

NBER WORKING PAPER SERIES

AN EXPERIMENT IN CANDIDATE SELECTION

Katherine Casey
Abou Bakarr Kamara
Niccoló Meriggi

Working Paper 26160
<http://www.nber.org/papers/w26160>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2019

We thank the All People's Congress, the Sierra Leone People's Party, Search for Common Ground, the Political Parties Registration Commission, and the National Electoral Commission of Sierra Leone for their collaboration. We thank Chiara Amato, Abdulai Bah, Arthur Baker, Fatu Conteh, Arja Dayal, Yuen Ho, Fatoma Momoh, Alfred Rasheed, Andreas Rodriguez, Marcos Salgado, Korinayo Thompson, and the entire Innovations for Poverty Action team in Freetown for excellent research assistance and fieldwork. We are especially indebted to Rachel Glennerster for her comments and support, as well as to Dave Baron, David Broockman, Steve Callander, Fred Finan, Emmanuel Gaima, Pedro Gardete, Guido Imbens, Saum Jha, Alhassan Kanu, Eddie Lazear, Herbert M'cleod and Stephan Seiler for insightful feedback. We gratefully acknowledge financial support from the Economic Development Initiative funded by UK Aid, the Governance Initiative at J-PAL, and the International Growth Centre. All errors are our own. This study was pre-registered on the AEA registry:

<https://www.socialscisceregistry.org/trials/2485/history/23509>

The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Katherine Casey, Abou Bakarr Kamara, and Niccoló Meriggi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

An Experiment in Candidate Selection
Katherine Casey, Abou Bakarr Kamara, and Niccoló Meriggi
NBER Working Paper No. 26160
August 2019
JEL No. D72,H1,P16

ABSTRACT

Are ordinary citizens or political party leaders better positioned to select candidates? While the direct vote primary system in the United States lets citizens choose, it is exceptional, as the vast majority of democracies rely instead on party officials to appoint or nominate candidates. Theoretically, the consequences of these distinct design choices on the selectivity of the overall electoral system are unclear: while party leaders may be better informed about candidate qualifications, they may value traits—like party loyalty or willingness to pay for the nomination—at odds with identifying the best performer. To make progress on this question, we partnered with both major political parties in Sierra Leone to experimentally vary how much say ordinary voters, as opposed to party officials, have in selecting Parliamentary candidates. We find evidence that more democratic selection procedures increase the likelihood that parties select the candidate most preferred by voters, favor candidates with stronger records of local public goods provision, and alter the allocation of payments from potential candidates to party officials.

Katherine Casey
Stanford University
Graduate School of Business
655 Knight Way
Stanford, CA 94305
and NBER
kecasey@stanford.edu

Niccoló Meriggi
International Growth Centre
niccolo.meriggi@theigc.org

Abou Bakarr Kamara
International Growth Centre
abou.kamara@theigc.org

I. Introduction

Are ordinary citizens or political party leaders better positioned to select candidates? While the direct vote primary system in the United States lets citizens choose, it is exceptional, as the vast majority of democracies rely instead on party officials to appoint or nominate candidates.¹ Despite this considerable heterogeneity, there is relatively little theoretical or empirical research on the consequences of these distinct design choices for the performance of the overall electoral system. While voter control straightforwardly facilitates representation, e.g. citizens get their most preferred candidate, it may come at the expense of selection on quality or competence, particularly if voters are poorly informed. Party leaders, by contrast, have the expertise and information needed to screen on quality, but may be out of step with voter preferences, either because their policy positions are shaped by their elite status, or because they reward candidates for traits unrelated to performance in office (like party loyalty or willingness to pay for the nomination). These tradeoffs are particularly consequential in partisan strongholds, as the process the dominant party uses to select its candidate is effectively the process that determines the identity of the elected politician.² Empirical work on these issues has been constrained by the fact that parties are generally loathe to vary how they choose candidates for anything but purely strategic (and thus endogenous) reasons.

This is the first paper to directly evaluate this question leveraging experimental variation in how much say ordinary voters, as compared to party officials, have in selecting candidates. To do so, we partnered with both major political parties in Sierra Leone to experiment with different methods of selecting candidates for the 2018 Parliamentary election. In the status quo, parties chose among potential candidates, referred to as “aspirants,” in a given constituency via recommendations from party officials at various levels, with no direct participation by voters. For a randomly selected subset of races, the parties implemented a new selection method with two components: i) a party convention where aspirants presented their qualifications to party officials and local residents, and engaged in informative policy-oriented debate that was broadcast over local radio; followed by ii) opinion polling, representative of all registered voters in the constituency, that elicited and aggregated voter preferences over aspirants, which was shared with

¹ And some argue that party leaders retain more influence in U.S. nominations than is popularly appreciated (e.g. Cohen et al. 2008).

² Strongholds, particularly for sub-national races, are common in the U.S.—e.g. a majority of House races over the past 60 years were decided by margins wider than 15 percentage points (Hirano and Snyder 2014)—and in much of the developing world, where ethnic- or caste-identities are correlated with both party loyalty and geographic sorting.

party officials via a one page report. Note that neither component is binding on the parties' ultimate choice of candidate. Yet if party officials followed the voter reports in all cases (which they did not), the intervention approximates a direct vote primary with mandatory turnout.³ The research team randomly assigned the intervention for each party independently, with 23 treated and 23 control races per party (for a total sample of 92 party-races nationwide); implemented the opinion polls; and collected rich data from voters, aspirants and party officials.

Analysis proceeds in two stages. First, we use our data to characterize candidate selection in the status quo. We document strong positive self-selection into politics on education and wealth: for example, while 43 percent of registered voters in Sierra Leone have no formal schooling, 80 percent of aspirants have been to university. Constituency-level party officials are situated in between the two, having on average completed 12 years of schooling, which suggests they may have an advantage over voters in screening aspirants on technical qualifications. At the same time, party officials appear poorly informed about voter preferences: they presume local voters share their first choice over aspirants in 90 percent of races, when voters in fact only agree with them in 55 percent; and in a third of races, not a single party official (among multiple surveyed per race) accurately guessed which aspirant ranked first among voters. Aspirants make financial contributions (in official application fees and unofficial payments) to the party to be considered for candidacy, data we elicited from aspirants via survey. These payments amount to an unadjusted mean of \$2,406, which is equivalent to 1.3 months of an MP's official salary and 34 times the monthly minimum wage. These facts make it unclear *ex ante* what the optimal allocation of control between party officials and voters might be in this setting. This allocation could be particularly impactful given that the majority of Parliamentary races are located in ethnicity-based, regional strongholds, where the dominant party's candidate is near guaranteed to win the general election.

Our second, and main, contribution is to estimate how the more democratic internal party process affects three key outcomes of interest, focusing on representation, selection on quality, and payments from aspirants to party officials. First, we find evidence for large positive treatment effects on representation, defined as whether the aspirant who ranked first among voters was awarded the party's "symbol" to compete in the general election. In control races, party officials selected the voters' first choice aspirant 37 percent of the time, which increased to 61 percent with

³ The leap from voluntary to mandatory turnout is not unreasonable in this setting, as voluntary turnout reached 87% in the general election studied. Turnout could, of course, be lower or more selective in the primary stage.

treatment, a large and highly statistically significant effect. To provide a sense of magnitude, this effect size suggests that party officials chose a different candidate than they would have otherwise in 11 races, and they thereby changed the identity of 6 elected Members of Parliament (MPs) (e.g. those from stronghold races where their party was likely to win the general election).

Second, regarding selection, we identify which aspirant traits predict their popularity among voters and party officials, respectively, and then estimate treatment effects on the characteristics of candidates eventually sent to the general election. *Ex ante*, we find little evidence for conflict in preferences: both voters and party officials prefer aspirants with a stronger record of having provided local public goods and those who are more conscientious. The former captures the aspirant's involvement in development-oriented projects, like small scale public infrastructure (constructing a bridge or community center), support to education (rehabilitating classrooms) and agriculture (procuring farm inputs like tools and tractors). The latter is a behavioral measure of how carefully the aspirant handled a financial reimbursement for transport expenses (described in detail in Section V.B.). *Ex post*, our treatment effect estimates suggest that the more democratic method led to the selection of candidates with stronger public goods records, meaning that aspirants who had provided more development projects in the past were more likely to be chosen to advance to the general election. To the extent that past provision predicts future provision, this is a cautiously optimistic result. Our data further suggest that the conventions and associated radio broadcasts enhanced voter knowledge of aspirant qualifications, which may have helped voters to select on quality.

Third, patterns of aspirant contributions in the data are consistent with their being tied to the expected returns to candidacy, and these returns, and hence payments, adjust under the experimental treatment. Expected returns are a function of the value of office, e.g. an MP's salary and promotion prospects; the probability the party wins the seat, which is higher in ethnic strongholds; and the likelihood of being selected as the candidate, which is affected by treatment. Regarding the second term, mean payments are three times higher in stronghold seats compared to weak seats, as expected. As to the third term, if the conventions and opinion polls reveal information about aspirant type (quality, local popularity) to party officials, then high (low) type aspirants become more (less) likely to be selected, and should thus increase (decrease) their contributions. Consistent with this, estimates suggest that selected candidates pay differentially more than unsuccessful aspirants in treated races. Lastly, we do not find much evidence that

advance announcement of the initiative, which was only partially implemented, induced aspirant entry.

Our analysis contributes to the relatively new literature on political selection,⁴ or the factors that determine the quality of elected politicians (see Dal Bó and Finan 2018 for review). There is emerging consensus that higher returns to holding office (Ferraz and Finan 2011, Gagliarducci and Nannicini 2013, Fisman et al. 2015) and greater political competition (Galasso and Nannicini 2011, De Paola and Scoppa 2011, Dal Bo et al 2017) facilitate positive selection on candidate quality.⁵ Much less is known about the influence of party leaders on selection, and the two most related studies provide contrasting results. Dal Bo et al (2017) point to the role of party leaders in making merit-based promotions (among other institutional factors) in contributing to Sweden’s “inclusive meritocracy,” one characterized by positive selection on candidate competence across the socioeconomic spectrum. By contrast, Besley et al (2017) use the same data and find evidence that party officials have incentives to select competent candidates, but only to the extent that they do not pose an internal leadership threat, a tension that induces lower competence leaders to field lower competence candidates. A similar theoretical result can be found in Mattozzi and Merlo (2007), who model party recruitment of candidates as an opportunity to maximize rents, with the effect of promoting “mediocracy.”

Historically, skepticism about the role party elites play in candidate selection was a key driver of the Progressive movement to adopt direct vote primaries in the United States. Hirano and Snyder (2019) describe the appeal of primaries as “a straightforward reform that would limit the ability of political and economic elites to manipulate and profit from the nomination process” (page 15). They find evidence via difference-in-differences analysis that primaries promote the selection of more competent candidates in safe, open seats (Hirano and Snyder 2014, 2019). This idea is supported by theoretical work arguing that American-style primaries produce higher quality candidates, where quality is defined as campaign skill (Adams and Merrill 2008), but potentially at the expense of ideological extremity (Serra 2011). More broadly, other models of primaries suggest they may be useful in regulating internal competition to induce effort in developing policy platforms (Caillaud and Tirole 2002); and in promoting the provision of broad based public goods

⁴ Besley (2005) argues that much of the modern political economy literature “has not only neglected the problem of political selection, it has been positively hostile to the topic” (page 44).

⁵ Noting the caveat that higher illicit (as opposed to official) returns may have the opposite effect (Brollo et al 2013).

as opposed to individualistic private transfers (Ting et al 2018). In the developing world, Carey and Polga-Hecimovich (2006) and Ichino and Nathan (2013) investigate the effects of primaries on general election vote shares, but do not look directly at the question of candidate selection. We follow the pioneering approach of Wantchekon (2003) in working directly with political party leaders on a randomized controlled trial.

Our paper brings these two disparate literatures together by directly estimating the effects of party leader versus voter control on candidate selection, and associated impacts on representation, payments, voter learning and aspirant entry. We structure our analysis to parse key tradeoffs between a poorly informed citizenry and a potentially misaligned political elite. In so doing, we extend consideration to the role of primaries in alleviating information constraints, which are pervasive in poor countries where transport and communication costs are high. They are arguably important in rich countries as well, where American voters, for example, are criticized for being poorly informed about politics (Delli Carpini and Keeter 1996, Achen and Bartels 2016) and local party leaders and representatives hold inaccurate views of public opinion (Butler and Nickerson 2011, Broockman et al. forthcoming). While Adams and Merrill (2008) and Folke et al (2016) share our view of primaries as a way for party leaders to learn about aspirant popularity, we go one step further and show that they also deliver performance-relevant information from voters to party officials. To the best of our knowledge, this is the first analysis in either economics or political science to exploit experimental variation in how parties select candidates. Overall, our results suggest that the more democratic selection method creates value for voters, in that they are more likely to get their preferred candidates, who have stronger public goods records.

The rest of this paper is organized as follows. Section II discusses variation in candidate selection processes around the world and introduces a simple framework to structure the analysis. Section III details the experimental design and interventions. Section IV characterizes candidate selection in the status quo. Section V presents estimates of how the more democratic selection process affects key outcomes of interest. Section VI concludes with policy considerations.

II. Candidate Selection in Perspective

II.A. Empirical Variation

There is substantial variation across country and over time in how political parties select candidates. To provide a sense of the dispersion in the relative control of party officials versus

voters, consider first the case of France: central party leaders historically chose all candidates and allocated them across space to populate sub-national races (Valen et al 1988). Parties in the United Kingdom have traditionally used a more decentralized approach, where the Labour party for instance delegates candidate selection to constituency-level party members. Historically, relatively high barriers to membership—via financial dues and time requirements—have meant that this group is quite narrow: data from the 1980’s suggests that on average 40 Labour Party members chose the candidate on behalf of some 70,000 constituents (Bochel and Denver 1983).

The U.S. anchors the other end of the spectrum, where all states now use some form of direct primary. Yet the direct vote phenomenon is relatively new: most states adopted mandatory primary laws between 1900 and 1920, with additional uptake staggered over subsequent decades (Hirano and Snyder 2019, p. 23). At the Presidential level, the outcome of state-level primaries in determining each party’s candidate only began to bind after the contentious 1968 Democratic convention in Chicago. The extent of voter control remains in flux today: the 2016 Presidential race divided the two major parties over whether there is currently “too little” or “too much” democracy. On one side, members of the Democratic National Convention moved to increase voter control by circumscribing the role of so-called “superdelegates,” or party elites not beholden to vote the way the primaries went in their respective states (Levy 2018). Across the aisle, some Republicans troubled by the prospect of Donald Trump’s candidacy called for a return to the historically stronger role for party elites as gatekeepers of the nominating process and as a check on the excesses of “hyperdemocracy” (Sullivan 2016). Globally, the demand for direct vote primaries is on the rise: primaries are becoming popular in Latin America (Carey and Polga-Hecimovich 2006); and for the first time in French history, both major political parties held direct vote primaries to select their Presidential candidates in 2016 (Briançon 2016).

The consequences of these disparate design choices on the performance of the electoral system in delivering high quality, representative candidates and elected politicians are poorly researched. This study is designed to address this gap. We explore selection of Parliamentary candidates in Sierra Leone, which in the status quo shares features of the traditional British and French approaches and is similar to many selection processes across Sub-Saharan Africa. The experimental treatment we evaluate moves candidate selection in the direction of an American-style process, by increasing the amount of say that ordinary voters have in selecting candidates, without getting all the way to a binding direct vote primary that currently reigns in the U.S.

As background, Parliament in Sierra Leone consists of 132 single-member jurisdictions, won by plurality. The general election of interest was held in March 2018 and declared largely free and fair by domestic and international observers.⁶ In the status quo, candidate selection is guided by the country's Constitution,⁷ which specifies eligibility requirements for becoming an MP, and the parties' own constitutions and regulations, which outline their internal procedures. In principle, both major parties begin the process with constituency-level executives screening candidates. These officials make recommendations to district- or regional-level executives, who in turn make recommendations to national executives, who have the final say. For the Sierra Leone People's Party (SLPP), the first, most local, step in this chain involves between three and ten constituency-level officials, compared to an average of 24,000 registered voters per constituency. When multiple aspirants are under consideration, this group is meant to vote among themselves to determine which aspirant to recommend upward. The All People's Congress (APC) party takes a more centralized approach. Its constitution enshrines the right to "elect or select" all candidates, for all levels of office, and mandates that the party's National Advisory Committee approve all candidates. Neither party has a clear mechanism in place to capture the preferences of voters or rank-and-file party members. At the time of the experiment (late 2017), the APC was the ruling party, however this changed when the SLPP won the Presidency in the election studied.

II.B. Conceptual Framework

There is no model in the literature that speaks directly to how a move from party leader control towards a more democratic process will affect candidate selection and associated outcomes of representation and payments to the party. A simple framework is thus useful to define key concepts, illuminate tradeoffs between a poorly informed electorate and a potentially misaligned political elite, and frame the experiment we study with respect to information constraints.

Set up: Suppose that each jurisdiction has a single representative voter and a party official, either of whom could select one candidate from a finite pool of aspirants. Aspirants are heterogeneous in quality, which is a vector of universally valued traits (like ability, intelligence, integrity) and match-specific traits associated with the jurisdiction (fluency in local languages,

⁶ See for example the report of the European Union Election Observation Mission, available at: https://eeas.europa.eu/sites/eeas/files/eu_eom_sl_2018_final_report_4.pdf

⁷ See the Constitution of Sierra Leone, 1991, available at <http://www.sierra-leone.org/Laws/constitution1991.pdf>.

knowledge of local priorities and resources). All quality traits contribute positively to a single dimension of performance in office. To fix ideas, define performance as the local population's valuation of the bundle of public goods that the candidate will produce if elected, which is a function of two traits: competence, or the volume of goods produced from a fixed public budget; and alignment, which captures the idea that local voters value a new school more than a clinic if there are currently few schools and many clinics nearby, or if they prefer education over health. There is a third factor, party loyalty, which is orthogonal to performance but potentially correlated with quality. Aspirants make payments to the party to be considered. Their willingness to pay is increasing in the expected returns to candidacy, which is a function of the returns to holding office (salary and other opportunities for remuneration), their likelihood of being selected as the candidate, and their party's likelihood of winning the seat in the general election.

Information: As is standard in principal agent models (see for example Banks and Sundaram 1993, Fearon 1999), aspirants have private information about their type, meaning that both the voter and party official select under information constraints.⁸ We extend consideration to the intuitive case where the voter may be relatively better informed about local alignment, and the party official relatively better informed about competence.

Outcomes: We are interested in how allocating control to the voter versus the party official affects three outcomes: i) representation, defined as the likelihood that the selected candidate is the voter's first choice, where her choice is conditioned on her information set; ii) selection on quality, which is the expected performance (value of public goods produced) of the candidate selected; and iii) the financial contributions aspirants make to the party.

Preference divergence: A standard concern is that only the party official values party loyalty and maximizes a combination of loyalty and quality that selects a loyal aspirant over a more competent one. If the voter and party official's information sets are similar, then transferring control to the voter straightforwardly enhances representation and selection on quality. This was a key argument made by Progressive reformers in the U.S. All else equal, however, the extent to which preference divergence compromises selection depends on the correlation between loyalty and quality: if sufficiently negative, the effects could be pernicious; however if sufficiently positive, they could be of little consequence.

⁸ This paper focuses on unknown type and the selection channel, noting that our ongoing data collection regarding the performance of general election winners will enable future exploration of the moral hazard channel.

Screening: A countervailing concern is that the voter is at an absolute screening disadvantage compared to the party official. This could be the case if her information set is strictly worse, or if the trait she can observe has a lower marginal product with respect to performance. If so, then reallocating control from the official to the voter could increase representation (she gets her most preferred aspirant), but at the cost of selection on quality (if she had the party official's information on competence, she would have chosen a different aspirant to maximize performance).

Experimental treatment: Under this framework, the intervention studied can be interpreted as alleviating information constraints: i) the conventions reveal information about competence to the party official (and indirectly to the voter), enhancing both of their ability to select on it; and ii) the opinion polling delivers the voter's information about alignment to the party official, enhancing his ability to select on it. As the polling data is delivered with free disposal, the party official retains control. While we do not have sharp predictions to take to the data, it seems plausible that under a reasonable set of parameters this treatment would: i) increase representation, but not to 100 percent, as the party official will deviate from the voter's first choice either to increase performance or for a more loyal aspirant; ii) enhance selection on quality, via both the competence and alignment channels, and hence increase expected performance; and iii) increase (decrease) payments by higher (lower) quality aspirants, as these are the ones for whom the information treatments increase (decrease) their likelihood of being selected as the candidate.

Other considerations: This simple framework abstracts away from aspirant entry, which is a key driver in the models of Dal Bó and Finan (2018) and Hirano and Snyder (2019). As entry is largely shut down in our context, we focus on how information and preferences affect the choice of the voter versus the party official, conditional on the aspirant pool. There is also a vast literature in political science about the interaction between primaries and ideology. As the political parties in Sierra Leone are not strongly differentiated by ideology, it allows us to better isolate factors that contribute to selection on quality. It is worth noting that the lack of strong ideological labeling is not unique to Sierra Leone: Cruz and Keefer (2015) aggregate data on parties in more than 100 countries and find that about half are “non-programmatic,” which means they cannot be classified on a left, center, right scale or other metric of economic policy.

III. Experimental Design

The experimental intervention studied has two main components, which were implemented in

tandem at the constituency-party level: i) a town hall-style convention, broadcast over local radio; and ii) representative opinion polling of registered voters, aggregated into a one page report and shared with party officials. These components were randomly assigned across constituencies (e.g. races) for each participating party independently. Figure 1 presents an overview of the design.

The offer to participate in the initiative and associated research was managed by the Political Parties Registration Commission of Sierra Leone, which has the constitutional mandate to register, supervise, and monitor the conduct of political parties. This includes monitoring the accountability of parties to their members and to the broader electorate.⁹ The PPRC extended the offer to participate to all registered political parties, whose leadership decided whether or not to opt in. For the parties that expressed interest in participating, the PPRC asked national party leaders for a list of constituencies that they were willing to include in the initiative, with the requirement that more than one aspirant be under consideration for each race.¹⁰

Search for Common Ground, a non-profit civil society organization, contributed to the design and implementation of the support provided to parties via the initiative. SFCG works in 35 countries with a mission of transforming the way the world deals with conflict. It established a Sierra Leone office in 2000 under the name Talking Drum Studio. SFCG works extensively on programs to promote transparency, accountability and communication on the political process in Sierra Leone, providing neutral and reliable content through a network of 27 local radio and television broadcasters. It played a leading role in coordinating broader civil society efforts to support the 2018 (and previous) elections.¹¹ The parties, PPRC, SFCG and the research team worked together to design and implement the initiative.

III.A. Party Conventions

Constituency-level party conventions provide an opportunity for aspirants to present their qualifications and debate each other on policy issues in front of an audience of party officials, rank-and-file party members, and local residents. SFCG led support for the parties' conventions via three stages: public mobilization, town hall implementation, and radio dissemination.

⁹ For more information, see <https://www.pprcsierraleone.org/>.

¹⁰ Three minor, or “emerging,” parties expressed interest in participating, however did not have any races with more than one aspirant under consideration, so did not proceed to the implementation stage. One more expressed interest after the implementation window had closed.

¹¹ For more information, see <https://www.sfcg.org/sierra-leone>.

SFCG field teams launched the mobilization process in each constituency a few days before each convention, encouraging residents to attend the town hall or listen to the convention replayed over radio. These promotional efforts included jingles broadcast on local radio and activation of the “town crier” network, which are local people associated with the traditional chieftaincy system whose role it is to inform communities of events and news. The data collection team also provided 25 voters with information about the broadcasts in pre-convention surveys (randomly sampled from three voter registration centers, see description of the V1 survey below). While the conventions were open to all interested party officials, a core set of three standard constituency-level positions—the party’s constituency chair, secretary and treasurer—were explicitly encouraged to attend their respective convention.

A typical town hall convention began with a trained moderator introducing the aspirants to the audience, and then posing a series of policy questions, allowing each aspirant two minutes to respond to each one and alternating who spoke first across questions. Five standard questions were: i) explain who you are and what qualifies you to be a good MP; ii) what is your first priority for additional government funding; iii) how would you spend the constituency facilitation fund, a pot of public money given annually to each elected MP; iv) what makes you a good representative of local people, including how you would know what local people want and how you would represent their interests in Parliament; and v) what is your strategy for improving teaching and educational outcomes? Two further questions followed, tailored to the area, covering topics like how to deal with local power supply constraints and allocating mining royalties.

Shortly after the convention finished, SFCG delivered audio recordings of the event to independent local radio stations that re-broadcast the convention multiple times over subsequent days. In post-convention surveys in treatment areas, 26 percent of voters reported having attended or listened to a convention. These rates are higher for the quarter of surveyed voters who were notified in advance about the conventions (46 percent versus 21 percent).

III.B. Voter Reports

A few days after the conventions and associated radio broadcasts, the research team fielded an opinion poll of registered voters in the constituency, visiting voters in-person at their residence. We sampled voters directly from the official registry of voters maintained by the National Electoral Commission, which includes names, demographics and home address. To ensure representation,

we first randomly selected 10 voter registration centers per constituency. As constituencies on average contain 25 centers, this means that field teams visited 40 percent of all centers in a given constituency. To facilitate rapid field work, we limited consideration to voters whose community of residence matched the community of the registration center, which excludes the 24 percent of voters who live in other nearby communities. We then randomly selected ten voters per registration center, stratifying on age and gender, for a target of 100 respondents per constituency. Each target respondent was accompanied by two potential replacements from the same demographic bin. The respondent hit rate was high: on average, 94 voters were surveyed per constituency, where 67 percent of those polled were the target respondent, 20 percent were the first replacement, and 13 percent were the second replacement. These surveys are thus substantially more representative than the telephone polls commonly conducted in the U.S. where, for example, the Gallup poll currently has a survey response rate of 7 percent (Marken 2018).

The research team aggregated this opinion poll data, weighted by demographics, into one page voter reports that displayed the share of votes each of the party's aspirants in the constituency received among poll respondents (see example in Appendix Figure A1). The top of the report read as follows: "The first choice of voters in this constituency for the [*party*] MP symbol is: [*name of top ranked aspirant*] who has [*X*]% of the popular vote. This is based on polling results from a representative sample of [*N*] registered voters living near 10 different polling centers in this constituency." The report includes two bar charts showing the vote shares each aspirant received, first among all voters surveyed, and second for self-reported party supporters only. Due to strong geographic sorting by party loyalty, these two tabulations rarely identified a different frontrunner. For analysis, we focus on the former. The two parties diverged in how many copies of these reports they requested printed: the research team delivered over two hundred copies to the SLPP for distribution to all affected constituency-, district- and national-level executives; and delivered 20 reports to national executives of the APC. This request reinforced the traditional perception that the SLPP is more decentralized and the APC more centralized of the two main parties.

III.C. Treatment Assignment

Random assignment proceeded in stages and was conducted independently for the two political parties. Each party first selected 46 constituencies from the universe of 132 where they were interested in piloting the initiative. The research team then randomly assigned, via computer

program, half of each party's constituency list to treatment and half to control, stratifying by small geographic bins. This generates an experimental sample of 92 party-races (see Figure 1, noting that it masks locations to protect anonymity¹²). As the two parties occasionally picked the same constituency for inclusion, these 92 party-race observations cover 80 unique constituencies.

Panel C of Figure 1 details the implementation timeline. The conventions launched in mid-November 2017 and all were completed and voter reports delivered before the parties submitted their official list of candidates to the National Electoral Commission in early January 2018. Final outcomes of interest relate to which aspirants were registered with NEC to represent their party as candidates in the March 2018 Parliamentary elections.¹³

While the original implementation plan further included advance announcement of the list of constituencies assigned to the new selection method, this was only partially implemented (and is thus bracketed in Figure 1). Note first that the country's constitution stipulates that MP candidates in public employment (including teachers) must vacate their post a year before the election, so this initiative was launched too late to affect the entry decisions of those potential aspirants. Two months before the conventions began, the SLPP publicized the list of its 23 treatment constituencies—via public announcement and paper leaflets—during their national convention of delegates to nominate their Presidential candidate. Their promotional materials, which dub the initiative “Aspirant Voice and People’s Choice,” describe the two components in detail and characterize the program as a “pilot” designed to “strengthen the internal democracy of our party” (see flyer in Appendix Figure A2). The APC party, on the other hand, joined the initiative later and therefore did not announce the program or targeted constituencies at its own national convention. While the party did inform all national executives and a cross-section of district executives about the initiative at a subsequent meeting, this came too late in the process to affect entry decisions. We thus consider the entry channel largely shut down in this context (see Section V.D. for treatment effect estimates), however expect that it could be quite consequential in other settings.

Data collection was implemented in parallel for both treatment and control constituencies, and included pre- and post-convention surveys of voters (labeled V1 and V2, respectively),

¹² Panels A and B are illustrative only and show arbitrary assignments using the old 2007 constituency boundaries in order to protect the anonymity of all research participants. Boundaries were redrawn before the 2018 election.

¹³ Available here: <http://necsl.org/PDF/Media/List-of%20Parliamentary-Candidates%20-%202018%20Elections.pdf>

aspirants (A1 and A2), and party officials (P1 and P2). Data collection rolled out sequentially across stratification bins, so the pre- and post-survey timing proceeded in lockstep for both treated and control constituencies in a given geographic area. This is important as the information environment, and thus voter and party official opinions, were evolving over time in the control areas, as well as under the information-rich treatments. On average there were 11 days between the pre- and post-convention voter surveys. The only exception to the symmetry of data collection is that P2 was not collected in controls, as it was socially awkward for enumerators to ask party officials an identical set of questions a few days apart, when no observable event had occurred in the interim (P2 is for descriptive purposes only and no outcomes depend on P2 data).

The voter reports are compiled based on post-convention (V2) polling data, the sampling of which is described above. Respondents in the pre-convention (V1) survey are a subset of the voters targeted for V2: specifically, V1 covered 25 voters registered to 3 of the 10 sampled registration centers in V2. For party leader surveys (P1 and P2), enumerators surveyed the same three constituency-level positions (constituency chair, secretary and treasurer) that were targeted for encouragement to attend the convention. These respondents were replaced as necessary with holders from similar constituency positions (e.g. deputy constituency chair) or higher level party officials (e.g. district chair). Overall, 81 percent of those surveyed hold constituency-level positions, 6 percent hold district-level positions, and 13 percent hold other positions.

IV. Characterizing Status Quo Selection

We first use our data to document four facts about candidate selection in the status quo: there is positive self-selection into politics on education and wealth; party officials face nontrivial information constraints in eliciting voter preferences; most Parliamentary races are located inside regional partisan strongholds where competition in the general election is weak; and party leaders were most interested in experimenting with the new selection process in their own strongholds.

IV.A. Selection into Politics

Table 1 compares the characteristics of voters, party officials and aspirants, and shows that entry into politics is strongly associated with higher levels of education and wealth. Voters (in the V2 survey) have on average completed 5 years of education, 43 percent of them have no formal schooling, and only 4 percent have some university education. Aspirants, by contrast, have

completed over 15 years of education, none lack formal schooling, and 80 percent have been to university. Party officials (most of whom hold constituency-level positions) sit in between the two, with 12 years of education, 5 percent without any schooling, and 34% with university exposure. The picture looks similar for proxies of wealth. Voters on average own fewer than 3 assets from a list of 11 common household items (e.g. a mobile phone, generator, radio) and only 11 percent of them have a formal bank account. Aspirants, by contrast, on average own 9.5 of these assets and nearly all of them (98 percent) have bank accounts. Party officials again fall in between: they on average own 6 assets and a majority (77 percent) has a bank account. In terms of demographics, politics is a male-dominated activity: 80 percent of party officials and 90 percent of aspirants are male, compared to 47 percent of registered voters. Politicians are roughly ten years older (at 46 for party officials and 48 for aspirants) than voters (at 37 years).

Such pronounced selection on education suggests that party officials might provide a screening service for voters in assessing aspirants on their technical merits. As a concrete example, the first formal step for potential aspirants is to file a standardized application form with the party, which covers items like eligibility requirements and their standing in the party. Most voters will find it difficult to even read the aspirants' applications, or review their curriculum vitae, while party officials will both be able to read the documents and use their knowledge of government to assess which qualifications are important for carrying out the duties of Parliament. Regarding the last point, 33 percent of voters (in the V1 survey) could not name a single MP job responsibility.

Note that incumbency is much less entrenched in Sierra Leone than in the U.S. There were 25 incumbents seeking re-nomination in our sample, which covers 27 percent of the races (reassuringly incumbency is balanced across party and treatment assignment). Roughly half of them secured re-nomination from their party. Compare this to the 2018 U.S. House elections, where 88 percent of incumbents sought re-nomination and only 1 percent lost their primary.

IV.B. Information Constraints

Sierra Leone is a very poor country with high transport and communication costs, which create pervasive information constraints throughout the electoral process. The logistical challenges and costs associated with screening aspirants in 132 separate constituency-level pools, which occurs

over a matter of weeks, are thus substantial.¹⁴ Party officials may as a result be poorly informed about the characteristics of available aspirants, an information constraint that could be alleviated by the town hall conventions. Perhaps more importantly, there is no large scale, cost-effective polling technology accessible in this market, implying that party officials are particularly constrained in attempting to elicit and aggregate local voter preferences over these 132 pools of aspirants. The polling intervention immediately alleviates this second constraint.

To gauge how important the latter constraint is in practice, before the conventions the research team asked party officials (via the P1 survey in both treated and control races) two questions about the local pool of aspirants under consideration by their party: i) “if the choice to award the symbol was up to you today, who would be your first choice?” and ii) “if the registered voters in this constituency voted directly today for the symbol, who do you think would get the most votes?” Answers to these questions reveal that party officials are imperfectly informed about voter preferences over aspirants. Specifically, 90 percent of party officials indicated that local voters shared their first preference, e.g. that the aspirant the leader himself preferred would win a local primary in that constituency. This is incorrect: only 56 percent of presumed shared preferences were in fact a match with the polling data on voter preferences (from the V2 polling data). Overall, party officials correctly guessed who would win a local primary only 53 percent of the time. This disconnect is severe: in 31 percent of races, no party leader surveyed correctly guessed which aspirant was the most popular with local voters. Reflecting back on the conceptual framework, these data suggest that the opinion polling component of the intervention could alleviate an information constraint in a way that facilitates representation.

IV.C. Partisan Strongholds

Political allegiances in Sierra Leone are rooted in long standing ties between different ethnic groups and the two major parties: the APC is historically associated with ethnic groups in the north of the country, including the Temne, and the SLPP is historically tied to groups in the south, including the Mende (Kandeh 1992). While national politics is quite competitive—as these two respective groupings are comparable in overall population size—geographic sorting means that

¹⁴ The constitution does not specify a timeframe for candidate selection. For this election, the sitting Parliament adjourned 3 months prior to the election and parties were required to submit their list of candidates to NEC 2 months prior to the election. Most of the candidate screening and selection activity occurred within this one month window.

most sub-national jurisdictions are located inside partisan strongholds, where the locally dominant party's candidate is near guaranteed to win the general election. As a benchmark, the average constituency in Sierra Leone has the same partisan leaning as the 5th Congressional district of Massachusetts, which covers suburban Boston, the 18th Congressional district in California, which contains Palo Alto, or the 4th district in Texas, covering counties northeast of Dallas.¹⁵

Figure 2 displays these allegiances graphically. For each constituency, we compute the difference in population shares of ethnic groups historically associated with the APC minus those associated with the SLPP. The darkest red shading indicates a constituency populated almost entirely by APC-affiliated groups (the color choice reflects the party's symbol of a rising sun); and the darkest green indicates a constituency populated wholly by SLPP-affiliated groups (that party's symbol is a palm tree). Using data from the 2007 election, we regress the constituency-level difference in vote shares for the APC minus the SLPP Parliamentary candidate on the measure of ethnic population differences displayed in Figure 2. This yields an R^2 of 0.92, implying that differences in ethnic population shares alone explain 92 percent of the variation in Parliamentary vote shares for that election. The estimated coefficient on the ethnicity-based measure is, as expected, positive, large in magnitude (0.76, noting that perfect positive correlation would generate a coefficient of 1.0), and precisely estimated (standard error of 0.02).

Thus for the majority of subnational races, the locally dominant party's candidate is delivered on the strength of ethnicity-based ties by a large margin. This underscores the importance of internal party selection, as the process the party uses to choose a candidate effectively determines the identity of the elected MP. Note that strongholds are not unique to Sierra Leone, nor a curiosity of the developing world: Hirano and Snyder (2014) calculate that a minority (44 percent) of U.S. House of Representatives races from 1952 to 2010 were decided by fewer than 15 percentage points, which is a fairly lax standard for competitiveness.¹⁶

IV.D. Where Parties Chose to Experiment

Budget constraints on how many selection interventions could be supported presented parties with a decision problem of choosing 46 races from the population of 132 Parliamentary constituencies

¹⁵ Based on 2018 House Elections in the U.S. and the 2007 Parliamentary Elections in Sierra Leone (pooling votes for the SLPP splinter party, the PMDC, with the SLPP vote).

¹⁶ Caste loyalties create strongholds in India, where the literature is mixed as to whether this facilitates (Munshi and Rosenzweig 2016) or hinders (Banerjee and Pande 2009) selection on quality for the locally dominant caste.

nationwide for inclusion in the research sample and thus the lottery that assigned the initiative. Panel A of Table 2 summarizes their choices with respect to how competitive the general election race was likely to be. The panel doubles over the constituency map to accommodate both parties simultaneously, and classifies each constituency as either a likely safe seat for the party in question, a swing race, or a weak seat (e.g. a safe seat for the rival party). Note that since the constituency boundaries were redrawn for this election, we classified competitiveness at the district level (the next higher administrative unit).

Nationwide, column 2 shows that 36 percent of races are expected to be safe for a given party, 28 percent are swing, and 36 percent are weak seats. Column 4 shows that parties disproportionately selected their own stronghold races for inclusion in the initiative: safe seats compose 52 percent of the experimental sample, which reflects statistically distinct positive selection. Parties demonstrate neutral selection for swing seats: the proportion in the experimental sample (30 percent) is not statistically distinct from the national proportion. This leaves strong negative selection out of weak seats, which constitute only 17 percent of the experimental sample, and suggests that parties did not see much value in experimenting with selection where they were very likely to lose the general election. In thinking about the experimental results, this first stage of selective inclusion is important to keep in mind: results are representative of races in which the parties were willing to experiment, where the modal race is located in the target party’s stronghold. It further suggests that parties did not view the initiative primarily as a way to garner general election votes, as they were near guaranteed to win these races.

V. Treatment Effect Estimates

We organize our experimental results with reference to three main outcomes of interest: representation, candidate selection, and payments. We also present results on voter learning from the conventions and radio broadcasts and on aspirant entry.

V.A. Representation

Our first experimental estimates capture how the new model of candidate selection affects representation, using the following model:

$$Y_{ipc} = \beta + \beta_T T_{pc} + \tau_{pc} + \varepsilon_{icp} \quad (1)$$

where outcome Y is an indicator variable equal to one if candidate i that represented party p in the

general election for constituency c ranked first among voters in the primary opinion polls; treatment indicator T signals assignment to the more democratic selection model, which recall was assigned at the party-constituency level; τ are fixed effects for 26 party-region strata used in the random assignments; and ε is an idiosyncratic error. This specification tests whether treatment increases the likelihood that party officials select the aspirant who local voters want to represent them. In the data, outcome Y is coded to one if the candidate listed in NEC official registration data is the same aspirant who ranked first with voters in the V2 polls. We estimate intention-to-treat effects, where the parties complied with treatment in 43 of 46 races (93.5 percent) and there was no non-compliance in controls (generating a scaling factor of 1.07 for treatment-on-the-treated estimates). In our view, this high rate of compliance is quite extraordinary and attests to the willingness of the political parties to experiment at the frontier of democratic practice.

In the raw data, the two political parties respectively selected the locally preferred aspirant in 30 and 43 percent of control races. This rate increased to 61 percent for both parties in treatment races. Panel A of Table 3 presents the regression analogue, where the control mean in the pooled sample is 37 percent and the estimated coefficient on treatment is 23.9 (standard error 10.5). The estimated treatment effect corresponds to a 65 percent increase in representation (as we have defined it), which is materially substantive and statistically distinct from zero at the 95 percent confidence level. To put this in perspective, it suggests that parties responded to the information provided by picking a different candidate than they would have in the status quo for 11 races (e.g. 0.239×46 treatment group races). Given that 52 percent of these races were in the parties' respective strongholds, the party thereby likely changed the identity of 6 elected MPs. Column 2 finds little evidence for heterogeneous treatment effects by party: the coefficient on the interaction between treatment and ruling party (the APC at the time) is not statistically different from zero. As we find little evidence for heterogeneous response by party for any of the outcomes of interest, we focus on pooled estimates in all following tables.

Linking back to earlier evidence on where the parties chose to experiment, Table 2 Panel B assesses heterogeneity in responsiveness to treatment by the likely level of competition in the general election. For safe, swing and weak seats separately, columns 1 and 2 report the baseline rate of selecting the voters' preferred aspirant for control and treatment races, respectively; column 3 reports the difference in these rates; and column 4 shows the associated p -values from t -tests of the equality of rates across treatment assignment. It is evident that the treatment effect on

representation is driven by responsiveness in safe and weak seats, where representation increases by 29 and 50 percentage points, respectively. There is no apparent response to treatment for swing seats, where the likelihood of selecting the local voters' preferred aspirant is 50 percent in both treated and control races. In the regression analogue, the treatment effect in non-competitive general election races (safe pooled with weak seats) is a highly significant 34.38 (standard error 11.67). The coefficient on the interaction between treatment and swing seat is an equally sized negative term (-34.38 with standard error 22.74), which falls just below the 90 percent confidence level (p -value of 0.13). In interpreting these results, note that the 50 percent baseline rate (e.g. for control races) in swing seats is substantially higher than that in safe (29 percent) and weak seat (38 percent) races. This pattern is consistent with the parties already investing more resources in determining who is locally popular for more competitive general election races.

One immediate follow-on question is whether this representation effect in the candidate selection stage has downstream effects on the party's general election vote share. Estimates in Panel B of Table 4 find little evidence that it does: the estimated coefficient of the candidate selection treatment on the general election vote share of the targeted party is small in magnitude and imprecisely estimated (-0.48 with standard error 2.96). Column 2 breaks this out by the level of competition, where we see that moving from a swing to safe seat increases a party's general election vote share by 28 percentage points. Moving from a swing to weak seat decreases this expected share by 26 percentage points. None of the estimated coefficients on treatment status and its interactions with competitiveness are statistically distinct from zero. This null result makes sense in light of the accumulated estimates thus far: Figure 2 reveals that the general election is strongly determined by ethnicity-party ties, limiting the scope for the candidate selection treatment to affect cross-party vote choices in the general election, at least for partisan strongholds. While there may be more scope for a downstream effect in swing areas, recall that swing seats comprise only 30 percent of the experimental sample, and more importantly, recall from Table 2 that there is no "first stage" of the experiment in swing races, as parties were already substantially more likely to pick the local favorite in the status quo and this does not vary by treatment status.

This null result aligns with historical evidence from the U.S. While Ware (2002) suggests that incumbent party elites expected an electoral benefit from introducing primaries in the U.S., Hirano and Snyder (2019) find no evidence that primaries affected the general election vote share of the advantaged party (page 37). It runs counter to more recent evidence from elsewhere, which

suggests that primaries can boost general elections vote shares, at least for “underdog” parties in Latin American Presidential races (Carey and Polga-Hecimovich 2006) and opposition parties in legislative elections in Ghana (Ichino and Nathan 2013).

V.B. Candidate Selection

We approach the question of how treatment might affect which types of candidates are ultimately selected to represent the party in the general election via two steps. We first investigate the degree of potential conflict in preferences between voters and party officials by comparing which characteristics of aspirants appear to make them popular with voters, and how these compare to the characteristics that make them popular with party officials. We then estimate treatment effects on the characteristics of the candidates who were selected to compete in the general election, to see if the more democratic process selected candidates of different type. As this problem has elements of both prediction (what traits predict preferences?) and inference (does the new process affect selection on these traits?), we use a combination of machine learning and traditional econometric methods (see Athey and Imbens 2019 for discussion).

Analysis leverages rich data on aspirant characteristics collected during in-person interviews (via the A1 and A2 surveys). We organized survey questions into eight categories: i) professional qualifications, including items like education, incumbency and previous elected office experience; ii) wealth, including assets and reported income; iii) economic development record, including the number and value of local public goods the aspirant was involved in providing in the constituency over the previous five years; iv) cognitive ability, based on a series of election-oriented questions that involve numeric computations; v) party loyalty, which includes the history of membership, leadership positions and service to the party; vi) public service motivation, using a series of questions adapted from Perry (1996); vii) local networks, including membership in constituency social and occupational groups (like saving clubs); and viii) campaign effort (like community visits) and expenditure during the primary stage.

We complement this survey data with one measure of directly observed behavior designed to capture conscientiousness and attention to detail. Since field enumerators requested that aspirants meet them in the constituency headquarter town to conduct the interviews, they followed local practice and reimbursed aspirants a set fee to cover their travel expenses. After the survey, enumerators handed aspirants an envelope explaining that they were giving them 150,000 Leones

(approximately US\$20) to cover their travel costs and asked them to verify that the money was correct. Inside each envelope were eighteen, not fifteen, 10,000 Leone notes. The measure of conscientiousness is whether aspirants detected and returned any of the extra three bills. As nearly all who gave back any money returned all three notes, we focus on the binary measure of whether any money was returned. Overall, 46 percent of aspirants returned some money.

The first prediction step identifies which of the many aspirant characteristics collected appear to be valued by voters, and whether these are the same traits that party officials value. To do so, we estimate variants of the linear model:

$$V_{ipc}^s = \alpha + \sum_{k=1}^K x_{ipc}^k \beta^{sk} + v_{ipc} \quad (2)$$

where V_{ipc}^s is the vote share aspirant i who is vying to represent party p in constituency c received in among selectors s , where $s \in \{v, l\}$ denotes voters and party leaders, respectively; x^k is one of K aspirant traits collected in the data (like education or wealth); and v is an idiosyncratic error term. To identify which traits predict preferences, we test null hypotheses of the form $\hat{\beta}^k = 0$, and to assess whether voters and party officials have common preferences over traits, we test $\hat{\beta}^{vk} = \hat{\beta}^{lk}$. Since the voter data was collected after the conventions and radio broadcasts (in V2), we estimate a second version of (2) that includes interaction terms between treatment assignment and the vector of traits to capture whether voter learning about aspirant characteristics affects their rankings. For completeness, we present the same interaction model for party officials, noting that this data is pre-convention (collected via P1) so estimates for these terms should be null.

The second step tests whether varying how the primary process is run has a causal impact on the types of candidates thereby selected. We estimate a series of regressions of the form:

$$x_{pc}^k = \gamma^k + \gamma_T^k T_{pc} + \tau_{pc} + \eta_{pc} \quad (3)$$

where x^k is a characteristic of the candidate selected to compete in the general election on behalf of party p in constituency c ; T_{pc} and τ_{pc} are the treatment assignment and randomization strata as defined for Equation (1); and η_{pc} is an error term.

As we collected data on a large number of characteristics relative to a modestly sized sample of aspirants ($N = 370$), candidate selection is a high dimensional problem. This generates two interrelated challenges. For the first prediction step, we expect that many of the measures we collected are in fact irrelevant to voters and party officials, however we do not know *ex ante* which ones these are. This immediately suggests that regularization may be useful to avoid overfitting

the model and draw meaningful conclusions about preferences (Belloni et al 2014). For the second inferential step, the existence of many potential traits that might be affected by the candidate selection intervention creates a temptation to “cherry pick” or “*p*-hack” those associated with statistically significant estimates (Simmons et al 2011, Casey et al 2012, Brodeur et al 2016). Here standard approaches to adjust for multiple inference over all traits (Anderson 2008) will quickly erode statistical power, as we have even fewer observations of selected candidates ($N=92$) in the second stage. This suggests a combination of approaches, which is what we use, and triangulate results across them, gaining confidence if different methods produce similar estimates.

We first implement traditional approaches to reducing the number of statistical tests without culling any traits. We roll up the K characteristics into standardized indices for each of the eight survey areas outlined above (following Kling, Liebman and Katz 2007) and enter the indices, along with our single behavioral measure, into the unified regression of Equation (2). We carry forward all indices and the behavioral measure into step two, and adjust standard errors on treatment effect estimates to control the false discovery rate (FDR) across the nine regressions. Second, we complement this with regularized regression methods to select a subset of traits with the greatest predictive power in step one (Zou and Hastie 2005), and carry only those traits forward to step two, where we estimate treatment effects and again implement FDR adjustments. This affords flexibility in searching over all K traits (instead of eight indices) and retaining only those found to be relevant. To stabilize our estimates, we run the regularization technique 200 times for voters and party officials, respectively, keeping track of the total number and list of specific traits selected in each iteration.

Table 4 presents the index-level results on which aspirant traits predict voter and party leader preferences. For voters, the two strongest drivers are the aspirant’s record of having provided local public goods and other economic development projects in the constituency, and their conscientiousness with respect to returning the extra transport allowance. The estimated coefficients are positive and statistically significant in both column 1, which pools all voters together, and column 2, which splits voters out by treatment assignment. In terms of magnitude, the coefficient on development record of 0.04 (standard error 0.01) implies that a one standard deviation unit increase in public goods provision is associated with a 4 percentage point increase in support among voters in the opinion polls. For conscientiousness, the estimate suggests that returning the extra money is associated with a six percentage point increase in popularity among

voters. There is also somewhat weaker evidence that voters prefer aspirants with higher levels of wealth, public service motivation and connectivity to local networks, as the associated coefficients are positive and at least marginally statistically significant.

Similar characteristics appear to also drive party leader preferences over aspirants. Coefficients in column 3 are again positive and at least marginally statistically significant for an aspirant's local economic development record and their conscientiousness with the transport allowance, as well as for wealth. The one area where there is evidence of a potential divergence in preferences regards professional qualifications: aspirant qualifications positively and significantly predict party official appraisals, but do not register among voters. The associated p -value in column 5, which tests for equivalence of the coefficients estimated for voters and party officials in control races (columns 2 versus 4), rejects equality at the 98 percent confidence level. This is consistent with the human capital advantage party officials have compared to voters in screening on technical merits, and could reflect differences in their respective information sets, a question we return to in Section V.D. The lower half of the table reassuringly finds little evidence for differences in party officials' preferences in treated versus control races, as recall this is estimated on pre-treatment data (from the P1 survey).

Table 5 turns to causal inference, testing whether the more democratic process selects candidates that differ on observables compared to those chosen via status quo methods. Building on what Table 4 reveals about preferences, it is noteworthy that the candidates selected by the more democratic process on average have stronger records of having previously provided public goods in the constituency. The estimated treatment effect for this economic development index is 0.29 standard deviation units (standard error 0.13), which is a materially large and highly significant effect on a naïve, or per comparison, basis (p -value = 0.03). There is also a positive treatment effect estimate for conscientiousness, of 0.15 standard deviation units, but it does not reach significance at conventional levels (p -value = 0.16). There is further a positive and significant effect on selecting aspirants with stronger local networks (including being born in the constituency and a member of local groups), which is unexpected since this did not rank strongly as a predictor of preferences in Table 4. Implementing FDR adjustments over all nine regressions sends these estimates below standard significance levels: the corresponding q -values for the two most precise estimates fall to 0.22. This is perhaps not surprising given the strains on statistical power in stage two, where there are only 92 observations of selected candidates. While not a main

focus, note further that we do not find treatment effects on demographics: 88 percent of selected candidates are male, average age 46.5, and neither varies significantly with treatment assignment.

Reassuringly, the regularization methods produce a similar pattern of results. Appendix Table A2 shows that the specific measure of how many development projects the aspirant provided in the constituency (an item in the economic development index) stands out as the most consistent predictor of both voter and party leader preferences in step 1. This trait was selected in 192 iterations for party officials and all 200 iterations for voters. Incumbency (an item in the professional qualifications index) also registers frequently for both voters and party leaders, and is the only additional trait that consistently survives penalization for party officials). For voters, notably conscientiousness with respect to returning the extra transport allowance is one of four additional traits consistently identified as relevant across iterations (the other three are aspirant education and two measures of party loyalty). As such, allowing the data to choose which specific measures have predictive power identifies items in the same three broad areas (development record, qualifications and conscientiousness) that the index level approach deemed important.

Table 6 presents the post-regularization prediction results for the union of these six traits. The number of development projects and incumbency are the strongest predictors of voter and party official preferences, entering positively at above 95 percent confidence. The other measures are less strongly predictive: education registers positively for officials but not voters, and two measures of party loyalty (willingness to spend money on the party's versus the aspirant's own campaign and number of relatives in party leadership) register at least weakly positively for both.

Returning to causal inference, Panel B of Table 5 estimates treatment effects on these six specific traits. The only estimate that is statistically distinct from zero is the large positive effect on the number of development projects (0.50 with standard error 0.21). In terms of magnitude, it implies that candidates selected via the more democratic primary process on average had been involved in providing half an additional local public good or other development project. Given the control group mean of 2.1 projects, this effect constitutes a 24 percent increase in such provision. As seen already, the estimate for conscientiousness is positive in sign but not quite significant at conventional levels (p -value = 0.16). In column 4, the FDR adjustments reduce the number of development projects result to the 86 percent confidence level.

As the estimated treatment effect on development record is a key result, it is worth providing additional context from the data. The number of projects provided over the past three

years ranges from zero to three, and 81 percent of aspirants have been involved in at least one project. The modal project is construction of small scale public infrastructure (roads, bridges, community centers), followed by support to education (classroom construction and rehabilitation), agriculture (provision of farming inputs, like seeds, tools and tractors), and healthcare (clinic construction and rehabilitation). While incumbents have provided more projects than others (a mean of 2.5 versus 1.9 projects, a highly statistically significant difference), incumbency does not drive these results. Specifically, all estimates indicating that an aspirant's economic development record (via the index or the specific number of projects variable) predicts voter preferences in tables 4 and 6, do so while controlling for incumbency (which voters also value). Moreover, there is evidence for a positive treatment effect on selecting high development record candidates, accompanied by a null result for (re-)selecting incumbents, in Table 5.

Overall, the headline results that come through this triangulation approach are first, that both voters and party officials value aspirants who have a demonstrated record of providing public goods in the constituency. And second, the more democratic selection methods increase the likelihood that aspirants who are strong in this regard are selected to then compete in the general election. Results for conscientiousness are directionally similar, yet estimated with lower precision. To the extent that past public goods provision predicts future provision, which is a question that ongoing data collection seeks to address, these are cautiously optimistic results.

V.C. Financial Payments

Recall that aspirants in control group races contribute more than two thousand dollars on average to be considered for candidacy. One way to interpret these payments is with reference to the expected return to candidacy, which is increasing in three multiplicative terms: the returns to office, the probability the party wins the seat in the general election, and the likelihood of being selected as candidate. This section explores how aspirant willingness to pay responds to each of these factors in the data, and to the experimental treatment.

Regarding the first term, returns to office include the salary of an MP, scope for promotion (e.g. to a Ministerial post), and other remunerative opportunities that positions in government afford access to. It is likely that there are also returns to candidacy even for those who lose the general election, particularly for members of the ruling party, as candidacy opens up avenues to public sector employment and positions within the party organization. In light of this, both major

parties were able to recruit aspirants and collect financial contributions in all 132 Parliamentary races nationwide.

To measure these contributions, enumerators asked aspirants shortly after candidate selections were formally announced (in January 2018) how much they had paid to the party in official fees and other payments. Note that these are distinct from the aspirant's own campaign expenses, which were recorded separately. The specific question asked, "How much contribution in total have you given to the party leaders for acquiring the party symbol from the start of your campaign? This includes registration/application fees, tips, small or big token to all the party leaders, kola [gratitude] money and transportation reimbursement tips." Enumerators further asked unsuccessful aspirants how much they thought the selected candidate in their pool had paid (see Appendix B for survey questions and prompts).

Regarding baseline rates of willingness to pay, self-reports from control group races suggest that nearly all aspirants made payments: 89.7 percent of aspirants report non-zero contributions, with an unadjusted mean converted to US dollars of \$2,406. To put this amount in perspective, it is equivalent to 1.3 months official salary of an elected MP. This looks modest compared the 17 months of salary that public healthcare workers paid to secure a promotion in another low income country (Weaver 2018), but lies well out of reach for ordinary Sierra Leoneans, where for instance the monthly minimum wage is \$71. Selected candidates self-reported the highest contribution in their pool in 35 percent of control races. This rate increases to 54 percent when substituting out low candidate self-reports with the average of other-reports on the candidate by unsuccessful aspirants.

As to the second term, the likelihood that the party wins the seat in the general election is increasing in the ethnic-party allegiance of voters in the constituency. This implies that willingness to pay should be higher in stronghold races, which is exactly what we see in the control group data: panel C of Table 2 shows that mean payments increase from \$1,089 in weak seats, to \$1,482 in swing seats, to \$2,983 in safe seats. Appendix C discusses how this increase in expected returns in turn affects aspirant entry in the status quo, and finds evidence that there are more aspirants in the pool, who are of higher average quality, in stronghold races, which is consistent with findings from Brazil (Ferraz and Finan 2011), Italy (Gagliarducci and Nannicini 2013), Sweden (Dal Bó et al 2017) and the U.S. (Hirano and Snyder 2019).

Regarding the third factor, willingness to pay should be higher for aspirants with a better

chance of being selected as the candidate, which is an object that the experimental treatment might affect. To see this, recall the assumption in the conceptual framework that aspirants have private information about their type, which the conventions and opinion polling deliver (at least in part) to party officials. For high types, this implies that their probability of being selected increases, and thus their willingness to pay should as well.¹⁷ For low types, the opposite should hold. A simple way to capture this in the data is to test whether selected candidates paid differentially more than unsuccessful aspirants in treated versus control races.

As a start, summary statistics in Table 2 are broadly consistent with this idea: while average payments across the whole aspirant pool are slightly higher in treated races, none of these differences are statistically significant. Narrowing consideration to selected candidates, however, reveals an increase in payments associated with treatment in safe and swing seats (Table 7 presents regression analogues, winsorizing payments at the 95th percentile). The increase is most pronounced in safe seats, where the unadjusted mean payment for selected candidates in control races is \$1,390 compared to \$5,771 in treated races. Note that mechanically, a given change in the probability of being selected will induce a larger nominal increase in payments where the probability of winning the general election is higher, implying that any effects of treatment on payments should be larger in safe versus weak seats.

Table 7 column 3 presents the more direct test of the hypothesis. It pools all aspirant data together and regresses contributions on whether the aspirant was selected to be the candidate, treatment assignment, and the interaction of the two terms. The positive and significant coefficient on the interaction between being selected and treatment implies that the divide between how much successful and unsuccessful aspirants paid is wider in treatment races. This divide is even more pronounced in column 4, which restricts attention to stronghold races. Overall, these results are consistent with aspirant willingness to pay being tied to the expected returns to candidacy, and thus for those whose expected returns increase with treatment—because they are more likely to be selected under the new process—their contributions increase as well.

V.D. Additional Treatment Effects

This final section presents treatment effect estimates on voter learning and aspirant entry to provide

¹⁷ This is somewhat analogous to tournament-based pay, where a worker's investment in productivity is increasing in the likelihood that productivity increases her chance of winning the tournament (Lazear and Rosen 1981).

additional insight into which channels are likely at work (and not) in the experiment.

A central concern about primaries is that if voters are poorly informed about the professional qualifications of aspirants, then giving them control over candidate selection increases representation at the cost of screening on quality. Publicizing the town hall conventions over radio could mitigate this concern if it improves voter knowledge of aspirants. To test this, Appendix Table A2 presents treatment effect estimates on what voters know about aspirants, using the post-convention V2 data. It shows that voters in treated races are more likely to accurately identify which aspirant in the local pool has the most education (by fifteen percentage points), the one with most public office experience (by 16 percentage points), and the one with the strongest record of previously providing local public goods (by 10 percentage points). Given baseline rates for correct answers between 30 and 36 percent, these estimates suggest that voter knowledge of aspirant qualifications increased by 42 percent on average compared to control races.

Referring back to Table 4, there is suggestive evidence that this knowledge gain increased the weight voters place on qualifications in ranking aspirants. In the bottom half of column 2, the interaction between treatment and an aspirant's professional qualification index is positive and marginally statistically significant (0.10 with standard error 0.06). This indicates that more qualified aspirants were somewhat more likely to be chosen by voters as their first preference for candidacy in treated versus control races. More broadly, the learning results are consistent with earlier work showing that publicizing policy-oriented debates between candidates in the general election is an effective way to build voter knowledge in this context (Bidwell et al 2019), and resonates with evidence from Uganda that debates build voter knowledge in both the primary and general election stages (Platas and Raffler forthcoming). It also aligns with survey evidence from the U.S. showing that voters learn about candidate policy positions and move into stronger ideological congruence as the primary campaign progresses (Hirano et al 2014).

Regarding aspirant entry, recall that only one of the two parties, the SLPP, announced in advance which constituencies would participate in the new selection process, which could have altered the entry decision of potential aspirants. Appendix Table A3 presents treatment effect estimates for the total number of aspirants considered by each party per constituency. The first two columns use administrative data from the Secretaries General of the parties. Estimates in column 1 show that just under three aspirants on average competed for the SLPP's symbol in control races. The estimated treatment effect is 1.00 (standard error 0.59), which is a marginally

significant increase in entry to just under four aspirants per race. Column 2 shows that the APC on average considered 4 aspirants per constituency, and the estimated treatment effect is small in magnitude and not statistically distinct from zero (0.20 with standard error 0.59), as one would expect in the absence of advance notification. This data reflect the number of aspirants who were under official consideration at the time candidate selections were made, at the end of the process. We can compare this to the number of aspirants who were surveyed earlier as part of the research, which captures those who were under consideration by the party at the time of the conventions, so could be a larger number if some aspirants dropped out or were disqualified in the interim. Estimates for the APC are nearly identical. For the SLPP, the mean number of aspirants increases to 4 per race and there is no evidence of a treatment effect. Putting the survey and administrative estimates together suggests that the initiative may have helped some SLPP aspirants stay under consideration longer.

In light of this, our results are best considered as partial equilibrium effects holding the pool of aspirants largely fixed. If the new selection process were to become internalized by parties and scaled up, we anticipate that the entry channel could become important in general equilibrium.¹⁸ We leave this question for future research.

VI. Conclusion

Elections are large public investments: the United Nations Development Programme, the largest international donor in the electoral space, expended more than three billion dollars to support elections in poor countries over the past fifteen years (UNDP 2019). Sierra Leone was the ninth largest recipient of these funds, despite having a population of only three million registered voters. A key point of this paper is that the efficacy of such investments in delivering competent elected officials depends critically on how candidates are selected. If party officials select candidates with little input from voters, citizens may well be perfectly enfranchised on paper—entitled to participate in free and fair general elections—but wholly irrelevant in practice, at least for partisan strongholds. This problem is largely absent from debates about foreign aid, for example as late as 2004 there was no explicit reference to political parties anywhere in the UNDP’s own multi-year funding framework (UNDP 2006). This neglect arises in part from concerns about impartiality

¹⁸ See for example Gulzar and Khan (2018) for experimental evidence from Pakistan on how ordinary citizens can be mobilized to run for local public office.

that lead donors to shy away from direct engagement with political parties.

The two major political parties in Sierra Leone demonstrate that this need not remain the case. The experimental platform was offered in a party-neutral, equitable fashion, and relied on party leadership to determine whether it was in their party's interest to opt in. In a revealed preference argument, the fact that both major parties participated and selected their stronghold races for inclusion suggests they saw value in the initiative. The fact that they (for the most part) complied with the treatment assignment attests to their willingness to experiment at the frontier of democratic practice. Party officials responded to the information relayed by the party conventions and opinion polls by selecting different types of candidates, and the fact that these candidates had stronger records of local public goods provision is a cautiously optimistic result.

Analysis in this paper has focused on the question of selection. Ongoing data collection turns to the question of accountability. Through a combination of Parliamentary administrative records and field audits of public spending by elected MPs, we will soon have evidence on the question of whether these candidates who were selected via more democratic primary processes behave any more or less accountably once they get into public office.

References

- Achen, Christopher and Larry Bartels. *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton: Princeton University Press, 2016.
- Adams J, Merrill III S. "Candidate and Party Strategies in Two-Stage Elections Beginning with a Primary." *American Journal of Political Science* 52:2(2008), 344-359.
- Anderson, Michael L., "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedaian, Perry Preschool, and Early Training Projects," *Journal of the American Statistical Association*, 103:484(2008),1481-1495.
- Athey, Susan, and Guido W. Imbens. "Machine Learning Methods That Economists Should Know About," *Annual Review of Economics* 11 (2019), 685-725.
- Banerjee, Abhijit V., and Rohini Pande. "Parochial Politics: Ethnic Preferences and Politician Corruption," unpublished manuscript, Harvard Kennedy School, 2009.
- Banks, Jeffrey S. and Rangarajan K. Sundaram, "Moral Hazard and Adverse Selection in a Model of Repeated Elections" in *Political Economy: Institutions, Information Competition, and Representation* (W. Barnett, et al., Eds.), Cambridge Univ. Press, Cambridge, UK, 1993.
- Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen. 2014. "High-Dimensional Methods and Inference on Structural and Treatment Effects." *Journal of Economic Perspectives*, 28 (2): 29-50.
- Benjamini, Y., Krieger, A., and Yekutieli, D., "Adaptive Linear Step-Up Procedures That Control the False Discovery Rate," *Biometrika*, 93(2006), 491-507.
- Besley Timothy, "Political Selection," *Journal of Economic Perspectives* 19:3(2005): 43-60.
- Besley, Timothy, Olle Folke, Torsten Persson, and Johanna Rickne. "Gender Quotas and the Crisis of the Mediocre Man: Theory and Evidence from Sweden," *American Economic Review*, 107:8(2017), 2204-42.
- Bidwell, Kelly, Katherine Casey and Rachel Glennerster, "Debates: Voting and Expenditure Responses to Political Communication," conditionally accepted, *Journal of Political Economy*.
- Bochel, John and David Denver, "Candidate Selection in the Labour Party: What the Selectors Seek," *British Journal of Political Science*, 13:1(1983), 45-69.
- Briançon, Pierre, "The French (primary) revolution," Politico, 9/26/2016, <https://www.politico.eu/article/the-french-primary-presidential-election-2017-nicolas-sarkozy-francois-hollande-les-republicains-socialist/> accessed 12/5/2018.
- Broockman, David E. Nicholas Carnes, Melody Crowder-Meyer and Christopher Skovron, "Having Their Cake and Eating It, Too: Why Local Party Leaders Don't Support Nominating Centrists," *British Journal of Political Science*, forthcoming.
- Butler, Daniel M. and David W. Nickerson, "Can Learning Constituency Opinion Affect How Legislators Vote? Results from a Field Experiment", *Quarterly Journal of Political Science*, 6:1(2011), 55-83.
- Caillaud B, Tirole J. 2002. "Parties as Political Intermediaries." *The Quarterly Journal of*

Economics: 1453-1489.

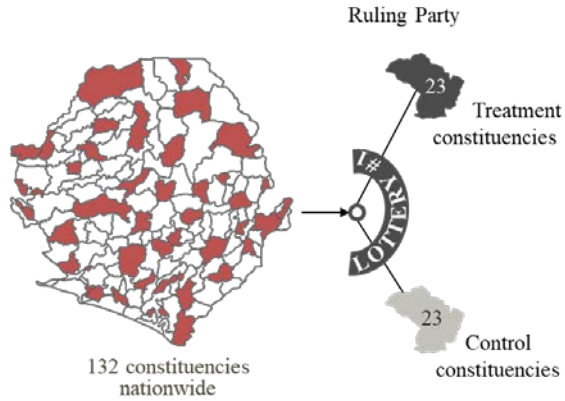
- Carey JM, Polga-Hecimovich J. 2006. "Primary Elections and Candidate Strength in Latin America." *The Journal of Politics* 68 (3): 530-543.
- Casey, Katherine, "Crossing Party Lines: The Effects of Information on Redistributive Politics," *American Economic Review*, 105:8(2015), 2410-2448.
- Cohen, Marty, David Karol, Hans Noel and John Zaller. *The Party Decides: Presidential Nominations Before and After Reform*, Chicago: University of Chicago Press, 2008.
- Cruz, Cesi and Philip Keefer, "Political Parties, Clientelism, and Bureaucratic Reform," *Comparative Political Studies*, 48:14(2015), 1942-1973.
- Dal Bó, Ernesto and Frederico Finan, "Progress and Perspectives in the Study of Political Selection," *Annual Review of Economics*, 10(2018), 541-575.
- Dal Bó, Ernesto, Frederico Finan, Olle Folke, Torsten Persson and Johanna Rickne, "Who Becomes A Politician?" *The Quarterly Journal of Economics*, 132:4 (2017), 1877-1914.
- Delli Carpini, Michael X. and Scott Keeter. *What Americans Know About Politics and Why It Matters*. New Haven: Yale University Press, 1996.
- De Paola, Maria and Vincenzo Scoppa, "Political competition and politician quality: evidence from Italian municipalities," *Public Choice* 148(2011): 547-559.
- Fearon, James D. "Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance," in *Democracy, Accountability, and Representation* (S. Przeworski et al., Eds.) Cambridge, UK: Cambridge Univ. Press 1999, 55–97.
- Ferraz, Claudio and Frederico Finan, "Exposing Corrupt Politicians: The Effects of Brazil's Publicly Released Audits on Electoral Outcomes," *Quarterly Journal of Economics*, 123:2(2008), 703-745.
- Ferraz, Claudio and Frederico Finan, "Motivating Politicians: The Impacts of Monetary Incentives on Quality and Performance," unpublished manuscript, University of California, Berkeley, 2011.
- Folke, Olle, Torsten Persson and Johanna Rickne, "The Primary Effect: Preference Votes and Politician Promotions," *American Political Science Review*, 110:3(2016), 559-578.
- Gagliarducci Stefano and Tommaso Nannicini, "Do Better Paid Politicians Perform Better? Disentangling Incentives from Selection." *Journal of the European Economic Association*, 11:2(2013), 369–98.
- Gulzar, Saad and Muhammad Yasir Khan, "Motivating Political Candidacy and Performance: Experimental Evidence from Pakistan" unpublished manuscript, Stanford University, 2018.
- Hirano, Shigeo, Gabriel S. Lenz, Maksim Pinkovskiy and James M. Snyder, Jr. "Voter Learning in State Primary Elections," *American Journal of Political Science*, 59:1(2014), 91-108.
- Hirano, Shigeo and James M. Snyder Jr., "Primary Elections and the Quality of Elected Officials," *Quarterly Journal of Political Science*, 9:4(2014), 473-500.
- Hirano, Shigeo and James M. Snyder Jr., Primary Elections in the United States. Cambridge:

- Cambridge University Press. 2019.
- Ichino N, Nathan NL. 2013. "Do Primaries Improve Electoral Performance? Clientelism and Intra-Party Conflict in Ghana." *American Journal of Political Science* 57 (2):428-441.
- Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz, "Experimental Analysis of Neighborhood Effects," *Econometrica*, 75:1(2007), 83-119.
- Lazear, Edward P. and Sherwin Rosen. "Rank-Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy*, 89:5(1981), 841-864.
- Levy, Adam, "DNC changes superdelegate rules in presidential nomination process," accessed: <https://www.cnn.com/2018/08/25/politics/democrats-superdelegates-voting-changes/index.html>
- Mattozzi, Andrea and Antonio Merlo, "Mediocracy," *Journal of Public Economics* 130 (2007), 32-44.
- Marken, Stephanie, "Still Listening: The State of Telephone Surveys," 2018, <https://news.gallup.com/opinion/methodology/225143/listening-state-telephone-surveys.aspx?version=print>
- Munshi Kaivan and Mark Rosenzweig, "Insiders and Outsiders: Local Ethnic Politics and Public Good Provision," unpublished manuscript, Yale University, New Haven, CT, 2016.
- Perry, James. 1996. "Measuring Public Service Motivation: An Assessment of Construct Reliability and Validity." *Journal of Public Administration Research and Theory: J-PART*, 6(1): 5-22.
- Platas, Melina and Pia Raffler, "Meet the Candidates: Field Experimental Evidence on Learning from Politician Debates in Uganda," forthcoming in Dunning, Grossman, Humphreys, Hyde, McIntosh, and Nellis (eds.): *Information, Accountability, and Cumulative Learning. Lessons from Metaketa I*. Cambridge University Press.
- Serra G. "Why Primaries? The Party's Tradeoff Between Policy and Valence." *Journal of Theoretical Politics* 23:1 (2011), 21-51.
- Simmons, Joseph P., Leif D. Nelson and Uri Simonsohn. "False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allow Presenting Anything as Significant", *Psychological Science*, 22:11(2011), 1359-1366.
- Sullivan, Andrew, "America Has Never Been So Ripe for Tyranny," *New York Magazine*, 1 May 2016: <http://nymag.com/intelligencer/2016/04/america-tyranny-donald-trump.html?gtm=bottom>
- Ting, Michael M., James M. Snyder, and Shigeo Hirano. "Primary Elections and the Provision of Public Goods," *Journal of Politics* 80:2 (2018), 647-661.
- Turner, Julius, "Primary Elections as the Alternative to Party Competition in 'Safe' Districts," *The Journal of Politics*, 15:2 (1953), 197-210.
- United Nations Development Programme, "A Handbook on Working with Political Parties," New York: 07 Jan 2006.
- United Nations Development Programme, "Funding Analysis of UNDP Electoral Assistance 2004-2018," internal memorandum, 4 January 2019.


- Valen, Henry, Michael Gallagher and Michael March, *Candidate Selection in Comparative Perspective: The Secret Garden of Politics*, London: Sage Publications, 1988.
- Wantchekon, Leonard, "Clientelism and Voting Behavior: Evidence from a Field Experiment in Benin," *World Politics* 55(2003), 399-422.
- Ware, Alan. *The American Direct Primary*. Cambridge: Cambridge University Press. 2002.
- Weaver, Jeffrey, "Jobs for Sale: Bribery and Misallocation in Hiring," manuscript, University of Southern California, 2018.
- Zou, Hui, and Trevor Hastie. "Regularization and Variable Selection via the Elastic Net," *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 67:2 (2005), 301-20.

Figure 1: Experimental design

Panel A: Ruling party (APC) constituency selection and assignment (locations masked)



Panel B: Opposition party (SLPP) constituency selection and assignment (locations masked)



Panel C: Implementation timeline

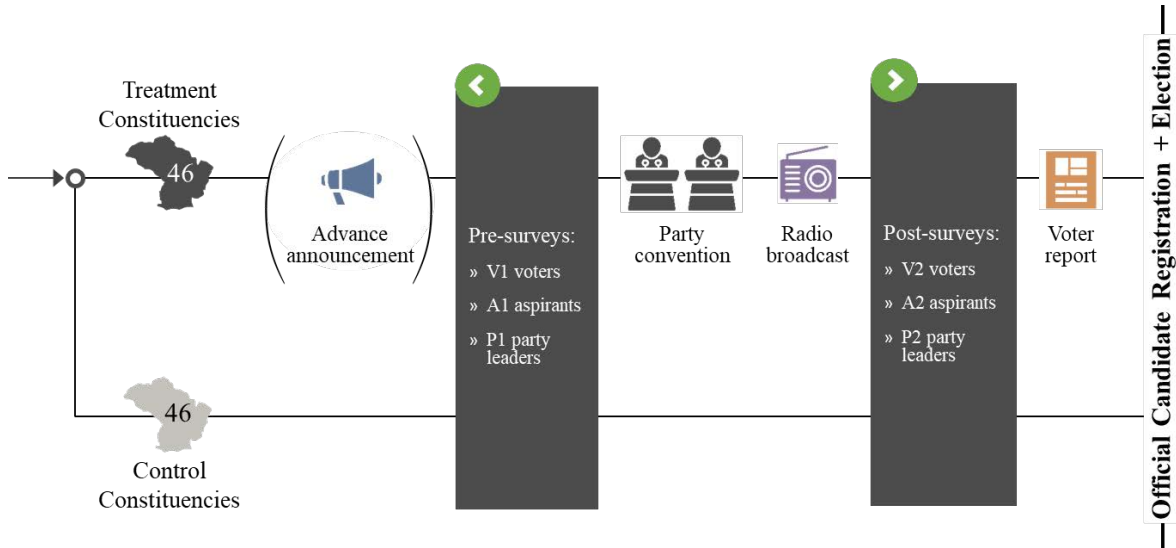
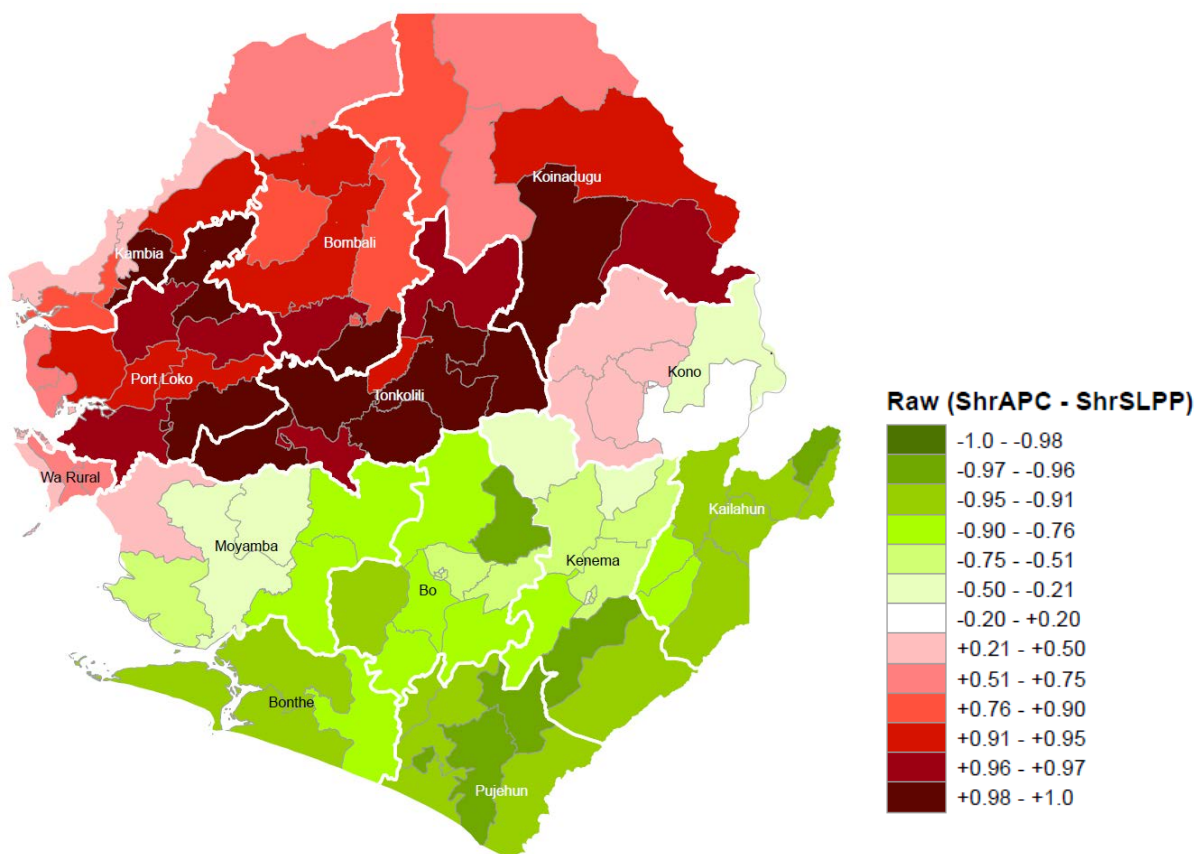


Figure 2: Ethnic-party strongholds



Notes: this figure shows the distribution and intensity of ethnicity-based ties to the two major political parties for Parliamentary constituencies in Sierra Leone. For each constituency, we compute the difference in population shares of ethnic groups historically associated with the All People's Congress minus those of ethnic groups associated with the Sierra Leone People's Party. Darker red shading indicates a constituency-level ethnic-party bias closer to 1 (e.g. where 1 indicates that the constituency is 100 percent populated by APC-affiliated ethnic groups) and darker green implies closer to -1 (e.g. a 100 percent SLPP-affiliated population). Mappings between ethnic groups and parties are from Kandeh (1992) and Casey (2005). Ethnicity data is from the 2004 census, mapped into contemporary constituency boundaries for the 2007 Parliamentary election (boundaries were redrawn before the 2018 election studied here).

Table 1: Selection into Politics

	Mean, voters	Mean, party officials	Mean, aspirants	<i>p</i> -value on (1) versus (2 and 3)
	(1)	(2)	(3)	(4)
Years of education	4.86	12.14	15.28	<0.001
Percent with no formal schooling	43%	5%	0%	<0.001
Percent with some university education	4%	34%	80%	<0.001
Asset ownership (of 11 household items)	2.66	6.44	9.49	<0.001
Proportion that have a bank account	0.11	0.77	0.98	<0.001
Proportion male	0.47	0.80	0.90	<0.001
Years of age	37.37	46.09	47.50	<0.001
Observations	7,544	245	433	

Notes: i) this table compares characteristics of voters (from survey V2), party officials (from survey P1) and aspirants (from survey A1); ii) p-values in column 4 refer to t-tests rejecting equality of means for voters as compared to party officials and aspirants pooled together; iii) the list of assets includes radio, personal computer, mobile phone, DVD player, refrigerator, bicycle, motor vehicle, generator, television, electric fan, and flashlight; and iv) bank account includes either domestic or foreign accounts.

Table 2: Heterogeneous Responses by Competition in the General Election

Panel A: Party Selection of Races to Include in the Experiment						
	Races nationwide		Races in experiment		Selection into experiment	
	<i>N</i>	Percent	<i>N</i>	Percent	Difference	<i>p</i> -value
	(1)	(2)	(3)	(4)	(5)	(6)
Safe seat constituencies	95	36%	48	52%	+14 pts	0.00
Swing constituencies	74	28%	28	30%	+2 pts	0.60
Weak seat constituencies	95	36%	16	17%	-20 pts	0.00
Observations (party-races)	264		92			

Panel B: Heterogeneous Treatment Effects on Representation					
	Selected candidate is the aspirant most preferred by voters				
	Mean controls	Mean treatment	Difference	<i>p</i> -value	<i>N</i>
	(1)	(2)	(3)	(4)	(5)
Safe seat constituencies	29.17	58.33	29.17	0.04	48
Swing constituencies	50.00	50.00	0.00	0.99	28
Weak seat constituencies	37.50	87.50	50.00	0.04	16

Panel C: Heterogeneity in Mean Payments from Aspirants to Party Officials						
	All aspirants			Candidates only		
	Control	Treatment	<i>p</i> -value	Control	Treatment	<i>p</i> -value
Competitiveness	(1)	(2)	(3)	(4)	(5)	(6)
Safe seat constituencies	\$2,983	\$3,076	0.92	\$1,390	\$5,771	0.07
Swing constituencies	\$1,482	\$1,917	0.43	\$994	\$2,016	0.29
Weak seat constituencies	\$1,089	\$1,355	0.76	\$839	\$604	0.67
Observations (aspirants)	185		202		46	

Notes: i) this table shows heterogeneity for three outcomes--where parties choose to experiment, how responsive they are to the voter reports, and mean payments from aspirants to party officials--by the level of likely competition in the general election; ii) *p*-values are from *t*-tests on the equality of means across samples, which is for the nation population of races versus those included in the experimental sample for panel A and across treatment and control party-races for panels B and C.

Table 3: Estimated Treatment Effects on Representation

Panel A: Direct Effect on Representation		
	Selected candidate is voters' first choice	
	(1)	(2)
Treatment	23.91** (10.52)	30.43* (15.89)
Ruling party		31.52 (40.87)
Ruling party X Treatment		-13.04 (21.14)
Mean in controls	36.96	
Observations	92	92
Panel B: Indirect Effect on General Election Vote Shares		
	Party's vote share in the general election	
	(1)	(2)
Treatment	-0.48 (2.96)	-0.64 (3.40)
Safe seat		27.90*** (6.64)
Safe seat X Treatment		-2.01 (6.13)
Weak seat		-25.95*** (6.89)
Weak seat X Treatment		6.91 (5.54)
Mean in controls	44.98	
Observations	91	91

*Notes: i) significance levels indicated by * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$; ii) ordinary least squares regression with robust standard errors; iii) specifications include fixed effects for 26 party-region strata used in the random assignments; and iv) one race is missing in Panel B since the general election was disputed and is being resolved by the courts.*

Table 4: Voter and Party Official Preferences over Aspirant Characteristics

	Aspirant's share in voter polls		Aspirant's share in party official survey		<i>p</i> - value (2 vs 4)
	(1)	(2)	(3)	(4)	
Professional qualifications index	0.04 (0.03)	-0.01 (0.04)	0.11*** (0.04)	0.13*** (0.05)	0.02
Wealth index	0.03** (0.01)	0.04* (0.02)	0.05* (0.03)	0.06* (0.03)	0.60
Economic development record index	0.04*** (0.01)	0.04* (0.02)	0.04* (0.02)	0.05* (0.03)	0.75
Public service motivation (PSM) index	0.04* (0.02)	0.05* (0.03)	0.03 (0.03)	0.01 (0.04)	0.38
Party loyalty index	0.04 (0.03)	0.02 (0.06)	0.04 (0.06)	-0.04 (0.08)	0.44
Cognitive ability index	0.01 (0.02)	-0.00 (0.03)	0.02 (0.03)	-0.02 (0.04)	0.65
Local network index	0.03 (0.03)	0.10** (0.04)	0.05 (0.05)	0.13** (0.07)	0.69
Campaign effort and expenditure index	-0.03 (0.02)	-0.01 (0.03)	-0.02 (0.04)	0.03 (0.04)	0.48
Conscientiousness indicator	0.06** (0.02)	0.07** (0.03)	0.09*** (0.03)	0.09** (0.04)	0.53
Professional qualifications X Treatment		0.10* (0.06)		-0.04 (0.08)	
Wealth X Treatment		-0.02 (0.03)		-0.02 (0.05)	
Development record X Treatment		0.01 (0.03)		-0.00 (0.05)	
PSM X Treatment		-0.02 (0.04)		0.04 (0.06)	
Party loyalty X Treatment		0.05 (0.07)		0.15 (0.12)	
Cognitive ability X Treatment		0.01 (0.05)		0.07 (0.06)	
Local network X Treatment		-0.14** (0.06)		-0.15 (0.09)	
Campaign X Treatment		-0.03 (0.04)		-0.09 (0.07)	
Conscientiousness X Treatment		-0.02 (0.03)		0.00 (0.05)	
Observations	370	370	369	369	368

Notes: i) columns 1 and 2 (3 and 4) use aspirant characteristics to predict their popularity among voters (party officials) in the V2 opinion polls (P1 leader surveys); ii) column 5 tests for differences in preferences between voters and party officials in control races, reporting the *p*-value from chi-squared tests of equality of coefficients across estimates in columns 2 and 4 from a seemingly unrelated regression framework; iii) significance levels indicated by * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$; iv) robust standard errors clustered by party-constituency; v) specifications include fixed effects for 24 party-region randomization strata; and vi) independent variables are eight indices of aspirant traits expressed in standard deviation units plus the binary behavioral measure of conscientiousness.

Table 5: Estimated Treatment Effects on Candidate Selection

	Treatment effect	Standard error	Naïve p -value	FDR q -value
	(1)	(2)	(3)	(4)
Panel A: Index-level regression estimates				
Personal qualifications index	-0.03	(0.10)	0.75	0.75
Wealth index	-0.12	(0.13)	0.37	0.56
Economic development index	0.29**	(0.13)	0.03	0.22
Public service motivation index	-0.08	(0.15)	0.60	0.75
Party loyalty index	0.04	(0.09)	0.62	0.75
Cognitive ability index	-0.13	(0.12)	0.29	0.53
Local networks index	0.19**	(0.09)	0.04	0.22
Campaign expenditure index	-0.16	(0.12)	0.19	0.49
Conscientiousness indicator	0.15	(0.10)	0.16	0.35
Panel B: Regularized regression estimates				
Number of development projects	0.50**	(0.21)	0.02	0.14
Incumbent MP	-0.09	(0.07)	0.23	0.57
Years of schooling	-0.17	(0.23)	0.45	0.72
Party versus own campaign expenditure	-0.13	(0.19)	0.50	0.72
Number of relatives in party leadership	0.24	(0.22)	0.29	0.57
Conscientiousness indicator	0.15	(0.10)	0.16	0.57
Observations	92			

Notes: i) this table reports treatment effect estimates on the characteristics of selected candidates for 9 indices of traits in Panel A and for the 6 individual traits selected via regularized regression in Panel B; ii) significance levels indicated by * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$ on a per comparison basis (p -values in column 3); iii) each row reports results from a separate OLS regression with robust standard errors that includes fixed effects for 26 randomization strata; iv) in Panel A, all indices are equally weighted sums of underlying traits expressed in standard deviation units (following Kling, Liebman and Katz 2007) and conscientiousness is a binary indicator; v) in Panel B all traits are in natural units; vi) party versus own expenditure indicates an affirmative response to the question "Are you more willing to spend money on your party's campaign versus your own?;" and vii) column 4 presents false discovery rate (FDR)-sharpened q -values that adjust for multiple inference over all estimates by panel, following Benjamini, Krieger and Yekutieli (2006) and Anderson (2008).

Table 6: Voter and Party Official Preferences in Post-Regularization Regressions

	Aspirant's share in voter polls	Aspirant's share in party official survey
	(1)	(2)
Number of development projects	0.04*** (0.01)	0.04*** (0.01)
Incumbent MP	0.16** (0.06)	0.21** (0.09)
Years of schooling	-0.01 (0.01)	0.03** (0.01)
Party versus own campaign expenditure	0.04* (0.02)	0.08** (0.03)
Number of relatives in party leadership	0.03** (0.02)	0.04* (0.02)
Conscientiousness indicator	-0.02* (0.01)	-0.03 (0.02)
Observations	370	369

*Notes: i) this table presents post-regularization estimates that retain only those aspirant traits with the greatest predictive power selected via 200 iterations of k-fold elastic net procedures; ii) significance levels indicated by * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$; iii) robust standard errors clustered by party-constituency; iv) specifications include 26 randomization strata; v) all traits expressed in natural units; and vi) party versus own expenditure indicates an affirmative response to the question "Are you more willing to spend money on your party's campaign versus your own?"*

Table 7: Estimated Treatment Effects on Payments to Party Officials

	All aspirants	Selected candidates	All aspirants	All aspirants, stronghold races
	(1)	(2)	(3)	(4)
Treatment	-67 (247)	809* (458)	-357 (310)	-777* (408)
Selected candidate			-607* (352)	-821 (502)
Selected X Treatment			1166** (555)	1932** (818)
Observations	385	92	385	237

Notes: i) this table estimates how payments (demarcated in US\$) from aspirants to party leaders are affected by the experimental treatment; ii) significance levels indicated by * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$; iii) ordinary least square regression with robust standard errors clustered by party-race; iv) specifications include fixed effects for 26 party-region strata used in the random assignments; and v) payments are winsorized at the 95th percentile.

SUPPLEMENTAL ONLINE MATERIALS

“An Experiment in Candidate Selection”

By Casey, Kamara and Meriggi

Contents

- Appendix A: Supplemental figures and tables
 - Figure A1: Example voter report
 - Figure A2: SLPP advance announcement flyer
 - Table A1: Aspirant traits selected via regularized regression
 - Table A2: Voter learning from conventions and broadcasts
 - Table A3: Treatment effects on aspirant entry
- Appendix B: Payment survey script and questions
- Appendix C: Status quo self-selection in safe versus weak seats
 - Table A4: Aspirant characteristics by general election competitiveness

Figure A1: Example voter report (aspirant names redacted)





Figure A2: SLPP advance announcement flyer (constituency list redacted)

Sierra Leone People's Party (SLPP)
National Headquarter and Western Region Office
Address: 15 Wallace Johnson Street, Freetown



Aspirant Voice and People's Choice

As the Party of the People, we are proud to announce a new pilot program that will complement existing procedures for awarding MP symbols. It is hoped that this programme will strengthen the internal democracy of our party. It will be tested in 23 constituencies. In these constituencies, this program will do two things:

1. Aspirant voice: The Party will host town hall debates amongst all aspirants for the party symbol inside the constituency. We will ask aspirants to stand before party leadership and community members and tell us why they are qualified to be an Honorable, what their policy positions are, and how they will represent the will of the local people in Parliament. Everybody can listen to these debates on local radio.
2. People's choice: The Party will ask the local people directly which aspirant has their support to become the symbol bearer. We will do this by polling voters directly via survey. The Party will seriously consider the local people's choice in deciding whom to award the symbol to.

We will implement this new program on a pilot basis in partnership with the Political Parties Registration Commission (PPRC) in 23 constituencies as follows (see next page).

Table A1: Aspirant Traits Selected via Regularized Regression

Panel A: Aspirant vote share in leader survey		Panel B: Aspirant vote share in voter polls	
Variable	Frequency	Variable	Frequency
No. of Development Projects	192	No. of Development Projects	200
Incumbent MP	103	Incumbent MP	200
Number of Meetings with Party Leaders	80	No. of Relatives in Party Leadership	200
No. of Relatives in Party Leadership	31	Education Level	199
Years Spent Serving in Elected Office	17	Party Versus Own Campaign Expenditure	181
Conscientiousness	17	Conscientiousness	100
Party Versus Own Campaign Expenditure	3	Time Spent as Member of Party	37
Has Previously Run for Public Office within Party	3	Number of Rallies Hosted in Constituency	7
Lives in Home Constituency	2	Number of Iterations	200
From a Chief/Ruling Family	2		
Education Level	1		
Number of Visits to Constituency in Past 6 Weeks	1		
Has Received in-kind Support from Party	1		
Log of Development	1		
Number of Different People Candidate Has Held Meetin	1		
Number of Rallies Hosted in Constituency	1		
Votes in Own Constituency	1		
Has a White-Collar Job	1		
Number of Iterations	200		

Notes: i) this table ranks aspirant traits by the number of times each was selected across 200 iterations of regularized regression; ii) the dashed line indicates the median number of traits selected over the 200 iterations, where traits above this frequency are carried forward into the post-regularization regressions of main text Table 6; iii) to tune the penalization parameters, each iteration uses k -fold cross validation, making ten random subsets of the data, using nine to train the model and the tenth as the validation sample; and iv) with an eye toward sparsity, we instruct the algorithm to search for optimal α values in the range $(0.5, 1)$, where $\alpha=1$ corresponds to LASSO with zero traits retained and $\alpha=0$ corresponds to ridge regression with all traits retained.

Table A2: Voter Learning from Conventions and Broadcasts

	Mean in controls	Treatment effect	Std. error
Dependent variables:	(1)	(2)	(3)
Identify most educated aspirant	0.36	0.15***	(0.05)
Identify most public office experience	0.30	0.16***	(0.06)
Identify strongest development record	0.32	0.10**	(0.05)
Observations	4,097	8,961	

*Notes: i) significance levels indicated by * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$; ii) each row reports results from a separate ordinary least squares regression using V2 data with robust standard errors clustered by party-race; and iii) specifications include fixed effects for 26 party-region strata used in the random assignments.*

Table A3: Treatment Effects on Aspirant Entry

	Party administrative data		Research survey data	
	SLPP	APC	SLPP	APC
	(1)	(2)	(3)	(4)
Treatment	1.00*	0.20	0.30	0.52
	(0.59)	(0.59)	(0.70)	(0.66)
Control mean	2.68	4.00	4.00	3.96
Observations (races)	45	38	46	46

Notes: i) this table estimates treatment effects on the total number of aspirants considered per party-race; ii) significance levels indicated by * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$; iii) ordinary least squares regression with robust standard errors; iv) specifications include fixed effects for each party's respective randomization strata; and v) columns 1 and 2 use administrative data from each party's Secretary General, columns 3 and 4 use the number of aspirants surveyed by the research team.

Appendix B: Payment survey script and questions

ENUMERATOR PROMPT: “Now I would like to ask you some questions about the campaign for the symbol. We know it's expensive to campaign and we know parties are short of funding.”
[Enumerator: Remind everyone that their answers will be kept secret.]

EXPENSE QUESTION: “How much contribution in total have you given to the party leaders for acquiring the party symbol from the start of your campaign? This includes registration/ application fees, tips, small or big token to all the party leaders, kola money and transportation reimbursement tips.”

[Enumerator: Please list the total amount in Leones. For example, if the amount is 100,000 Leones (one hundred thousand) enter 100000. If none, please enter ZERO.]

WINNER EXPENSE QUESTION: “To the best of your knowledge, how much contribution do you think the aspirant who received the symbol has given in total to the party leadership for receiving the party symbol?”

[Enumerator: Please list the total amount in Leones. For example, if the amount is 100,000 Leones (one hundred thousand) enter 100000. If none, please enter ZERO.]

Appendix C: Status Quo Self-Selection in Strong versus Weak Seats

A key idea in the existing literature is that higher returns to office induce positive selection into the pool of aspirants, which has been documented empirically with data from Brazil (Ferraz and Finan 2011), Italy (Gagliarducci and Nannicini 2013), Sweden (Dal Bó et al 2017) and the U.S. (Hirano and Snyder 2019). We can leverage the fact that differences in ethnic composition in Sierra Leone directly affect how likely the party is to win the Parliamentary seat, which in turn determines the expected returns to candidacy. Under this view, in the status quo we would expect there to be more aspirants, and of higher average quality, in stronghold races.

To explore this idea, Appendix Table A3 presents characteristics of the aspirant pool by the expected level of general election competition. It shows that there are more aspirants under consideration for stronghold safe seats, 5 on average, which is statistically distinct from the 4 vying for the nomination in swing seats and the just over 2 in weak seats. Aspirants on average appear to be of higher quality in stronghold areas: they on average have completed an additional half year of education (15.6 years versus 15.0 in swing and 14.7 in weak) and are more likely to have some university training (86 percent versus 74 and 62). They also appear a bit wealthier, although the differences are only significant comparing safe to weak seat aspirants, where the mean number of assets owned is 10.9 versus 9.2 (of 11). No clear pattern emerges with respect to demographics.

As mentioned in the main text, it is clear from the last row of estimates that payments to the party are also increasing in the likelihood that the aspirant's party will win the general election. Together these estimates are consistent with the probability of winning driving much of the expected return to office and thereby inducing positive selection into the aspirant pool.

Table A4: Aspirant Characteristics by General Election Competitiveness

	Mean, safe seats (1)	Mean, swing seats (2)	Mean, weak seats (3)	<i>p</i> -value (1) vs (2) (4)	<i>p</i> -value (2) vs (3) (5)	<i>p</i> -value (1) vs (3) (6)
Number of aspirants	4.94	4.04	2.31	0.09	<0.01	<0.01
Years of education	15.58	14.95	14.70	<0.01	0.51	<0.01
Percent with some university education	0.86	0.74	0.62	0.01	0.18	<0.01
Asset ownership (of 11 household items)	10.88	10.82	9.24	0.73	<0.01	<0.01
Proportion that have a bank account	0.97	0.96	0.92	0.91	0.27	0.17
Proportion male	0.90	0.86	0.95	0.20	0.15	0.40
Years of age	48.21	45.23	49.76	0.01	0.02	0.39
Average payment to the party (controls only)	\$2,983	\$1,482	\$1,089	0.12	0.64	0.25
Observations (party-races)	48	28	16			
Observations (all aspirants)	237	113	37			
Observations (aspirants, control races only)	118	51	16			

Notes: i) this table compares characteristics of aspirants across races where the general election is expected to be a safe, swing or weak seat for the aspirant's party; ii) p-values refer to t-tests rejecting equality of means across columns; iii) the list of assets includes radio, personal computer, mobile phone, DVD player, refrigerator, bicycle, motor vehicle, generator, television, electric fan, and flashlight; iv) bank account includes either domestic or foreign accounts; and v) payment refers to self-reported official and unofficial fees paid by aspirants to party leaders in control group races only.