be not much of a problem in practice. Appendix Exhibit A4 details this part of the empirical strategy and shows evidence that spillover estimates are null.

III. Research Design

III.A. Sample

To evaluate the interventions, which were collectively referred to among our partners as the "Local Champions Initiative," we first selected an experimental sample of wards within the district councils. The sample excludes wards where the intervention was piloted and those whose boundaries cut across chiefdom boundaries (to avoid confusion since chiefdom divisions may be more salient to citizens than wards). This generates a sample of 250 wards, located in 92 distinct Parliamentary constituencies.

Since the initiative was done in collaboration with the two major parties, each sampled unit was available to both parties, which yields a full national sample of 500 party-wards and 184 party-constituencies. Considering clear regional patterns of support for each major party, we delineate a stronghold sample (250 wards, 92 constituencies) that assigns wards to the party that historically dominates politics in that locality (see Figure 1). Strongholds arise from allegiances between ethnic groups and political parties that date back to the post-Independence period (Kandeh, 1992). As community nominees are most interested in contesting for the dominant party and individuals placed on (the top half of) its candidate list have a high chance of getting elected, we focus much of our discussion on this sample. We present corresponding estimates for the full 500 party-wards in the Appendix.

III.B. Treatment Assignment

At a high level, there were two rounds of treatment assignment (see Figure 2). The first, at the ward level, randomly allocated the 250 sample wards to the nominations and screening procedure (T1) or a pure control group (C), stratified by district and partisan competitiveness (as measured by 2018 electoral returns). This was completed under the presumption that the preexisting majoritarian electoral system would remain intact. Field teams conducted nominations and screening procedures in T1 wards and 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 round, at the constituency level, randomly assigned the profile sharing component (T2) among the T1 units, which created three groups (T1 only, T1+T2, and C). This was completed after it became clear that the switch to PR would go through, and in response, uses a saturation design to measure potential spillovers across wards within a given district-level candidate list. To do so, it grouped wards into their 92 distinct host constituencies and assigned constituencies to three levels of treatment saturation, stratified by district: (i) zero saturation constituencies (no information sharing in any wards), (ii) low saturation constituencies (treat one third of wards with information sharing), and (iii) high saturation constituencies (treat two thirds of wards). Assignment of profile sharing was done independently for the two parties.

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 (40 percent) than the low saturation arm (33 percent), and more to the low saturation arm than to pure controls (27 percent). Appendix Figure A3 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] Appendix Table A1 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.

Once the constituency-level saturation was assigned, we randomly assigned the required number of T1 wards within each party-constituency to the T2 information sharing condition. Given the effective saturation levels, we end up with 40 percent of the party-wards in our sample receiving both T1 and T2. For the stronghold sample this number is 99 (T1+T2) wards, alongside 87 nominations-only wards (T1 only), and 64 pure control (C) wards. Appendix Table A2 shows

[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.

that assignments to nominations and profile sharing are reasonably balanced across population, population density and 2018 electoral outcomes at the ward level.

IV. Implementation

Fieldwork for the main T1 arm proceeded in a sequence of structured community nominations followed by technocratic screening. The T2 profiles collated and shared information from this two-step process in a subset of T1 wards.

IV.A. Representation via Community Nominations

A key aspect of representation is the support of local voters. To assess this, field teams elicited the names of popular individuals via household visits in each of the three largest communities per ward. Field teams divided each community into geographic sectors, with enumerators conducting random walks to visit households along both main and side roads to ensure coverage of both central and peripheral areas. Sampling protocols required one respondent per dwelling, balancing across gender, age cohorts (18-35 versus above 35 years), and ethnic groups where relevant. This approach was designed to achieve representative demographic coverage while minimizing geographic clustering in the nomination process. Coarse estimates using census data suggest that the household sample is representative of roughly 29 to 37 percent of the total ward population.[5]

At each household, enumerators first explained the skills and competencies that are needed to be an effective Local Councillor (see script in Appendix Exhibit A1). They then asked respondents to list up to three individuals from their ward that they thought would make good candidates (unconstrained round). Respondents were then asked to think of up to three additional names of female nominees in particular (second gender round). Teams implemented these elicitations with, on average, 84 individuals per ward. To facilitate aggregation across households, respondents were asked to provide relevant nicknames and contact details for each person they nominated.

Field teams also targeted a standard set of local leaders at the community level. In each community they interviewed the town chief and the women's leader, as well as five respondents from the following list of positions: ward development committee member, youth leader, master farmer, religious leader, traditional society head, community health worker, traditional birth attendant, head teacher or principal of a local school, and town speaker. Some individuals from the household

[5]More specifically, teams visited the headquarter community of each of the 3 largest chiefdom sections per ward (there are on average 3.3 sections per ward). A section is a geographic sub-unit within the boundaries of the traditional chieftaincy system that is salient to voters. In 2015 census data, a headquarter community contains about a third of its section's population (median 32 percent, mean 41 percent). Thus, the share of the total ward population represented by our household sample is: 3/3.3 sections x 32 (41) percent median (mean) share = 29 (37).

elicitation also held leadership positions. Overall, on average 27 respondents per ward hold leadership positions.

IV.B. Quality via Technocratic Screening

Field Protocols

The research team selected up to ten of the highest ranked nominees by frequency of household nominations for private screening. This search was independent of party affiliation and included provisions to include the most popular women even if they obtained slightly fewer votes than some of the nominated men. On average, enumerators successfully contacted 86 percent of highest ranked nominees by frequency of household nominations. In the stronghold data, where we focus exclusively on nominees from the party most likely to win the general election, our data includes on average 7 nominees for the dominant party per ward.

The enumeration script for the screening session includes information that could plausibly encourage entry (see Appendix Exhibit A2). When approaching a nominee, field enumerators explained that part of the purpose of the research program was to help political parties locate high-quality candidates. They told the nominee the number of household nominations they had personally received from residents of their ward. The enumerator also read aloud reasons household respondents had given for why they thought that individual would be a good Local Councillor, text that was preloaded directly from the nominations dataset. Such social recognition could motivate these individuals to enter politics. For example, Gulzar & Khan (2025) show that messaging matters for political entry and that pro-social phrasing about the benefits of serving in politics is more effective than emphasizing private returns in encouraging Pakistani men to stand for local elections.

Nominees were informed that there was a chance, but no guarantee, that some of the information from the screening would be shared with political parties. All nominees were asked if they consented to have their profiles shared with party leaders and if so, with which parties they were willing to share. Overall willingness was high: 85 percent (89 percent) of all (top scoring) nominees were willing to share their profile. For female nominees, these numbers were 80 (88) percent. In practice, nearly all nominees (95 percent) who consented listed only one party, and the vast majority (78 percent) picked the locally dominant party. After processing the screening results, the research team identified the top two highest scorers in each ward. If one of these had not consented to have their profile shared, they were replaced by the third highest scorer unless the drop in scores was too severe, in which case no replacement was made. We refer to this popular, high-scoring, consenting set simply as "top nominees."

Enumerators also asked nominees whether they were interested in running for Local Council. Stated interest was high: among top nominees, 89 percent said they were "interested" or "very interested"

in running. This number was 85 percent among top female nominees. Those who ranked lower on the quality screening expressed less interest: only 67 percent said they were "interested" or "very interested" in running. Their lower intent to run foreshadows the positive selection on quality for filing candidate applications that analysis below documents.

At the end of the screening, nominees were clearly informed that the research team would not share all profiles, that their particular profile may not be shared with party leaders, and that even if it were, the parties were under no obligation to use the profile. The script closes with "We provide information to the parties. We cannot control what parties do with that information."

Quality Measures

To pin down quality in a meaningful way, we selected measures that have been empirically validated in related work to positively correlate with observable outputs of politician performance in Sierra Leone. Specifically, we pulled measures from three separate validation exercises.

We first used data from Casey, Kamara, & Meriggi (2021) to select measures of Parliamentary candidates that positively predict the subsequent public spending of elected Members of Parliament (MPs), as verified by field audits. Second, we use measures from a screening instrument in Casey, Glennerster, Miguel, & Voors (2023) that identifies high-skill members of rural communities. These individuals were involved in drafting proposals for local infrastructure that were submitted to a district government grants competition, and we use measures that positively correlate with proposal quality and winning a grant. Third, during the piloting for the current research, 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 to their constituents) and anonymous peer ratings of competence in office and lack of corruption. We select measures that positively predict either type of outcome.

To avoid overfitting our prediction of quality, particularly with variables that are biased toward groups over-represented among elected politicians (namely older, elite males), we add measures that are widely perceived to be associated with quality and human capital, like years of education, professional experience, and IQ proxies. Our main empirical specifications roll the full set of indicators up into an equally weighted index expressed in standard deviation units. We deploy the same instrument to collect data on status quo applicants and incumbents to fully characterize the potential candidate pool. The text box below lists the individual component measures with superscripts linking them to the validation exercises.

Because a key part of our argument rests on quality measures reflecting real differences in political competencies, let us flesh out a few examples of how index components map to performance data. Starting with conscientiousness, we adapt the transport subsidy exercise deployed among

Text Box 1: List of Quality Measures

1. Human capital:
    a. Years of education
    b. Indicator for literacy*,+
    c. Score in numeracy test*,+
    d. Score in Raven's Test
    e. Score in digit memorization test
2. Work experience:
    a. Wage in previous job
    b. Number of jobs in sectors like health, education, NGOs among last five jobs*
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. Number of jobs with experience managing projects (e.g. budgets, supervising people) among last five jobs+
5. Conscientiousness
    a. Returned extra transport subsidy†

* predicts the performance of sitting Local Councillors in self- or peer-evaluations
† predicts the public spending performance of elected MPs in audit data
+ predicts the quality of local infrastructure grant proposals developed by community members

Parliamentary candidates in the previous election (Casey, Kamara, & Meriggi, 2021). As enumerators asked respondents to meet them in the nearest headquarter town to conduct the survey, it is common practice to offer a transport reimbursement and ask the individual to sign for receipt. In that study, the reimbursement was 150 thousand Leones (roughly $20 at the time), yet in the envelope, there were eighteen instead of fifteen 10,000 Leone notes. The conscientiousness measure indicates whether the individual noticed and returned the extra bills, where the mean return rate was 46 percent. The subsample of candidates who won seats was later audited on how they spent the constituency facilitation fund, which is public money entrusted to them for transport to and development projects in their constituencies. Correlations over time suggest that returning the extra notes as candidates is associated with 60 percent more verified development expenditures by elected MPs in the audit data. For the current study, where enumerators met respondents closer to their home or workplace, cash notes were replaced by mobile telephone top-up cards. Instructions stated the respondent would receive four 10,000 Leone cards (worth $8 at the time), but there were six cards in the envelope.

For managerial capital, we draw from instruments used and validated by Casey, Glennerster, Miguel, & Voors (2023). Individual community members were asked a battery of questions and exercises to assess their managerial ability as it relates to local public infrastructure provision. This included their experience managing others or coordinating local projects, and a simulated grant writing exercise for agricultural infrastructure that required knowledge about project management, procurement and budgeting. Communities later submitted actual proposals to their relevant Local Council who evaluated them and competitively allocated infrastructure grants to fund the best submissions. In that study, a one standard deviation increase in the managerial capital test score of community members involved in the proposal process correlates with a gain in measured proposal quality of 0.29 standard deviation units (standard error 0.05) and an increased probability of winning a grant by 4.1 percentage points (standard error 1.8). Breaking out these correlations for each of eight core competencies covered by the test, project management experience, literacy and numeracy have the most predictive power for proposal quality.

For human capital, the literacy of incumbent Local Councillors was assessed in the current study by asking them to read a short passage aloud, and the enumerator judged whether they read it with ease or struggled. While only a small subset of incumbents were judged illiterate (4.3 percent in the research sample), that literacy divide is highly predictive of anonymous peer ratings. Literate incumbent Councillors on average score 16 percent higher (standard error 3.4) on peers' review of their effectiveness in office.

IV.C. Information Provision via Profile Sharing

The two-track process above generates lists of representative, high-quality potential candidates. For a randomly selected subset of these nominations and screening (T1) wards, the research team compiled short profiles of top nominees to share with political party leaders for their consideration as they screened applicants and awarded party symbols (T2). Profiles present information under free disposal, with no obligation on parties to select any of the nominees.

The profiles include the nominee's name and photo, along with basic information on their age, education, occupation and residence (see Figure 3 for an example). To summarize relative quality, each profile includes one chart that compares the nominee's score on the quality screening instrument to the average score of sitting incumbents in that district on the same instrument.[6] For popularity, it explains how many households were surveyed in the ward, how many unique names they generated, and that this nominee received more votes than $[X]$ percent of nominees in their ward. The profile further lists relevant work experience and skills the nominee possesses and highlights areas

[6]Scores are presented at the decile level, so each individual scores a 1, 2, 3, . . . , 10. The average for the district is the average of these deciles for all incumbents serving on the relevant Council, all of which was explained to the party officials receiving the profiles.

within the screening instrument where the nominee excelled. It closes with a quotation from the household nominations data, labeled "community testimonial." The testimonial in Figure 3 reads, "He is helpful to the community, bold, educated, can speak well, developmentally oriented, and has humanitarian feeling for his community." Delivery rates of generated profiles to party leaders were high, at 92 percent for constituency-level officials (signed on average by at least three different leaders) and 82 percent for district-level officials (signed on average by two).

While the sharing of profiles focuses on the demand-side of parties selecting candidates, it also includes a supply-side "nudge" to top nominees. The nudge was a phone call to top nominees in T2 wards informing them of their strong performance in the screening process and explaining that their profile would be shared with party leaders. The call included a reminder to file a candidate application with their preferred party by the relevant deadline, if they had not already done so. The research team successfully contacted 84 percent of the top nominees assigned to T2. This nudge operates in addition to the social recognition that all nominees received when originally contacted for screening (in T1), both of which might encourage individuals to enter the applicant pool.

The timeline of implementation proceeded roughly as follows. Sitting incumbents were surveyed from October to November of 2022. Household nominations were elicited between mid-February and end-March 2023, with the implementation of private screening slightly trailing as teams roved across wards (T1). Top nominee phone calls (T2) were made between March 29 and April 4. Profile packets (T2) were delivered to party officials April 11 to 17. Parties had to submit final candidate lists to the National Electoral Commission by April 28. Status quo applicants were surveyed between mid-May and mid-June. Election Day was June 24, 2023. "Mop up" surveys of missing status quo applicants and newly elected Councillors were completed over the following months.

V. Identifying New Political Talent

To tackle the first hurdle for political renewal, namely finding new talent, this section explores whether there exists a pool of high-quality potential candidates, who enjoy broad support among local voters, that can be tapped for public service. Analysis assesses the intervention's efficacy in identifying such new talent and compares their characteristics to those who enter politics via the status quo channels.

V.A. Who Runs in the Status Quo?

Compared to the average voter in rural districts, individuals who run for local office, and those who win, tend to be drawn from the socioeconomic elite. Table 1 shows that applicants for local

candidacy are more likely to be male (at 73 percent versus 50 percent for voters), somewhat older, and wealthier as measured by various proxies including asset counts, monthly wages and being less likely to work in agriculture (26 versus 63 percent). They are also positively selected on educational attainment, including higher literacy rates (82 versus 49 percent) and higher primary school completion (92 versus 37 percent). Sitting incumbents are somewhat more strongly selected on most of these dimensions, including male gender, age, assets and educational attainment.

Such positive selection on education does not imply no room for improvement. In the applicant pool, only about half have pursued education beyond secondary school and 18 percent are illiterate. There is certainly scope to improve descriptive representation for women and those from lower socioeconomic status.

Previewing differences that we unpack below, top nominees appear similar to status quo applicants on the core demographic and economic indicators. They diverge by being better educated: 14 percentage points more likely to work in education, all are literate, and 11 percentage points more likely to have completed post-secondary schooling. The emergence of teachers through the nomination system resonates with the motivation to find high human capital individuals who are well known in their communities and involved in public service provision. Top nominees are much less likely to be registered with a party (by 22 percentage points), signaling their relative newness to politics.

Strikingly, nominees are more heavily drawn from traditional chieftaincy lineages. The share from ruling families rises at each step in Table 1: from 16 percent of voters, to 21 percent of status quo applicants, 36 percent of incumbents and 63 percent of top nominees. Note that nominees are not traditional leaders themselves, as only 4 percent of them report being a section or town chief, but rather their younger relatives. This resonates with the positive views of traditional authorities among rural voters, who see them as less corrupt and more trustworthy than Local Councillors (Afrobarometer, 2024). Overall, these differences characterize top nominees as alternative elites who possess high human capital, strong standing in their communities, and connections with the traditional governance system that runs largely in parallel to elected government.

These descriptive patterns speak to a broad literature that paints a nuanced picture of how traditional institutions coexist with the state and with electoral politics (Baldwin, 2025). Chiefs can act as complements to the state, coordinating with it to mobilize collective action, deliver services and state transfers, and collect taxes (Balán, Bergeron, Tourek, & Weigel, 2022; Basurto, Dupas, & Robinson, 2020). They can also act as substitutes or rivals, with tensions over authority and revenue (Acemoglu, Reed, & Robinson, 2014; Henn, 2023). Traditional authorities also shape electoral outcomes in many settings, for instance by coordinating voters behind particular parties or candidates, with ambiguous implications for the accountability of elected representatives (Baldwin,

2025; de Kadt & Larreguy, 2018).[7] Coming from a ruling family thus places top nominees within a select elite embedded in both systems.

V.B. Assessing Relative Quality

Quality analysis focuses on whether the intervention identified individuals who are stronger on dimensions that are empirically correlated with performance in office than others whom parties consider for candidacy. Estimation compares the characteristics of top community nominees (CN) to incumbent Local Councillors (INC) and status quo applicants (the omitted group), using the following specification:

$$Q_{iw} = \beta_0 + \beta_1 CN_{iw} + \beta_2 INC_{iw} + \mathbf{s}_w + \varepsilon_{iw} \tag{1}$$

where Q is the index of quality measures (from Text Box 1), $\mathbf{s}$ is a vector of randomization strata (by district and competitiveness), and ε is an idiosyncratic error term. Construction of the index follows Kling, Liebman, & Katz (2007) by orienting each variable so that higher values indicate more favorable outcomes, transforming them into standard deviation units by subtracting the mean and dividing by the standard deviation of the omitted group, imputing missing values at the mean of the omitted group (so they take a standardized value of zero), and weighting each variable equally. The estimand of interest is β_1, where $\beta_1 \geq 0$ indicates that the intervention identified nominees who were at least as high quality as status quo applicants. We also test $\beta_1 \geq \beta_2$, which indicates that the intervention identified nominees that were at least as high quality as sitting incumbents at their *previous* election.[8] The sample includes all top nominees from the 186 wards assigned to community nominations and screening (T1), and incumbents and status quo applicants from the 250 wards in the research sample.[9]

Table 2 shows that individuals identified by community nominations and private screening score substantially higher on the technocratic quality index than both status quo applicants and incumbent Councillors. The first row of Table 2 shows that on average top nominees scored 0.275 standard deviation units (0.023 standard error) higher than status quo applicants on the index in the full national sample. Incumbents also scored higher, by 0.172 standard deviation units (0.030 standard error) than status quo applicants. Post-estimation tests of the equality of these two coefficients strongly reject (with associated p-value of 0.001), indicating that top nominees are higher quality

[7] See Baldwin & Raffler (2019) on how chiefs' role in service delivery shapes electoral accountability. Revenue collection illustrates both coordination and tension: chiefdom authorities collect local taxes, market dues and natural-resource royalties (Glennerster, Miguel, & Rothenberg, 2013), and share part of the local tax with district councils.

[8] While most of the quality measures are time invariant (like highest completed education), for indicators that vary over time, like political experience, we measure them for incumbents at the point of candidacy for the previous election.

[9] Results are similar when we expand the comparison sample to include incumbents and status quo applicants from all rural wards in the study districts, or restrict it to the same 186 wards from which nominees are drawn.

than the incumbents they seek to unseat. Column 2 repeats these estimates restricting consideration to the stronghold sample. Top nominees again strongly outperform status quo applicants, by 0.200 standard deviation units (0.027 standard error) and sitting incumbents (p-value rejecting equality < 0.001). Incumbents' scores are higher but not statistically distinct from status quo applicants' scores in the strongholds (coefficient of 0.016, standard error 0.033).

Figure 4 breaks the index into its five clusters of components and shows that top nominees dominate status quo applicants across the board. The mean score of status quo applicants in the full national sample is standardized to zero, as indicated by the red dashed vertical line, and the coefficient and confidence interval estimating how top nominees' scores differ is shown by the blue horizontal lines. Top nominees display higher levels of human capital, relevant work and local experience, managerial ability, and conscientiousness in the top-up card behavioral exercise. Incumbent scores (depicted by red horizontal lines) are higher than status quo applicants for human capital and local experience, and indistinguishable for work history and managerial capital. Note that conscientiousness was not measured for incumbents due to our data collection partner's reluctance to provide any token compensation to elected public officials for survey participation. Removing conscientiousness from the quality index does not alter estimates of top nominees' relative advantage by much: in the stronghold sample, they retain a 0.191 unit (standard error 0.029) edge over applicants and 0.181 unit (p-value < 0.001) over incumbents (in Appendix Table A12). Corresponding regression estimates for index components and individual variables are in Appendix Tables A3 and A4.

Are the nominees new entrants to politics?

Figure 5 substantiates the claim that top community nominees are new to politics, and not simply the same set of people who would have run for local office in the absence of the intervention. As the most direct measure, only 28 percent of top nominees have held elected office before, which is less than status quo applicants (at 41 percent) and incumbents (100 percent of whom hold office now, and 39 percent did prior to their previous election). Nominees are also less connected to existing political networks, as they are less likely to be registered with a political party, have less political experience, and are less likely to have a family member working in a party office.

A further test of whether these are truly a different set of individuals choosing to run leverages the pure control sample, where no interventions were implemented. We compare the best people who emerged organically and applied to political parties in control wards to the best people identified by the intervention in treatment wards. Panel A of Table 3 compares the characteristics of the top nominees in T1 wards to the two highest quality status quo applicants found in control wards in the stronghold sample. Estimates show that top nominees are more likely to be female and from a ruling family, are less wealthy and politically experienced, and are overall higher quality (by a statistically significant 0.102 standard deviation units, standard error 0.042). Panel B takes one step

further, retaining only top nominees who formally entered politics (i.e., filed a candidate application with a party) in T1 wards. The patterns remain similar: nominee entrants are more likely to be female, from a ruling family, and higher quality (by 0.151 standard deviation units, standard error 0.060). Appendix Table A5 shows that in the full sample of 500 party-ward units, top nominees remain more likely to be from a ruling family, less politically experienced, and higher quality (with a positive and highly significant advantage on the quality index), although the differences in gender and wealth do not hold. These results argue against the notion that those identified by community nominations and screening would have run anyway, by showing that nominees who did in fact run differ on observables from the best people who ran in pure control wards.

Heterogeneity

We explore whether there is variation in how well the intervention worked in identifying high-quality nominees along prespecified dimensions of gender, party and competitiveness. Table 4 does so by comparing average quality scores among community nominees (excluding incumbents and status quo applicants) in the full national sample. The first column is largely mechanical and shows that top nominees are on average 0.474 standard deviation units (0.022 standard error) higher quality than lower ranked nominees, which is consistent with technical screening winnowing consideration to high scorers.

Column 2 shows that lower ranked female nominees are substantially lower quality, by 0.259 standard deviation units (0.026 standard error), than their lower ranked male counterparts. The screening process reduces this gender disadvantage by half, as the coefficient on the interaction between female and top nominee is a significant 0.129 standard deviation units (0.052 standard error). This highlights potential for the intervention to be particularly useful in identifying high-quality women at a time when parties were working to greatly expand the number of female candidates they select in response to the introduction of the quota.

Column 3 shows that lower ranked nominees are of similar quality across the two major parties, but the quality premium of top nominees is larger for the ruling party, by 0.124 standard deviation units (0.044 standard error). Top nominees nonetheless substantially outperform lower ranked peers for both parties. Column 4 shows that the average quality of lower ranked nominees does not differ in more electorally competitive wards (coefficient −0.019, standard error 0.043), where competitiveness is defined at the ward-level based on previous electoral results. The quality premium of top nominees is somewhat smaller in these areas, by 0.097 standard deviation units (0.051 standard error), but remains large at 0.399 units. Similarly, lower ranked nominees are of similar quality in the stronghold sample and in the parties' respective weakholds (0.057, standard error 0.037), while top nominees outperform their lower ranked peers by more in strongholds (0.534 versus 0.354 units, with an interaction of 0.180, standard error 0.051). This means that screening

adds the most value in party strongholds, which is where parties are most likely to win seats and where much of our analysis focuses.

V.C. Representation

Analysis next explores whether the intervention successfully identified potential candidates with broad support among local voters. While the answer to this question is largely, "yes, by design," we can probe further in two ways. First, we characterize the amount of local support evidenced in the community nominations data, and second, we draw comparisons between nominees and other applicants, with some caveats.

Nominations Data

There is large dispersion of nomination "vote" shares in Figure 6. On average, 47 unique names were generated per ward. Aggregating across households, the nomination shares range from 1 to over 70 percent. Panel A shows the distribution for all community nominees (shaded in gray), where it is evident that the vast majority are nominated by only one or two individuals. The average nomination share is 5 percent. We selected up to the ten most popular individuals to carry forward to the private screening component. The average nomination share for this group (in red outline) is 13 percent, which is mechanically higher. Using the screening data, we selected the top scoring of these popular nominees to potentially carry forward to parties via the profiles. For these top nominees, their average nomination share (in blue outline) is 16 percent. This additional increase is not mechanical and suggests that those who score higher on merit are more popular, which is consistent with households possessing information about who holds the relevant skills that field enumerators emphasized were useful for the work of Local Council.

Panel B limits the distribution to female nominees. There are many fewer women than men in the sample because households nominated fewer women: 69 percent do not name a woman before being prompted to do so in the second elicitation round, and 50 percent do not name any in the second round, either. As mentioned above, women performed less well on the quality screening, so the share of top nominees falls further. Overall, women compose 30 percent of community nominees and 20 percent of top nominees, indicating that identifying popular women with strong local support proved challenging. That said, once these women were identified, they held similar average nomination shares: among all female nominees the average is 4 percent, which rises to 12 percent for the set carried forward to screening, and 15 percent for the top scorers on the quality instrument.

Relative Comparisons

To draw comparisons between household support for nominees and other potential candidates, we estimate the following regression:

$$VS_{iw} = \beta_0 + \beta_1 CN_{iw} + \beta_2 INC_{iw} + \mathbf{s}_w + \varepsilon_{iw} \tag{2}$$

where VS is the "vote" share of household nominations for individual i in ward w and all other terms remain as defined above for Equation (1). This is estimated for the 186 wards where community nominations were conducted.

An important caveat is in order before exploring results: estimates will be biased in favor of community nominees given the difficulty of matching other applicants back (in time and name) to the household nominations data. This is unavoidable as the set of status quo applicants could not be known until the candidate application deadline, which was well past the conclusion of the community nominations process. *Ex post* matching of names can be difficult, especially as many politicians are known by various nicknames among voters. To facilitate accurate matches, we used text-similarity algorithms that were manually checked by a research assistant. A second issue is that we sampled voters in the largest communities in the ward, so if someone had a base of support across many more rural villages, we would not pick this up (see footnote 5). Both issues are somewhat less of a problem for incumbents, as they were known at the time of the nominations process and represent the entire ward. Nevertheless, consider the following estimates as suggestive in nature only.

Table 5 shows that community nominees have higher shares of household nominations than incumbents, who in turn have higher shares than status quo applicants. This advantage is apparent for measures of the share of household nominations, the total number of unique nominations, and the extensive margin of whether any household nominated that individual. For example, status quo applicants in the full national sample have on average a recovered household nomination share of 2.4 percent in column 1, compared to incumbents whose vote share is 4.2 percentage points larger and community nominees whose vote share is 13.4 points larger (both differences are significant at the 99 percent confidence level). Similar extensive and intensive margin differences in magnitude and significance obtain in the stronghold sample (columns 4 to 6). Post-estimation tests of equality across coefficients for nominees versus incumbents reject at 99 percent confidence for all six estimates.

Figure 7 breaks out differences in the number of household nominations by type of respondent, relative to status quo applicants (again depicted by the red dashed vertical line at zero). Interestingly, the popularity advantage for top community nominees is apparent among all types of voters: female and younger voters (defined as aged 18 to 35 years), community leaders, chiefs and

elders. For all respondent types, vote counts for incumbents fall in between the vote counts of top nominees and those of status quo applicants, and are statistically distinct from both at the 99 percent confidence level. The data provide a consistent ranking where nominees receive more nominations than incumbents, who in turn receive more than other applicants, across all specifications and subsamples.

VI. Entry into Politics

The second hurdle to clear in improving political selection is the willingness of new talent to enter politics. We define entry as filing an application with a political party to be considered as a candidate. For status quo applicants, this rate by construction is 100 percent as the data on who applied comes from the parties themselves. What this number might be for the community nominees was unknown *ex ante*, particularly as there are few empirical estimates of inducing political entry in the literature. It was also unclear how many incumbents would appear in the applicant pool, as the shift to proportional representation could influence their re-run decisions.

VI.A. Entry of Community Nominees

The revealed baseline entry for lower ranked (outside the top two) nominees in the full national sample is 12.9 percent in Table 6. Interestingly, there is evidence of strong self-selection on quality into entry, as top nominees were 9.6 percentage points (standard error 3.0) more likely to enter than their lower ranked peers, a difference that is significant at 99 percent confidence. As a sanity check, profile sharing had no impact on the entry rates of lower ranked nominees, who did not receive any nudge (estimate of -1.4 points, standard error 2.1). The nudge, which recall was a phone call to a subset of top nominees explaining that they performed best on the screening and would have their profiles shared with parties, did increase entry of top nominees by a sizeable magnitude (4.6 percentage points); however, it is not statistically significant at conventional levels (standard error 3.8). While imprecise, these estimates are causally identified, as recall that T2 was randomly assigned among T1 wards.

Heterogeneity

Column 2 shows that female nominees enter at lower rates than males, by 5.0 percentage points, a difference that is significant at 95 percent confidence (standard error 2.0). Column 3 limits the sample to female nominees and shows that they too exhibit strong positive self-selection on quality, where top women were 11.0 percentage points more likely to enter than lower ranked women, a

difference that is significant at 90 percent confidence (standard error 5.9). The randomized nudge again led to large but imprecisely estimated impacts on entry: the phone call increased entry by 7.4 percentage points (standard error 8.1), which accounts for a 37 percent increase over the rate observed for top women who were not contacted (20.1 percent).

The last set of estimates compares entry rates within and outside party strongholds. Entry of lower ranked nominees does not vary along this margin in column 4, while top nominees are somewhat less likely to enter in strongholds (by 6.0 percentage points, standard error 4.3), so that the self-selection gap between top and lower ranked nominees is smaller there (by 10.3 percentage points, significant at 95 percent confidence). Importantly, though, there are vastly more nominees in strongholds than weakholds (recall that 78 percent of nominees indicated that they were only interested in running for the locally dominant party). Thus estimates in column 5 that restrict attention to the stronghold sample (retaining 1,251 of the total 1,597 individual observations) are very similar to column 1: baseline entry of 12.7 percent for lower ranked nominees and strong positive self-selection on quality of 10.0 percentage points (3.7 standard error) for top nominees. Note that there is no discernible response to the profile nudge among top picks in the stronghold sample (estimate of -0.2 percentage points, 4.5 standard error).

Relation to the Literature

Are these entry rates, which range from an average of 9.1 percent for lower ranked women to 25.4 percent for top males, low or high? For starters they are much lower than the community nominees' stated interest in running for local office, which recall was 89 percent for top nominees. They are in the broad range of estimates in the literature, notwithstanding stark differences in context and intervention. An NGO targeting young progressives for local U.S. races saw a conversion rate from website intake forms through to candidacy of 11.6 percent (Green, Conroy, & Hammond, 2024). A candidate training program in Brazil saw that application to the program accounted for a 16 percent entry rate, which doubles with participation in the training (Dal Bó, Ferraz, Finan, & Pessoa, 2024).

Does a nudge effect of 5 to 7 percentage points align with expectations? The phone call informing top nominees of their status increased entry from 22.5 to 27.1 percent on average in the national sample, which is a 20 percent increase. In Pakistan, Gulzar & Khan (2025) find that using a pro-social versus neutral prompt increased the likelihood that male household heads stood in nonpartisan village council elections from 3 to 4 percent, or a 33 percent increase on base. By comparison, baseline entry rates in Sierra Leone are higher, which is consistent with leveraging community knowledge to target more likely prospects than randomly selected citizens. The percent increase is lower, perhaps reflecting the lighter touch nature of a phone call versus in-person promotion of public service, which was used in the Pakistan study and worked only when delivered in a group setting.

Self-selection on Quality

The observed self-selection on quality measures is a striking and potentially optimistic result. Figure 8 unpacks this result and shows that all five component areas of the quality index positively predict entry among nominees. In addition, those who have more local support in the nominations data are also more likely to run (recall that all nominees were informed when contacted for screening of the number of households who had put their name forward), as are those who expressed stronger interest in running when asked by the survey screening team (measured on a 5-point Likert scale). Nominees from ruling families were significantly less likely to enter, which could align with them realizing that in the existing political equilibrium, ruling house descendants constitute a minority of those selected.

Panel B shows that other characteristics that predict entry suggest that costs matter. For example, wealthier nominees and those already registered with a party (so do not need to pay registration fees or catch up on past membership dues) are more likely to enter. Also, those already plugged into politics, either because they have some political experience or family members in party office, were more likely to enter.

VI.B. Re-running by Incumbents

Historically, there has been very high political turnover from year to year in Sierra Leone. Based on calculations using publicly available data from the National Electoral Commission, of those elected to Council in 2012, only 14 percent (56 of 394 incumbents nationwide) appear again as candidates for the 2018 elections. Of those who re-ran, only 25 were elected for a second time, which implies near total churn: only 6.3 percent of the 2012 cohort served again in the 2018 elected cohort. The switch to closed-list proportional representation in 2023 will encourage more incumbents to run again if they perceive this as tipping the balance from voters to parties in candidate selection, and they are in relatively good standing with their party's leaders.

Re-run rates are indeed much higher for the 2023 elections studied. Homing in on rural areas for comparability with the study sample, Table 7 estimates these differences in incumbent behavior using regression analysis with district fixed effects.[10] The baseline rate in rural areas for those elected in 2012 and running again in 2018 is 14.0 percent. The re-run rate more than tripled for the 2023 election, jumping by a highly statistically significant 34.1 percentage points (standard error 3.5 points) for those elected in 2018. This implies that in roughly half of stronghold wards there is

[10]There are more ward observations in the estimation (293) versus research (250) sample as the latter excludes wards that cross chiefdom boundaries and where piloting activities took place.

an incumbent seeking re-election.[11]

VI.C. From Entry to Compliance

As the experimental design is anchored to ward-based geographical assignment, it is useful to explore how individual entry rates map to ward-level measures of treatment compliance. This matters because individual entry decisions in the 9 to 25 percent range aggregate up to ward-level compliance far below 100 percent, with implications for statistical power under an intention-to-treat design, the severity of which varies by sample and treatment arm.

This conception of compliance goes beyond the intervention happening where it should (which is near perfect) to it generating new applicants for parties to consider, as otherwise it is unclear what exactly there is for parties to respond to in selecting candidates. For the full sample, 133 of the 522 identified top nominees applied for candidacy (or 25 percent). This maps to nominees filing an application (with either party) in 103 of the 186 wards where nominations and screening occurred, for a 55 percent "full" compliance rate for T1 wards. Compliance is 41 percent in the stronghold sample and 33 percent in the full national party-ward sample. This immediately implies that estimation will be higher powered for the stronghold sample, as compliance in the full national sample is diluted by observations in the respective weakholds.

Impacts on power will be more consequential for the profile sharing arm as the starting sample for T2 is smaller than that for T1 and compliance is somewhat lower. The T2 design aimed to share profiles in 195 party-wards in the full national sample, but because there were often no nominees interested in sharing their information with the weakhold party, in practice profiles were available in only 120 party-wards (62 percent of T2 wards). Crossing this with applications data, we observe 72 applications from profiled individuals in 60 distinct wards. This implies an effective "full" compliance rate of 31 percent for T2 wards. Compliance in the stronghold sample is only slightly higher, at 38 percent. Low levels of realized compliance on the smaller T2 sample suggest that statistical power to detect party leader responses to profiles will be limited as receiving a profile of someone who did not apply for candidacy is of little use.

VII. Impacts on the Applicant Pool

This section investigates how the entry of community nominees affects the pool of applicants available for parties to consider when assembling their candidate lists. We first look at ward-level

[11] We reference the stronghold sample because incumbents naturally come from the locally dominant party. In the research sample, incumbents of either party applied for re-selection in 137 of the 250 wards (55 percent). Restricting to the stronghold party, 117 of 152 sitting Councillors re-applied (77 percent).

measures of the size of the pool, where estimates point in the expected direction but are low power. We gain statistical power when turning to individual measures of quality within the pools, where estimates suggest positive treatment effects.

VII.A. Geographic Representation

To see how the observed individual-level entry rates impact the size and quality of the applicant pool, we start with ward-level estimates of the form:

$$Y_w = \beta_0 + \beta_1 T_w^1 + \beta_2 T_w^2 + \mathbf{s}_w + \varepsilon_w \tag{3}$$

where Y is the number of candidate applications filed by individuals registered to vote in ward w; T^1 is an indicator equal to one for assignment to the nominations plus screening treatment arm; T^2 is an indicator for assignment to the profile sharing arm (the omitted category is pure control wards where no interventions were implemented); and $\mathbf{s}$, ε remain as defined above. We focus on the stronghold sample of 250 wards and report full sample analogues in Appendix Table A6. Hypothesis tests of interest are $\beta_1 \geq 0$, which indicates that the combination of community nominations and technical screening expands the applicant pool; and $\beta_2 \geq 0$ indicates that the encouragement aspect of the profile component has an additional marginal effect. For T^1, recall that the script of the private screening included data on how many people in the individual's ward had nominated them and quotations about why those voters thought the person would be a good candidate for Local Council, information that could encourage entry. For T^2, recall that a subset of top nominees received a phone call informing them that they were in fact the top scorers and that their profiles would be shared with parties, which could have an additional effect beyond their learning about voter support. For these entry outcomes, the T^2 nudge works only on the supply-side of encouraging applicants and does not reflect demand-side actions by party leaders, who did not receive the profiles until after the candidate application deadline.

The first tests concern geographic representation. On the extensive margin, the control mean in column 1 of Table 8 shows that nearly every ward (98.4 percent) had at least one local applicant. This leaves little scope for treatment effects, rendering this test moot. (Indeed, the estimated treatment effect for nominations is a null 1.4 percent, standard error 1.3.) Column 2 shows a noisily estimated treatment effect of community nominations and screening on the intensive margin, an increase of 0.22 people (0.25 standard error) on a base of 2.83 applicants per ward, which points in the expected direction but is not statistically distinct from zero. Realized entry rates limit statistical power for this test: the minimum detectable effect size is 0.71 people, so closing the gap between this and the estimated 0.22 would have required 0.49 more applicants per ward, or an average increase in entry of 7.3 percentage points across all nominees screened per ward. On average, about one nominee per

T1 ward filed an application (roughly 0.45 top nominees and 0.60 lower-ranked nominees). Taken at face value, the noisily estimated increase in pool size is smaller than this, so some of these people (particularly lower ranked nominees) may have applied anyway. However, as we show in the next subsection, the best applicants are higher quality in T1 compared to control wards, so our results are also consistent with community nominations and screening shifting who applied (bringing in individuals at the top end), and not only how many applied.

For women, there are again no effects of nominations and screening on the extensive margin in column 3, yet now in column 4 there is a positive and marginally statistically significant impact on the number of female applicants, of 0.22 women (standard error 0.13) on a base of 0.51 female applicants, or a 44 percent increase. This is an early indication that the intervention could operate as a complement to the introduction of the gender quota in helping parties recruit women. Estimates for the full sample are positive across all four outcomes but smaller in magnitude and not statistically significant (Appendix Table A6).

There is no evidence of marginal treatment effects for the phone call nudge at the level of the applicant pool: no estimates in the second row are statistically distinct from zero and three are negative in sign. The lack of effect is unsurprising given earlier null estimates in the stronghold sample for individual responses to the phone call in making entry decisions (shown in Table 6).

VII.B. Quality of the Applicant Pool

The second area of exploration is whether the interventions enhanced the quality of the applicant pool. These analyses take the form:

$$Q_w = \beta_0 + \beta_1 T_w^1 + \beta_2 T_w^2 + \mathbf{s}_w + \varepsilon_w \tag{4}$$

where Q_w is the ward-level measure of the average or maximum quality observed in the applicant pool; and all other terms remain as defined above. We standardize the quality index with respect to the mean and standard deviation of applicant quality in pure control wards.

Table 9 presents evidence of positive treatment effects of the nominations arm on quality, using first the observed and then higher simulated entry rates. The first two columns show standard treatment effect estimates for the realized level of entry by community nominees. Column 1 reveals a positive, but noisy estimate of a 0.052 standard deviation unit (0.044 standard error) increase in the average quality of applicants in wards assigned to nominations and screening. Estimates in column 2 show an improvement in the maximum quality of applicant observed equal to 0.092 standard deviation units, which is significant above the 90 percent confidence level. To gauge magnitude, the control mean implies that the best applicant observed in control wards scored 0.267 standard deviation

units higher than the mean for all applicants in control wards. Relative to this wedge, the estimated treatment effect suggests that nominations and screening increased the quality of the top scorer in T1 wards by roughly 34 percent. As expected based on null results in Tables 6 and 8, there is no evidence for treatment effects of profile nudges.

A useful feature of the experimental design and data collection is that we can simulate treatment effects that would have obtained if more community nominees had filed candidate applications. We conduct a Monte Carlo simulation with 500 replications, drawing a new random sample of simulated entrants in each replication. Columns 3 through 6 do this by retaining the top nominees who filed applications in practice and, for wards where no top nominee applied, we take the highest quality top nominee who was interested in running and boost their entrance randomly at different rates. Columns 3 and 4 add these top nominees with a probability of 25 percent. This leads to limited changes in mean quality, but substantial and precisely estimated increases in the maximum quality observed in the pool, equivalent to 0.121 standard deviation units (0.053 standard error). Columns 5 and 6 boost entry by 50 percent and uncover positive and statistically significant impacts on both the mean and maximum quality of the ward-level pools. These estimates suggest an increase in the mean of 0.074 standard deviation units (standard error 0.044) and an increase in the maximum of 0.149 standard deviation units (standard error 0.053). Analogous simulations for the full 500 party-ward sample are all positive but smaller in magnitude and not statistically significant (Appendix Table A7).

The simulations reveal untapped potential that could be unlocked through greater investment in converting high-quality nominees to applicants. This resonates with the evidence presented in Table 3, which shows that the top nominees who did enter are of measurably higher quality than the best status quo applicants found in pure control wards. Greater effort to convert just one top pick in each ward identified by communities would have provided party leaders with more high-quality options to choose from and increased precision of treatment effect estimates. The encouragement component of phone calls to top nominees was the lightest touch aspect of the intervention. Evaluating what types of conversion effort or resources would be more effective is a priority for future research.

VIII. Party Gatekeeping and Candidate Selection

The translation of upstream impacts on identifying new talent and encouraging entry, to potential downstream impacts on candidacy and elected seats, depends crucially on the third and final issue of party gatekeeping. This marks the transition from studying candidate entry issues (a supply-side feature) to candidate selection (on the demand side). To better assess the scope for such translation, we first use the data to characterize party strategy in candidate selection and how it may have responded to the switch to PR, and then consider their responsiveness to the interventions.

VIII.A. Parties' Geographic Strategy

An informative place to start is some basic statistics about geographic representation in the raw data. The first thing to investigate is whether the switch to PR skewed representation, on candidate lists and elected seats, away from uniform representation across wards as mandated by the old majoritarian structure.

Figure 9 suggests that parties largely adhered to the old ward-based system. Using National Electoral Commission data, covering both parties across all rural districts nationwide, panel A shows that the modal ward has two candidates on each party's district list. There are wards with only one candidate and others with two or more, but overall, there are very few wards (only 6.2 percent) with no representation on a district list of one party or the other. This is consistent with parties taking a straightforward approach to the electoral reform by simply doubling the number of candidates they selected from each ward.

Panel B shows a similar pattern for ward representation in elected seats. This panel takes the National Electoral Commission's official results on everyone who won a seat in a Local Council, regardless of their party, and traces these elected officials back to the ward where they are registered to vote. It shows that the vast majority of wards can be linked to exactly one elected official. This pattern is broadly consistent with parties ordering their respective lists so that each ward has one candidate in the top half of the list (which covers the maximum number of seats they could possibly win), and one in the bottom half of the list (which comprises the waitlist to replace anyone who leaves office during their term). Nonetheless, a nontrivial minority of rural wards nationwide (22 percent) lost direct geographic representation in their elected Local Council.

Econometrically, the fact that parties largely adhered to the old ward-based system suggests that potential interference between units for T2—i.e., selecting profiled candidates from treatment wards directly at the expense of status quo candidates in control wards—is unlikely to present a major challenge for estimating potential treatment effects. But we can test this formally, leveraging the constituency-level saturation design in the assignment of profile sharing. Appendix Exhibit A4 provides details on the motivation and structure of the empirical specifications. As expected, estimates for spillover effects on geographic representation are null (in Appendix Table A17).

VIII.B. Strongest Predictors of Party Selections

We next examine individual characteristics that predict whom parties select from the applicant pool to advance to candidacy on a district list. These estimates are run on the realized applicant pools, which include all status quo applicants, incumbents who ran for re-election and community nominees who filed a candidate application. For simplicity we focus on regression estimates of how different traits influence the probability of candidacy within our randomization strata (i.e.,

fixed effects by district and competitiveness).

The two characteristics that stand out across the panels of Figure 10 are incumbency and female gender, which are the strongest positive predictors of candidacy. Gender straightforwardly responds to the quota, and the new policy achieved an overall share of elected female Councillors just shy of its 30 percent target. Incumbency preference breaks with past precedent.

Party leaders' preference for incumbents is observable in the raw data tracking how the three different groups of potential candidates progress through the selection process. Figure 11 shows that conditional on filing an application, sitting Councillors are the most likely to get selected onto a candidate list (at 82 percent in panel A), followed by top nominees (69 percent) and status quo applicants (59 percent). Ordering these groups by size shows top nominees (84) and incumbents (117) are both well behind status quo applicants (415). Conditional on candidacy, incumbents are vastly more likely to win a seat (at 67 percent), compared to nominees and status quo applicants (both 33 percent). This is consistent with party leaders stacking the top of candidate lists with incumbents.

Placing this in historical perspective, the shares of both those who reached a second candidacy and re-won seats increased dramatically compared to the 2012 to 2018 political cycle (see Appendix Figure A7). In regression analysis, conditional on running again, incumbents increased their re-election rates by 21.6 percentage points (standard error 8.7 points) on a base rate of 36.6 percent (in Table 7). Overall, this increases the stickiness of representation from an average of 5.1 percent incumbent carryover in 2018 to 30.7 percent in 2023. This reduction in political renewal substantially narrows the space for new entrants, including both community nominees and those challenging incumbents in the status quo.

Increased privileging of incumbents is consistent with a broader literature showing that closed-list proportional representation systems concentrate power in party leadership, benefiting incumbents through party-controlled list placement at the expense of voters. This provides an explanation for why the incumbency advantage is substantially larger under closed-list versus open-list systems (Fiva & Røhr, 2018; Fiva & Smith, 2018). The lack of entry and incumbent preference can further reduce individual accountability and explain why closed-list proportional representation systems tend to exhibit higher levels of rent-seeking relative to plurality systems (Kunicová & Rose-Ackerman, 2005). This is of particular concern in weaker states that struggle to control corruption more generally. It is in these same settings that one hopes that higher quality leaders might compensate for lower quality institutions (Carreri, 2021), compounding the political selection problem.

VIII.C. Selection on Quality

Returning to Figure 10, parties do not demonstrate a strong overall pattern of selection with respect to our quality measures. Parties appear to positively select on local knowledge and elected political experience, and somewhat on popularity as measured by the community nominations data. Parties appear to negatively select on human capital and work experience, and somewhat so on managerial capital, all of which are areas that the intervention sought to identify people with strength in.

Tracking selection on quality throughout the pipeline of entry, candidacy and winning a seat, Figure 12 illustrates the degree of observed selection on the quality index for each type of applicant. For entry, the first set of columns compares the mean quality score for each type with reference to the average score among all status quo applicants (shown by the first green bar which is standardized to zero). The positive bars for incumbents who ran for re-election (red bar) and top community nominees who filed applications (blue bar) suggest that they are on average higher quality than status quo entrants. The negative (yellow) bar for lower ranked nominees (on the screening instrument) suggests that nominees outside the top picks who filed applications are relatively lower quality than those filing in the status quo.

The next set of columns reflects candidate selections by parties. The heights of all four bars remain similar to their counterparts at the application stage, implying neutral selection by parties with respect to quality in choosing candidates to place on their district lists. The final set of columns is for election, which is driven by how parties ranked individuals on the list. The heights of the blue and green bars increase, indicating that parties positively select on quality in ordering status quo and top nominee candidates toward the top of their lists. The red and yellow bars hold their heights, suggesting quality-neutral ordering of incumbents and low-ranked applicants on the candidate lists.

VIII.D. Treatment Effect Estimates on Candidate Selection

To evaluate how parties responded to treatment, we start with geographic representation and test whether the intervention impacted the likelihood that a given ward had a candidate on the party's district-level list, and whether anyone from that ward won a seat. Analysis takes the form of Equation (3). A key difference in interpretation is that the β_2 estimated for T^2 now captures whether the information shared with parties impacts the probability that they select a profiled nominee.

As presaged in the raw data on ward representation from Figure 9, baseline rates of geographic representation on the dominant party's list are high, leaving little scope for treatment effects. The control mean in Table 10 column 1 shows that 93.7 percent of wards in this sample did indeed have at least one local candidate on the dominant party's list. The estimated treatment effect for nominations and screening, which this high base renders largely moot, is correspondingly small in

magnitude and imprecisely estimated (0.9, standard error 3.7).

For geographic representation by female candidates, estimates in column 2 suggest that the nominations and screening may have been helpful in this regard. In control group wards, the baseline rate of having a female candidate from the ward on the dominant party's district list was 34.9 percent, which increases with nominations by 13.8 percentage points (standard error 7.8 points). This effect is marginally statistically significant and concentrated in areas where no profiles were shared.

There is no evidence that nominations and screening impacted either the rank on the candidate list (see also Appendix Table A10) or the likelihood that someone registered to vote in the ward got elected. Estimates in columns 3 and 4 are small in magnitude and not statistically distinct from zero. This is perhaps unsurprising in light of parties' newly increased preference for re-running incumbents.

For profile sharing, the one notable estimate in Table 10 is a positive and statistically significant treatment effect on the inverse rank of candidates from profile sharing wards on the list. Among wards with at least one candidate on the list, candidates from T2 wards are on average ranked 4.4 positions higher (as in closer to the first rank on a descending list) on their district lists relative to candidates from wards with nominations but no profile sharing. Note that this estimate is relative to nominations-only wards, where candidates were ranked 2.8 positions lower than in control wards (not statistically significant). Relative to pure control wards, candidates from profile sharing wards rank 1.7 positions higher, a difference that is not statistically significant (p-value 0.32).

Given power issues, it is useful to complement this with a look at the raw data. Figure 13 compares outcomes for top nominees in the T1 only arm (in red) to those in the T1+T2 arms (in blue), which recall was randomly assigned. The share of top nominees is higher at each stage of entry, candidacy and election, although these differences are not statistically significant. Appendix Figure A8 maps the probability density functions of candidates' numerical ranks on party lists separately for wards randomly assigned to profile sharing (T2) versus those assigned to T1 only. The dashed vertical line at rank 22 denotes the average last rank on the dominant party's list to win a seat in its regional stronghold. Profile wards have more density to the left of this line (or ranks closer to the top of the list), suggesting that parties may have been open to information in ranking candidates. Appendix Table A10, which tracks whether a ward has a candidate in each quartile of the party list, similarly provides suggestive evidence of mild effects on candidate positioning.

Turning to the quality of the candidate pool, effects in Table 11 mirror those seen earlier for the applicant pool. The nominations intervention shifted the maximum quality observed among candidates on a list upwards, by 0.117 standard deviation units of the quality index, which is marginally statistically significant (standard error 0.061). Treatment effect estimates on the mean quality of candidates are positive but imprecise. Coefficient estimates for impacts on the mean and maximum quality of elected politicians from treatment wards are positive in sign but not statistically

distinct from zero. These last estimates also face the issue that quality is not observed for wards with no winning candidate (roughly 30 percent of the sample), so is run on a selected subset of wards with winners. There is no evidence that profile sharing had any additional marginal effects on the quality of selected candidates or elected officials, as all four corresponding estimates are small in magnitude, negative in sign and imprecisely estimated. Estimates for the full sample are null (Appendix Tables A8 and A9), as are estimates that adjust for multiple inference testing across all prespecified hypotheses in both the stronghold and full samples (Appendix Table A11 and Exhibit A3).

VIII.E. Interpretations of Party Gatekeeping

There are multiple ways to interpret the fact that parties privileged incumbents over top community nominees and status quo challengers in ordering their candidate lists. The most immediate explanation is that the switch to PR created institutional incentives to advance incumbents. An interesting question is how this reprioritization may have impacted the quality profile of elected Councillors. While multiple factors are at play, and statistical power to estimate treatment effects is limited, several empirical patterns help inform what potential consequences might be for the productivity of elected Councils.

First, it is clear that parties passed over higher quality community nominees, as measured by our screening protocol, to advantage incumbents. When we restrict estimates of relative quality to include only nominees who applied for candidacy and incumbents who ran for re-election, top nominees retain a large and highly significant quality advantage (of 0.207 standard deviation units, p-value < 0.001 in Appendix Table A12). This quality advantage survives when further limiting to wards where nominees and incumbents ran against each other (of 0.161 units, p-value 0.050).[12] In these head-to-head races, top nominees also have significantly higher vote shares in the household nominations data (by 14.6 percentage points, p-value < 0.001 in Appendix Table A13). Nor does ethnic politics appear to explain this choice: the ethnic composition of top nominees closely resembles that of incumbents, and 65 percent of top nominees (341 of 522) share the ethnicity of the incumbent in their ward, a share that is similar in the stronghold sample and among those who applied.

Second, parties also passed over higher quality status quo applicants to favor incumbents. To see this, Appendix Figure A9 compares the observed shares of elected Councillors by type, with a simulation of the shares that would obtain under a strict rank ordering based on the quality index. If parties had ranked applicants by the quality index and ordered their lists accordingly (while respecting the gender quota), the share of incumbents making it high enough on the list to win a

[12] All estimates in Appendix Table A12 exclude the conscientiousness component as it was not measured for incumbents.

seat would fall from 30 to 15 percent (dropping 46 incumbents). These winning ranks would be replaced by more top nominees (increasing from 8 to 18 percent for a gain of 28 individuals) and high-quality status quo challengers (whose share would increase from 51 to 61 percent, for 30 more people). Lower ranked nominees would decrease somewhat, from a low base (10 to 6 percent).

To the extent that the quality components we are able to measure robustly predict Council performance, this suggests negative impacts on productivity for the elected group. Note that parties could not readily observe all the quality measures we collected, apart from the information shared in the one-page profiles of some top nominees. Yet parties might have their own information that is unobservable to the research team, in particular about the returns to experience on performance, which could explain their choices and net out positively for Council productivity.

To assess the scope for this, we focus on what we all can observe: years in office, which is what incumbency most directly implies, and how well it predicts politicians' views of each other's performance in our incumbent survey data. For outcomes, we use first an average of the anonymous peer reviews on effectiveness and honesty; and second, self-reported metrics of performance in office, namely engagement with voters and the number and average value of development projects they mobilized. We benchmark years in office (which ranges from less than one year to 16 years in the sitting incumbent data) to the individual components of the quality index. We add to this the two other visible predictors of parties' candidate selections from Figure 10, namely gender and political experience. The goal is a series of "horse races" between the quality index components and what we observe parties selecting on in terms of predicting performance.

Appendix Tables A14 and A15 report results for the peer review and self-reported indices, respectively. Column 1 reports how often an elastic-net regression, which drops weak predictors, retains each variable when run on many random halves of the data and across a range of increasingly strict penalties (Meinshausen & Bühlmann, 2010; Shah & Samworth, 2013). A high share means the variable predicts performance consistently, and not only in one draw of the data or one choice of model. Column 2 reports the share of 300 elastic-net runs in which each variable is selected, with the penalty chosen by cross-validation using newly drawn folds in each run. Because penalties with similar predictive performance can select different sets of variables, these selection frequencies are sensitive to the random partition of the data.[13] Columns 3 and 4 report bivariate and multivariate coefficients to gauge magnitudes.

Several components of the quality index are more consistent predictors of performance than years in Council. Specifically, 9 of 13 index components rank higher with positive and at least marginally statistically significant coefficient estimates in one table or the other, and two (project proposal score and development projects) do so in both. Estimated coefficients for years in Council are small and not statistically distinct from zero in every specification. Political experience fares no better. The differences in magnitude are large: a one standard deviation increase in the project proposal score is

[13] See Ritzwoller & Romano (2025) on this instability and why we prefer the approach in column 1.

associated with a 0.206 standard deviation increase in the peer review index (standard error 0.050), compared to 0.021 (standard error 0.053) for years in Council. As an alternative to incumbency, the project proposal measure offers an easy-to-implement screening device: it is straightforward to administer, is a hypothetical that applies to everyone equally, and is clearly related to the daily work of managing local infrastructure projects.

Appendix Table A16 sharpens the comparison by including the quality index and years in Council in the same regression. For peer reviews, the coefficient on the quality index is roughly 20 times larger than that on years in Council (0.199 versus 0.011 without controls), and the latter turns negative once controls and district fixed effects are added. For self reports, the quality index coefficient is about three times larger (0.179 versus 0.051), and only the quality index is statistically significant across specifications. At least based on this descriptive exercise, it seems unlikely that the returns to experience in office are high enough relative to our quality metrics to rationalize parties' choice to prioritize incumbents over higher quality challengers. Parties' choices thus appear more consistent with the incentives created by the switch to PR than with an assessment of who would perform best in office.

IX. Discussion and Conclusion

Widespread discontent with politicians and their performance calls for creative solutions to bring new talent into public office. While appealing in theory, there are reasons to doubt whether this can be accomplished in practice. It may be difficult to identify individuals with the right skills and sufficient public support; promising recruits might be unwilling to endure the hardships of campaigning or frustrations of government service; and even if they were, political parties may not be interested in sponsoring them as candidates. Breakdowns at any point in this chain will undermine efforts to improve political selection.

The intervention studied in Sierra Leone successfully cleared the first two supply-side hurdles. It identified popular, high-quality individuals and induced many of them to file an application to be formally considered by a party for candidacy. Structured household consultations surfaced names of individuals with broad support for their fitness for office, lists that were honed on quality using measures that correlate with on-the-job performance validated in related studies. With screening data on all party applicants—including those associated with the intervention, re-running incumbents, and others who applied in the status quo—analysis shows that these new entrants were on average higher quality than their competitors. While they were also more representative in the sense that a larger share of household respondents said they would like the person to represent them in local government, they were nonetheless socioeconomic elites, strongly selected from traditional chiefly lineages.

Patterns of entry among community nominees reveal strong self-selection on measures of quality, with those who scored highest on technocratic screening entering at rates more than 50 percent higher than others. If we were to run this study again, we would invest more resources in converting the best quality nominees into applicants. This part of the intervention was light touch: consisting of providing nominees with information about their popularity and testimonials from community members about their fitness for office, as well as an encouragement phone call (to a subset). Using data on all actual nominees, simulations show how boosting their entry rates even by a little yields large and precisely estimated impacts on the quality of applicants available for parties to choose in selecting candidates.

The Government of Sierra Leone implemented a gender quota for the first time in these 2023 elections, which raises the importance of finding strong women to run. While women faced higher barriers to entry than men, those who cleared them performed relatively well. Voters showed bias against women in nominating fewer of them, but those women whose names surfaced enjoyed broad support. Female nominees on average performed less well than their male peers on the merit-based screening, but those who ranked highly on merit closed much of this gender performance gap compared to top ranking men. There is suggestive evidence that nominations and screening helped women from treatment wards get onto party lists. These results collectively highlight potential for recruitment efforts to complement gender quotas, as parties need to source many more female candidates than they are used to while respecting tight electoral calendars.

The impacts of community nominations, screening and profile sharing largely stalled when it came to securing a rank at the top of party lists, the essential final step in getting elected. This demand-side breakdown reflects intensified party gatekeeping that favored incumbents in response to the abrupt shift from majoritarian to proportional representation seen in the election studied. By increasing the rate at which incumbents both sought and won re-election (from 5 to 31 percent), political renewal dropped dramatically in the space of a single electoral cycle. This reduced scope for challengers—from both within and beyond the intervention studied—to gain a foothold in elected office, including those of demonstrably higher quality. Another potentially pernicious impact of the switch to proportional representation is that 21 percent of wards in the research sample lost geographic representation in that they no longer have any elected official from their ward serving in their Local Council.

In sum, gathering local information on who is best fit for public office from community nominations, and sharpening it with merit-based screening, proved highly effective in identifying new political talent. A quarter of the top nominees filed candidate applications with parties, providing proof of concept that there is new talent out there that can be brought into politics. Getting them elected, however, requires a different approach to party gatekeeping that favors incumbents, and particularly so for electoral systems like PR that privilege parties over voters in selecting individual candidates.

References

Acemoglu, D., Reed, T., & Robinson, J. A. (2014). Chiefs: Economic development and elite control of civil society in Sierra Leone. *Journal of Political Economy*, *122*(2), 319–368.

Acemoglu, D., & Robinson, J. A. (2012). *Why nations fail: The origins of power, prosperity, and poverty*. Crown Business.

Afrobarometer. (2024). *Afrobarometer round 9 survey in Sierra Leone, 2022 [Data set]*. http://www.afrobarometer.org.

Ahuja, A., & McClendon, G. (2026). The effects of exposure to new electoral rules: Field experimental evidence from Sierra Leone. *American Political Science Review*, 1–20.

Anderson, M. 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), 1355–1394.

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

Balán, P., Bergeron, A., Tourek, G., & Weigel, J. L. (2022). Local elites as state capacity: How city chiefs use local information to increase tax compliance in the Democratic Republic of the Congo. *American Economic Review*, *112*(3), 762–797.

Baldwin, K. (2016). *The paradox of traditional chiefs in democratic Africa*. Cambridge University Press.

Baldwin, K. (2025). Traditional authorities around the world: Toward a better understanding of their de facto power, downward accountability, and state recognition. *Annual Review of Political Science*, *28*, 95–114.

Baldwin, K., & Raffler, P. (2019). Traditional leaders, service delivery, and electoral accountability. In J. A. Rodden & E. Wibbels (Eds.), *Decentralized governance and accountability: Academic research and the future of donor programming* (pp. 61–90). Cambridge University Press.

Basurto, M. P., Dupas, P., & Robinson, J. (2020). Decentralization and efficiency of subsidy targeting: Evidence from chiefs in rural Malawi. *Journal of Public Economics*, *185*, 104047.

Bazzi, S., Hilmy, M., Marx, B., Shaukat, M., & Stegmann, A. (2026). *It takes a village election: Turnover and performance in local bureaucracies.* (Working Paper)

Benjamini, Y., Krieger, A. M., & Yekutieli, D. (2006). Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*, *93*(3), 491–507.

Besley, T., Folke, O., Persson, T., & Rickne, J. (2017). Gender quotas and the crisis of the mediocre man: Theory and evidence from Sweden. *American Economic Review*, *107*(8), 2204–2242.

Besley, T., Montalvo, J. G., & Reynal-Querol, M. (2011). Do educated leaders matter? *Economic*

Journal, *121*(554), F205–F227.

Besley, T., Pande, R., & Rao, V. (2005). *Political selection and the quality of government: Evidence from South India* (Working Paper No. 921). Economic Growth Center, Yale University.

Bowles, J., & Marx, B. (2023). *Turnover and accountability in Africa's parliaments.* (Working Paper)

Carreri, M. (2021). Can good politicians compensate for bad institutions? Evidence from an original survey of Italian mayors. *The Journal of Politics*, *83*(4), 1229–1245.

Casey, K., Glennerster, R., Miguel, E., & Voors, M. (2023). Skill versus voice in local development. *Review of Economics and Statistics*, *105*(2), 311–326.

Casey, K., Kamara, A. B., & Meriggi, N. F. (2021). An experiment in candidate selection. *American Economic Review*, *111*(5), 1575–1612.

Chang, E. C., & Golden, M. A. (2007). Electoral systems, district magnitude and corruption. *British Journal of Political Science*, *37*(1), 115–137.

Chattopadhyay, R., & Duflo, E. (2004). Women as policy makers: Evidence from a randomized policy experiment in India. *Econometrica*, *72*(5), 1409–1443.

Cheng, C., & Tavits, M. (2011). Informal influences in selecting female political candidates. *Political Research Quarterly*, *64*(2), 460–471.

Clayton, A. (2021). How do electoral gender quotas affect policy? *Annual Review of Political Science*, *24*, 235–252.

Colonnelli, E., Prem, M., & Teso, E. (2020). Patronage and selection in public sector organizations. *American Economic Review*, *110*(10), 3071–3099.

Cruz, C., Keefer, P., & Scartascini, C. (2021). *The database of political institutions 2020 (DPI2020)* (Tech. Rep.). Inter-American Development Bank, Department of Research and Chief Economist.

Dal Bó, E., Ferraz, C., Finan, F., & Pessoa, P. (2024). *When democracy refuses to die: Evaluating a training program for new politicians* (NBER Working Paper No. 33251). National Bureau of Economic Research.

Dal Bó, E., & Finan, F. (2018). Progress and perspectives in the study of political selection. *Annual Review of Economics*, *10*, 541–575.

Dal Bó, E., Finan, F., Folke, O., Persson, T., & Rickne, J. (2017). Who becomes a politician? *Quarterly Journal of Economics*, *132*(4), 1877–1914.

Dal Bó, E., Finan, F., & Rossi, M. A. (2013). Strengthening state capabilities: The role of financial incentives in the call to public service. *Quarterly Journal of Economics*, *128*(3), 1169–1218.

de Kadt, D., & Larreguy, H. A. (2018). Agents of the regime? Traditional leaders and electoral politics in South Africa. *Journal of Politics*, *80*(2), 382–399.

Fiva, J. H., & Røhr, H. L. (2018). Climbing the ranks: Incumbency effects in party-list systems. *European Economic Review*, *101*, 142–156.

Fiva, J. H., & Smith, D. M. (2018). Political dynasties and the incumbency advantage in party-centered environments. *American Political Science Review*, *112*(3), 706–712.

Fujiwara, T., Hilbig, H., & Raffler, P. (2025). Biased party nominations as a source of women's electoral underperformance. *American Journal of Political Science*. (Forthcoming)

Gagliarducci, S., & Nannicini, T. (2013). Do better paid politicians perform better? Disentangling incentives from selection. *Journal of the European Economic Association*, *11*(2), 369–398.

Glennerster, R., Miguel, E., & Rothenberg, A. D. (2013). Collective action in diverse Sierra Leone communities. *Economic Journal*, *123*(568), 285–316.

Government of Sierra Leone. (1991). *The constitution of Sierra Leone.* Act No. 6 of 1991.

Government of Sierra Leone. (2022). *The local government act.* Supplement to the Sierra Leone Extraordinary Gazette Vol. CXLXIII, No. 91 dated 30th December 2022.

Green, J., Conroy, M., & Hammond, C. (2024). Something to run for: Stated motives as indicators of candidate emergence. *Political Behavior*, *46*(2), 1281–1301.

Gulzar, S. (2021). Who enters politics and why? *Annual Review of Political Science*, *24*, 253–275.

Gulzar, S., & Khan, M. Y. (2025). Good politicians: Experimental evidence on motivations for political candidacy and government performance. *Review of Economic Studies*, *92*(1), 339–364.

Gulzar, S., Pathak, D., Thompson, S., & Tóth, A. (2025). Can party elites shape the rank and file? Evidence from a recruitment campaign in India. *American Political Science Review*, *119*(2), 812–831.

He, G., & Wang, S. (2017). Do college graduates serving as village officials help rural China? *American Economic Journal: Applied Economics*, *9*(4), 186–215.

Henn, S. J. (2023). Complements or substitutes? How institutional arrangements bind traditional authorities and the state in Africa. *American Political Science Review*, *117*(3), 871–890.

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.

Jones, B. F., & Olken, B. A. (2005). Do leaders matter? National leadership and growth since World War II. *The Quarterly Journal of Economics*, *120*(3), 835–864.

Kandeh, J. D. (1992). Politicization of ethnic identities in Sierra Leone. *African Studies Review*, *35*(1), 81–99.

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.

Kling, J. R., Liebman, J. B., & Katz, L. F. (2007). Experimental analysis of neighborhood effects. *Econometrica*, *75*(1), 83–119.

Kunicová, J., & Rose-Ackerman, S. (2005). Electoral rules and constitutional structures as constraints on corruption. *British Journal of Political Science*, *35*(4), 573–606.

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), 2213–2246.

Mahmoud, Y. (2014). *Pervasive disconnect between state and citizen offers two paths: Promise or peril.* IPI Global Observatory. https://theglobalobservatory.org/2014/06/pervasive-disconnect-between-state-citizen-offers-paths-promise-peril/. (June 16)

Martínez-Bravo, M. (2017). The local political economy effects of school construction in Indonesia. *American Economic Journal: Applied Economics*, *9*(2), 256–289.

Marx, B., Pons, V., & Rollet, V. (2025). Electoral turnovers. *Review of Economic Studies*, *92*(5), 3306–3339.

Meinshausen, N., & Bühlmann, P. (2010). Stability selection. *Journal of the Royal Statistical Society, Series B*, *72*(4), 417–473.

Mettler, S. (2018). *The government-citizen disconnect*. Russell Sage Foundation.

OECD. (2024). *OECD survey on drivers of trust in public institutions – results: Building trust in a complex policy environment.* OECD publishing. https://doi.org/10.1787/9a20554b-en.

Pande, R. (2003). Can mandated political representation increase policy influence for disadvantaged minorities? Theory and evidence from India. *American Economic Review*, *93*(4), 1132–1151.

Persson, T., Tabellini, G., & Trebbi, F. (2003). Electoral rules and corruption. *Journal of the European Economic Association*, *1*(4), 958–989.

Pew Research Center. (2023). *Americans' dismal views of the nation's politics.* https://www.pewresearch.org/politics/2023/09/19/americans-dismal-views-of-the-nations-politics/.

Querubín, P. (2016). Family and politics: Dynastic persistence in the Philippines. *Quarterly Journal of Political Science*, *11*(2), 151–181.

Ravanilla, N. (2016). *Attracting good people into public service: Evidence from a field experiment in the Philippines.* (Working Paper)

Riaño, J. F. (2023). *Bureaucratic nepotism.* (Available at SSRN 3995589)

Ritzwoller, D. M., & Romano, J. P. (2025). *Reproducible aggregation of sample-split statistics.* Working Paper.

Schäfer, A., & Zürn, M. (2023). *The democratic regression: The political causes of authoritarian populism*. Polity Press. (S. Curtis, Trans.)

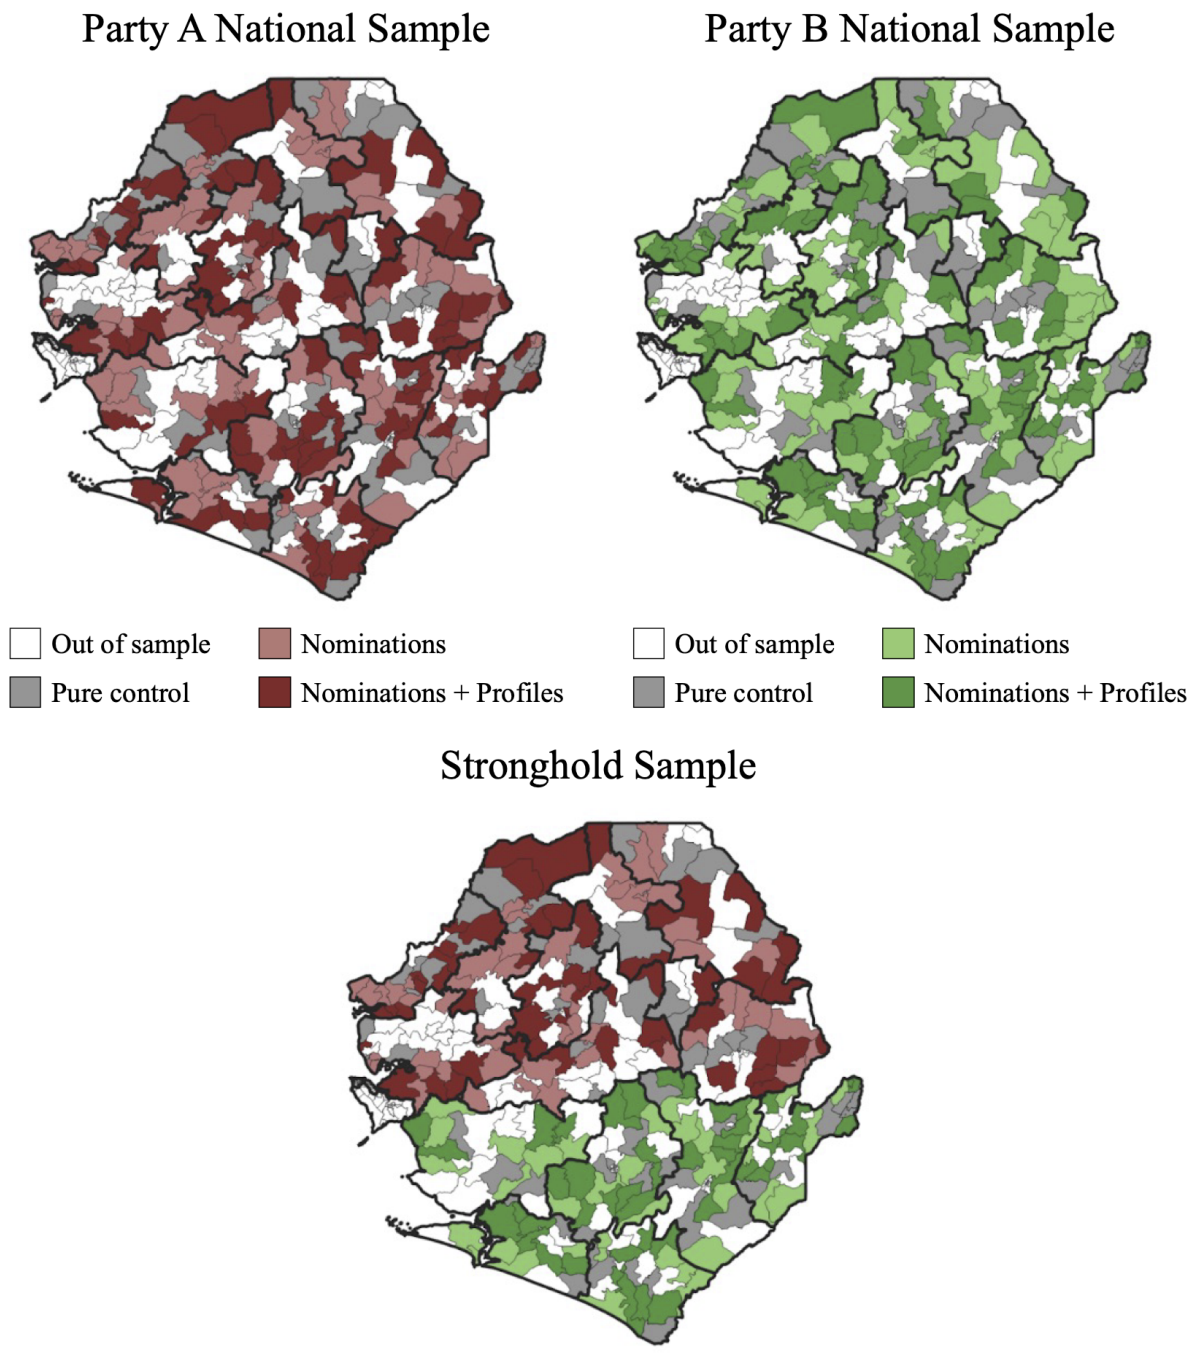
Shah, R. D., & Samworth, R. J. (2013). Variable selection with error control: Another look at stability selection. *Journal of the Royal Statistical Society, Series B*, *75*(1), 55–80.

Valgarðsson, V. O., Jennings, W., Stoker, G., Bunting, H., Devine, D., McKay, L., & Klassen, A. (2025). A crisis of political trust? Global trends in institutional trust from 1958 to 2019. *British Journal of Political Science*, *55*, e15.

World Bank. (2024). *Worldwide governance indicators [Data set].* http://www.govindicators.org. (Accessed 4/30/2025)

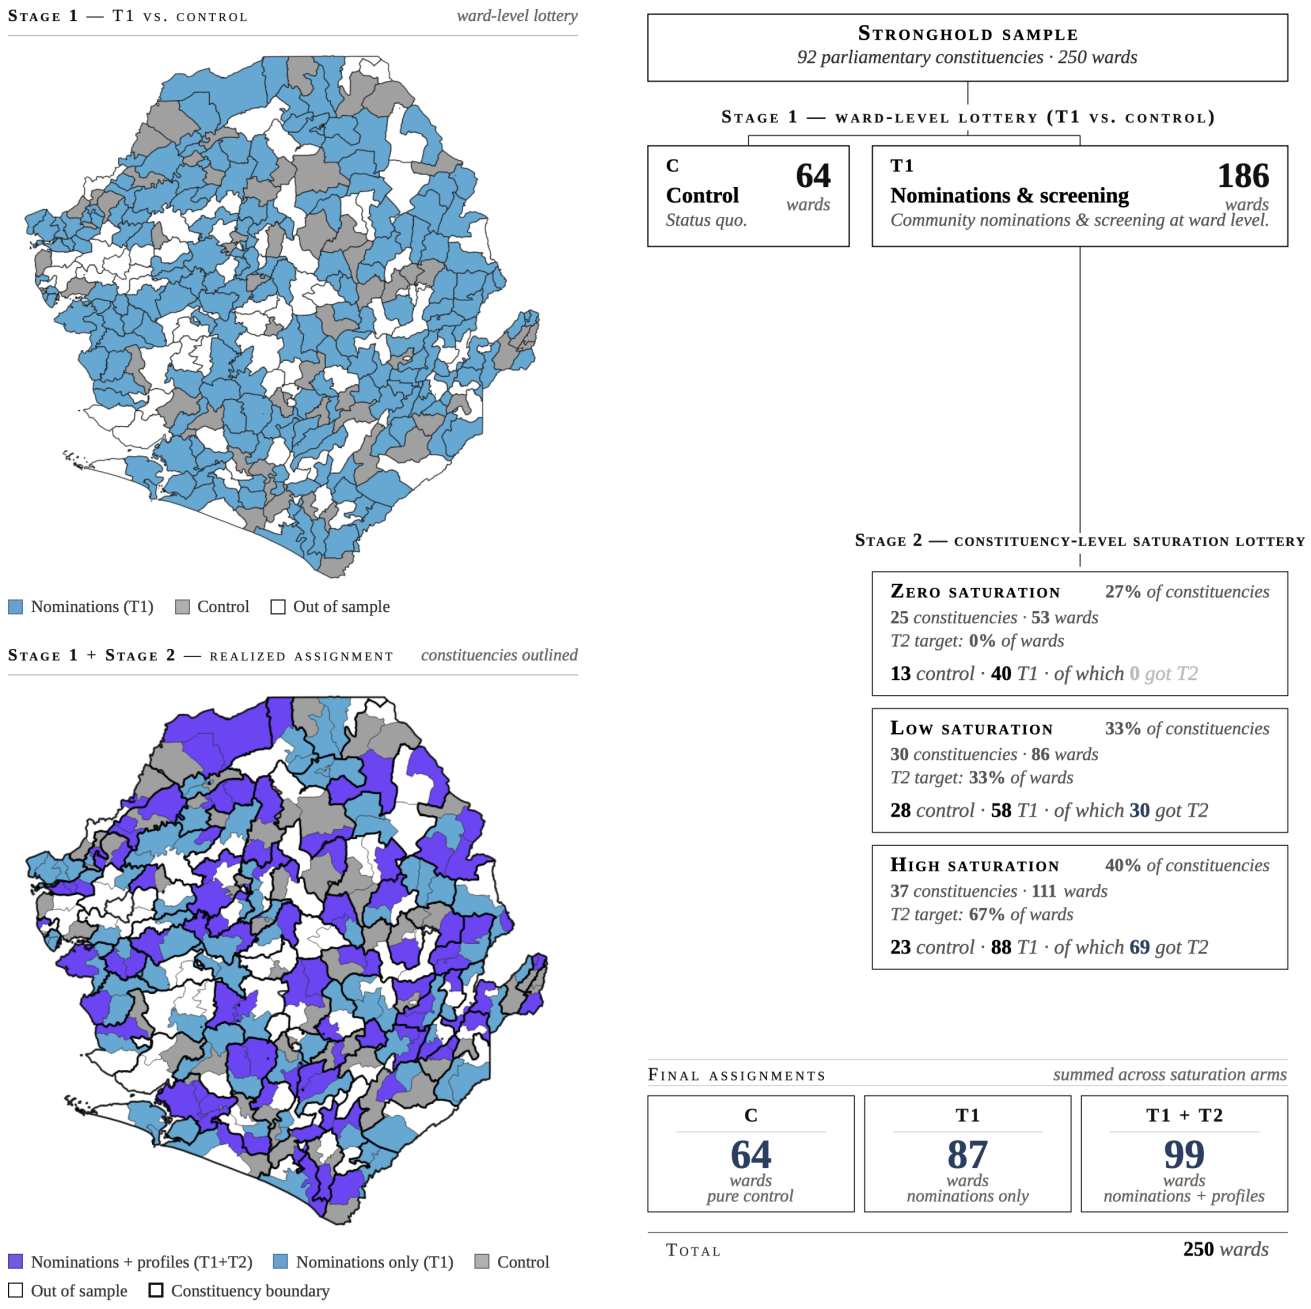
Figures and Tables

Figure 1: National Party Samples and the Stronghold Sample



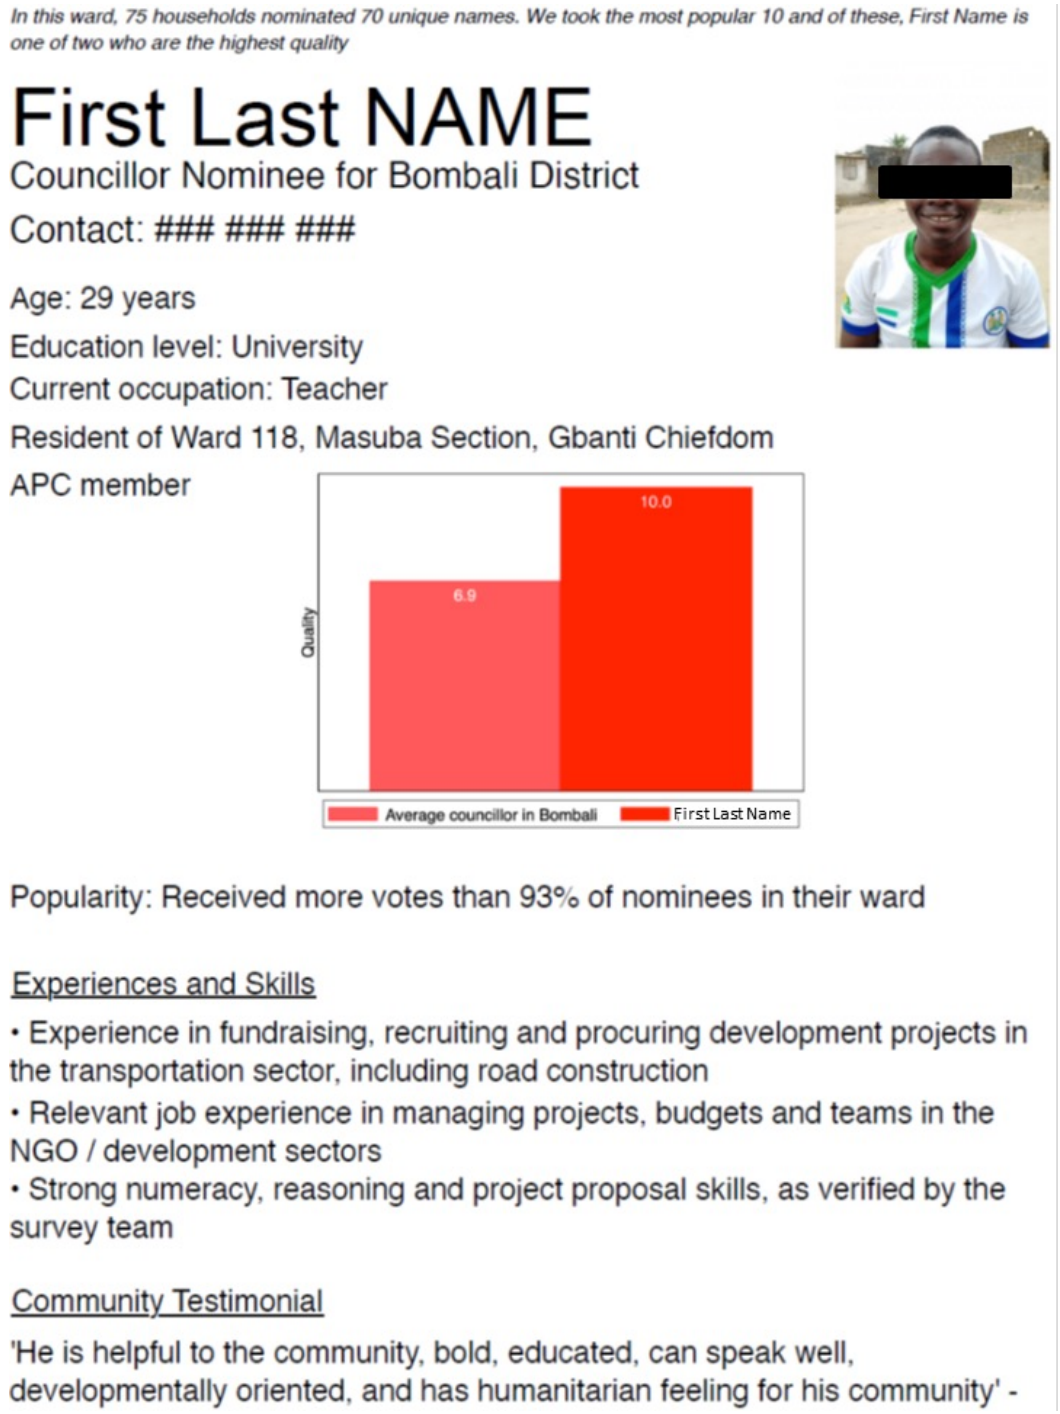
Notes: This figure maps the geographic scope of the study. The top panel shows wards assigned to treatment for both parties in the full sample. This comprises the full sample of 500 party-ward units. The bottom panel delineates the stronghold sample of 250 wards, which assigns each ward only to one party, the locally dominant one based on 2018 district-level electoral data.

Figure 2: Experimental Design for the Stronghold Sample



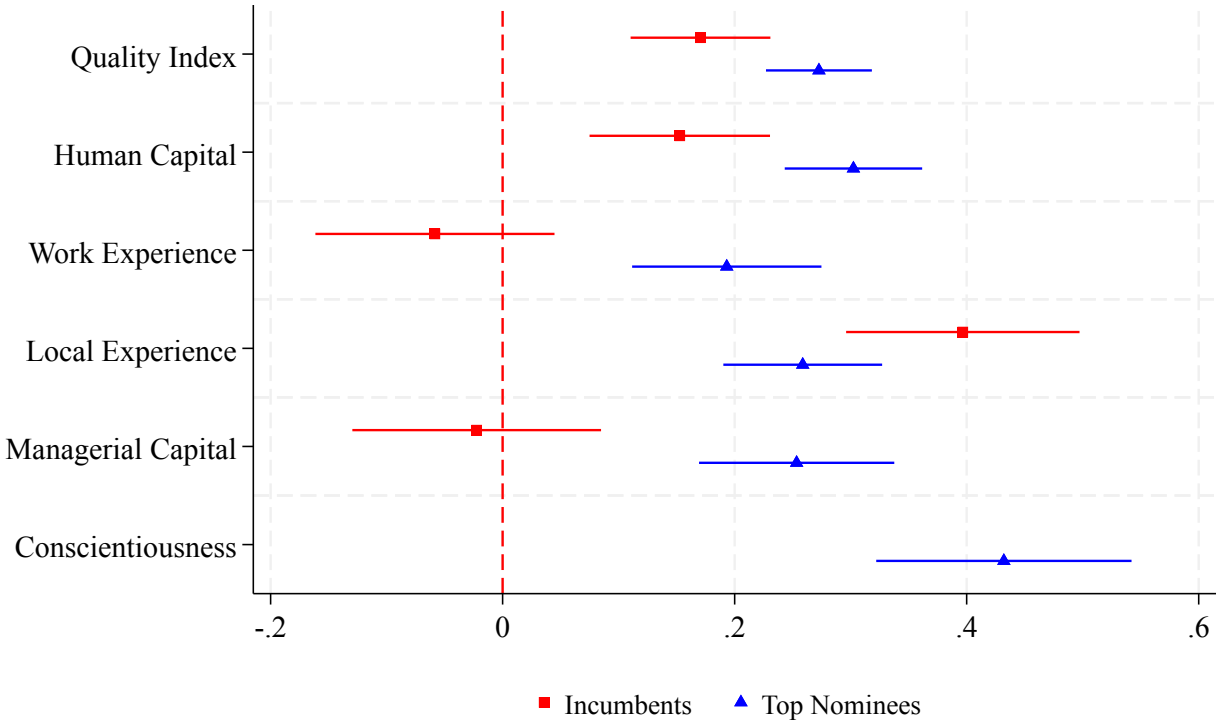
Notes: Maps (left) show the realized ward-level assignment across Sierra Leone for Stage 1 (T1 vs. control) and the combined Stage 1 + Stage 2 outcome (T1 only vs. T1+T2, with host constituencies outlined). Diagram (right) shows how this multi-stage randomization was implemented: a first lottery randomizes the 250 wards (for the stronghold sample) to control (C) or to community nominations & screening (T1); a second lottery then randomizes the 92 host constituencies to one of three saturation arms, which sets the fraction of wards in each constituency to be randomized into profile sharing (T2).

Figure 3: Profile Example (redacted)



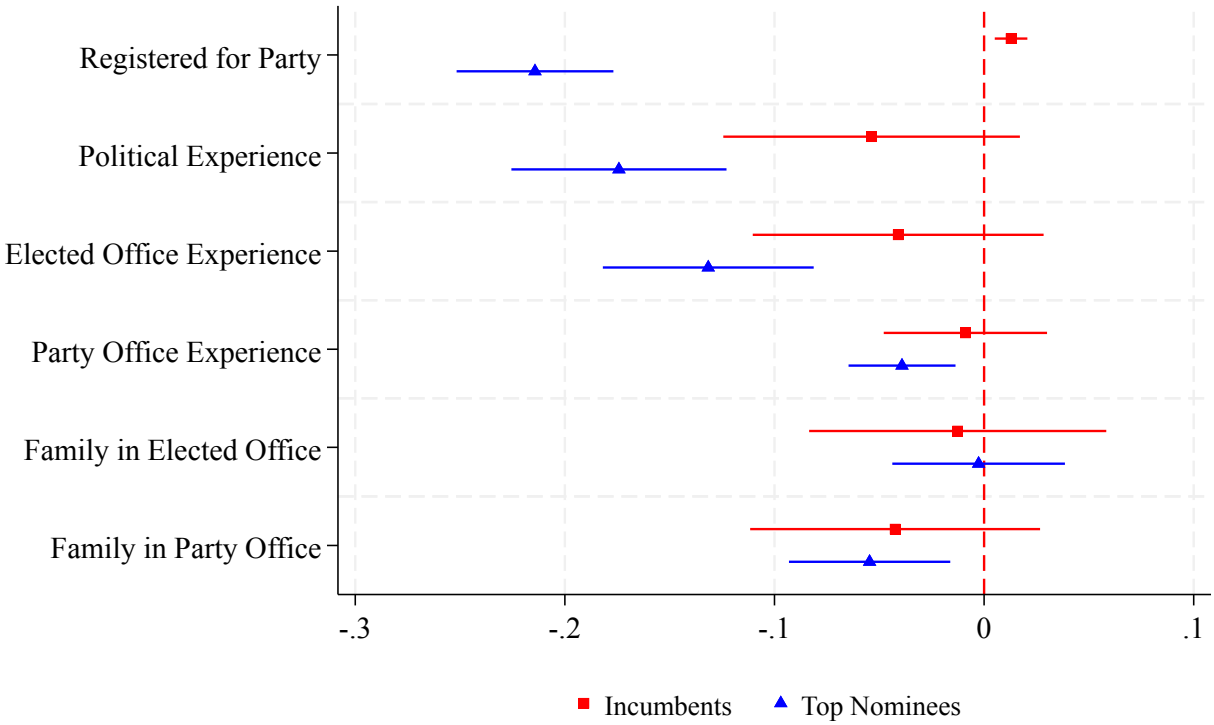
Notes: This figure presents a redacted example of the candidate profiles shared with party leaders in wards assigned to the profile sharing treatment (T2).

Figure 4: Comparisons of Quality Sub-components



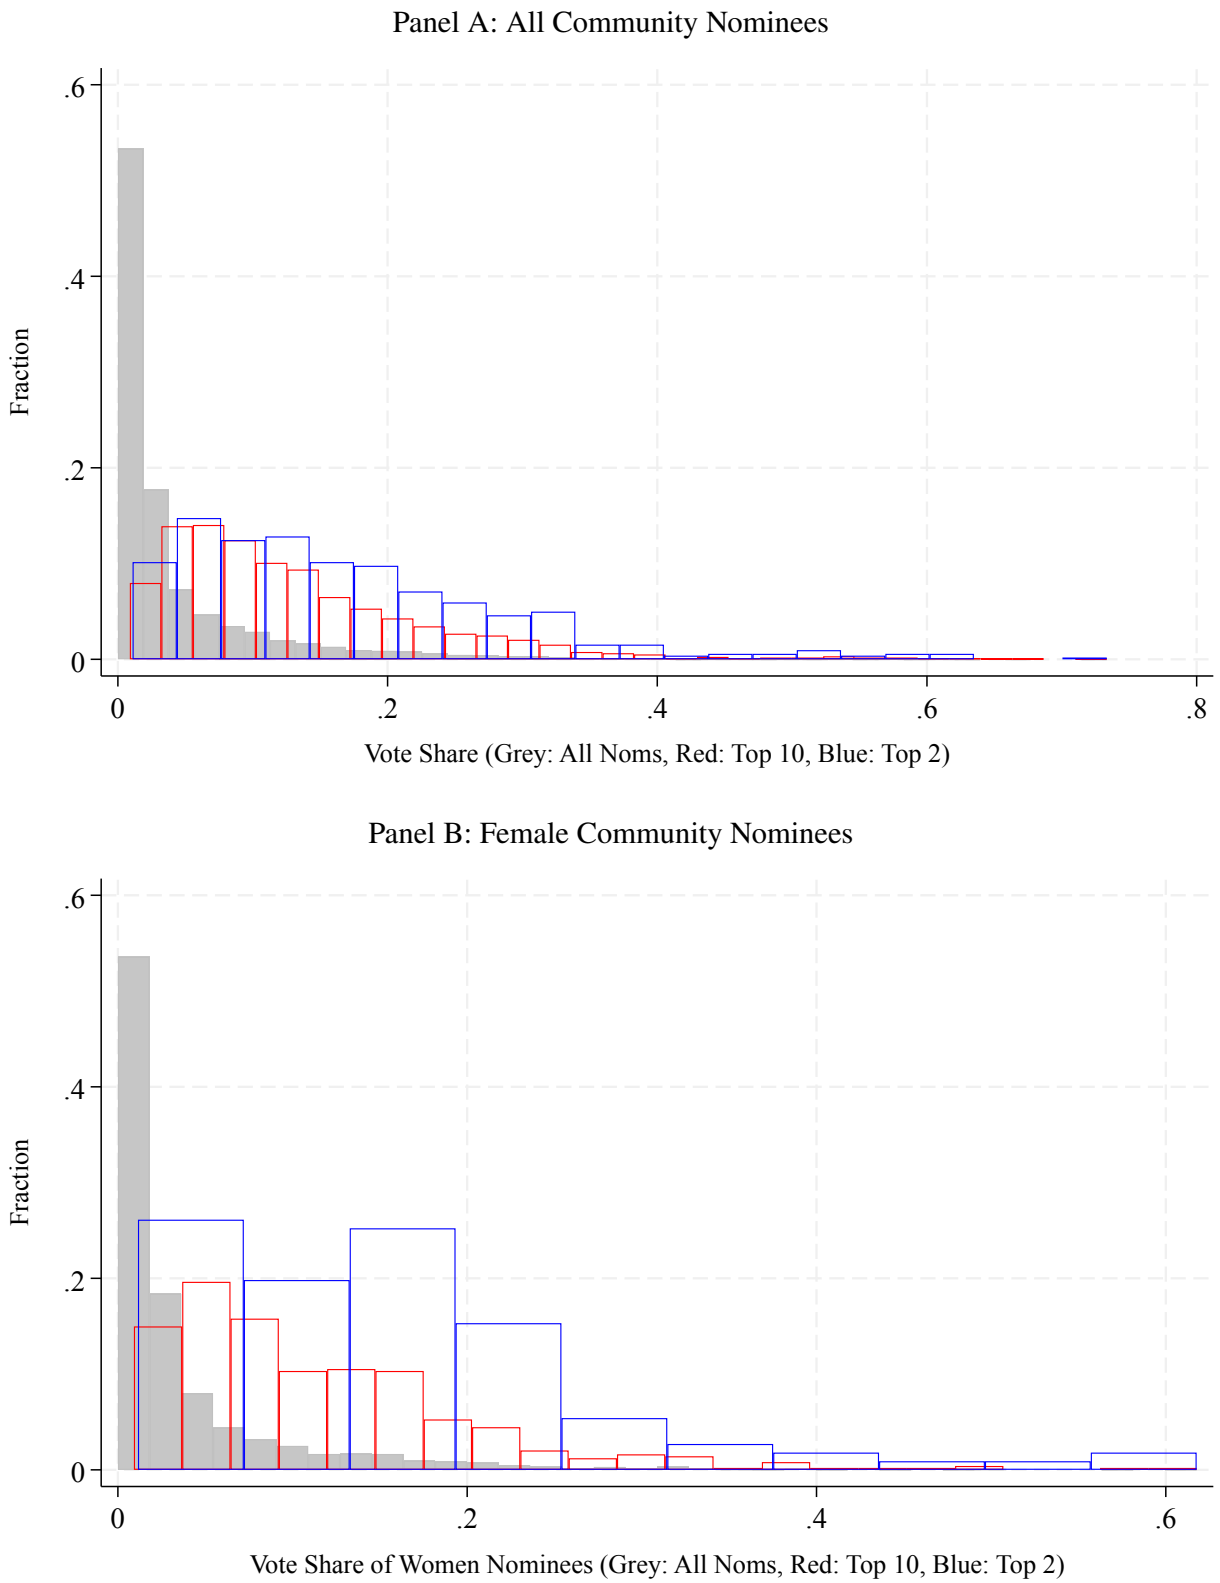
Notes: This figure shows regression coefficients with 95% CIs comparing incumbents and top nominees to status quo applicants in the full sample. The first outcome is the quality index and the other outcomes are indices created with the same methodology grouping variables by category. The mean score of status quo applicants is standardized to zero and used as the control group, indicated by the red dashed vertical line. Regressions control for geographic strata used in the randomization of T1.

Figure 5: Indicators that Nominees are New to Politics



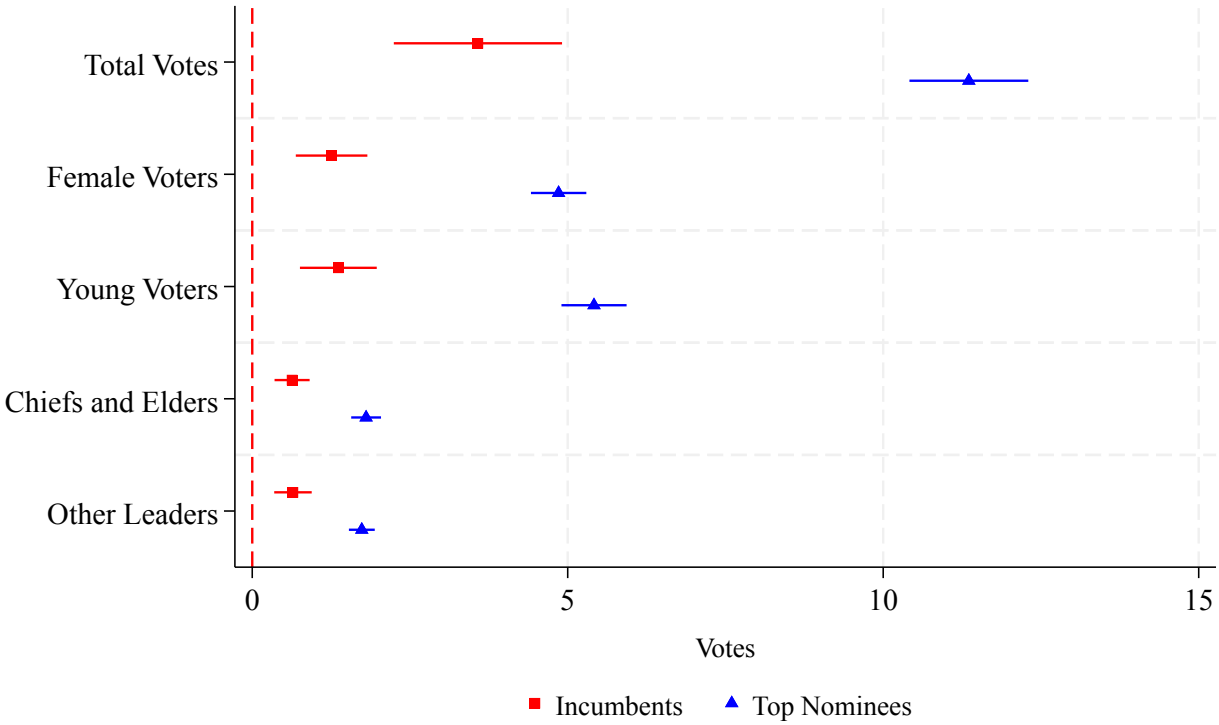
Notes: This figure shows regression coefficients with 95% CIs comparing incumbents and top nominees to status quo applicants in the full sample. All outcomes are binary variables. Regressions control for geographic strata used in the randomization of T1.

Figure 6: Distribution of Community Nomination Shares



Notes: This figure plots the distribution of community nomination shares garnered by each individual who was nominated by at least one respondent (Panel A) and separately for female nominees only (in Panel B).

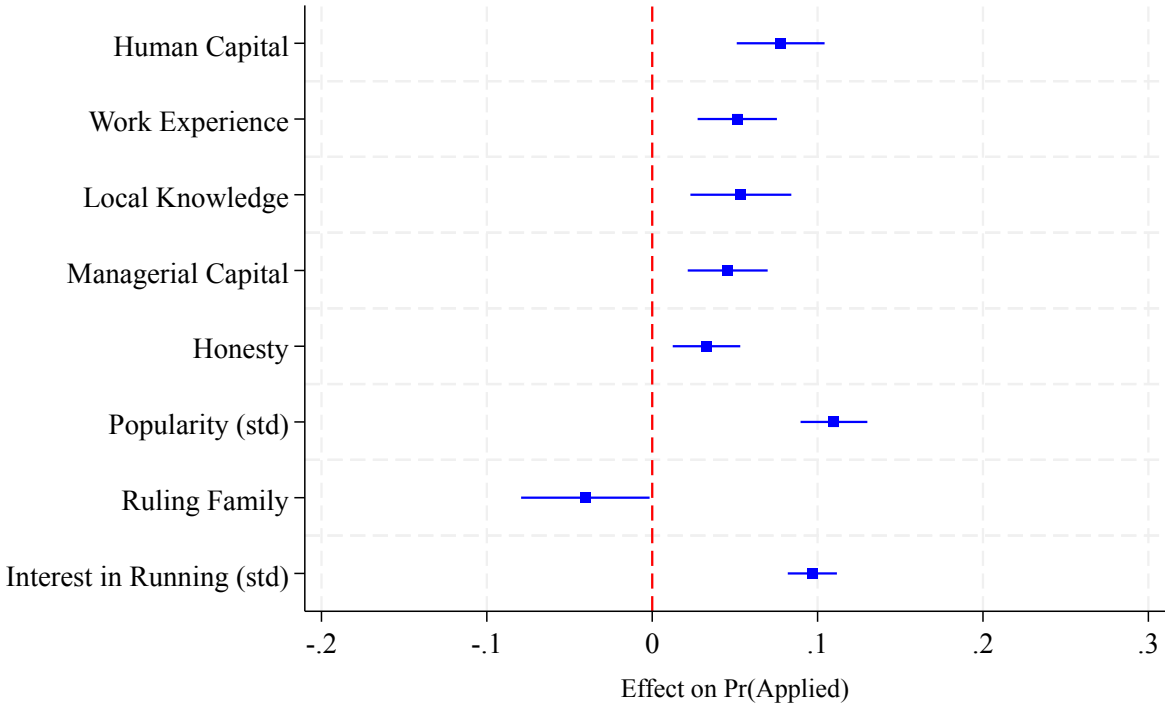
Figure 7: Differences in Local Support by Household Respondent Type



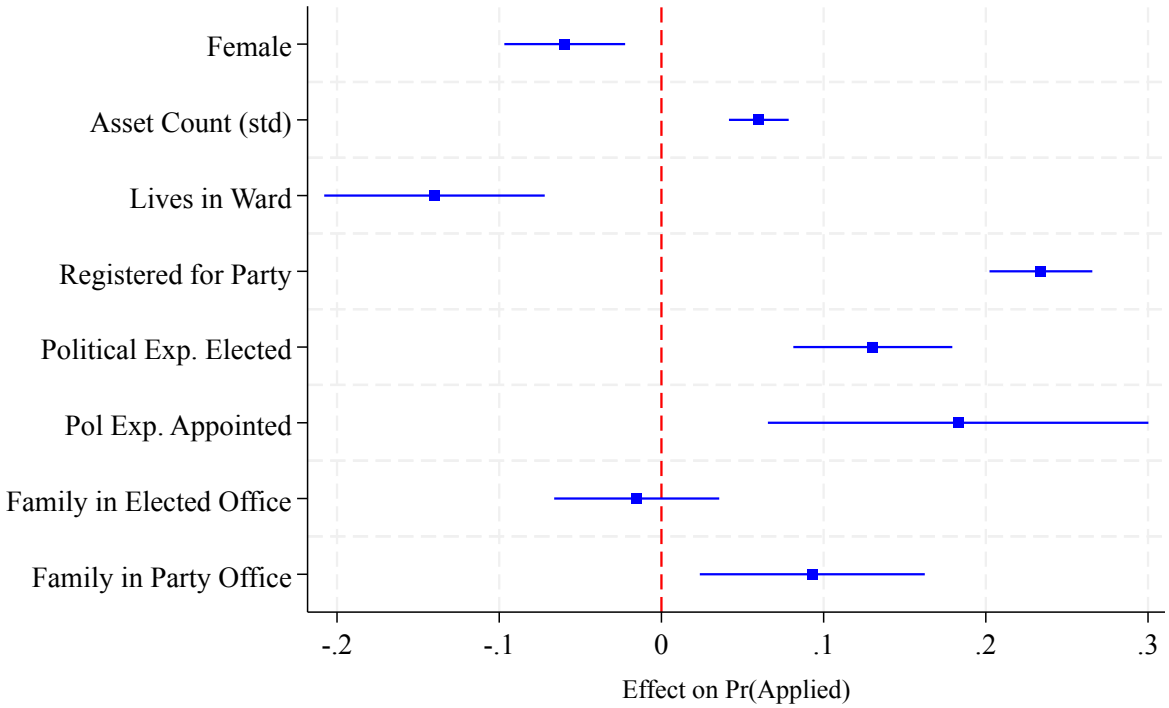
Notes: This figure shows regression coefficients with 95% CIs comparing incumbents and top nominees to status quo applicants in the full sample. All outcomes are the raw number of votes from private nominations. Regressions control for geographic strata used in the randomization of T1.

Figure 8: Drivers of Self-selection into Entry among Nominees

Panel A: Components of quality, popularity, ruling house membership and interest

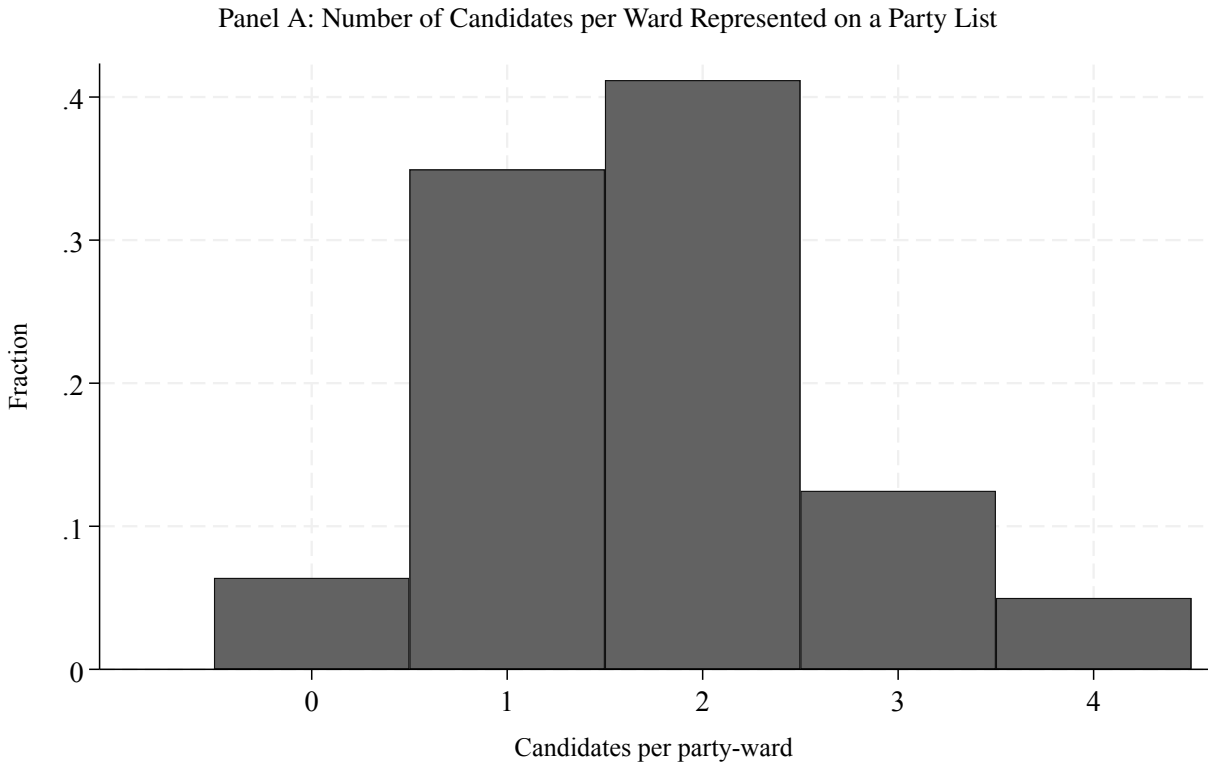


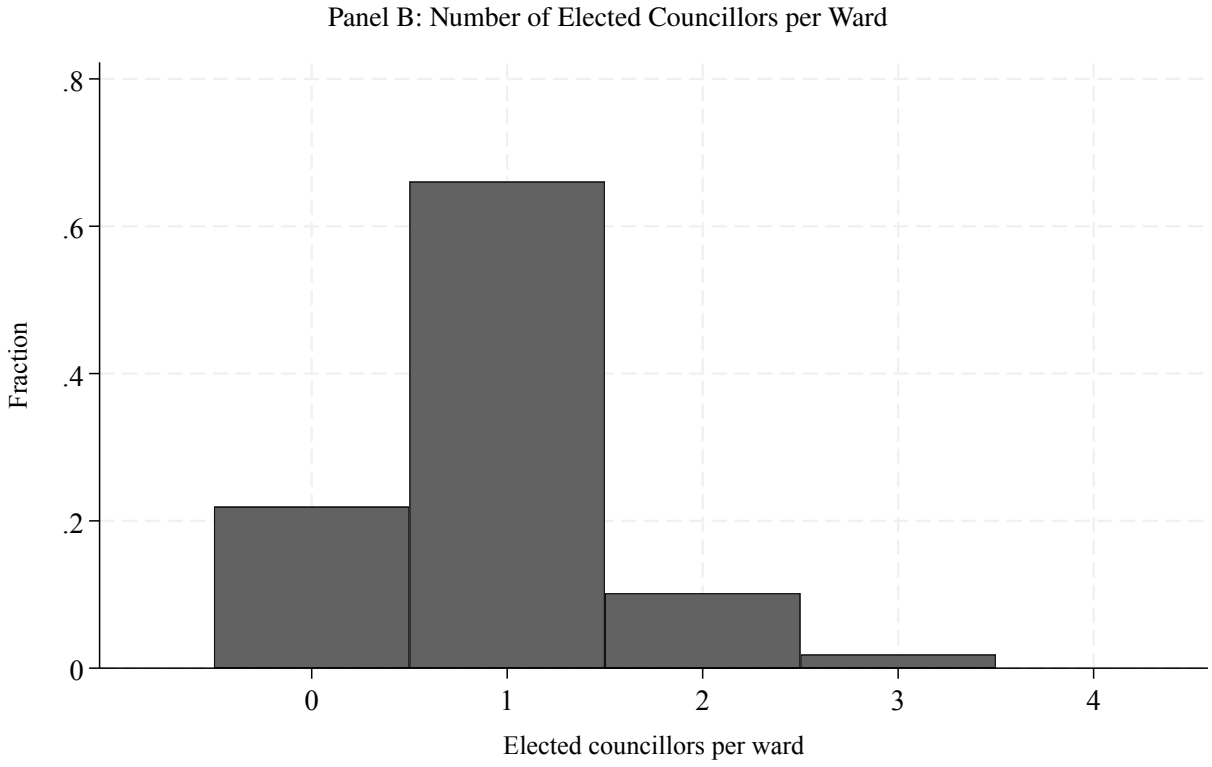
Panel B: Measures associated with costs of entry or political connections



Notes: Figure shows simple regressions of how different variables predict which nominees in the full sample apply to become councillors. Regressions control for geographic strata used in the randomization of T1.

Figure 9: Parties Largely Adhered to the Old Ward-based System

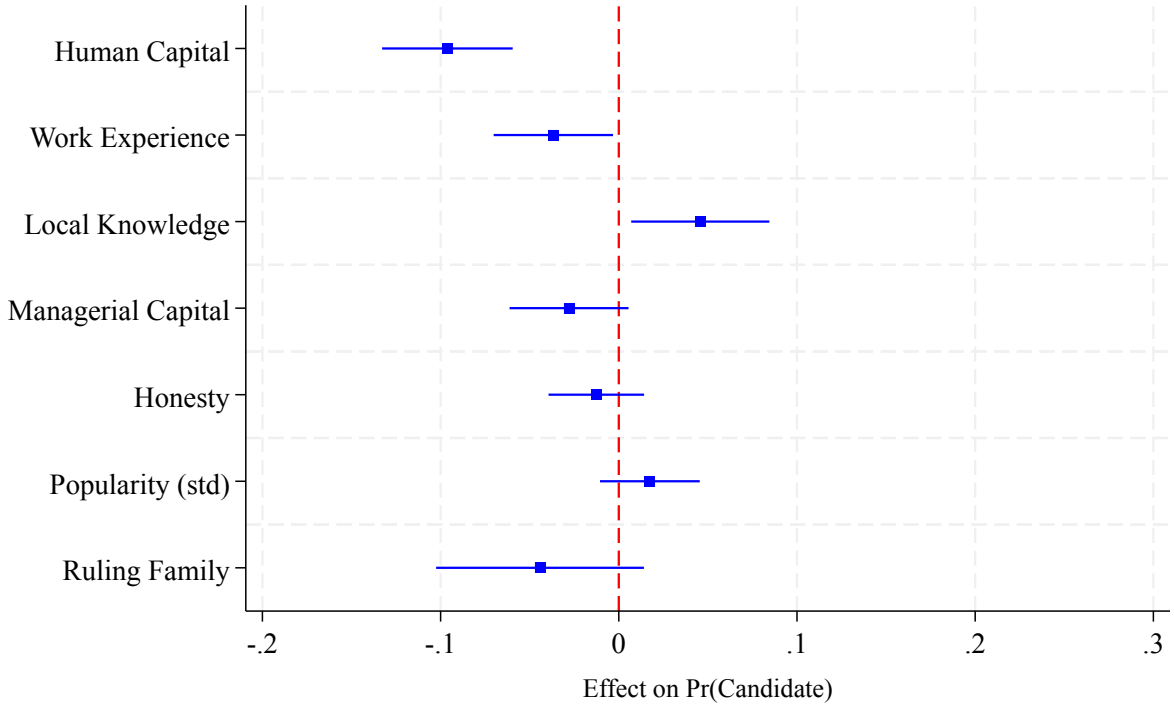




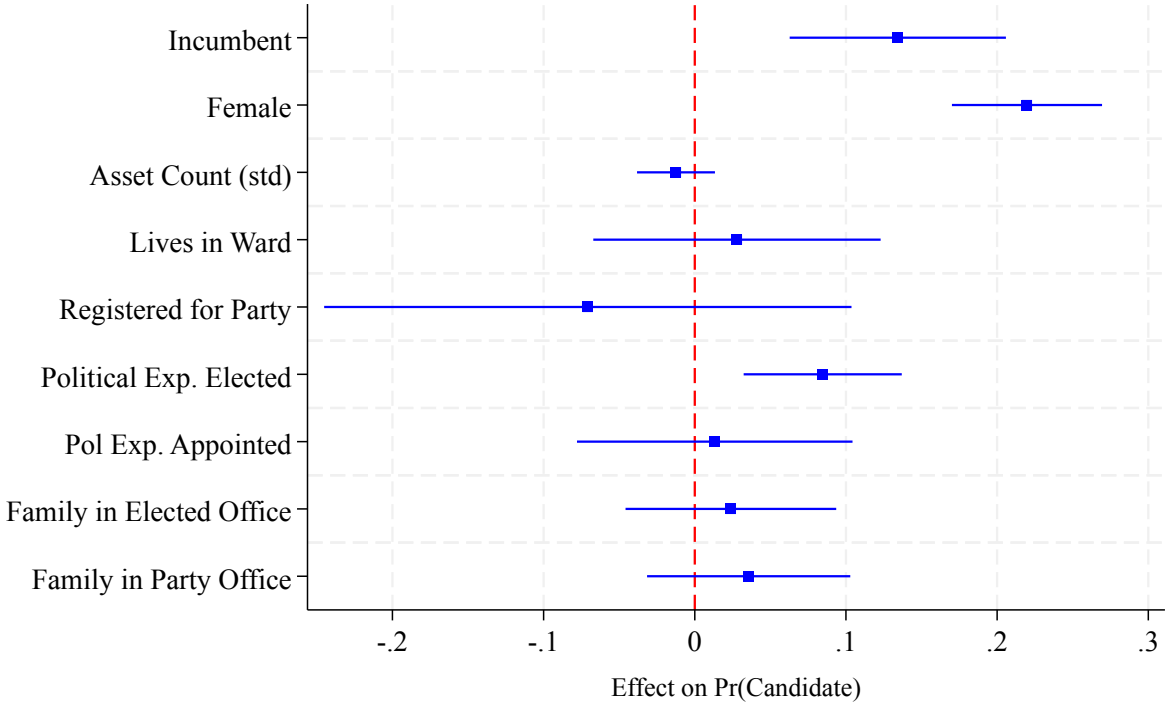
Notes: This figure uses electoral data for both major parties across rural districts that are part of our research sample. Panel A plots the distribution of the number of candidates per ward that appear on a party's district candidate list. Panel B shows the distribution of the number of elected councillors per ward.

Figure 10: Predictors of Party Selection of Candidates

Panel A: Components of quality, popularity and ruling house membership



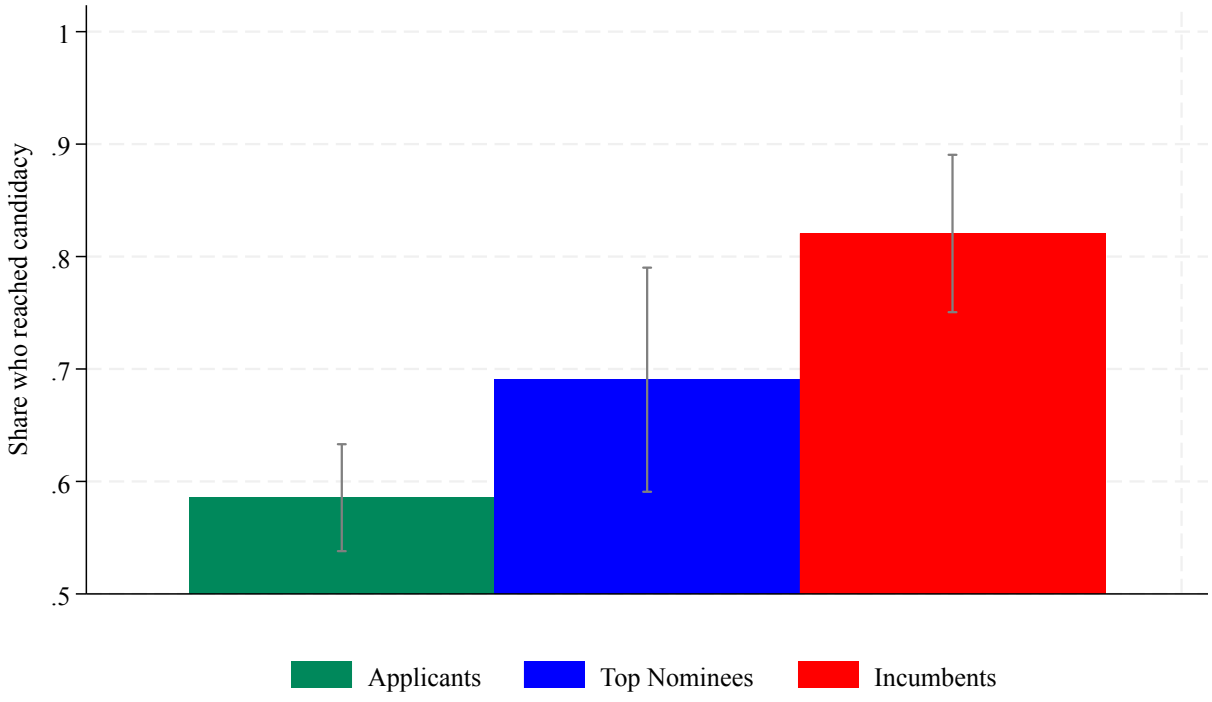
Panel B: Measures associated with costs of entry or political connections



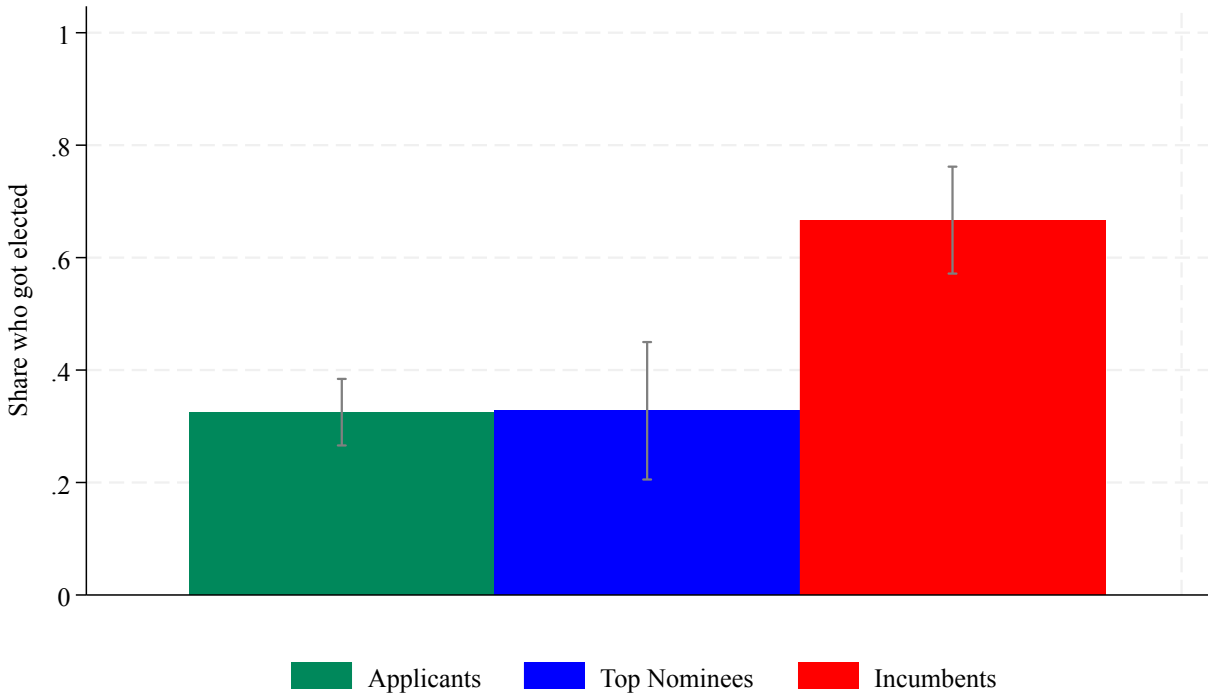
Notes: Figure shows simple regressions of how different variables predict which individuals in the full sample who applied to become councillors go on to be selected as candidates. For incumbents, time-varying characteristics (political experience and family members in office) are measured as of their candidacy in the 2018 election. Regressions control for geographic strata used in the randomization of T1.

Figure 11: Raw Data on Progression from Entry to Election by Type

Panel A: Conditional Likelihood of Candidacy Selection by Applicant Type

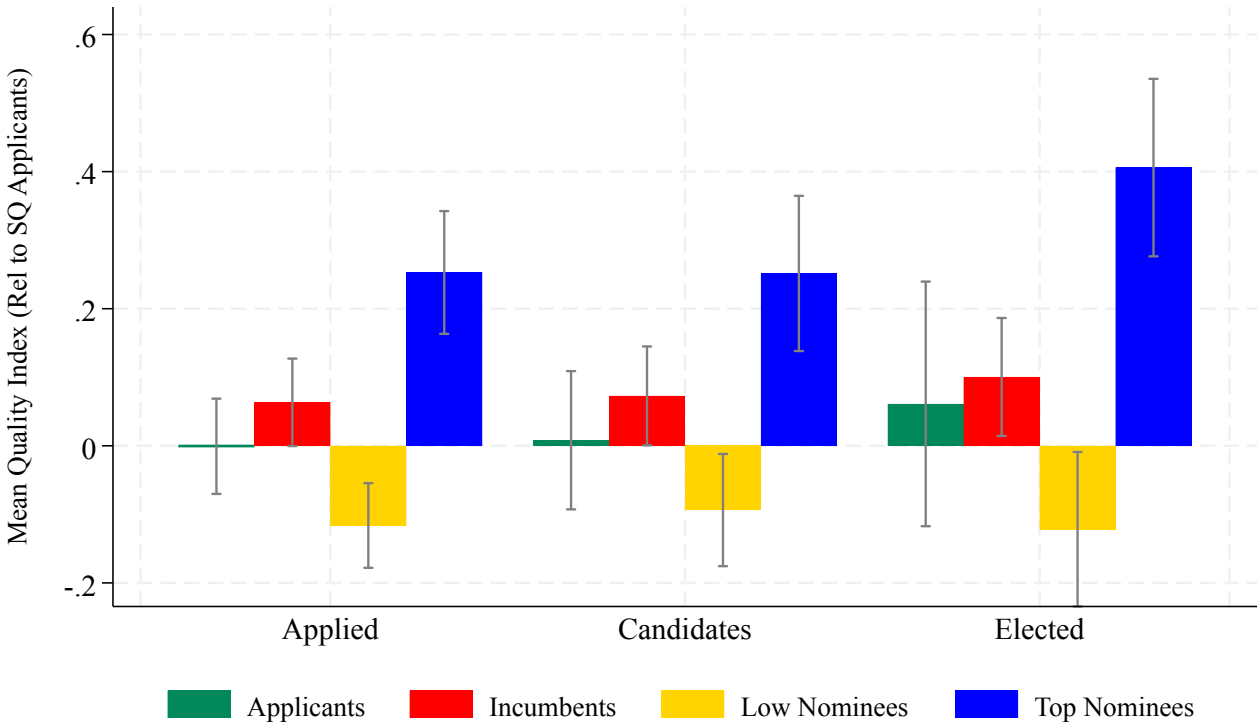


Panel B: Conditional Likelihood of Election by Candidate Type



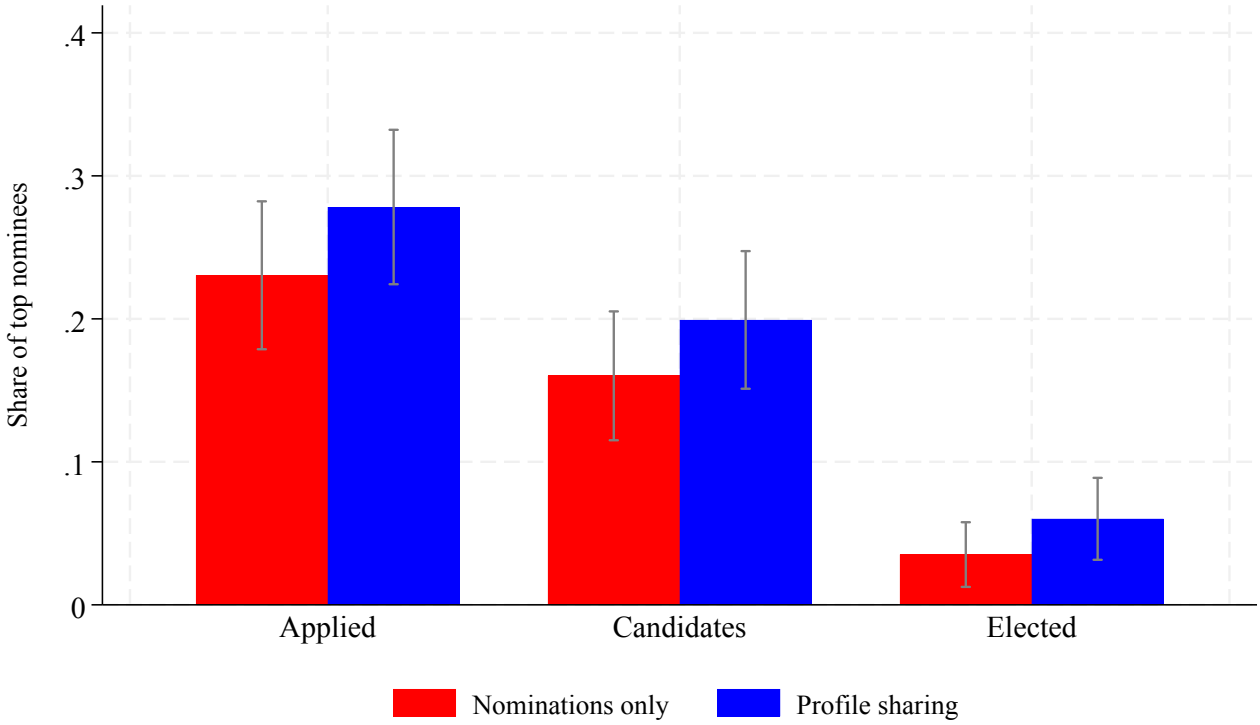
Notes: This figure plots the conditional likelihood of progressing to candidacy and being elected for the potential candidates in the stronghold sample of our study: incumbents, top nominees, and status quo applicants. Panel A shows the share of each type that, conditional on filing an application, was selected onto a party's candidate list. Panel B shows the share of each type that, conditional on being placed on a candidate list, won a council seat. Panel A includes 117 incumbents, 84 top nominees and 415 status quo applicants who filed applications; panel B includes the 96, 58 and 243 of them, respectively, who were selected as candidates.

Figure 12: Party Selection on Quality



Notes: This figure illustrates the degree of selection on the quality index at each stage of the political entry process, for each type of potential applicant in the stronghold sample. Each column is the average quality index within each group relative to status quo applicants (first applicant bar in green is 0 by construction). Figure uses 95% confidence intervals.

Figure 13: Raw Data on Profile Sharing Impacts



Notes: This figure presents evidence on the impact of the profile sharing treatment (T2) using raw electoral data for the stronghold sample. It shows the share of top nominees who reached each stage of the process: applied to a party, were selected as candidates, and won a seat. It compares wards assigned to nominations only (T1 only, in red) and wards assigned to both nominations and profile sharing (T1+T2, in blue).

Table 1: Characteristics of Potential Candidates

	(1) Voters	(2) Applicants (N=1,218)	(3) Incumbents (N=293)	(4) Top Nominees (N=522)
Male %	50.1	72.8	86.4	78.7
Age	35.4	38.9	45.8	39.7
Asset Count	3.8	8.9	10.2	9.3
Monthly Wage (USD)	17.5	51.2	47.1	46.1
Household Size	7.7	8.8	10.7	9.3
Works in Education %	6.0	21.1	29.8	35.2
Works in Business %	8.5	23.7	15.8	19.7
Works in Farming %	62.6	26.2	31.6	24.7
Literacy %	49.0	82.0	95.6	100
Primary Completed %	36.9	91.8	98.6	99.0
Any Post-Secondary %	6.2	53.3	71.9	64.3
Registered to a Party%	17.8	98.8	100	76.8
Ruling Family %	15.7	20.5	35.5	62.6

Notes: All data in columns 2, 3, and 4 was collected via surveys for this experiment. Column 3 includes all sitting Councillors surveyed in rural districts. Column 4 includes top nominees in the full sample. Data for voters in column 1 is obtained from multiple sources. Sex, age, and party registration data come from the households surveyed during the private nominations (N=10,365). Assets, wages, household size, and ruling family data come from a separate study conducted by one of the coauthors covering 6 rural districts during the same period of time as this project (N=1,974). The data on sector of employment and education completion come from Round 9 of the Afrobarometer Survey done in Sierra Leone, restricting to respondents in rural areas and approximating occupations based on their classification (N=1,200). Literacy rate among adults is from 2022 and was obtained from the World Bank Data.

Table 2: Quality Index by Type of Potential Candidates

	(1)	(2)
	Outcome: Quality Index	
	Full Sample	Stronghold
Top Nominees	0.275***	0.200***
	(0.023)	(0.027)
Incumbents	0.172***	0.016
	(0.030)	(0.033)
Observations	1,670	979
P-value β_1=β_2	0.001	<0.001

Notes: Column 1 reports estimates for the full sample and column 2 for the stronghold sample. Top nominees are drawn from the 186 wards assigned to community nominations and screening (T1); incumbents and status quo applicants are drawn from the 250 wards in the research sample. Regression outcome is standardized relative to the omitted group in each regression, which is status quo applicants. Regressions include fixed effects for geographic strata used in the randomization of T1.

Table 3: Top Nominees Differ from the Best Applicants in Control Wards

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Female	Asset Count	Ruling Family	Political Experience	Family in Politics	Family in Party	Quality Index
Panel A: Top nominees vs best two status quo applicants in Control wards							
Top Nominees	0.096**	-0.261**	0.461***	-0.265**	0.008	-0.013	0.102**
	(0.039)	(0.113)	(0.046)	(0.110)	(0.041)	(0.039)	(0.042)
Observations	448	448	448	448	448	448	448
Panel B: Top nominees that applied vs best two status quo applicants in Control wards							
Top Nominees	0.121**	0.053	0.367***	0.056	-0.024	0.038	0.151**
	(0.059)	(0.140)	(0.067)	(0.148)	(0.054)	(0.050)	(0.060)
Observations	185	185	185	185	185	185	185

Notes: Regressions include all observations in the stronghold sample and include fixed effects for geographic strata used in the randomization of T1.

Table 4: Heterogeneity in Identifying High Quality Nominees

	(1)	(2)	(3)	(4)	(5)
	Outcome: Quality Index				
Top Nominee	0.474***	0.405***	0.409***	0.496***	0.354***
	(0.022)	(0.024)	(0.032)	(0.025)	(0.044)
Female		-0.259***			
		(0.026)			
Female × Top Nominee		0.129**			
		(0.052)			
Ruling Party			-0.026		
			(0.033)		
Ruling Party × Top Nominee			0.124***		
			(0.044)		
Competitive Ward				-0.019	
				(0.043)	
Competitive Ward × Top Nominee				-0.097*	
				(0.051)	
Stronghold					0.057
					(0.037)
Stronghold × Top Nominee					0.180***
					(0.051)
Observations	1,597	1,597	1,597	1,597	1,597
Share in Het. Category		32%	52%	20%	78%

Notes: Regressions compare all top nominees to lower ranked nominees in the full sample. Share in Het. Category reports the share of nominees in the category indicated by each column. Regressions include fixed effects for geographic strata used in the randomization of T1.

Table 5: Measures of Local Support by Type of Potential Candidate

	(1)	(2)	(3)	(4)	(5)	(6)
	Full Sample			*Stronghold*		
	Vote Share	Tot. Votes	Any Votes	Vote Share	Tot. Votes	Any Votes
Top Nominees	0.134***	11.320***	0.737***	0.130***	11.013***	0.677***
	(0.006)	(0.486)	(0.018)	(0.008)	(0.670)	(0.029)
Incumbents	0.042***	3.715***	0.301***	0.030***	2.710***	0.225***
	(0.008)	(0.704)	(0.042)	(0.010)	(0.841)	(0.049)
Observations	1,297	1,297	1,297	769	769	769
P-value β_1=β_2	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001
Control Mean	0.024	2.036	0.259	0.031	2.602	0.322

Notes: Regressions cover all T1 wards where we conducted the nominations exercise. We used name matching algorithms to search *ex post* for the names of status quo applicants and incumbents in the nominations data. The omitted category is status quo applicants. Regressions include fixed effects for geographic strata used in the randomization of T1.

Table 6: Formal Entry Rates of Community Nominees

	(1)	(2)	(3)	(4)	(5)
	Outcome: Applied to Party				
Top Nominee	0.096***	0.111***	0.110*	0.204***	0.100***
	(0.030)	(0.025)	(0.059)	(0.043)	(0.037)
Profile Sharing	-0.014		-0.019		-0.004
	(0.021)		(0.031)		(0.024)
Top Nominee × Profile Sharing	0.060		0.093		0.002
	(0.043)		(0.086)		(0.050)
Female		-0.050**			
		(0.020)			
Top Nominee × Female		0.038			
		(0.050)			
Stronghold				0.044	
				(0.027)	
Top Nominee × Stronghold				-0.103**	
				(0.049)	
Observations	1,597	1,597	510	1,597	1,251
Sample	All Nominees	All Nominees	Female Nominees	All Nominees	Stronghold
Control Mean	0.129	0.143	0.091	0.104	0.127

Notes: Regressions use nominees in the full sample (female nominees only in column 3; stronghold sample in column 5) and include fixed effects for geographic strata used in the randomization of T1. Control Mean is the entry rate of the omitted group: lower-ranked nominees (outside the top two) in wards without profile sharing in columns 1, 3 and 5; lower-ranked male nominees in column 2; and lower-ranked nominees outside party strongholds in column 4.

Table 7: Incumbency Advantage After PR

	(1)	(2)	(3)
	Re-run	Re-elected given re-run	Re-elected
After PR	0.341***	0.216**	0.256***
	(0.035)	(0.087)	(0.029)
Observations	586	182	586
Control Mean	.140	.366	.051

Notes: Regression includes incumbent councillors elected both in 2012 and 2018, meaning they were up for re-election in 2018 and 2023 respectively. The main outcomes are whether each incumbent re-ran in the subsequent election and whether they were re-elected. Only councillors from the same 293 wards are included for both election cycles. This covers wards in rural districts where the experiment happened. Regression includes district fixed effects.

Table 8: Estimated Impacts on the Size of the Ward-level Applicant Pool

	(1)	(2)	(3)	(4)
	Any Applicant	Num. of Applicants	Any Women Applicant	Num. of Women Applicants
T1 Nominations	0.014	0.222	0.125	0.222*
	(0.013)	(0.252)	(0.083)	(0.133)
T2 Profiles	0.004	-0.221	-0.121	-0.087
	(0.004)	(0.228)	(0.078)	(0.137)
Observations	250	250	250	250
Control Mean	0.984	2.825	0.397	0.508

Notes: Regressions estimated with observations collapsed at the ward level using the stronghold sample and include fixed effects for geographic strata used in the randomization of T1.

Table 9: Estimated and Simulated Impacts on the Quality of the Applicant Pool

	(1)	(2)	(3)	(4)	(5)	(6)
	Applicant Pool		+25% App. Rate		+50% App. Rate	
	Mean	Max	Mean	Max	Mean	Max
T1 Nominations	0.052	0.092*	0.063	0.121**	0.074*	0.149***
	(0.044)	(0.051)	(0.044)	(0.053)	(0.044)	(0.053)
T2 Profiles	0.002	-0.040	0.008	-0.034	0.013	-0.027
	(0.038)	(0.048)	(0.038)	(0.051)	(0.038)	(0.051)
Observations	249	249	249	249	249	249
Control Mean	0.018	0.267	0.018	0.267	0.018	0.267
P-value $\beta_1 = 0$	0.234	0.075	0.148	0.021	0.090	0.005
Individuals in Sample	728	728	753	753	778	778
Top Noms in Sample %	11.5	11.5	14.5	14.5	17.2	17.2
Ward Compliance %	40.9	40.9	54.3	54.3	67.7	67.7

Notes: Regressions estimated with observations collapsed at the ward level using the stronghold sample and include fixed effects for geographic strata used in the randomization of T1. Main outcome is quality index aggregated in different ways for each ward observation: Average across all applicants and maximum quality of applicant pool. Columns 3 to 6 report results from a Monte Carlo simulation with 500 draws. In each draw, for wards where we did not observe a top nominee entering, we select the highest quality top nominee who expressed interest in running but did not apply and we boost their entry, first adding them to the pool with a 25% probability (columns 3 and 4) and then with a 50% probability (columns 5 and 6). Coefficients are pooled means across simulations. Standard errors and p-values reported combine within-simulation sampling uncertainty and across-simulation variance. Individuals in sample, top nominees in sample, and ward compliance are averages across simulations.

Table 10: Estimated Impacts on the Selected and Elected Candidates

	(1)	(2)	(3)	(4)
	Any Candidate	Any Women Candidate	Inv. Candidate Ranking	Any Elected
T1 Nominations	0.009	0.138*	-2.751	-0.035
	(0.037)	(0.078)	(1.694)	(0.068)
T2 Profiles	0.026	-0.122*	4.406***	0.044
	(0.027)	(0.073)	(1.421)	(0.058)
Observations	250	250	235	250
Control Mean	0.937	0.349	-21.525	0.714
P-value $\beta_1 + \beta_2 = 0$	0.286	0.828	0.319	0.884

Notes: Regressions estimated with observations collapsed at the ward level using the stronghold sample and include fixed effects for geographic strata used in the randomization of T1. Column 3 is estimated on the 235 wards with at least one candidate whose list position we observe in the party ranking data (positions are missing for 10 candidates from 4 wards). Because treatment may affect whether a ward has a candidate, this estimate conditions on a post-treatment outcome and should be interpreted with caution.

Table 11: Estimated Impacts on the Quality of the Candidate and Elected Pool

	(1)	(2)	(3)	(4)
	Candidate Lists		Elected Officials	
	Mean Q	Max Q	Mean Q	Max Q
T1 Nominations	0.076	0.117*	0.105	0.108
	(0.053)	(0.061)	(0.077)	(0.078)
T2 Profiles	-0.001	-0.077	-0.030	-0.043
	(0.042)	(0.048)	(0.057)	(0.057)
Observations	239	239	175	175
Control Mean	-0.004	0.158	-0.001	0.020
P-value $\beta_1 + \beta_2 = 0$	0.146	0.504	0.296	0.373

Notes: Regressions estimated with observations collapsed at the ward level using the stronghold sample and include fixed effects for geographic strata used in the randomization of T1.

Online Appendix

Contents

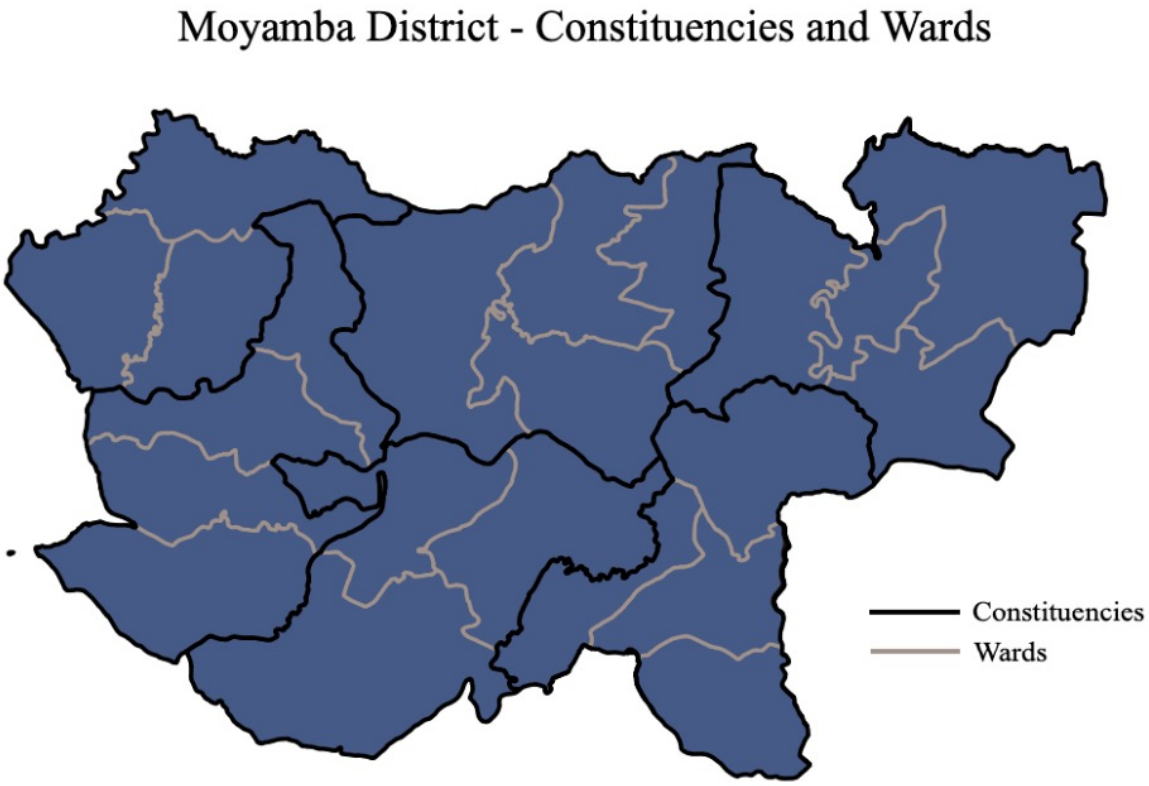
1 Supplemental Figures **A1**

2 Supplemental Tables **A10**

3 Supplemental Exhibits **A21**

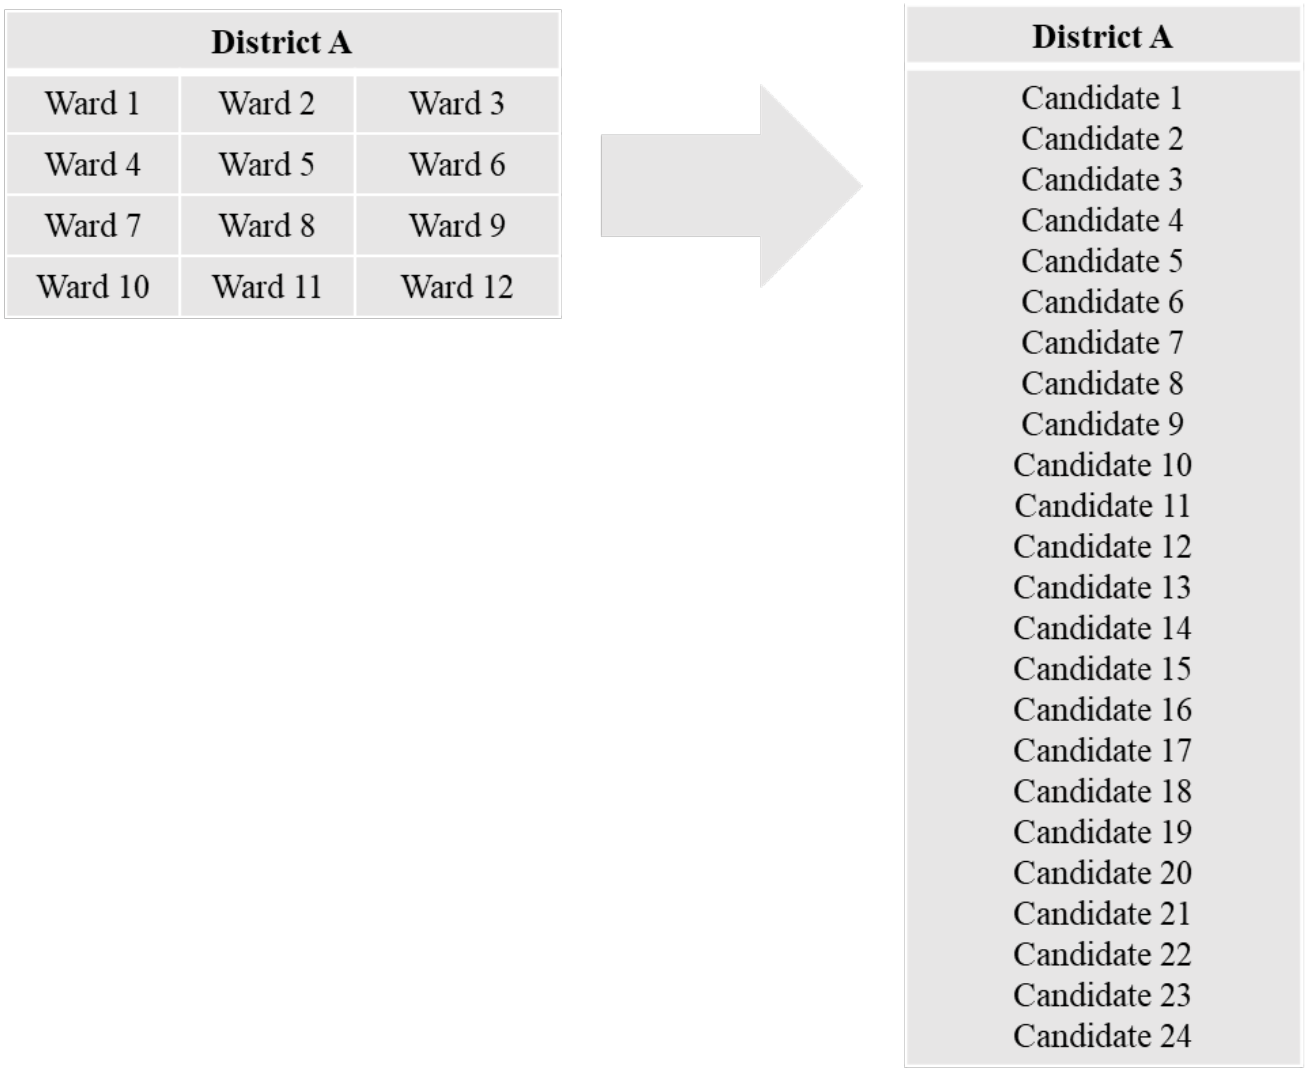
1 Supplemental Figures

Figure A1: Map of Nested Wards and Constituencies for Moyamba District



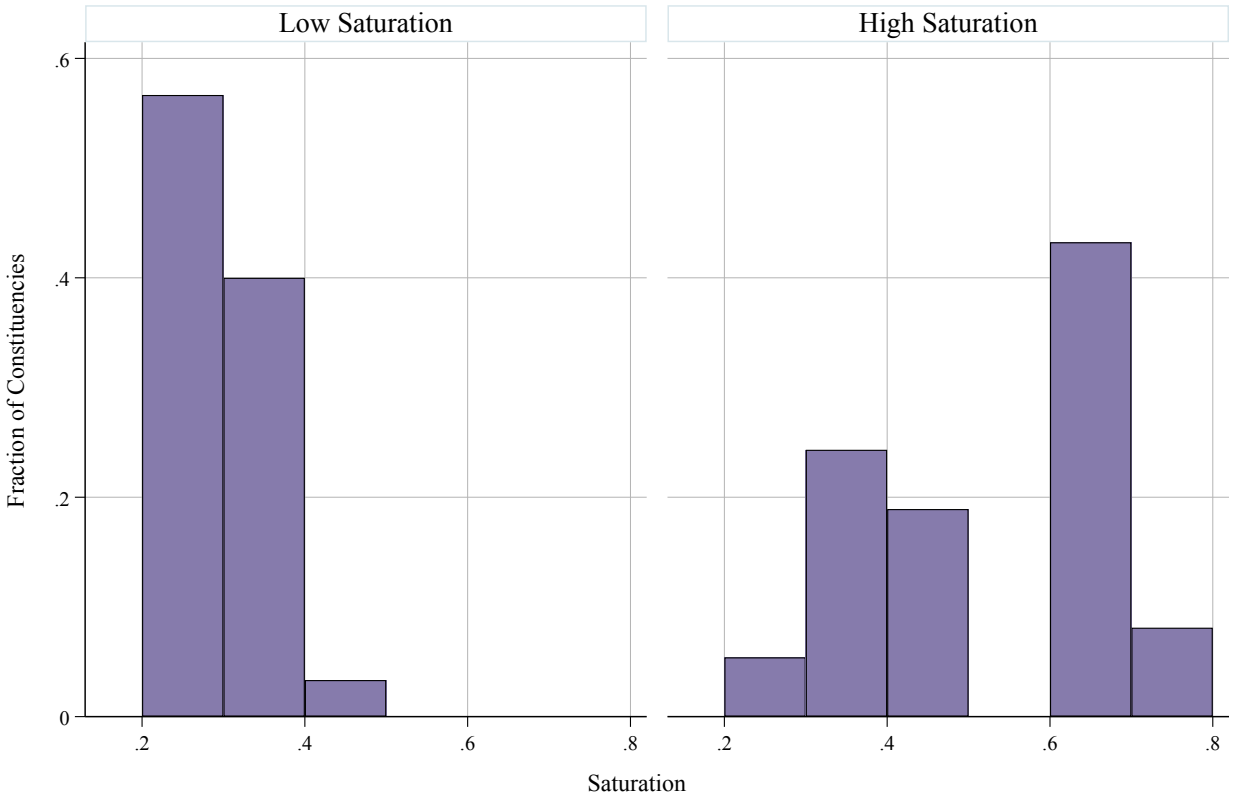
Notes: This figure shows how Local Council wards nest neatly inside Parliamentary constituencies, which in turn map neatly into district boundaries. This neat mapping, with roughly 3 to 4 wards per constituency, motivates our choice of constituency as the level to assign the profile sharing treatment (T2) as a reasonable level to capture potential interference between ward units in party selections of candidates to put on their respective district-level candidate lists. Moyamba District Council is one of the Local Councils in the research sample.

Figure A2: Illustration of Electoral Reform



Notes: This figure illustrates the electoral change from first-past-the-post majoritarian single-member jurisdictions (image on left) to district-block proportional representation (image on the right) that was enacted shortly before the 2023 elections studied. In the hypothetical district presented, there were originally twelve wards, each of which elected a single Local Councillor to represent them in the district council. After the electoral change, political parties prepared candidate lists that were twice as long as the original number of wards, so a list of 24 in this example. These are closed lists, implying that on election day voters simply choose between the party lists, without seeing or ranking individual candidates.

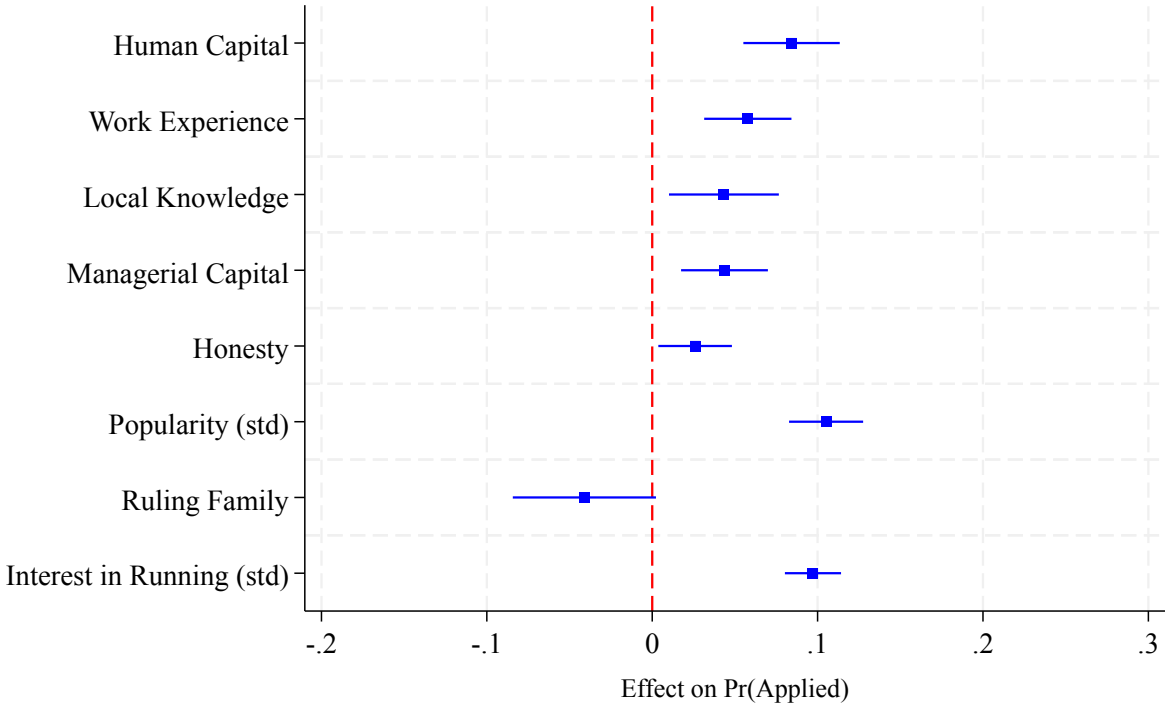
Figure A3: Effective Saturation Levels Across Constituencies



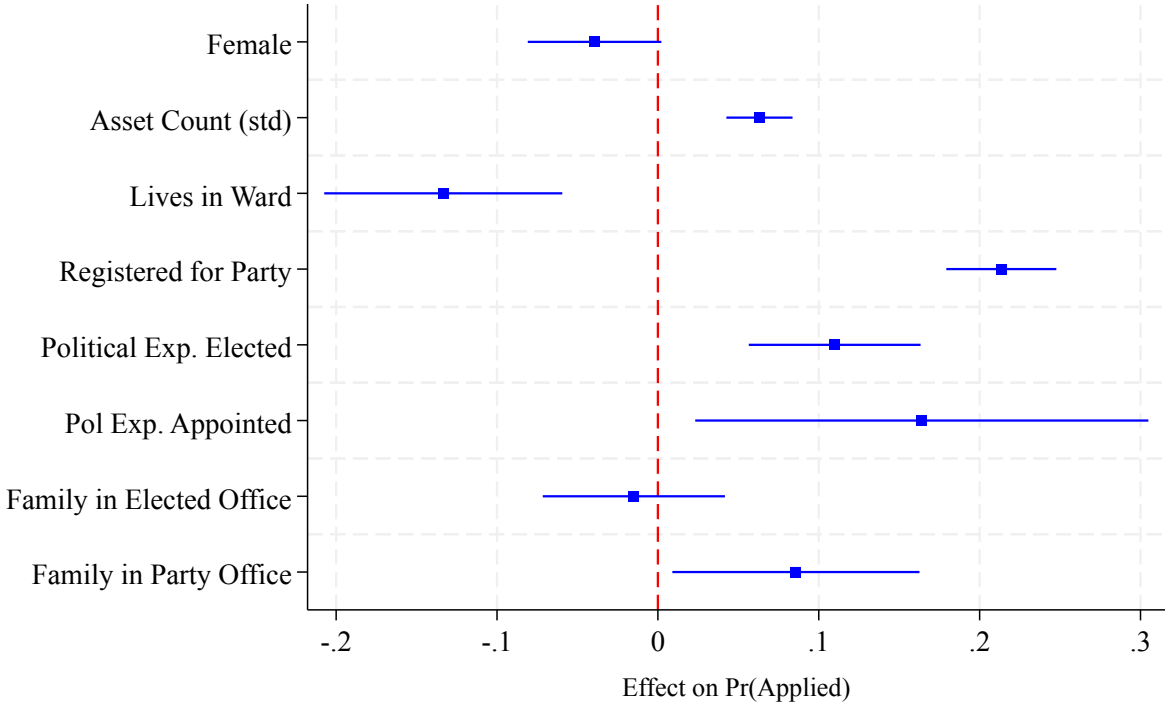
Notes: This figure plots the distribution of effective saturation levels of T2 treatment wards across constituencies. The X-axis shows the share of wards located within a given constituency that were assigned both the nominations and screening process in the first round of ward-level randomization and then assigned to the profile sharing treatment in the second round of ward-level randomization. The group of zero saturation constituencies is omitted from the figure as they contain no T2 treatment wards (saturation share = 0.0).

Figure A4: Drivers of Self-selection into Entry among Nominees

Panel A: Components of quality, popularity, ruling house membership and interest



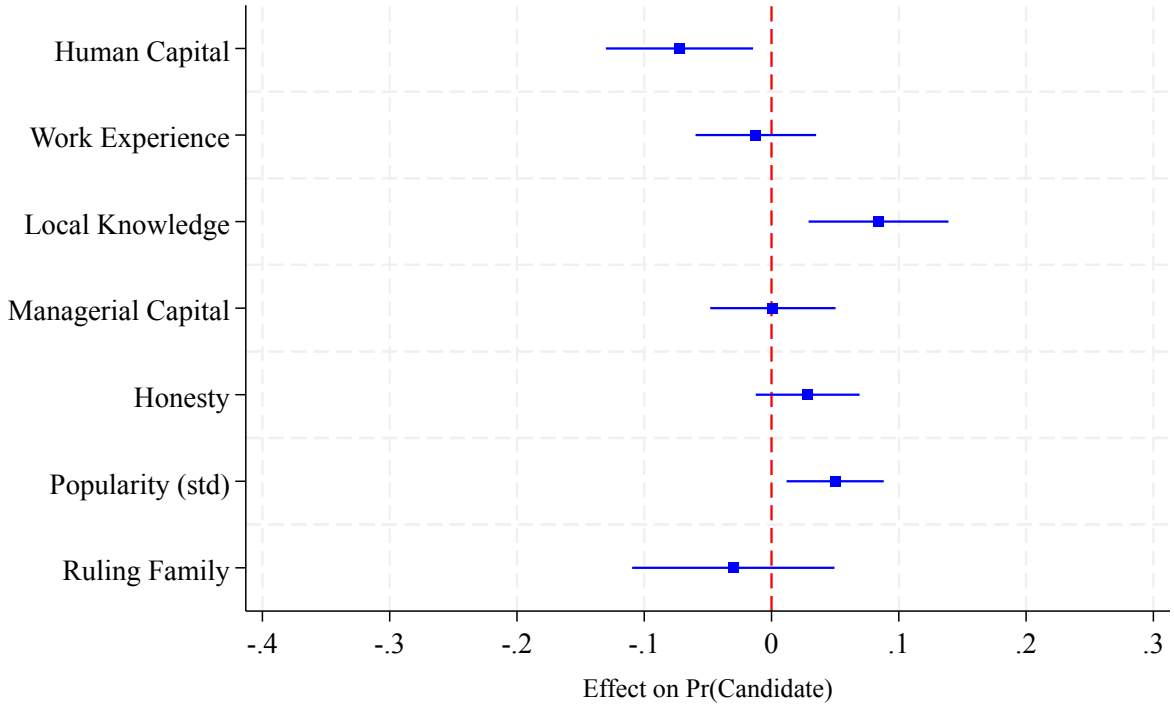
Panel B: Measures associated with costs of entry or political connections



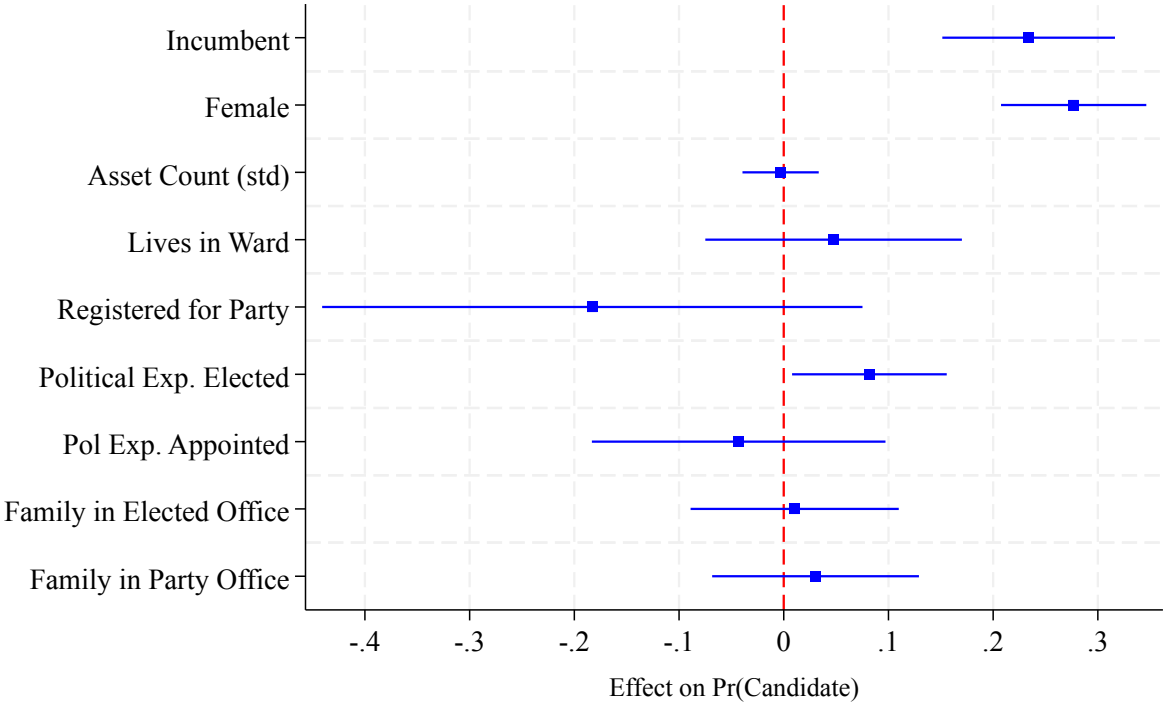
Notes: Figure shows simple regressions of how different variables predict which nominees in the stronghold sample apply to become councillors. Regressions control for geographic strata used in the randomization of T1.

Figure A5: Predictors of Party Selection of Candidates

Panel A: Components of quality, popularity and ruling house membership



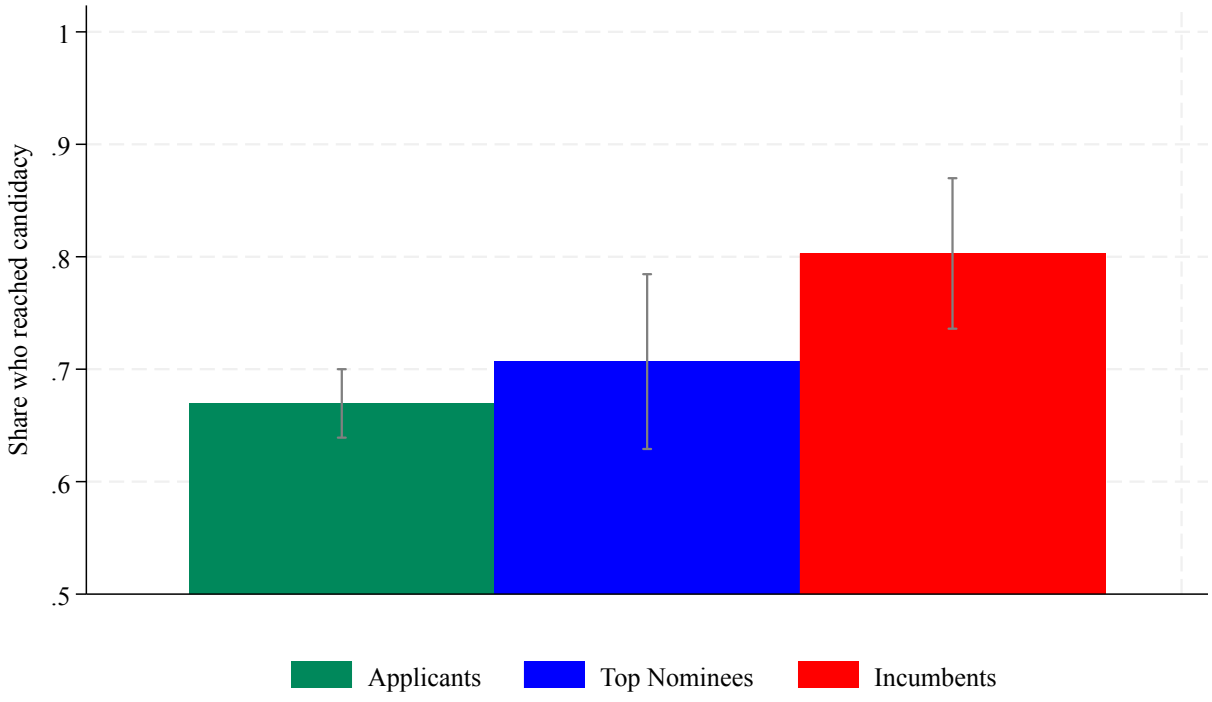
Panel B: Measures associated with costs of entry or political connections



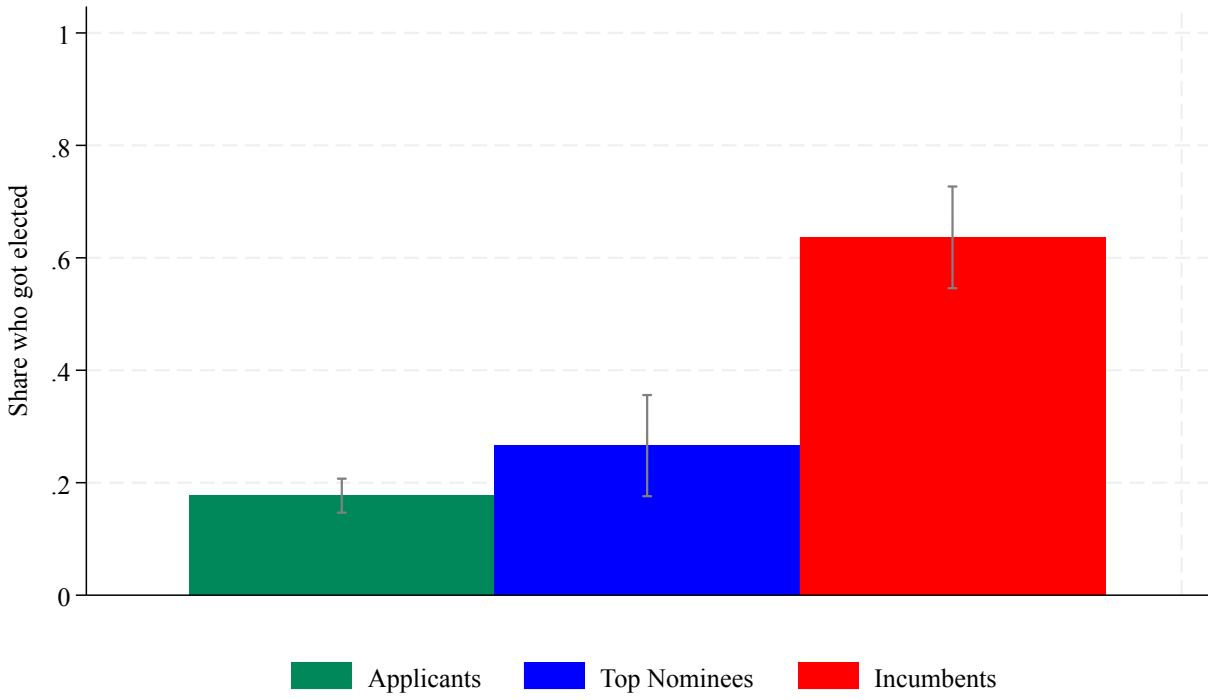
Notes: Figure shows simple regressions of how different variables predict which individuals in the stronghold sample who applied to become councillors go on to be selected as candidates. For incumbents, time-varying characteristics (political experience and family members in office) are measured as of their candidacy in the 2018 election. Regressions control for geographic strata used in the randomization of T1.

Figure A6: Raw Data on Progression from Entry to Election by Type

Panel A: Conditional Likelihood of Candidacy Selection by Applicant Type

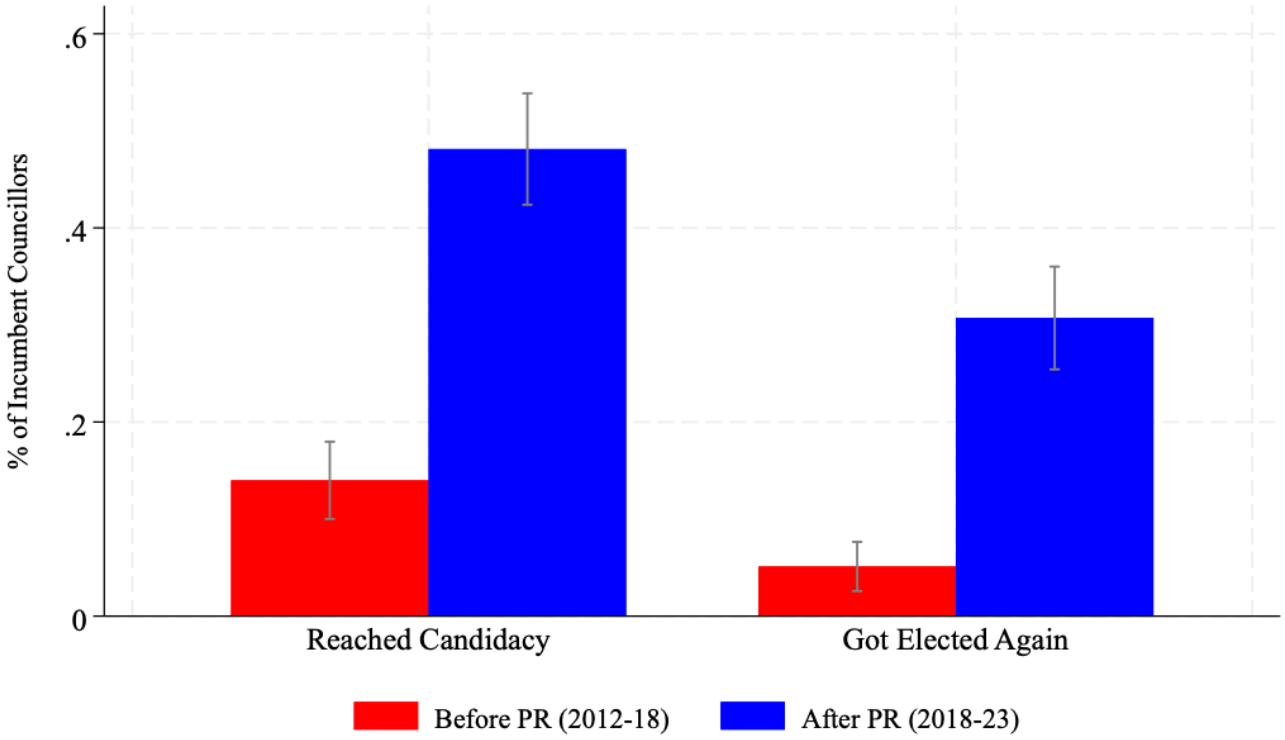


Panel B: Conditional Likelihood of Election by Candidate Type



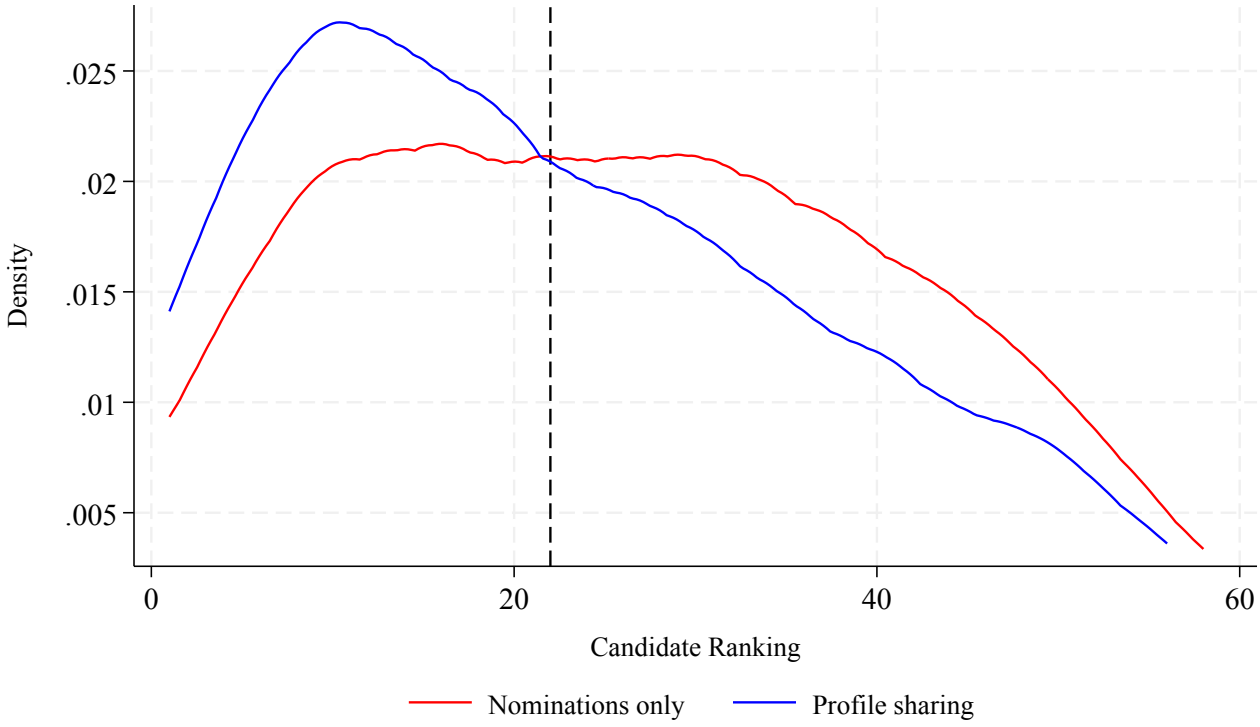
Notes: This figure plots the conditional likelihood of progressing to candidacy and being elected for the potential candidates in the full sample of our study: incumbents, top nominees, and status quo applicants. Panel A shows the share of each type that, conditional on filing an application, was selected onto a party's candidate list. Panel B shows the share of each type that, conditional on being placed on a candidate list, won a council seat.

Figure A7: Incumbency Advantage After PR



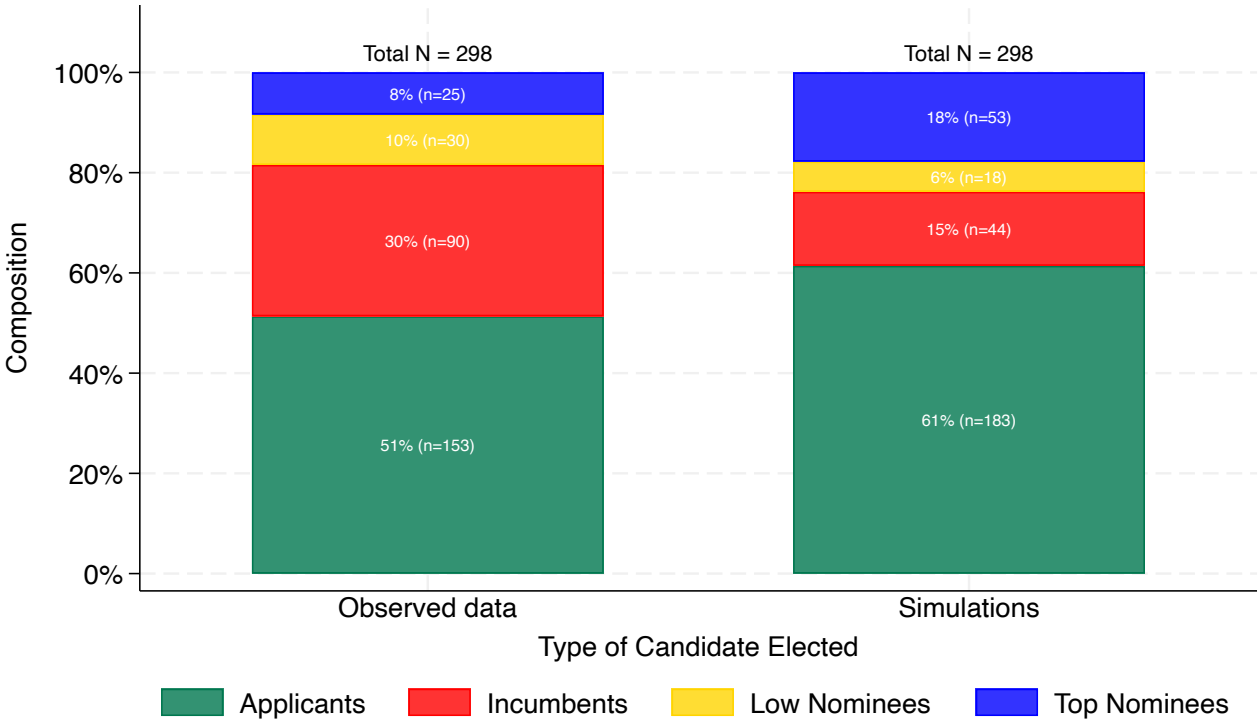
Notes: This figure shows raw number of incumbent councillors and their performance in the next elections. It shows the percentage of them, across 293 wards, that appear as candidates and that win again.

Figure A8: Distribution of Ward Ranks on Candidate Lists by T2 Status



Notes: This figure presents evidence on the impact of the profile sharing treatment (T2) using raw electoral data for the stronghold sample. It compares wards assigned to nominations only (T1 only) and wards assigned to both nominations and profile sharing (T1+T2), plotting the probability density function of candidates' numerical ranks on party candidate lists. The dashed vertical line at rank 22 denotes the average last winning rank on the dominant party's list in its regional stronghold.

Figure A9: Observed and Simulated Elected Councillors by Type



Notes: This figure shows the composition of elected Councillors among the different types of candidates, in contrast to the composition we would have observed if candidate lists were made based on our quality index. Simulated lists rank each party's applicants in a district by the quality index, placing the highest-ranked woman available in every third slot to respect the gender quota, and keep as many candidates as the party won seats. The sample includes all elected Councillors in the 14 study districts (N = 298).

2 Supplemental Tables

Table A1: Constituency Measures of Balance Across Saturation Assignment

	(1)	(2)	(3)	(4)	(5)
	Population	Pop. Density	Turnout	Win Margin	Num. Wards
Low Saturation	-1,261	-23.82	-968.0	-0.014	-0.027
	(1,417)	(76.38)	(768.9)	(0.032)	(0.104)
High Saturation	-2,314*	-68.26	-509.7	-0.003	-0.066
	(1,379)	(74.35)	(748.5)	(0.032)	(0.101)
Observations	184	184	184	184	184
Control Mean	54341	223.3	16993	0.610	3.520
F-Test p-value	0.246	0.632	0.453	0.892	0.797

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.

Table A2: Ward-level Measures of Balance Across Treatment Assignment

	(1)	(2)	(3)	(4)
	Population	Pop. Density	Turnout	Win Margin
T1	529.9*	-105.6	27.53	-0.034*
	(272.3)	(101.1)	(158.02)	(0.019)
T1 + T2	254.1	-109.0	-82.52	-0.019
	(242.5)	(90.0)	(140.70)	(0.017)
Observations	500	500	500	500
Control Mean	14661	309.5	4,593	0.585
F-Test p-value	0.143	0.458	0.630	0.199

Notes: This table tests for balance on observable characteristics across treatment arms. Nominations and screening (T1) were assigned once per ward and apply to both major parties; profile sharing (T2) was then assigned separately for each party among T1 wards. Pooling the two parties yields 500 party-ward observations, corresponding to 250 ward-level T1 assignments. Population data comes from the 2015 census. Turnout and Win margin are based on the presidential race of 2018.

Table A3: Quality Sub-Components by Type of Potential Candidates

	(1)	(2)	(3)	(4)	(5)	(6)
	Quality Index	Human Capital	Work Experience	Local Experience	Managerial Capital	Conscientiousness
Top Nominees	0.275***	0.305***	0.197***	0.260***	0.257***	0.432***
	(0.023)	(0.030)	(0.041)	(0.035)	(0.043)	(0.056)
Incumbents	0.173***	0.157***	-0.061	0.394***	-0.016	
	(0.030)	(0.038)	(0.052)	(0.051)	(0.053)	
Observations	1,670	1,670	1,670	1,670	1,670	1,433
P-value $\beta_1=\beta_2$	0.001	<0.001	<0.001	0.016	<0.001	

Notes: Regression outcomes are standardized relative to the omitted group in each regression, which is status quo applicants. Regressions include fixed effects for geographic strata used in the randomization of T1.

Table A4: Quality Variables by Type of Potential Candidates

Outcome	Nominees	Incumbents	p-value ($\beta_1 = \beta_2$)	Control Mean	N
Years of Education	1.490***	1.326***	0.256	11.435	1,669
	(0.138)	(0.153)			
Literate	0.174***	0.133***	0.006	0.821	1,670
	(0.012)	(0.019)			
Number Score	0.071*	0.080*	0.854	2.073	1,670
	(0.038)	(0.047)			
Ravens Score	0.544***	0.249	0.081	3.133	1,669
	(0.114)	(0.154)			
Digits Score	0.332***	-0.357***	<0.001	3.337	1,670
	(0.090)	(0.109)			
Monthly Wage (USD)	-5.385*	-4.245	0.784	61.263	1,598
	(3.038)	(3.874)			
Exp. in Social Sector	0.507***	-0.082	<0.001	1.275	1,670
	(0.062)	(0.075)			
Exp. with Development Projects	0.219***	0.205***	0.869	0.926	1,670
	(0.060)	(0.078)			
Exp. with Community Activities	-0.059	0.114*	0.009	1.317	1,670
	(0.049)	(0.061)			
Leadership Roles	0.309***	0.278***	0.699	1.002	1,670
	(0.054)	(0.071)			
NGO Knowledge	0.359***	0.600***	0.004	0.247	1,670
	(0.043)	(0.078)			
Project Proposal Score	0.339***	-0.005	0.005	3.602	1,670
	(0.095)	(0.109)			
Managerial Experience	0.317***	-0.028	<0.001	1.745	1,670
	(0.058)	(0.073)			
Conscientiousness	0.140***	–	–	0.138	1,433
	(0.023)				

Notes: Outcomes presented in raw units. Robust standard errors in parentheses. Equality p-value tests whether the Nominees and Incumbents coefficients are equal. The omitted group in each regression is status quo applicants. Regressions include fixed effects for geographic strata used in the randomization of T1.

Table A5: Top Nominees Differ from the Best Applicants in Control Wards - Full Sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Female	Asset Count	Ruling Family	Political Experience	Family in Politics	Family in Party	Quality Index
Panel A: Top nominees vs best two status quo applicants in Control wards							
Top Nominees	-0.002	0.054	0.448***	-0.310***	0.037	-0.034	0.176***
	(0.033)	(0.081)	(0.034)	(0.078)	(0.029)	(0.029)	(0.034)
Observations	719	719	719	719	719	719	719
Panel B: Top nominees that applied vs best two status quo applicants in Control wards							
Top Nominees	0.003	0.218**	0.381***	-0.068	0.003	0.005	0.201***
	(0.045)	(0.108)	(0.049)	(0.106)	(0.038)	(0.040)	(0.046)
Observations	347	347	347	347	347	347	347

Notes: Regressions include the full sample and include fixed effects for geographic strata used in the randomization of T1.

Table A6: Estimated Impacts on the Size of the Applicant Pool - Full Sample

	(1)	(2)	(3)	(4)
	Any Applicant	Num. of Applicants	Any Women Applicant	Num. of Women Applicants
T1 Nominations	0.006	0.011	0.030	0.076
	(0.017)	(0.156)	(0.058)	(0.088)
T2 Profiles	0.014	-0.053	-0.071	-0.063
	(0.011)	(0.137)	(0.052)	(0.083)
Observations	500	500	500	500
Control Mean	0.976	2.651	0.492	0.635

Notes: Regressions estimated with observations collapsed at the party-ward level using the full sample and include fixed effects for geographic strata used in the randomization of T1.

Table A7: Estimated and Simulated Impacts on the Quality of the Applicant Pool - Full Sample

	(1)	(2)	(3)	(4)	(5)	(6)
	Applicant Pool		+25% App. Rate		+50% App. Rate	
	Mean	Max	Mean	Max	Mean	Max
T1 Nominations	0.015	0.025	0.022	0.041	0.029	0.056
	(0.037)	(0.043)	(0.038)	(0.044)	(0.038)	(0.044)
T2 Profiles	0.019	-0.007	0.019	-0.006	0.020	-0.004
	(0.033)	(0.040)	(0.033)	(0.041)	(0.033)	(0.041)
Observations	493	493	494	494	495	495
Control Mean	0.006	0.274	0.006	0.274	0.006	0.274
P-value $\beta_1 = 0$	0.695	0.564	0.558	0.350	0.450	0.205
Individuals in Sample	1321	1321	1351	1351	1381	1381
Top Noms in Sample %	10.1	10.1	12.1	12.1	14	14
Ward Compliance %	32.5	32.5	40.6	40.6	48.7	48.7

Notes: Regressions estimated with observations collapsed at the party-ward level using the full sample and include fixed effects for geographic strata used in the randomization of T1. Main outcome is the quality index aggregated in different ways for each ward observation: average across all applicants and maximum quality of the applicant pool. Columns 3 to 6 report results from a Monte Carlo simulation with 500 draws. In each draw, for wards where we did not observe a top nominee entering, we select the highest quality top nominee who expressed interest in running but did not apply and we boost their entry, first adding them to the pool with a 25% probability (columns 3 and 4) and then with a 50% probability (columns 5 and 6). Coefficients are pooled means across simulations. Standard errors and p-values combine within-simulation sampling uncertainty and across-simulation variance. Individuals in sample, top nominees in sample, and ward compliance are averages across simulations.

Table A8: Estimated Impacts on the Selected and Elected Candidates - Full Sample

	(1)	(2)	(3)	(4)
	Any Candidate	Any Women Candidate	Inv. Candidate Ranking	Any Elected
T1 Nominations	-0.002	0.064	-0.666	-0.039
	(0.029)	(0.058)	(1.265)	(0.058)
T2 Profiles	0.011	-0.073	0.947	-0.014
	(0.024)	(0.051)	(1.161)	(0.052)
Observations	498	498	463	498
Control Mean	0.935	0.419	-22.537	0.468
P-value $\beta_1 + \beta_2 = 0$	0.720	0.869	0.823	0.350

Notes: Regressions estimated with observations collapsed at the party-ward level using the full sample and include fixed effects for geographic strata used in the randomization of T1. Column 3 is estimated on the 463 party-wards with at least one candidate on the party list. Because treatment may affect whether a ward has a candidate, this estimate conditions on a post-treatment outcome and should be interpreted with caution.

Table A9: Estimated Impacts on the Quality of the Candidate and Elected Pool - Full Sample

	(1)	(2)	(3)	(4)
	Candidate Lists		Elected Officials	
	Mean Q	Max Q	Mean Q	Max Q
T1 Nominations	0.016	0.013	0.065	0.073
	(0.043)	(0.048)	(0.068)	(0.069)
T2 Profiles	0.028	-0.024	0.004	-0.013
	(0.038)	(0.043)	(0.057)	(0.056)
Observations	468	468	217	217
Control Mean	0.005	0.201	0.000	0.017
P-value $\beta_1 + \beta_2 = 0$	0.302	0.817	0.275	0.349

Notes: Regressions estimated with observations collapsed at the party-ward level using the full sample and include fixed effects for geographic strata used in the randomization of T1.

Table A10: Estimated Impacts on Candidate List Positions

	(1)	(2)	(3)	(4)
	Any Candidate in Party List Quartile			
	Top	Second	Third	Bottom
T1 Nominations	-0.000	0.080	-0.033	0.134
	(0.085)	(0.085)	(0.085)	(0.088)
T2 Profiles	0.090	-0.040	0.009	-0.204***
	(0.075)	(0.076)	(0.074)	(0.075)
Observations	250	250	250	250
Control Mean	0.365	0.429	0.349	0.397
P-value $\beta_1+\beta_2 = 0$	0.261	0.620	0.769	0.395

Notes: Regressions estimated with observations collapsed at the ward level using the stronghold sample and include fixed effects for geographic strata used in the randomization of T1. Outcomes indicate whether the ward has at least one candidate in the given quartile of the party list, where quartiles are defined relative to list length, which varies with the number of wards in the district.

Table A11: Multiple Hypothesis Testing Corrections

	p-value	Sharpened q-value
Panel A: Stronghold Sample		
Number of applicants — Nominations	0.189	0.608
Number of applicants — Profiles	0.834	0.716
Avg. quality applicants — Nominations	0.119	0.608
Avg. quality applicants — Profiles	0.479	0.608
Any candidate — Nominations	0.404	0.608
Any candidate — Profiles	0.168	0.608
Avg. quality candidates — Nominations	0.076	0.608
Avg. quality candidates — Profiles	0.510	0.608
Panel B: Full Sample		
Number of applicants — Nominations	0.472	1.000
Number of applicants — Profiles	0.651	1.000
Avg. quality applicants — Nominations	0.343	1.000
Avg. quality applicants — Profiles	0.283	1.000
Any candidate — Nominations	0.528	1.000
Any candidate — Profiles	0.324	1.000
Avg. quality candidates — Nominations	0.355	1.000
Avg. quality candidates — Profiles	0.231	1.000

Notes: p-values are from the primary regressions and are one-sided, testing for positive treatment effects as prespecified in the pre-analysis plan. Sharpened q-values are false discovery rate (FDR) adjusted following Anderson (2008) and Benjamini, Krieger, & Yekutieli (2006). The adjustment is applied jointly across the 8 estimates in each sample (two coefficients, Nominations and Profiles, for each of four outcomes: number of applicants, average quality of applicants, any candidate, and average quality of candidates)

Table A12: Quality Index by Type of Potential Candidates Who Applied

	(1)	(2)	(3)	(4)	(5)
	Outcome: Quality Index (excluding conscientiousness)				
	Full Sample	Stronghold	Only Applicants	Applicants in Wards Incumbent Applies	Applicants in Wards Both Apply
Top Nominees	0.266***	0.191***	0.201***	0.224***	0.244***
	(0.024)	(0.029)	(0.050)	(0.072)	(0.085)
Incumbents	0.170***	0.010	-0.006	0.020	0.083
	(0.030)	(0.034)	(0.040)	(0.045)	(0.078)
Observations	1,670	979	616	322	106
P-value β_1=β_2	0.002	<0.001	<0.001	0.005	0.050
N Wards	250	250	242	117	31

Notes: Column 1 uses the full sample; columns 2 to 5 use the stronghold sample. Top nominees are drawn from the 186 wards assigned to community nominations and screening (T1); incumbents and status quo applicants are drawn from the 250 wards in the research sample. Regression outcome is standardized relative to the omitted group in each regression, which is status quo applicants. The quality index excludes the conscientiousness component, which was not measured for incumbents. Regressions include fixed effects for geographic strata used in the randomization of T1. Column 3 restricts to Top Nominees and Incumbents who apply to parties. Column 4 only includes wards where an incumbent re-applies. Column 5 uses only wards where an incumbent re-applies and a top nominee from that ward also applied.

Table A13: Local Support by Type of Potential Candidates Who Applied

	(1)	(2)	(3)	(4)
	Outcome: Vote Share in Nominations			
	Stronghold	Only Applicants	Applicants in Wards Incumbent Applies	Applicants in Wards Both Apply
Top Nominees	0.130***	0.187***	0.187***	0.187***
	(0.008)	(0.016)	(0.026)	(0.025)
Incumbents	0.030***	0.050***	0.059***	0.041*
	(0.010)	(0.014)	(0.014)	(0.023)
Observations	769	428	208	106
P-value β_1=β_2	<0.001	<0.001	<0.001	<0.001
N Wards	185	178	83	31

Notes: All columns use the stronghold sample. Regression outcome is vote share in the nominations exercise, which implies only individuals from T1 wards are included. Regressions include fixed effects for geographic strata used in the randomization of T1. Column 2 restricts to Top Nominees and Incumbents who apply to parties. Column 3 only includes wards where an incumbent re-applies. Column 4 uses only wards where an incumbent re-applies and a top nominee from that ward also applied.

Table A14: Predictors of Incumbent Performance: Index of Peer Ratings

	Selection frequency		Standardized coefficient	
	(1)	(2)	(3)	(4)
Predictor	Prediction stability	CV-selected	Bivariate	Multivariate
Literacy	0.904	1.000	0.216***	0.202***
			(0.043)	(0.053)
Project proposal score	0.888	1.000	0.206***	0.185***
			(0.050)	(0.058)
Woman	0.761	1.000	0.153***	0.176***
			(0.047)	(0.048)
Development projects	0.320	1.000	0.059	0.093*
			(0.049)	(0.052)
Digits memory test	0.244	0.627	0.096*	0.054
			(0.049)	(0.054)
Community projects	0.240	0.627	-0.040	-0.061
			(0.054)	(0.054)
Numeracy score	0.163	0.000	0.085**	0.013
			(0.043)	(0.045)
Wages	0.148	0.000	0.041	0.000
			(0.053)	(0.057)
Years in Council	0.143	0.000	0.021	-0.029
			(0.053)	(0.053)
Political experience	0.143	0.000	0.043	0.033
			(0.047)	(0.047)
Experience in social sector	0.142	0.000	0.047	-0.003
			(0.053)	(0.064)
Leadership positions	0.140	0.000	0.069	0.007
			(0.047)	(0.051)
Years of education	0.128	0.000	0.099*	-0.011
			(0.053)	(0.051)
Management experience	0.096	0.000	0.043	-0.026
			(0.046)	(0.059)
Raven's score	0.086	0.000	0.066	-0.021
			(0.047)	(0.048)
NGO knowledge	0.080	0.000	0.030	-0.003
			(0.043)	(0.044)
R^2 (full model)		0.116		

Notes: Outcome averages standardized peer ratings (by other councillors) on effectiveness in office and honesty. The sample includes 436 councillors from all district and city councils except Freetown City Council, representing 398 wards (some city wards elect multiple councillors). All variables are standardized. Predictors are the quality-index components plus three characteristics that predict party selection (Figure 10), in bold: years in Council, gender, and political experience. Rows are sorted by column (1). Column (2) reports the share of 300 elastic-net runs that select the variable, each run choosing its penalty by cross-validation on newly drawn folds. In these data many penalties predict about equally well but retain very different numbers of variables, so the penalty chosen in each run is close to arbitrary. The share in column (2) therefore reflects that choice as much as how well a variable predicts; Ritzwoller & Romano (2025) document the instability of this method using a very similar dataset from Casey, Kamara, & Meriggi (2021). Column (1) does not depend on a single penalty. We take a grid of relevant penalty values and, at each one, fit the elastic net on 250 pairs of complementary random half-samples (Meinshausen & Bühlmann, 2010; Shah & Samworth, 2013). The column reports the share of fits that select the variable, averaged across half-samples, which captures whether it predicts performance in many subsets of the data, and across the penalty grid, which captures whether it survives stricter, more heavily penalized models; while also improving the reproducibility of the method (Ritzwoller & Romano, 2025). Columns (3) and (4) report bivariate and multivariate coefficients. Robust standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Table A15: Predictors of Incumbent Performance: Index of Self Reported Proxies

	Selection frequency		Standardized coefficient	
	(1)	(2)	(3)	(4)
Predictor	Prediction stability	CV-selected	Bivariate	Multivariate
Project proposal score	0.765	1.000	0.183***	0.151**
			(0.049)	(0.061)
NGO knowledge	0.670	1.000	0.155***	0.114**
			(0.048)	(0.049)
Community projects	0.522	1.000	0.132**	0.089*
			(0.051)	(0.049)
Development projects	0.496	1.000	0.133***	0.073
			(0.050)	(0.053)
Raven's score	0.366	1.000	0.114***	0.101**
			(0.043)	(0.050)
Woman	0.338	1.000	-0.106**	-0.072
			(0.048)	(0.049)
Management experience	0.196	0.340	0.084*	0.079
			(0.047)	(0.054)
Wages	0.189	0.340	-0.044	-0.056
			(0.062)	(0.061)
Numeracy score	0.112	0.060	-0.013	-0.044
			(0.052)	(0.054)
Experience in social sector	0.106	0.000	0.067	-0.017
			(0.047)	(0.056)
Political experience	0.100	0.057	0.049	0.027
			(0.047)	(0.045)
Years in Council	0.099	0.057	0.034	0.033
			(0.052)	(0.053)
Leadership positions	0.074	0.057	0.010	-0.062
			(0.046)	(0.050)
Literacy	0.056	0.007	0.033	0.018
			(0.044)	(0.048)
Digits memory test	0.056	0.057	0.010	-0.061
			(0.047)	(0.054)
Years of education	0.046	0.000	0.006	-0.006
			(0.047)	(0.052)
R^2 (full model)		0.096		

Notes: Outcome averages standardized self reported proxies of performance in office: number of development projects implemented, average cost of such projects, and number of meetings with constituents. The sample includes 436 councillors from all district and city councils except Freetown City Council, representing 398 wards (some city wards elect multiple councillors). All variables are standardized. Predictors are the quality-index components plus three characteristics that predict party selection (Figure 10), in bold: years in Council, gender, and political experience. Rows are sorted by column (1). Column (2) reports the share of 300 elastic-net runs that select the variable, each run choosing its penalty by cross-validation on newly drawn folds. In these data many penalties predict about equally well but retain very different numbers of variables, so the penalty chosen in each run is close to arbitrary. The share in column (2) therefore reflects that choice as much as how well a variable predicts; Ritzwoller & Romano (2025) document the instability of this method using a very similar dataset from Casey, Kamara, & Meriggi (2021). Column (1) does not depend on a single penalty. We take a grid of relevant penalty values and, at each one, fit the elastic net on 250 pairs of complementary random half-samples (Meinshausen & Bühlmann, 2010; Shah & Samworth, 2013). The column reports the share of fits that select the variable, averaged across half-samples, which captures whether it predicts performance in many subsets of the data, and across the penalty grid, which captures whether it survives stricter, more heavily penalized models; while also improving the reproducibility of the method (Ritzwoller & Romano, 2025). Columns (3) and (4) report bivariate and multivariate coefficients. Robust standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Table A16: Quality Index and Experience Predicting Incumbent Performance

	(1)	(2)	(3)
Panel A. Outcome: Performance Index by Peer Reviews			
Quality Index	0.199***	0.225***	0.271***
	(0.045)	(0.044)	(0.052)
Years in Council	0.011	-0.017	-0.016
	(0.051)	(0.052)	(0.053)
Observations	436	436	436
Panel B. Outcome: Performance Index by Self Reports			
Quality Index	0.179***	0.179***	0.211***
	(0.044)	(0.044)	(0.049)
Years in Council	0.051	0.067	0.063
	(0.049)	(0.053)	(0.050)
Observations	436	436	436
Controls	No	Yes	Yes
District fixed effects	No	No	Yes

Notes: Sample includes sitting councillors before 2023 elections. Outcomes are standardized performance indices, Panel A uses the anonymous peer reviews among councillors (on effectiveness and honesty), while Panel B proxies performance with self reports of the number of development projects during office, the average cost of such projects, and number of meetings with constituents. Controls included in regressions are age, household size, a wealth index, political experience, ruling party membership, ruling family membership and gender. Robust standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

3 Supplemental Exhibits

Exhibit A1: Enumerator script for household elicitation

We are helping all political parties find the best candidates for local council by including the people in this process of nominations. However, the political parties are not required to accept the nominations; they will simply be encouraged to consider the people's choices.
Good Councillors have the following skills: • Hardworking • Honest • Have administrative skills • Listen to people • Represent their community well on council. The job of councillor involves the following: • Deciding how to use council resources • Councillor is a part-time job • Councillors are not paid (other than sitting fees and transport), so people keep other jobs while being a Councillor.
Every adult that lives in this ward regardless of their sex, ethnicity, or other characteristics are able to fill this role. Think of people in the whole ward (not just your own town). Also think about sons or daughters of the soil, even if they are not currently living in the ward. Councillors must be able to read and write English and speak Krio.
I am now going to ask you to identify people in your community who would be good candidates to run for local council. Can you think of any such community members? **Only identify people who can read English and speak Krio**
What is the full name of this person? *ENUMERATOR NOTE: Enter their name as it would appear on a **legal government document** in all capital letters. Enter any nicknames in the next question.*
Does this person have any nicknames?
What is this person's sex?
What is this person's phone number?
Which village / town does this person live in right now?
Where could we find this person? *ENUMERATOR NOTE: We want instructions on how to find this person, this person's house, this person's job, etc. Also record other useful information, such as if this person is out of town. This is the information you will use to locate nominees to survey.*
Why do you think this person would be a good candidate for councillor? *ENUMERATOR NOTE: Summarize what the respondent said to 1 sentence*
**** REPEAT UNCONSTRAINED NOMINATION UP TO 3 NAMES ****

I am now going to ask you to identify women in your community who would be good candidates to run for local council. Can you think of any such community members? **Only identify women who can read English and speak Krio**
**** IDENTICAL DATA COLLECTION ON NOMINEE ****
**** REPEAT CONSTRAINED NOMINATION UP TO 3 NAMES ****

Exhibit A2: Excerpt of enumerator script for community nominees

We are working with political parties to identify high quality candidates. We are talking to you because at least [number_nominations] person(s) in your community said you would make a good Councillor *ENUMERATOR NOTE: You cannot tell the respondent who nominated them*
Specifically, people mentioned the following as why you would make a good Councillor: *ENUMERATOR NOTE: Summarize the following and read to respondent:* "[nominee_why_prefill]" *ENUMERATOR NOTE: You cannot tell the respondent who nominated them, or gave this reason for nominating them*
Can we give your name and profile to party leadership as a candidate for Councillor in this area? If you say yes, you would consent to us sharing information we collect in this interview with party leadership. If you say you are not interested in sharing this information, it will all remain a secret. Regardless of your answer the interview will continue as normal.
Which parties would you be ok with receiving this information?
Can I take a headshot photo of you? We will include this photo when we send your photo to the political parties *ENUMERATOR NOTE: Take a photo that would be used for an ID card, only showing their face and shoulders*
Thank you for your interest. Please know the following: 1) We will not share the information of everyone we interview 2) There is a chance your information will not be shared with any party 3) If your information is shared, the parties are not required to consider your application, advance you in the selection process, or select you as a candidate 4) We provide information to the parties. We cannot control what parties do with that information.

Exhibit A3: A Note on Preregistration

We submitted our pre-analysis plan for this paper to the Journal of Development Economics Registered Reports on Feb 15, 2024. At the moment of pre-registration the research team had only partial access to the data used in this paper and could not construct the main outcomes needed for the analysis. This was possible because Innovations for Poverty Action in Sierra Leone committed to withhold key data from the research team until February 21, 2024 as part of our pre-registration. The data withheld included ward-level identifiers and survey data needed for quality metrics for all aspirants and candidates in their research sample. The table below maps the questions and hypotheses from our PAP to the Figures and Tables in the paper.

Test proposed in the Pre-analysis Plan	**Corresponding table or figure in this paper**
Q1: Does the intervention identify high quality potential candidates?	Figure 4, Table 2
Q2: Does the intervention identify popular potential candidates?	Figure 7, Table 5
Q3: Does the intervention identify new entrants into politics?	Figure 5
Q4: Do the nominees enter politics?	Table 6
Q5: Did party leaders accept the nominee profiles?	Text passage in Section VI.C
H1: The intervention expands the aspirant pool	Table 8
H2: The intervention enhances the quality of the aspirant pool	Table 9
H3: The intervention enhances geographic representation on the candidate lists	Table 10
H4: The intervention increased the quality of selected candidates	Table 11
H5: The intervention enhanced geographic representation in the elected Local Council	Table 10
H6: The intervention increased the quality of elected Councillors	Table 11
Estimation of potential spillover effects	Table A17
Adjustments for multiple inference	Table A11
Dimensions for potential heterogeneous effects	Table 4
Full sample analogs of stronghold results	Tables A5–A9

Exhibit A4: Discussion and Analysis of Spillovers

A key issue with the switch from ward-level majoritarian elections to district-level proportional representation is that 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. How much of a risk this posed to analysis could not be known until political parties submitted their official candidate lists to the National Electoral Commission, which is after all interventions were delivered. In anticipation of, and to mitigate any realized interference, we use a saturation strategy for profile sharing at the constituency level. As it turned out, parties largely stuck to the old system, allocating two candidates from each ward onto the list, making this analysis largely moot. We include below our strategy from our pre-analysis plan and formal estimation of spillovers here for transparency and completeness.

To see how bias could arise, consider first estimates for the number of candidates from a ward selected for the district candidate list. 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 estimated treatment effect on the number of candidates per ward 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. Extending the scenario above, suppose that in the absence of treatment, parties select one high quality (H) and one low quality (L) 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 $1/2$ to $3/4$).

To explore the relevance of potential interference, we leverage the saturation design at the constituency level. We illustrate the method for 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 hypotheses. Moreover, if the candidate lists submitted by parties to the National Electoral Commission adhere closely to two candidates per ward, which they in fact did, spillovers are unlikely to matter. Yet if we had found evidence of interference, a similar specification could have been applied to the other hypotheses.

To do so, we follow the approach in Baird, Bohren, McIntosh, & Özler (2018) where our randomized saturation design can be characterized by saturation vector $\Pi = \{0, 0.33, 0.67\}$ 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 was not assigned to profile sharing. 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 ($p = 0.67$).
To assess whether there are spillovers from treated to control wards, we estimate:

$$Y_{wcs} = \beta_0 + \beta_{1L} Profiles_{wcs} L_c + \beta_{1H} Profiles_{wcs} H_c + \gamma_{1L} S_{wc} L_c + \gamma_{1H} S_{wc} H_c + \alpha_s + \varepsilon_{wcs} \quad \text{(A1)}$$

where Y_{wcs} is the number of selected candidates from ward w in constituency cluster c; $S_{wc} = 1$ indicates a

within-cluster control, or ward in a cluster with positive treatment saturation that was not assigned to profile sharing; $L_c = 1$ indicates a constituency cluster allocated to low saturation treatment ($p = 0.33$), and $H_c = 1$ a cluster allocated to high saturation ($p = 0.67$). Hypotheses of interest are: $\gamma_{1L} \neq 0$ or $\gamma_{1H} \neq 0$ which 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 estimates of interest for policymakers are 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.
Appendix Table A17 presents results, where we see that spillover estimates for the stronghold sample are null. The point estimates are small in magnitude, vary in sign across outcomes, and none are statistically significant at conventional levels. This aligns with the evidence in the raw data presented in the main text Figure 9 that shows how parties largely stuck with the old system: panel A shows that the modal ward has 2 candidates on the list, and panel B shows that the vast majority of wards ended up with a single elected Local Councillor.

Table A17: Estimation of Spillover Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	No. Candidates		No. Elected		Quality Candidates		Quality Top Candidates	
Profile sharing	-0.221*		0.029		0.024		0.010	
	(0.129)		(0.074)		(0.046)		(0.064)	
Profiles (Low)		-0.304		0.079		0.115		0.157
		(0.210)		(0.111)		(0.081)		(0.110)
Profiles (High)		-0.245		0.098		-0.063		-0.066
		(0.168)		(0.108)		(0.046)		(0.069)
Spillover (Low)		-0.151		0.064		-0.060		0.012
		(0.198)		(0.109)		(0.065)		(0.091)
Spillover (High)		0.048		0.134		-0.031		-0.042
		(0.183)		(0.117)		(0.058)		(0.094)
Observations	250	250	250	250	239	239	196	196
Control Mean	1.967	2.057	0.742	0.679	0.166	0.176	0.221	0.224

Notes: This table estimates potential spillover effects from wards assigned to profile sharing (T2) onto neighboring control wards within the same constituency, using the saturation design at the constituency level. Columns 1–2 report effects on the number of candidates, columns 3–4 on the number of elected councillors, columns 5–6 on the mean quality of candidates, and columns 7–8 on the quality of the top candidate per ward. In each pair of columns, the first reports results for the intention-to-treat effect of assignment to profile sharing and the second estimates direct and indirect effects within constituencies. Regressions use the stronghold sample and follow the saturation specification detailed in Exhibit A4, including geographic strata fixed effects.