

An Experiment in Candidate Selection[†]

By KATHERINE CASEY, ABOU BAKARR KAMARA, AND NICCOLÓ F. MERIGGI*

Are ordinary citizens or political party leaders better positioned to select candidates? While the American primary system lets citizens choose, most democracies rely instead on party officials to appoint or nominate candidates. The consequences of these distinct design choices are unclear: while officials are often 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. We partnered with both major political parties in Sierra Leone to experimentally vary how much say voters have in selecting Parliamentary candidates. Estimates suggest that more democratic procedures increase the likelihood that parties select voters' most preferred candidates and favor candidates with stronger records of public goods provision. (JEL D72, H41, O17)

The competence and integrity of political leaders is a key determinant of government performance. James Madison (1788) went so far as to argue that the primary objective of any political constitution ought to be leadership selection, specifically to yield rulers with the wisdom and virtue to best pursue the common good. And yet despite the importance of selection becoming ever more apparent, Besley (2005, p. 44) writes that much of the modern political economy literature “has not only

*Casey: Stanford University and NBER (email: kecasey@stanford.edu); Kamara: International Growth Centre (email: abou.kamara@theigc.org); Meriggi: International Growth Centre and Wageningen University (email: niccolo.meriggi@theigc.org). Esther Duflo was the coeditor for this article. 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, Samuel Asher, Abdulai Bah, Arthur Baker, Tom Bangura, Chris Boyer, Fatu Conteh, Arja Dayal, Linda Gassama (in memoriam), Yuen Ho, Erin Iyigun, Fatoma Momoh, Alfred Rasheed, Andrés Rodríguez, Marcos Salgado, Rosemarie Sandino, Osman Siddiqi, 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, to the referees, as well as Dave Baron, David Broockman, Steve Callander, Fred Finan, Emmanuel Gaima, Pedro Gardete, Guido Imbens, Saum Jha, Alhassan Kanu, Eddie Lazear, Herbert M'cleod, Ben Olken, 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, the International Growth Centre and the Stanford Institute for Innovation in Developing Economies. The authors declare that we have no relevant or material financial interests that relate to the research described in this paper. This research received institutional review board approval from Stanford University (protocol 42946) and the Sierra Leone Ethics and Scientific Review Committee. This study was preregistered in the AEA registry of randomized controlled trials (<https://www.socialscienceregistry.org/trials/2485/history/23509>). All errors are our own.

[†]Go to <https://doi.org/10.1257/aer.20200125> to visit the article page for additional materials and author disclosure statements.

neglected the problem of political selection, it has been positively hostile to the topic.”¹

One critical component of selection is how political parties choose candidates. In most countries, this is by appointment or nomination by party elites. This contrasts sharply with the direct vote primary system in the United States, which devolves control to ordinary citizens. The divergence raises questions about which selection method works better, and what the consequences of voter versus party leader control might be for the overall electoral system.

It is hard to predict what would happen if the locus of control were to shift from party elites toward voters. To start, we do not know exactly which traits voters and party elites value in candidates, and whether their preferences are shared or in conflict with one another. We also do not know how relatively well informed they are about which candidates possess the traits they value, or how much they know or care about each other’s preferences. Both sets of agents may further act strategically when selecting candidates, weighing a broad array of factors that may be difficult to fully anticipate.

While generating precise predictions may thus be infeasible, there are some commonly held views about what a shift toward primaries might do. The first comes from primary skeptics, who are concerned that voters are poorly informed about politics. If so, then giving voters control straightforwardly delivers representation, i.e., citizens get their most preferred candidate, but it may come at the cost of selection on quality. And if severe enough, voters would be better off delegating the choice to party elites, who use their superior information and expertise to screen candidates’ on their technical merits and identify the best performers.

A countervailing concern is that elite choices may diverge from voter preferences. This happens, for instance, if elites’ political preferences are shaped by their privileged status, or they value candidate traits unrelated to performance in office (like party loyalty or willingness to pay for the nomination). They may also deploy more complicated strategies than voters, taking into account how selections in one race affect voting in other races, or how today’s selections affect the party’s options in future. And fundamentally, if no primary is held, party elites may have little idea what voter preferences are, leading the choices of even the best intentioned party delegates away from what voters are looking for in candidates.

There is scant evidence about how these trade-offs between a poorly informed citizenry and a potentially misaligned political elite are resolved in practice, and what their implications are for representation and the quality of selected candidates. Empirical progress has been constrained by the fact that political parties are generally loathe to vary how they choose candidates for anything but purely strategic (and thus endogenous) reasons.

This paper overcomes this identification challenge by partnering with both major political parties in Sierra Leone on a novel experiment that varied how much say registered voters, as compared to party officials, have in selecting candidates for the 2018 Parliamentary elections. In the status quo, parties chose among potential candidates, referred to as “aspirants,” in a given constituency via recommendations from

¹We review the more recent literature on political selection below.

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 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 party officials via a one page report. Neither component is binding on the party's ultimate choice of candidate, and both are best characterized as alleviating information constraints. Yet note that if party officials followed the voter reports in all cases, the intervention approximates a direct vote primary with mandatory turnout.² In practice, they increased the rate at which they nominated the voters' first choice candidate from 38 to 61 percent.

To join this initiative, interested parties were first asked to select 46 races from the population of 132 Parliamentary constituencies nationwide where they would be willing to experiment with the new selection mechanism. More than one-half of the races they chose were located in their respective regional strongholds, where they were very likely to win the general election. The research team then randomly assigned the intervention for each party independently, with 23 treated and 23 control races (for a total sample of 92 party-races); implemented the opinion polls; and collected rich data from voters, aspirants and party officials in both treated and control races. A well-known media group worked with the parties to moderate the party conventions and put them on the radio. We use the data to first characterize candidate selection in the status quo, and then estimate causal effects of the new selection method on key outcomes of interest, focused on representation and its relationship to selection on quality.

We find evidence that the status quo method of delegating candidate selection to party officials distorts choices away from voter preferences. This impedes representation, as defined as the candidate chosen to run in the general election being the aspirant who ranks first among voters. The baseline rate of selecting voters' first choice of 38 percent is low, given that there were on average only four aspirants under consideration per race. This rate increases by 23 percentage points with treatment, a large and highly statistically significant effect, which is driven by party responsiveness in safe and weak seats. Back-of-the-envelope calculations suggest that this positive effect on representation corresponds to party officials choosing a different candidate than they would have otherwise in 11 races, and thereby changing the identity of 6 elected Members of Parliament (i.e., those from treated stronghold races).

In designing this research, we thought that distortions away from voter preferences would likely be due to the fact that party officials and voters have very different ideas about what makes for a good candidate. This proved not to be the case empirically. To test this, we identify which aspirant characteristics predict their popularity among voters, and compare these to the traits that make them popular among constituency-level party officials. The data suggest that voters and party officials broadly agree: they both prefer aspirants with a stronger record of having provided

²The leap from voluntary to mandatory turnout is not unreasonable in this setting, as voluntary turnout reached 84 percent in the general election studied. Turnout could, of course, be lower or more selective in the primary stage.

local public goods, and those who are more conscientious. The former is measured by the aspirant's involvement in development projects, like small scale public infrastructure (constructing a bridge or community center), support to education (rehabilitating classrooms), and agriculture (procuring farm tools and tractors) over the past three years. The latter is measured by a behavioral indicator of how carefully the aspirant handled a financial reimbursement for transport expenses (described in detail in Section VA).

Given the apparent lack of preference divergence, what else might explain why party officials often fail to select the most popular candidate in the status quo? The data point strongly to 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 local party leaders and representatives in the United States have been shown to hold inaccurate views of public opinion (Butler and Nickerson 2011, Broockman et al. 2019). In Sierra Leone, 90 percent of constituency-level party officials presume local voters share their first choice over aspirants, when voters in fact only agree with them about one-half of the time; and in one-third of races, not a single party official (among multiple surveyed per race) accurately guessed which aspirant ranked first among voters. Their responsiveness to the voter reports reveals the usefulness of the primary stage in informing party officials about voter preferences. And since their own preferences were not in direct conflict with those of voters, there was substantial scope to accommodate voter opinion in choosing candidates.

Did the documented increase in representation come at the expense of selection on quality? To answer this, we compare the characteristics of candidates ultimately sent to the general election across treatment and control races. Experimental 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. Estimates for conscientiousness are directionally similar but imprecise. Voter learning from the conventions and radio broadcasts may have aided these positive results, as the data show that voters in treatment races knew more about aspirants than voters in control races. Our headline results thus demonstrate scope for party officials to incorporate the preferences of voters, even those who are very poor and lack formal education, in a way that facilitates representation without compromising the quality of selected candidates.

In thinking about generalizability, it is important to recall a few distinctive features of the research design. Our estimates are representative of races where the parties were willing to experiment, where the modal race is a safe seat where the party was near guaranteed to win the general election. Unlike in a binding direct vote primary, party leaders retained authority and discretion, as they were free to disregard the voter reports. Compiling these reports from a representative sample of voters further sidesteps some contentious issues in American primaries, like the relationship between selective turnout and the choice of ideologically extreme candidates.

We can largely rule out two alternative channels, aspirant entry and financial contributions, as drivers of the observed empirical patterns. Differential aspirant entry into treated races is unlikely to explain our experimental results, as we find

little evidence that advance announcement of the initiative, which was only partially implemented, induced entry. Average aspirant contributions (in official application fees and unofficial payments) amount to an unadjusted mean of \$2,477, which is equivalent to 1.3 months of an MP's salary and 34 times the monthly minimum wage (data we elicited via survey). While this presents a substantial barrier to entry, the data are not consistent with some of the more nefarious interpretations of these payments, like seats being sold to the highest bidder.³

Our analysis contributes to the relatively new literature on political selection (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 Bó et al. 2017) facilitate positive selection.⁴ Much less is known about the influence of party leaders on selection, and the two most related studies provide contrasting results. While Dal Bó et al. (2017) point to party leaders' use of merit-based promotions in helping make Sweden an "inclusive meritocracy," Besley et al. (2017) argue that their incentives to do so are tempered by a desire to avoid internal leadership threats. In a similar vein, Mattozzi and Merlo (2015) model party recruitment of candidates as a rent maximization opportunity that leads to "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, p. 21) 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." They find evidence that primaries promote the selection of competent candidates in safe, open seats (Hirano and Snyder 2014, 2019). This idea is supported by theoretical work arguing that primaries produce higher quality candidates (Adams and Merrill 2008), but potentially at the cost of ideological extremity (Serra 2011). More broadly, other models suggest that primaries regulate internal competition to induce greater effort developing policy (Caillaud and Tirole 2002) and promote the provision of public goods over private transfers (Ting, Snyder, and Hirano 2018).

Our paper brings these two literatures on political selection and primaries together by measuring what happens to representation and candidate quality when voters, as compared to party leaders, are afforded more say in selecting candidates. Treating primaries as a mechanism to alleviate information constraints distinguishes our approach from much of the literature: party officials and voters learn about aspirants from the conventions; party officials learn about voter valuations of aspirants from the reports; and pooling all this information together helps identify candidates who are strong on dimensions that both voters and party officials agree are important. Overall, our results suggest that the more democratic selection method creates value

³This concern is not limited to new or weak democracies, as the (now former) Illinois governor's attempt to sell President Obama's vacated Senate seat attests: see Monica Davey and Jack Healy, "Illinois Governor Charged in Scheme to Sell Obama's Seat," *New York Times*, December 9, 2008, <https://www.nytimes.com/2008/12/09/us/politics/10Illinois.html>.

⁴Noting the caveat that higher illicit (as opposed to official) returns may have the opposite effect (Brollo et al. 2013).

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 I discusses variation in candidate selection processes and introduces a simple framework to structure the analysis. Section II details the experimental design and interventions. Sections III through VI analyze key trade-offs in the status quo, and estimate how they are impacted by the new selection mechanism, focusing on representation, preference divergence, and selection on quality. Section VII explores alternative mechanisms related to aspirant entry and contributions. Section VIII concludes.

I. Candidate Selection in Perspective

A. Empirical Variation

There is substantial variation across country and over time in how political parties select candidates. Consider first the centralized approach practiced in France: central party leaders historically chose all candidates and allocated them across space to populate subnational races (Gallagher and March 1988). Parties in the United Kingdom have traditionally used a more decentralized method, 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 made this group quite narrow: data from the 1980s suggest that on average 40 Labour Party members chose the candidate on behalf of some 70,000 constituents (Bochel and Denver 1983).

The US 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, pp. 31–32). More recently, the 2016 presidential race divided the two major parties over whether there is now “too little” or “too much” democracy. The Democrats increased 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.⁵ Across the aisle, some pundits who were troubled by the prospect of Donald Trump’s candidacy called for a stronger role for party elites as gatekeepers of the nominating process, and as a check on the excesses of “hyperdemocracy.”⁶ 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 2016.⁷

The consequences of these disparate design choices on the performance of the electoral system in delivering high-quality, representative candidates and elected

⁵ See Adam Levy, “DNC Changes Superdelegate Rules in Presidential Nomination Process,” *CNN*, August 25, 2018, <https://www.cnn.com/2018/08/25/politics/democrats-superdelegates-voting-changes/index.html>.

⁶ See Andrew Sullivan, “America Has Never Been So Ripe for Tyranny,” *New York Magazine*, May 1, 2016, <http://nymag.com/intelligencer/2016/04/america-tyranny-donald-trump.html>.

⁷ See Pierre Briancon, “The French (Primary) Revolution,” *Politico*, September 26, 2016, <https://www.politico.eu/article/the-french-primary-presidential-election-2017-nicolas-sarkozy-francois-hollande-les-republicains-socialist>.

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. Both major parties rely on the recommendations of party officials at various levels, beginning with constituency-level officials, and neither party has a clear mechanism in place to capture the preferences of voters or rank-and-file party members. The experimental treatment we evaluate—conventions combined with polling—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 United States.

B. Conceptual Framework

There is no model in the literature that speaks directly to how a shift from party leader control to a process that gives more voice to voters affects candidate selection. A simple framework is thus useful to define key concepts, illuminate trade-offs between poorly informed voters and potentially misaligned elites, and frame the experiment studied with respect to information constraints.

Setup.—Suppose 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, integrity) and match-specific traits associated with the jurisdiction (fluency in local languages, knowledge of local priorities). 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 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, where local voters value a school more than a clinic if there are few schools and many clinics nearby, or if they prefer education over health. There is a third factor (such as party loyalty or willingness to pay for the nomination) that does not directly contribute to performance but is potentially correlated with quality.

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. We focus on the intuitive case where the voter is relatively better informed about local alignment, and the party official relatively better informed about competence.

We are interested in how allocating more say to the voter as compared to the party official affects two 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; and (ii) selection on quality, which is the expected performance (value of public goods produced) of the candidate selected. We seek to understand when improvements in representation come with consequences (positive or negative) for selection on quality, and the role of a nonbinding information treatment in this process.

Sources of Divergent Choices.—There are three main channels that would lead the voter and party official to select different candidates. The first is divergent preferences: a common concern about the status quo policy of delegating the choice to the party official is that only the official values the third factor (call it party loyalty). He thus maximizes a combination of loyalty and quality that at times selects a loyal aspirant over a more competent one. If the voter and party official's information sets are similar, then giving the voter more say straightforwardly enhances both representation and selection on quality. This was a key argument made by Progressive reformers in the United States. 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 positive, it could be of little consequence.

Disparate strategic objectives could similarly drive a wedge between the choices of the voter versus the party official. While the voter cares about getting the best MP for her jurisdiction, the party official may be playing a broader game that accounts for how the selection of that MP candidate affects other races. He might, for example, select a poorly aligned MP candidate who appeals to voters in nearby jurisdictions, in order to attract their votes for a higher level office. This could be relevant in our context as the Parliamentary races studied ran concurrently with a closely contested presidential race. As above, giving more say to the voter in this case enhances representation and selection on quality so long as her information set is not too inferior.

A third source of divergence arises when the voter is at an absolute screening disadvantage compared to the party official, which in contrast to the above cases, negatively affects selection when she has control. This happens when her information set is strictly worse, or if the trait she can observe relatively well (alignment) has a lower marginal product with respect to performance. If so, then giving the voter more voice 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 the voter (via radio), 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 to reward a more loyal aspirant, appeal to voters in nearby races, or choose a more competent aspirant; and (ii) enhance selection on quality, by pooling the information that the voter and party official respectively possess about competence and alignment, and hence increase expected performance.

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. This lack of strong ideological labeling is fairly common: Cruz, Keefer, and Scartascini (2016) find that roughly one-half of parties across 179 countries are “non-programmatic,” which means they cannot be classified on a left, center, right scale or other metric of economic policy.⁸

II. Experimental Design

This study explores how political parties in Sierra Leone chose candidates to compete in the 2018 Parliamentary elections. These are single-member jurisdictions, won by plurality, and the general election of interest was declared largely free and fair by domestic and international observers.⁹

Status quo selection in Sierra Leone is guided by the country’s Constitution,¹⁰ which specifies eligibility requirements for becoming a Member of Parliament (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 selection by three to 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 mechanism to systematically elicit or aggregate voter preferences.

For a randomly selected subset of races, the parties experimented with a new selection mechanism that had two components, implemented in tandem at the constituency-party level: (i) a town hall-style party convention, broadcast over local radio; and (ii) representative opinion polling of registered voters, aggregated into one page reports and shared with party officials. The offer to participate in the initiative and associated research was managed by the Political Parties Registration Commission (PPRC) of Sierra Leone, which has the constitutional mandate to register, supervise, and monitor the conduct of parties. Its remit includes monitoring the accountability of parties to their members and the broader electorate.¹¹ The PPRC

⁸In 2015 data, 54 percent of parties are non-programmatic (authors’ calculation). See also Cruz and Keefer 2015 (p. 1949).

⁹See, 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.

¹⁰See the Constitution of Sierra Leone (1991) available at <http://www.sierra-leone.org/Laws/constitution1991.pdf>.

¹¹For more information, see <https://pprc.gov.sl>.

extended the offer to participate to all registered parties, whose leadership decided whether to opt in.

The parties, the PPRC, the research team, and a civil society group called Search for Common Ground (SFCG) worked together to design and implement the two-pronged initiative. SFCG produces programs to promote transparency and accountability around the political process, providing neutral and reliable content through a network of local radio and television broadcasters. It played a leading role in coordinating broader civil society efforts to support the 2018 (and previous) elections.¹²

A. Party Conventions

The first component of treatment, a constituency-level party convention, provides 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. These town-hall style gatherings typically began with a moderator, trained by SFCG, introducing the aspirants to the audience, and then posing a series of policy questions. Standard questions included (i) explain who you are and what qualifies you to be a good MP; (ii) how would you spend the constituency facilitation fund, a pot of public money given annually to each elected MP; and (iii) what makes you a good representative of local people, including how you would know what local people want and represent their interests in Parliament?¹³ Additional questions followed, tailored to the local area, covering topics such as how to deal with local power supply constraints and allocating mining royalties. Shortly after each convention finished, SFCG delivered audio recordings of the event to local radio stations that rebroadcast the convention multiple times over subsequent days.

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. SFCG also publicized these events to local residents via radio jingles and community visits, encouraging them to attend the town hall or listen to the broadcasts. The data collection team further provided 25 voters per constituency with advance notice of the events via survey (see V1 survey described below). In post-convention surveys, 18 percent of voters in treatment areas reported having attended the town hall or listened to a broadcast. This rate is higher for those notified in advance (at 24 percent).

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 individuals from the official registry of voters maintained by the National Electoral Commission (NEC), which includes

¹²For more information, see <https://www.sfcg.org/sierra-leone>.

¹³Moderators gave each aspirant 2 minutes to respond to each question, alternating who spoke first across questions.

names, demographics, and home address. To ensure representation, we first randomly selected 10 voter registration centers (or 40 percent) per constituency, and then randomly selected 10 voters per center, stratifying on age and gender. The respondent contact 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.¹⁴ These surveys are thus substantially more representative than the telephone polls commonly conducted in the US where, for example, the Gallup poll currently has a survey response rate of 7 percent (Marken 2018).

Voters were asked the question, “If you could give the symbol¹⁵ directly, who would be your first choice?” and had the option to say “don’t know” if they had not heard of any of the aspirants. 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 received among poll respondents in the constituency (see example in online 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 partisan affiliation, these two tabulations rarely identified a different front-runner. 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 200 copies to the SLPP for distribution to all affected constituency-, district-, and national-level executives; and delivered 25 reports to national executives of the APC.

C. Treatment Assignment

For each party that responded positively to the initial invitation to participate, the PPRC asked its national leaders for a list of 46 constituencies from the universe of 132 nationwide where they were willing to experiment with the new candidate selection mechanism. The research team then randomly assigned, via computer program, one-half of each party’s constituency list to treatment and one-half to control, stratifying by small geographic bins.¹⁶ This generates an experimental sample of 92 party-races. As the two major parties occasionally picked the same constituency for inclusion, these 92 party-race observations cover 80 unique constituencies.

Figure 1 presents an overview of the random assignment and implementation timeline.¹⁷ The conventions launched in mid-November 2017 and all were completed and voter reports delivered before the parties submitted their official list of

¹⁴Each target respondent was accompanied by two potential replacements from the same demographic bin.

¹⁵In local parlance, “give the party symbol” means select the candidate to proceed to the general election.

¹⁶In addition to the two major political parties, 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.

¹⁷To protect the anonymity of research participations, panels A and B mask locations and show arbitrary assignments using the old 2007 constituency boundaries (the boundaries were redrawn before the 2018 election).

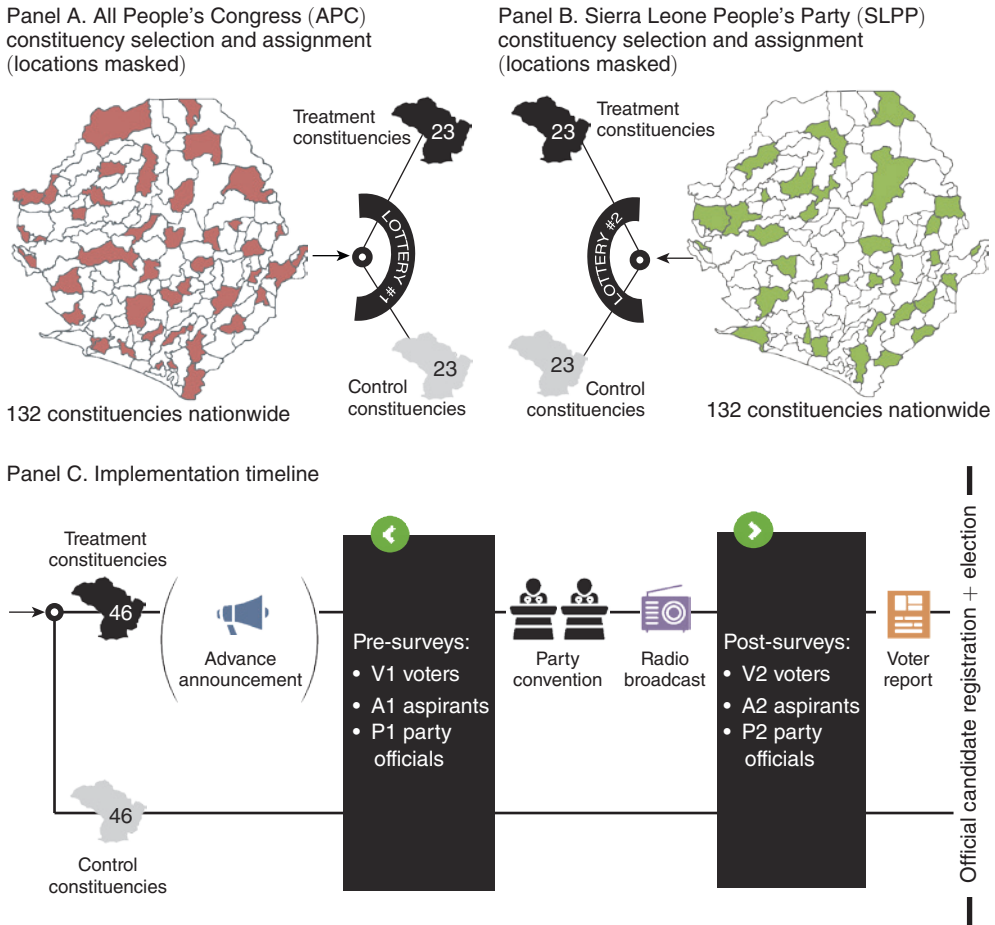


FIGURE 1. EXPERIMENTAL DESIGN

candidates to NEC 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.

While the original implementation plan further included advance announcement of which races would use the new selection method, this was only partially implemented (and is thus bracketed in Figure 1). This is for two reasons: (i) the country's constitution stipulates that MP candidates in public employment must vacate their post a year before the election, and this initiative launched too late to affect the entry decisions of those potential aspirants; and (ii) only one of the two participating parties did any advance announcement (the SLPP publicized the list of treated races and what treatment entailed two months before the conventions). We thus consider the entry channel largely shut down in this context (see Section VIIA for discussion), however expect that it could be consequential in other settings. Limited advance announcement further implies that party leaders were unlikely to know the list of races their rival party selected when choosing which of their own races to include in the experiment, which could be a strategic consideration in other contexts.

TABLE 1—SELF-SELECTION INTO POLITICS

	Mean, voters (1)	Mean, party officials (2)	Mean, aspirants (3)	<i>p</i> -value on (1) versus (2 and 3) (4)
Years of education	4.86	12.15	15.32	<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	6.45	9.52	<0.001
Proportion that have a bank account	0.11	0.77	0.97	<0.001
Proportion male	0.47	0.81	0.89	<0.001
Years of age	37.37	46.15	47.56	<0.001
Observations	7,544	239	390	

Notes: (i) This table explores self-selection into politics by comparing the characteristics of voters (from the V2 survey), party officials (from the P1 survey), and aspirants (from the A1 survey); (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.

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), aspirants (A1 and A2), and party officials (P1 and P2). Data collection rolled out sequentially across stratification bins and treatment assignment, which is important since the information environment was evolving over time in all races. 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). Anonymized data for the analysis in this article are available (see Casey, Kamara, and Meriggi 2021). See summary statistics in Table 1 (discussed further in Section IIIA).

Voter reports were 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 official surveys (P1 and P2), enumerators surveyed those holding the same standard constituency-level positions (constituency chair, secretary, and treasurer) who were encouraged to attend the conventions. These respondents were replaced as necessary with holders of similar constituency positions (e.g., deputy constituency chair) or higher level party officials (e.g., district chair). Overall, 73 percent of officials surveyed hold constituency-level positions, 5 percent hold district-level positions, and 22 percent hold other positions.

III. Characterizing Selection in the Status Quo

This section uses the data to explore three aspects of candidate selection in the status quo: (i) who self-selects into politics; (ii) which types of races party leaders chose to include in the experimental sample; and (iii) how frequently voters get their

most preferred candidates in control races where selection is fully delegated to party officials.

A. *Self-Selection into Politics*

Comparing the characteristics of voters, party officials and aspirants reveals strong positive self-selection into politics on education and wealth. It further suggests that trade-offs emphasized in the conceptual framework between a poorly informed electorate and a potentially misaligned political elite could be relevant in Sierra Leone.

Table 1 shows that registered voters have on average completed 5 years of education, 43 percent of them have no formal schooling, and only 4 percent have been to university. Aspirants, by contrast, have completed over 15 years of education, none lack formal schooling, and 80 percent have been to university. Party officials sit in between the two, with 12 years of education, 5 percent without schooling, and 34 percent with some university. Such pronounced selection on education suggests that party officials might be better able than voters to screen aspirants on their technical merits. As a concrete example, the first formal step for aspirants is to file an application with the party, which covers items like eligibility requirements and their standing in the party. Most voters would find it difficult to read the aspirants' applications, or review their curriculum vitae (even if they were publicly available), 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.¹⁸

The countervailing concern is that the socioeconomic divide between voters and party officials will lead the latter to select candidates who are not aligned with voter interests. This would obtain under a citizen candidate model (Osborne and Slivinski 1996, Besley and Coate 1997), for example, if the elite status of party officials and candidates shapes their policy preferences away from those of voters. Inequities apparent in proxies for wealth lend some credence to this concern: voters on average own fewer than 3 assets from a list of 11 household items (e.g., mobile phone, radio) and only 11 percent have a formal bank account. Aspirants, by contrast, on average own 9.5 of these assets and nearly all of them (97 percent) have an account. Party officials again fall in between: they on average own 6 assets and 77 percent have an account. As to demographics, politics is a male-dominated activity, 81 (89) percent of party officials (aspirants) are male, and politicians are roughly ten years older than the average voter. These summary statistics make it unclear *ex ante* whether voters or party officials are better positioned to select candidates, and point to potential tensions between representation and selection on quality.

B. *Where Parties Chose to Experiment*

Recall that participating parties selected 46 races from 132 Parliamentary constituencies nationwide for inclusion in the research sample and thus the lottery that assigned the initiative. It is instructive to classify their choices with respect to how

¹⁸ As a point of reference, 33 percent of voters (in the V1 survey) could not name a single MP job responsibility.

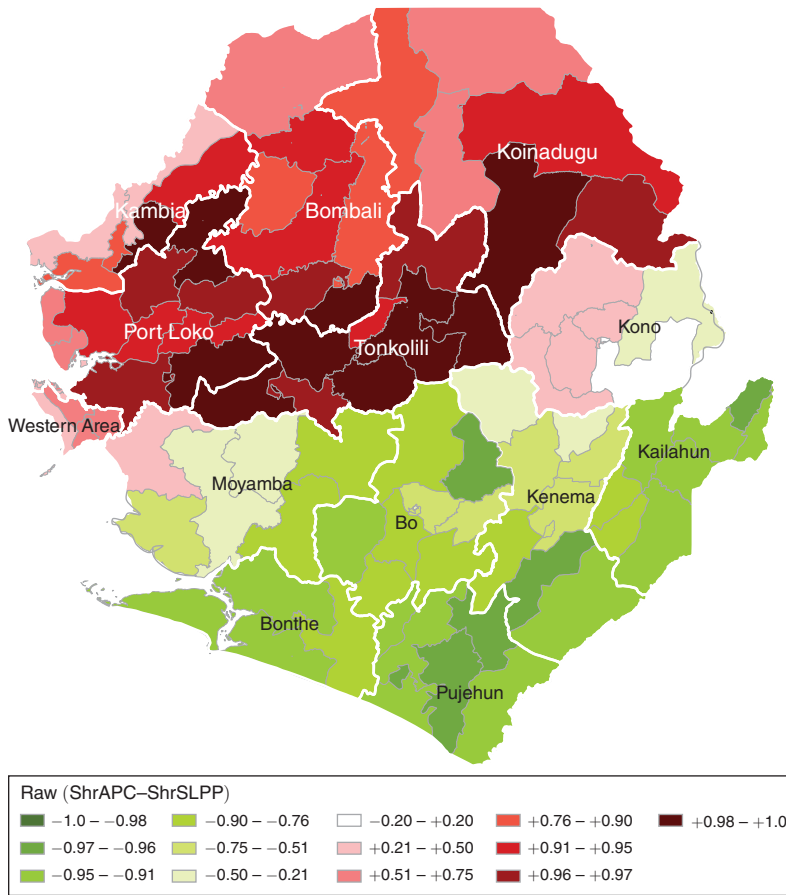


FIGURE 2. ETHNIC-PARTY STRONGHOLDS

Notes: This map shows the geographic 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 (APC) minus those of ethnic groups associated with the Sierra Leone People’s Party (SLPP). Darker red shading indicates a constituency-level ethnic-party bias closer to 1.0 (e.g. where 1.0 indicates that the constituency is 100 percent populated by APC-affiliated ethnic groups) and darker green implies closer to -1.0 (e.g. a 100 percent SLPP-affiliated population). Color choices reflect party symbols: the APC’s logo is a red rising sun and the SLPP’s is a green palm tree. Mappings between ethnic groups and parties are from Kandeh (1992) and Casey (2015). Ethnicity data is from the 2004 census (Statistics Sierra Leone 2004), mapped into constituency administrative boundaries for the 2007 Parliamentary elections (National Electoral Commission 2007). Note that constituency and district boundaries were redrawn before the 2018 elections studied here.

competitive the general election is likely to be, which can be done using census data on constituency-level ethnic composition.

As background, 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 most subnational jurisdictions are not, as they are located inside one of the two party’s regional strongholds. Figure 2 maps these allegiances. For each constituency, we compute the difference in population shares of ethnic groups historically associated with the APC versus the SLPP. The darkest

TABLE 2—PARTY CHOICE OF RACES AND RESPONSIVENESS TO THE PRIMARY INTERVENTION BY LEVEL OF GENERAL ELECTION COMPETITION

	<i>Panel A. Selection of races for the initiative</i>				<i>Panel B. Selected candidate is voters' first choice</i>			
	Proportion nationwide	Proportion in sample	Difference (2) versus (1)	<i>p</i> -value for (3)	Mean control	Mean treatment	Difference (6) versus (5)	<i>p</i> -value for (7)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Safe seat races	36.0	52.2	16.2	<0.01	30.4	58.3	27.9	0.06
Swing races	28.0	30.4	2.4	0.66	50.0	50.0	0.0	0.99
Weak seat races	36.0	17.4	-18.6	<0.01	37.5	87.5	50.0	0.04
Observations (party-races)	264	92						

Notes: (i) This table shows how the expected degree of general election competition affected party response to the primary stage intervention along two dimensions, namely which type of race party officials chose to include in the experimental sample (in panel A) and how responsive party officials were to the voter reports (in panel B); (ii) *p*-values are from *t*-tests on the equality of means across samples, which compares the national population of races to those included in the experimental sample for panel A, and across treatment and control group races for panel B.

red shading indicates a constituency populated almost entirely by APC-affiliated groups, and the darkest green indicates one populated wholly by SLPP-affiliated groups. To demonstrate how strong these allegiances are, regression analysis shows that differences in ethnic population shares explain 92 percent of the variation in the two-party vote for the 2007 Parliamentary elections.¹⁹

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 US House of Representatives races from 1952 to 2010 were decided by fewer than 15 percentage points, which is a fairly lax standard for competitiveness.^{20,21}

These stronghold races are the ones where party leaders were most interested in piloting the new selection initiative. Panel A of Table 2 shows that while 36 percent of races nationwide are expected to be safe for a given party, safe seats compose 52 percent of the experimental sample, reflecting statistically distinct positive selection by party leaders.²² Parties demonstrate neutral selection for swing seats: the proportion in the experimental sample is not statistically distinct from that nationwide (30 versus 28 percent). This leaves strong negative selection out of weak seats, which constitute

¹⁹Specifically, we regress the constituency-level difference in vote shares for the APC minus the SLPP Parliamentary candidate on the population share of APC-affiliated ethnic groups minus the population share of SLPP-affiliated groups (as displayed in Figure 2). This yields an R^2 of 0.92. The estimated coefficient on the ethnicity-based measure is, as expected, positive, large in magnitude (0.76), and precisely estimated (standard error 0.02). Note that in the regression we include any votes for the candidate from SLPP's splinter party (the PMDC) in the SLPP vote share.

²⁰Caste 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.

²¹As another benchmark, the average constituency in Sierra Leone has the same partisan leaning as the 18th Congressional district in California, which contains Palo Alto. This is based on 2018 US House Elections and 2007 Parliamentary Elections in Sierra Leone (pooling votes for the SLPP splinter party, the PMDC, with the SLPP vote).

²²We double over the map to accommodate both parties simultaneously, and classify competitiveness at the district level (the next higher administrative unit) since constituency boundaries were redrawn for this election.

only 17 percent of the sample, and suggests that parties did not see much value in experimenting with selection where they were likely to lose the general election. This first stage of selective inclusion is important for interpreting our experimental results: they are representative of races where party leaders were willing to experiment, where the modal race is in the party's respective 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.

C. *Status Quo Distortion*

We can use the polling data from control group races to assess how frequently the status quo model of delegating candidate selection to party officials gives voters their most preferred candidates.

The data reveal that this baseline rate is low: in the control group, the two political parties on average selected the aspirant who ranked first among local voters 38 percent of the time, which constitutes a clear and sizable distortion away from voter preferences. In about one-half of these deviations, the party selected the aspirant who ranked second among voters, while in one-third, the party selected someone ranked fourth or lower. This distortion is particularly important in party strongholds, as it means that party officials are not only picking candidates, but effectively future Members of Parliament, who diverge from voter preferences.

IV. Treatment Effects on Representation

The first prediction from the conceptual framework is that the new selection mechanism will reduce the status quo distortion away from voter preferences and thereby enhance representation, as defined as the selected candidate being the aspirant who ranks first among local voters. This section evaluates this hypothesis.

A. *Econometric Specification*

To capture how the new candidate selection mechanism affects representation, we estimate

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

where outcome Y_{ipc} is a variable equal to 100 (so that coefficients are expressed in percentage points) if candidate i that represented party p in the general election for constituency c ranked first among a plurality of voters in the V2 opinion polls (and zero otherwise); treatment indicator T_{pc} signals assignment to the more democratic selection model, which recall was assigned at the party-constituency level; τ_{pc} are fixed effects for 23 party-region strata used in the random assignments; and ε_{ipc} is an idiosyncratic error.²³ The parties held conventions and accepted voter reports in 43 of 46 treatment races.

²³ Selected candidates are those listed in the NEC official candidate registration data.

TABLE 3—EFFECTS OF THE PRIMARY INTERVENTION ON REPRESENTATION

Selected candidate is voters' first choice	(1)	(2)
<i>Panel A. Direct effect on representation</i>		
Primary intervention	22.96 (10.61)	28.66 (16.17)
Party 1		22.30 (39.50)
Party 1 × primary intervention		-11.27 (21.37)
Mean in controls	37.78	
Observations	91	91
<hr/>		
Party's vote share in the general election	(1)	(2)
<i>Panel B. Indirect effect on general election vote shares</i>		
Primary intervention	-0.48 (2.96)	-0.64 (3.40)
Safe seat		27.90 (6.64)
Safe seat × primary intervention		-2.01 (6.13)
Weak seat		-25.95 (6.89)
Weak seat × primary intervention		6.91 (5.54)
Mean in controls	44.98	
Observations	91	91

Notes: (i) This table reports the effects of the primary intervention on representation in the candidate selection stage and downstream effects on the party's vote share in the general election for Parliament; (ii) ordinary least squares regression with robust standard errors; (iii) specifications include fixed effects for 23 party-region strata used in the random assignments; and (iv) one race is missing in panel A because voters were surveyed about aspirants from the wrong party and one race is missing in panel B because the general election was disputed and resolved by the courts.

B. Treatment Effect Estimates

The new selection mechanism has a large positive impact on representation: the estimated treatment effect on selecting the voters' first choice aspirant is 23 percentage points, standard error 10.6 (panel A of Table 3). On the base rate of 38 percent in control races, this effect corresponds to a 61 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 via the conventions and voter reports by picking a different candidate than they would have otherwise for 11 races (i.e., 0.23×46 treatment group races). As 52 percent of the experimental sample is located in ethnic strongholds where the target party is likely to win the general election, this further implies that party officials thereby likely changed the identity of 5 to 6 elected MPs.

Column 2 finds little evidence for heterogeneous treatment effects by party: the coefficient on the interaction between treatment and an indicator to distinguish one

party from the other is not statistically different from zero. As we also find little evidence for heterogeneity by party for other outcomes, we focus on pooled estimates in all following tables. These results are robust to a variety of alternative specifications, including how ties are resolved, using pre-convention V1 data, and accounting for differing numbers of total primary “votes” cast (see discussion in Section VIC and online Appendix Table A8).

C. Heterogeneity in the Representation Effect

The treatment effect on representation is wholly concentrated in safe and weak seats, where competition in the general election was expected to be low. To see this, panel B of Table 2 reports the baseline rate of selecting the voters’ preferred aspirant in each type of race (safe, swing, and weak seats), the difference in these rates across treatment assignment, and the associated p -value that the difference is equal to zero. In safe and weak seats, representation increases by 28 and 50 percentage points, respectively. There is no apparent response to treatment for swing seats, where the likelihood of selecting voters’ preferred aspirant is 50 percent in both treated and control races.

Estimates from an analogous regression, which further includes the randomization strata, are comparable. Online Appendix Table A1 shows that the treatment effect in noncompetitive general election races (safe pooled with weak seats) is a highly significant 33 percentage points (standard error 12). The coefficient on the interaction between treatment and swing seat is an equally sized negative term (−33, standard error 25), however falls below the 90 percent confidence level (p -value = 0.18). In interpreting these results, note that the 50 percent baseline rate (i.e., for control races) in swing seats is substantially higher than that in safe (30 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.

D. Downstream Effects on General Election Vote Shares

Did this representation effect in the candidate selection stage lead to a downstream increase in the party’s general election vote share? We find little evidence that it did: the estimated effect of the new selection method on the general election vote share of the targeted party is small in magnitude and imprecisely estimated (−0.48 percentage points, standard error 2.96 in panel B of Table 3). 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 on average by 28 percentage points. Moving from a swing to weak seat decreases this expected share by 26 points. Both estimates are significant at 99 percent confidence. None of the estimated coefficients on treatment or 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 from Table 2 that there is no “first stage” of the experiment in swing races, as parties were already more likely to pick the local favorite in the status quo and this does not vary with treatment.

The null result further aligns with historical evidence regarding the introduction of direct vote primaries in the United States: while Ware (2002) suggests that incumbent party elites expected an electoral benefit, Hirano and Snyder (2019, p. 47) find no evidence that primaries affected the general election vote share of the advantaged party. This contrasts with contemporary evidence from other regions, where adopting primaries has been shown to boost general election prospects, at least for subgroups like “underdog” parties in Latin American presidential races (Carey and Polga-Hecimovich 2006) and opposition parties in Ghanaian legislative elections (Ichino and Nathan 2013). Similarly, Gulzar, Hai, and Paudel (forthcoming) show that party leaders in Nepal respond to polling data by changing which candidates they put forward onto a party ticket, which in turn increases their general election vote share. None of these latter papers investigate selection on quality.

V. Preference Divergence

What causes the status quo distortion in representation: is it due to a conflict in preferences between voters and party leaders, or something else? While we dedicated a substantial amount of data collection to measuring potential preference divergence, analysis in this section shows instead that voters and party officials broadly agree on the characteristics that make for a good candidate.

A. Data and Econometric Specifications

To gauge preferences, we gathered rich data on aspirants and analyze which specific characteristics predict their popularity among voters and party officials, respectively. We organize aspirant traits (collected in the A1 and A2 surveys) into eight categories: (i) professional qualifications, including 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 three years; (iv) cognitive ability, based on a series of election-oriented questions that involve numeric computations; (v) party loyalty, including history of membership and leadership positions; (vi) public service motivation (PSM), using questions adapted from Perry (1996); (vii) local networks, including membership in constituency-level social and occupational groups; and (viii) campaign effort and expenditure during the primary stage. See online Appendix Table A2 for complete variable list and summary statistics.

We complement these 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 (roughly US\$20

at the time) to cover their travel costs and asked them to verify that the money was correct. Inside each envelope were 18, not 15, 10,000 Leone notes. The measure of conscientiousness is whether the aspirant 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.

Analysis 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:

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

where V_{ipc}^s is the vote share aspirant i who is vying to represent party p in constituency c received among selectors s , where $s \in \{v, o\}$ denotes voters and party officials, respectively; x_{ipc}^k is one of K aspirant traits collected in the data (like education or wealth); τ_{pc} contains the 23 party-region strata from equation (1); and ν_{ipc} is an idiosyncratic error term. To identify which traits predict preferences, we test null hypotheses of the form $\hat{\beta}^{sk} = 0$, and to assess whether voters and party officials have common preferences over traits, we test $\hat{\beta}^{vk} = \hat{\beta}^{ok}$.

As we collected data on a large number of characteristics relative to a modestly sized sample of aspirants ($N = 390$), candidate selection is a high dimensional problem. We respond to this challenge in two ways. First, we 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 (as in Kling, Liebman, and Katz 2007) and enter the indices, along with our single behavioral measure, into the unified regression of equation (2). Second, we use regularized regression methods to select a subset of individual traits with the greatest predictive power (Zou and Hastie 2005). This affords flexibility in searching over all K traits and retaining only those found to be relevant, which is useful since some of the measures collected are likely extraneous in practice, but *ex ante* we do not know which ones these are (see discussion in Belloni, Chernozhukov, and Hansen 2014). To stabilize estimates, we run the regularization technique with k -fold validation 400 times each for voters and party officials, keeping track of which specific traits are selected in each iteration, and carry forward those selected above median frequency. We triangulate results across these two complementary approaches, gaining confidence if the distinct methods produce similar estimates.²⁴

B. Preference Estimates

Voters and party officials appear to value similar traits in aspirants. Starting with the index approach, the two strongest drivers of preferences are the aspirant's record of having provided local public goods and other economic development projects

²⁴This combination of machine learning and traditional econometric methods is useful given that analysis here and in Section VI has elements of both prediction (which traits predict preferences?) and inference (does the new process affect selection on these traits?). See Athey and Imbens (2019) for discussion.

TABLE 4—VOTER AND PARTY OFFICIAL PREFERENCES OVER ASPIRANT CHARACTERISTICS

	Aspirant's share in voter polls (1)	Aspirant's share in party official survey (2)	<i>p</i> -value on (1) versus (2) (3)
Professional qualifications index	0.05 (0.03)	0.11 (0.04)	0.09
Wealth index	0.03 (0.02)	0.03 (0.03)	0.84
Economic development record index	0.04 (0.01)	0.04 (0.02)	0.95
Public service motivation (PSM) index	0.04 (0.02)	0.03 (0.03)	0.84
Party loyalty index	0.01 (0.03)	0.02 (0.06)	0.76
Cognitive ability index	0.01 (0.03)	0.06 (0.04)	0.12
Local network index	0.02 (0.03)	0.04 (0.05)	0.55
Campaign effort and expenditure index	−0.00 (0.02)	−0.01 (0.04)	0.98
Conscientiousness indicator	0.05 (0.02)	0.09 (0.04)	0.17
R^2	0.24	0.17	
Observations	385	367	

Notes: (i) This table uses indices of aspirant characteristics to predict which categories of traits make them popular among voters in the V2 opinion polls (in column 1) and among party officials in the P1 survey (in column 2); (ii) column 3 tests whether voters and party officials have similar preferences over aspirants, by reporting *p*-values that correspond to χ^2 -tests rejecting the equality of coefficients in columns 1 and 2 estimated in a seemingly unrelated regression framework; (iii) robust standard errors clustered by party-constituency; (iv) specifications include fixed effects for 23 party-region randomization strata; and (v) the 8 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.

in the constituency, and their conscientiousness with respect to returning the extra transport allowance. The coefficient on development record in Table 4 is a precisely estimated 0.04 for both voters and party leaders (standard error 0.01 and 0.02, respectively), which implies that a one standard deviation unit increase in public goods provision is associated with a 4 percentage point gain in an aspirant's vote share in the V2 and P1 surveys. Estimates for the conscientiousness indicator suggest that returning the money is associated with a precisely estimated increase of 5 (9) percentage points in vote share among voters (party officials). Wealth and public service motivation also at least marginally predict voter preferences, and while not precise predictors for party officials, the estimated coefficients are positive and comparable in magnitude for both groups. Online Appendix Table A3 presents alternative specifications that limit consideration to control races, stronghold races, or pre-convention V1 data for voters. Development record stands out as the one index that registers positively and significantly for both voters and party officials across all of these specifications.

The one area where there is some evidence of a potential divergence in preferences regards professional qualifications: aspirant qualifications positively and significantly predict party official appraisals (with a coefficient of 0.11 and standard

error 0.04), but do not register among voters at conventional levels (0.05, standard error 0.03). The associated p -value in column 3, which tests for equivalence of coefficients estimated for voters and party officials, rejects equality at the 90 percent confidence level. This is consistent with the human capital advantage party officials have compared to voters in screening on technical merits, and could further reflect differences in their respective information sets, a question we return to in Section VIC.

Reassuringly, the regularization methods produce a similar pattern of results. Online Appendix Table A4 shows that the specific measure of how many development projects the aspirant provided in the constituency (an item in the economic development index) ranks first as the most consistent predictor of both voter and party leader preferences. This trait was selected in 395 iterations for voters and 287 iterations for party officials. Incumbency (an item in the professional qualifications index) is the next most frequently selected trait for voters and ranks highly for party officials as well. While conscientiousness with respect to returning the extra transport allowance registers frequently, it just misses the cutoff of being above median frequency to be carried forward. Five additional traits are retained for party officials, including years of elected office experience, three measures of party loyalty, and one PSM measure.

Table 5 presents the post-regularization prediction results for the union of the seven traits selected for voters or officials. The number of development projects and incumbency are the strongest predictors, with comparably sized positive coefficients for voters and party officials. No other measure strongly predicts preferences for both groups. As such, allowing the data to speak to which specific measures have predictive power identifies items in the same broad areas (development record, professional qualifications) that the index-level approach deemed important, and again shows fairly aligned preferences between voters and party officials.

In light of this alignment, why is it that, in the status quo, party leaders and voters so often disagree on the specific person to select, even though they broadly agree on the traits that make a good candidate? The conceptual framework emphasizes information constraints: both groups have imperfect information on aspirants, so both make errors. For voters, this could reflect challenges in screening on technical merits that leads them away from the most professionally qualified, not because they do not value qualifications, but because they lack the party official's information on competence. And for party officials, while they agree with voters that local alignment is important, they may not be well informed about all the community-level development projects that each aspirant has been involved in. Another explanation is residual conflict in preferences that we are not picking up in the data. The R^2 indicates that these regressions explain between 17 and 24 percent of the total variation in vote shares, suggesting that much of what is driving selections remains unobservable to the econometrician. So it could be the case that idiosyncratic factors break the tie between equally competent and aligned aspirants, and that voters and party officials diverge in how they assess these factors. This seems plausible in the case of education, for example: since 80 percent of aspirants have been to university, in many races all aspirants will be equally competent along this dimension and other factors will break the tie.

TABLE 5—VOTER AND PARTY OFFICIAL PREFERENCES IN POST-REGULARIZATION REGRESSIONS

	Aspirant's share in voter polls (1)	Aspirant's share in party official survey (2)	<i>p</i> -value on (1) versus (2) (3)
Number of development projects	0.03 (0.01)	0.04 (0.01)	0.51
Incumbent MP	0.14 (0.06)	0.13 (0.10)	0.87
Years of public office experience	0.00 (0.00)	0.01 (0.01)	0.34
Party versus own campaign expenditure	−0.02 (0.01)	−0.04 (0.02)	0.43
Number of relatives in party leadership	0.01 (0.01)	0.03 (0.02)	0.41
Number of meetings with party officials	0.00 (0.00)	0.00 (0.00)	0.42
PSM welfare of strangers question	0.00 (0.01)	0.02 (0.01)	0.02
R^2	0.24	0.18	
Observations	385	367	

Notes: (i) This table uses specific aspirant characteristics that were selected via regularization methods to predict which traits make them popular among voters in the V2 opinion polls (in column 1) and among party officials in the P1 survey (in column 2); (ii) these post-regularization estimates retain only those aspirant traits with the greatest predictive power selected via 400 iterations of *k*-fold elastic net procedures (see online Appendix Table A4 for details); (iii) column 3 tests whether voters and party officials have similar preferences over aspirants, by reporting *p*-values that correspond to χ^2 -tests rejecting the equality of coefficients in columns 1 and 2 estimated in a seemingly unrelated regression framework; (iv) robust standard errors clustered by party-constituency; (v) specifications include 23 party-region randomization strata; (vi) all traits expressed in natural units; (vii) *party versus own expenditure* indicates an affirmative response to the question “Are you willing to spend more money on your party’s campaign versus your own?”; and (viii) the *PSM welfare of strangers question* indicates strength of disagreement with the statement “I seldom think about the welfare of people whom I don’t know personally” with missing values imputed at the control group mean.

C. Information Constraints

A key part of the information story is what party officials know about voter preferences. If they know very little, then even well-aligned, well-intentioned party officials will frequently fail to select the most locally popular aspirants. The primary intervention was designed to alleviate this constraint, and the data suggest that it is indeed important in the status quo.

Sierra Leone is a poor country with high transport and communication costs, which create pervasive information asymmetries throughout the electoral process. There is no large-scale, cost-effective polling technology accessible in this market, implying that party officials are constrained in attempting to elicit and aggregate voter preferences over 132 local pools of aspirants. Quite simply, if there is no primary, party officials may have little idea which aspirant local voters prefer.

To gauge the empirical relevance of this constraint, 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 52 percent of presumed shared preferences were in fact a match with the polling data on voter preferences (from the V2 survey). Overall, party officials correctly guessed who would win a local primary only 50 percent of the time. This disconnect is severe: in 33 percent of races, no party leader (among multiple surveyed per race) correctly guessed which aspirant was the most popular with local voters.

This information asymmetry is not limited to poor countries. Folke, Persson, and Rickne (2016) study open-list proportional representation systems in Sweden and Brazil. Drawing an analogy to primaries, they argue that preference vote tallies reveal information about popularity to party leaders, who are then more likely to promote a first versus second place finisher of equal ability.²⁵ A natural next question is what this implies for selection: namely, when party officials accommodate voter preferences, and switch to more popular candidates, does it affect the quality of those chosen? The impact could be negative, if voters are poorly informed and less able to identify high performers. Or it could be positive, if voters possess complementary information on traits that contribute to performance (e.g., about local alignment from the conceptual framework).

VI. Impacts for Selection on Quality

Assessing whether the documented increase in representation affects selection on quality requires comparing the characteristics of candidates chosen via the new versus status quo mechanisms.

A. Econometric Specifications

To test 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:

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

where x_{pc}^k is a characteristic (e.g., education or wealth, as in equation (2)) 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.

We link this estimation directly to the two-pronged approach used to explore preferences in Section V. First, we carry forward all eight indices of aspirant traits and the behavioral measure used previously. For each one, we estimate whether candidates selected under the new mechanism are more or less likely to be strong

²⁵By honing in on the discontinuity between close winners and runners up, their empirical strategy by design identifies a preference for popularity holding candidate quality fixed (in expectation).

TABLE 6—EFFECTS OF THE PRIMARY INTERVENTION ON CANDIDATE SELECTION

	Primary intervention effect (1)	Standard error (2)	Naïve <i>p</i> -value (3)	FDR <i>q</i> -value (4)
<i>Panel A. Primary effects on indices of candidate traits</i>				
Personal qualifications index	−0.01	(0.10)	0.92	0.99
Wealth index	−0.13	(0.13)	0.32	0.51
Economic development index	0.29	(0.14)	0.04	0.22
Public service motivation index	−0.02	(0.15)	0.89	0.99
Party loyalty index	0.03	(0.09)	0.74	0.99
Cognitive ability index	−0.17	(0.12)	0.15	0.36
Local networks index	0.19	(0.09)	0.04	0.22
Campaign expenditure index	−0.08	(0.12)	0.52	0.77
Conscientiousness indicator	0.15	(0.10)	0.13	0.36
<i>Panel B. Primary effects on candidate traits identified by regularization methods</i>				
Number of development projects	0.46	(0.21)	0.03	0.27
Incumbent MP	−0.09	(0.07)	0.23	0.62
Years of public office experience	0.76	(1.34)	0.57	0.99
Party versus own campaign expenditure	−0.09	(0.19)	0.66	0.99
Number of relatives in party leadership	0.15	(0.30)	0.61	0.99
Number of meetings with party officials	−11.33	(5.92)	0.06	0.27
PSM welfare of strangers question	−0.02	(0.33)	0.95	0.99
Observations	92			

Notes: (i) This table explores whether the primary intervention selected candidates with different characteristics to proceed to the general election; (ii) panel A reports treatment effect estimates for indices of traits and panel B reports effects for the 7 specific traits selected via regularized regression methods; (iii) each row reports results from a separate OLS regression with robust standard errors that includes fixed effects for 23 party-region randomization strata; (iv) the eight indices in panel A are equally weighted sums of underlying traits expressed in standard deviation units (following Kling, Liebman, and Katz 2007), conscientiousness is a binary indicator, and all traits in panel B are in natural units; (v) *party versus own expenditure* indicates an affirmative response to the question “Are you willing to spend more money on your party’s campaign versus your own?” (vi) the *PSM welfare of strangers question* indicates strength of disagreement with the statement “I seldom think about the welfare of people whom I don’t know personally,” with missing values imputed at the control group mean; 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).

along that dimension. As this requires nine distinct iterations of equation (3), we adjust standard errors on these estimates to control the false discovery rate (FDR) across the nine regressions (following Benjamini, Krieger, and Yekutieli 2006 and Anderson 2008). Second, we carry forward the union of seven specific traits identified as important under the regularization approach. We again estimate treatment effects for each one, and implement FDR adjustments. A caveat for both approaches is that statistical power to support such multiple inference adjustments is strained, as at this stage, estimation operates over the pool of selected candidates, where there are only 92 observations.

B. Treatment Effects on Selection

Estimates suggest that candidates selected under the more democratic mechanism look somewhat stronger on their observable characteristics as compared to those chosen via the status quo method. Candidates selected by the new process on average have stronger records of having previously provided public goods in the constituency (in panel A of Table 6). The estimated treatment effect for the

economic development index is 0.29 standard deviation units (standard error 0.14), which is a materially large and highly significant effect on a naïve, or per comparison, basis (p -value = 0.04). There is also a positive treatment effect estimate for conscientiousness, of 15 percentage points, but it does not quite reach significance at conventional levels (p -value = 0.13). Positive effects for these particular traits are noteworthy in light of the fact that they are strong drivers of voter preferences (in Table 4). There is further a positive and significant effect on selecting aspirants with stronger local networks (e.g., membership in local groups), which intuitively aligns with discussion in the conceptual framework, but is empirically unexpected since it did not rank as consistently as a predictor of voter preferences.


Implementing FDR adjustments over all nine index-level regressions sends these estimates below standard significance levels: the corresponding q -values for the two most precise estimates inflate to 0.22. This is perhaps not surprising given the strains on statistical power. While not a main focus, note further that we do not find treatment effects on demographics: 85 percent of selected candidates are male, average age 46, neither of which varies significantly with treatment assignment (online Appendix Table A5).

Encouragingly, the regularization approach again provides similar results. Panel B of Table 6 estimates treatment effects on the union of traits that were identified as being predictive of voter or party official preference rankings over aspirants. The most striking estimate is the large positive effect on the number of development projects (0.46 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 one-half an additional local public good or other development project in the past three years. Given the control group mean of 2.07 projects, this effect constitutes a 22 percent increase in such provision. There is a marginally significant negative impact on the number of meetings with party officials, which perhaps suggests a move away from the most loyal members. In column 4, the FDR adjustments again reduce significance levels, with an adjusted q -value for the number of development projects equal to 0.27.

As the estimated treatment effect on economic development record is a key result, it is worth providing additional context from the data. Overall, 82 percent of aspirants have been involved in at least one project. Enumerators collected detailed data (project location, time-frame, expenditure, sources of funds, status of completion, etc.) for up to three most relevant projects. In this project dataset, 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 health care (clinic construction and rehabilitation).

Does incumbency drive these public goods results? While incumbents have provided more projects than others (a statistically significant difference of 2.5 versus 1.9 projects on average), incumbency does not account for these findings. Specifically, all estimates that show how an aspirant's economic development record predicts voter preferences in Tables 4 and 5, 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 6. Note further that incumbency is much less

Panel A. Positive effect of the primary intervention on representation



Panel B. Positive effect of the primary intervention for selection on quality

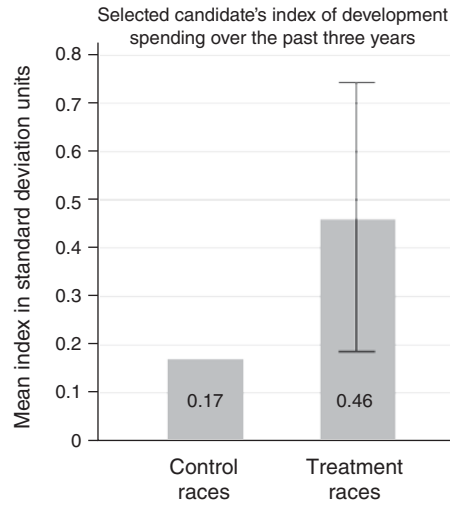


FIGURE 3. OVERVIEW OF MAIN RESULTS

Notes: This figure presents treatment effect estimates of the primary intervention on key outcomes of interest, where the error bars correspond to the 95 percent confidence interval as estimated via ordinary least squares regression that includes the 23 party-region randomization strata and robust standard errors. In panel B, the development spending index compiles multiple measures of a candidate's provision of local public goods over the previous three years, which is normalized with respect to all aspirants in control group races and expressed in standard deviation units.

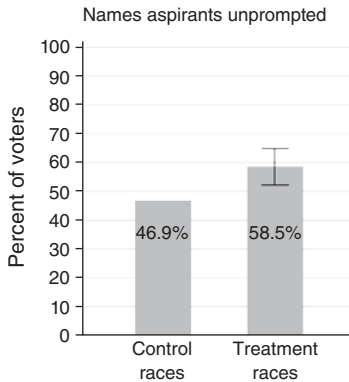
entrenched in Sierra Leone than in the United States. There were 25 incumbents seeking re-nomination in our sample, which covers 27 percent of the races (reassuringly incumbency is balanced across party and treatment assignment). Roughly one-half of them secured re-nomination from their party. Compare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

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 method increases 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 less precision. To the extent that past public goods provision predicts future provision, a question we begin to explore in Section VID, these are guardedly optimistic results. Figures 3 and 4 summarize main results and mechanisms.

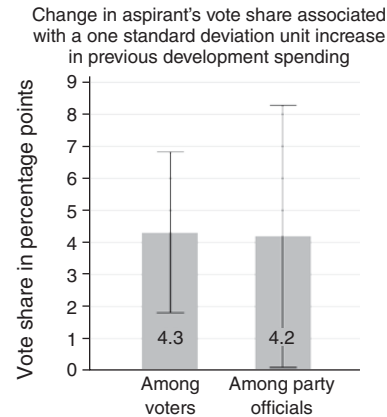
C. Voter Learning about Aspirants

The results of this experiment do not bear out concerns of primary skeptics that giving more say to poorly informed voters increases representation at the cost of worse screening on quality. This downside risk appears to have in part been mitigated by voter learning about aspirants from the town hall conventions and associated radio broadcasts.

Panel A. Positive effect of the primary intervention on voter knowledge



Panel B. Voters and party officials have aligned preferences over aspirant traits



Panel C. Null effect of the primary intervention on aspirant entry

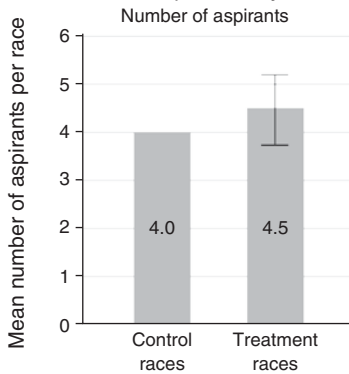


FIGURE 4. MECHANISMS TESTED

Notes: This figure explores mechanisms, where the error bars correspond to the 95 percent confidence interval as estimated via ordinary least squares regression that includes the 23 party-region randomization strata and robust standard errors (clustered by race for panels A and B). In panel B, the development spending index compiles multiple measures of a candidate's provision of local public goods over the previous three years, which is normalized with respect to all aspirants in control group races and expressed in standard deviation units. The specification in panel B further includes the eight other indices of related aspirant traits and the binary measure of conscientiousness (analogous to the specification in Table 4).

To substantiate this channel, Figure 4, panel A shows that voters in treated races are more likely to be able to name aspirants unprompted by 11.6 percentage points in the V2 survey, constituting a 25 percent increase on the base rate of 47 percent in control races, which is statistically significant at the 99 percent confidence level. We further find a precisely estimated null effect of treatment on naming aspirants in the V1 data, which was collected prior to the conventions and broadcasts, supporting the hypothesis that the observed difference in knowledge is a result of the primary intervention (see online Appendix Table A6). The positive learning effect is estimated with comparable magnitude across more and less competitive races, however, constitutes a larger percentage increase in weak and swing seats where baseline voter knowledge was lower (in online Appendix Table A6, panel B).

We have more speculative evidence that this general learning effect enhanced voter ability to identify the most qualified aspirants on specific metrics (in online Appendix Table A7). Due to an inconsistently applied skip pattern that created systematically more missing values for control races, these estimates are limited to strata where this skip pattern was not imposed and nearly all respondents were asked questions about which aspirant in the local pool was most qualified. In this subset, voters in treatment races are more likely to accurately identify the aspirant with the most professional experience by 18.9 percentage points (standard error 8.3) and the strongest record of development spending by 13.5 points (standard error 9.4). The effect on identifying the most educated is null. A joint test that all three estimates equal zero rejects at 99 percent confidence.

One consequence of this increase in knowledge and party activity around the conventions is that voters become more likely to form and express a preference over aspirants, which means that there are more primary “votes” cast in the treatment group (online Appendix Table A6, panel C). These increases in knowledge and voting are themselves expressions of enhanced representation. Yet this imbalance does raise a question about interpretation for our main measure of representation, namely whether the selected candidate is the one who ranked first among voters, as econometrically we may be more likely to make mistakes about who ranked first in control races. To ensure this imbalance is not driving our results, we show that our estimates for primary impacts on representation are robust to (i) limiting the sample to strongholds, where voter knowledge and engagement with the selection process was high in all races (online Appendix Table A8, panel A); (ii) setting a minimum total vote threshold, where we drop all races in a randomization strata where any single race had fewer votes than the threshold, and vary the threshold in increments of ten from at least 25 to more than 75 votes (online Appendix Table A8, panel B); and (iii) taking an adversarial approach (in the spirit of Manski bounds)²⁶ by recoding the representation outcome to voters getting their first choice in all control races with fewer votes than the smallest total in treated races (which is 32 votes), which generates a treatment effect estimate of 14.2 (standard error 10.5).

We can further rule out that treated voters are simply acting as party stooges who learn from the conventions which aspirant party leaders endorse and adopt their point of view. To assess this, we estimate the treatment effect on representation using pre-convention V1 polling data, which effectively shuts down the voter learning channel and fixes their preferences at baseline. We again find a positive and highly significant effect of the primary initiative on the likelihood that voters get their most preferred candidates, equal to 31 percentage points (online Appendix Table A8, panel A).

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, Casey, and Glennerster 2020), and resonates with evidence from Uganda that debates build voter knowledge in both the primary and general election stages (Platas and Raffler 2019). It also aligns with survey evidence from the United States showing

²⁶See Horowitz and Manski (2000).

that voters learn about candidate policy positions and move into stronger ideological congruence as the primary campaign progresses (Hirano et al. 2014).

D. Treatment Effects on the Performance of Elected MPs in Strongholds

This section narrows consideration to stronghold races, where selected candidates were very likely to become elected MPs, and traces the effects of the primary intervention on their subsequent performance in office. We first discuss how strongholds differ from the other races in the sample, show that our headline results hold for this subset of races, and then present treatment effect estimates for performance in office. Unfortunately, limited statistical power renders these latter tests inconclusive.

All actors—party officials, aspirants, and voters—engaged more intensively with the candidate selection process in stronghold races. As we have seen, these were the races where party leaders were most interested in experimenting with the new mechanism and where the stakes were highest, as selecting the candidate *de facto* means selecting the winner. These races also attracted more aspirants, who were of higher quality on average, than those in other areas. Specifically, online Appendix Table A9 shows that the mean number of aspirants increases from 2 in weak seats, to 4 in swing seats, to 5 in stronghold races. Aspirants in strongholds had completed an additional year of education (15.6 versus 14.7 in weak seats), were more likely to have been to university (86 versus 62 percent), and owned more assets (9.7 versus 8.2). Voters in turn were more likely to know who the aspirants were and have an opinion about which one should be selected: focusing on control races, voters were more than twice as likely to be able to name aspirants unprompted (59 percent in strong versus 23 in weak seats) and to report which one they would select (71 versus 34 “votes” cast per control race in strong versus weak seats, see online Appendix Table A6).

In light of these differences, it is informative to show that our headline results hold for this subsample of high engagement races. First, the treatment effect of the primary intervention on voters getting their most preferred candidate is 27.4 percentage points (standard error 14.7, in online Appendix Table A8). This is reassuring, as increased representation in choosing candidates has less impact in weak seats where few voters know or care about who wins the primary.²⁷ Second, voters in stronghold races value the same traits as voters elsewhere. Online Appendix Table A3 shows that stronghold voters prefer aspirants with better economic development records, and interestingly, further share party officials’ preference for those with more professional qualifications. And third, the primary intervention similarly facilitated the selection of candidates with stronger local networks, conscientiousness and economic development records, although the latter estimates are less precise: the treatment effect for the development index is 0.24 standard deviation units

²⁷This is equally true of primaries in high-income countries. To provide a US analogue, in the conservative 11th district of the House of Representatives in Texas, which has open primaries meaning that voters from any party can participate, there were nearly nine times as many votes cast in the Republican as compared to the Democratic primary. The Republican candidate went on to win the 2018 general election with 80 percent of the vote. See “2018 Primary Election Results,” *Texas Tribune*, <https://apps.texastribune.org/elections/2018/primary-election-results/>.

(standard error 0.18) and for the number of projects is 0.38 (standard error 0.28) in online Appendix Table A10.

To test whether the primary intervention had effects on subsequent performance in office, we contacted elected MPs from sampled stronghold races 18 months after the general election. We are interested in their utilization of the constituency facilitation fund, which is public money given to MPs to support development projects in their home constituency, establishment of a local office, and constituent outreach activities. We focus on three measures: (i) their total expenditure on development projects, as verified by field audits; (ii) the existence of a functional office accessible to citizens, verified by audit; and (iii) how many public meetings they held with citizens in their constituency, as reported by four key informants. Statistical power is constrained by two factors: this subsample, by design, is limited to the 48 winners from stronghold races; and disruption caused by the COVID19 pandemic shut down data collection before it was complete, reducing the effective sample by 13 percent.

Treatment effect estimates for all three performance outcomes are noisy and not statistically distinguishable from zero (in online Appendix Table A11). It is nonetheless worth noting that the estimated effect on development spending is positive in sign and large in magnitude (at 32 M Leones), which corresponds to a 23 percent increase over the verified spending of MPs from control races. As low power renders this analysis inconclusive, we see robustly measuring potential effects of the primary selection stage on performance as an important area for future research.

VII. Alternative Explanations

This section investigates two alternative channels, aspirant entry and financial contributions, and finds little evidence that either explain much of the observed variation in the data.

A. Aspirant Entry

Differential aspirant entry into treated races is unlikely to explain these results, largely because the advance announcement needed to induce entry was only partially implemented. Recall that only one of the two parties announced in advance which constituencies would participate in the new selection process. 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 delegates' convention to nominate their presidential candidate. Their promotional materials, which dub the initiative "Aspirant Voice and People's Choice," describe the two components and characterize them as a "pilot" designed to "strengthen the internal democracy of our party" (see flyer in online Appendix Figure A2). This could have altered the entry decision of potential SLPP aspirants. The APC, 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.

To nonetheless test whether the aspirant entry channel is operational in this experiment, online Appendix Table A12 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 (0.27, standard error 0.62), as one would expect in the absence of advance notification. These data reflect the number of aspirants who were under consideration at the end of the process when candidate selections were made.

We can compare this to the number of aspirants who were surveyed earlier in the process as part of the research, which captures those who were under consideration by the party at the time of the conventions. This could be a larger number if some aspirants dropped out or were disqualified in the interim. Estimates for the APC are comparable. 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. Online Appendix Table A13 reinforces this interpretation by showing that average characteristics of the aspirant pool are largely balanced across treatment assignment. Notice, for example, that the one index that is not balanced (specifically, mean scores on the cognitive ability questions were somewhat lower for aspirants in treated races) is not an index where we see positive treatment effects of the primary intervention on selected candidates. 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. We leave this question for future research.

B. Financial Contributions

Another factor that could influence candidate selection, both in the status quo and in response to treatment, is financial contributions to secure the nomination. To measure 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.²⁸ Note that these are distinct from the aspirant's own campaign expenses, which were recorded separately. Enumerators further asked unsuccessful aspirants how much they thought the selected candidate in their pool had contributed.

Self-reports from control group races suggest that nearly all aspirants (89.7 percent) made nonzero contributions, with an unadjusted mean converted to US dollars of \$2,392, which is equivalent to more than one month's official salary of an elected MP. This looks modest compared to the 17 months of salary that public health-care workers paid to secure a promotion in another low-income country (Weaver

²⁸The 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."

2018), but lies well out of reach for ordinary Sierra Leoneans, where for instance the monthly minimum wage is \$71. These official fees and other contributions thus constitute a substantial barrier to entry into politics.

Two of the more nefarious interpretations of these contributions are not consistent with the empirical patterns we find. First, the data do not suggest that candidacy is being sold to the highest bidder in the status quo. Selected candidates self-reported the highest contribution in their pool in only 36 percent of control races. While this top payer rate increases to 51 percent when substituting out low candidate self-reports with the average report of unsuccessful aspirants (about how much they think the selected candidate contributed), it remains well below 100 percent. Second, in treated races, the data are not consistent with contributions being used to compensate for low popularity levels, i.e., to buy off party officials so that they deviate from the voter reports. Specifically, selected candidates do not contribute differentially more than unsuccessful aspirants in treatment races where the party ultimately chose someone other than the voters' first choice compared to races where party officials went with the voters' choice (results not shown).

An alternative way to conceptualize these payments is with respect to the expected return to candidacy, which is increasing in three multiplicative components: the returns to office, the probability the party wins the seat in the general election, and the likelihood of being selected as candidate. The first component includes 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 returns 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 (see Colonnelli, Prem, and Teso 2020 for estimates in Brazil). In light of this, both major parties were able to recruit aspirants in all 132 Parliamentary races nationwide.

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 contribute should be higher in stronghold races, which is exactly what we see in the control group data: mean contributions increase from \$1,151 in weak seats, to \$1,527 in swing seats, to \$2,930 in safe seats (online Appendix Table A9). We can link this increase in the expected return back to status quo entry decisions, where Section VID shows that there are more aspirants, who are of higher average quality, in stronghold races, which is consistent with findings from elsewhere (Ferraz and Finan 2011, Gagliarducci and Nannicini 2013, Dal Bó et al. 2017, Hirano and Snyder 2019).²⁹

Regarding the third factor, willingness to contribute 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. For low types, the opposite should hold. A

²⁹ While interesting, note that this does not affect our experimental estimates as treatment assignment is stratified within small geographic areas and we include fixed effects for these strata in all relevant specifications.

simple way to capture this in the data is to test whether selected candidates contribute differentially more than unsuccessful aspirants in treated versus control races.

Online Appendix Table A14 tests this hypothesis by regressing 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 spread between how much successful and unsuccessful aspirants contributed is indeed wider in treatment races. This spread is more pronounced in column 4, which restricts attention to stronghold races. These estimates suggest that unsuccessful aspirants in treated races contributed less (by \$819 on average) and selected candidates contributed more (by \$2,256) than their counterparts in control races.

Overall, these results are consistent with aspirant willingness to pay being tied to the expected return to candidacy, and thus for those whose expected return increases with treatment, because they are more likely to be selected under the new process, their contribution increases as well. The offsetting reduction by aspirants who become less likely to be selected generates no net effect of treatment on mean contributions for the aspirant pool overall (online Appendix Table A14, column 1). This is one speculative interpretation of the data, and there may well be others that align with this constellation of results.

VIII. 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 15 years (UNDP 2019). The efficacy of such investments in delivering representative and competent elected politicians 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 discussions about the design of donor support for elections: to illustrate, as late as 2004 there was no explicit reference to political parties anywhere in the UNDP's own multi-year funding framework (UNDP 2006, p. 11). This neglect arises in part from concerns about impartiality that lead donors to shy away from direct engagement with parties.

The two major political parties in Sierra Leone demonstrate that there are practical ways to improve candidate selection. An alternative approach to selection 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. By revealed preference, the fact that both parties participated and put forward their stronghold races for inclusion suggests that they saw value in the initiative. Their openness and responsiveness to the new mechanism further attests to their willingness to experiment at the frontier of democratic practice.

The experiment demonstrates that the primary selection stage plays an important role in reducing information asymmetries. Primary conventions provide information about potential candidates to both voters and party officials, and opinion polling informs party officials about voter preferences. The finding that party officials

responded to the information relayed by selecting different types of candidates, and that these candidates have stronger records of local public goods provision, is a cautiously optimistic result.

In thinking about external validity, a few factors about the context and intervention are worth emphasizing. A concern about the US primary system is that it leads to the selection of candidates who are more ideologically extreme than the general population. This is less of an issue in our context as parties in Sierra Leone are largely non-ideological, as are about one-half of all parties worldwide (Cruz, Keefer, and Scartascini 2016). Moreover, in this study, preferences over candidates were gathered from a representative sample of voters via home visits. This again reduces the risk that those with more extreme views, or particular socioeconomic groups, wield disproportionate influence. In addition, weak advance announcement, combined with the nonbinding nature of the treatment, dampened incentives for new or different types of candidates to come forward. This effectively closes an entry channel that is likely to be consequential in other settings. While these factors help us isolate the role of information revelation in the candidate selection stage, they limit our ability to extrapolate from these results to what might happen when a well-publicized, binding, direct vote primary is introduced. In these differences lies a broad and rich agenda for future research.

REFERENCES

- Adams, James, and Samuel Merrill, III.** 2008. "Candidate and Party Strategies in Two-Stage Elections Beginning with a Primary." *American Journal of Political Science* 52 (2): 344–59.
- Anderson, Michael L.** 2008. "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American Statistical Association* 103 (484): 1481–95.
- Athey, Susan, and Guido W. Imbens.** 2019. "Machine Learning Methods that Economists Should Know about." *Annual Review of Economics* 11: 685–725.
- Banerjee, Abhijit V., and Rohini Pande.** 2009. "Parochial Politics: Ethnic Preferences and Politician Corruption." Unpublished.
- Banks, Jeffrey S., and Rangarajan K. Sundaram.** 1993. "Adverse Selection and Moral Hazard in a Repeated Elections Model." In *Political Economy: Institutions, Competition, and Representation*, edited by William A. Barnett, Melvin J. Hinich, and Norman J. Schofield, 295–311. Cambridge: Cambridge University Press.
- 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, Yoav, Abba M. Krieger, and Daniel Yekutieli.** 2006. "Adaptive Linear Step-Up Procedures that Control the False Discovery Rate." *Biometrika* 93 (3): 491–507.
- Besley, Timothy.** 2005. "Political Selection." *Journal of Economic Perspectives* 19 (3): 43–60.
- Besley, Timothy, and Stephen Coate.** 1997. "An Economic Model of Representative Democracy." *Quarterly Journal of Economics* 112 (1): 85–114.
- Besley, Timothy, Olle Folke, Torsten Persson, and Johanna Rickne.** 2017. "Gender Quotas and the Crisis of the Mediocre Man: Theory and Evidence from Sweden." *American Economic Review* 107 (8): 2204–42.
- Bidwell, Kelly, Katherine Casey, and Rachel Glennerster.** 2020. "Debates: Voting and Expenditure Responses to Political Communication." *Journal of Political Economy* 128 (8): 2880–924.
- Bochel, John, and David Denver.** 1983. "Candidate Selection in the Labour Party: What the Selectors Seek." *British Journal of Political Science* 13 (1): 45–69.
- Brollo, Fernanda, Tommaso Nannicini, Roberto Perotti, and Guido Tabellini.** 2013. "The Political Resource Curse." *American Economic Review* 103 (5): 1759–96.
- Brockman, David E., Nicholas Carnes, Melody Crowder-Meyer, and Christopher Skovron.** 2019. "Why Local Party Leaders Don't Support Nominating Centrists." *British Journal of Political Science*: 1–26.

- Butler, Daniel M., and David W. Nickerson.** 2011. "Can Learning Constituency Opinion Affect How Legislators Vote? Results from a Field Experiment." *Quarterly Journal of Political Science* 6 (1): 55–83.
- Caillaud, Bernard, and Jean Tirole.** 2002. "Parties as Political Intermediaries." *Quarterly Journal of Economics* 117 (4): 1453–89.
- Carey, John M., and John Polga-Hecimovich.** 2006. "Primary Elections and Candidate Strength in Latin America." *Journal of Politics* 68 (3): 530–43.
- Casey, Katherine.** 2015. "Crossing Party Lines: The Effects of Information on Redistributive Politics." *American Economic Review* 105 (8): 2410–48.
- Casey, Katherine, and Rachel Glennerster.** 2017. "The Impact of Primary Conventions." AEA RCT Registry. November 26. <https://doi.org/10.1257/rct.2485-1.0>.
- Casey, Katherine, Abou Bakarr Kamara, and Niccoló F. Meriggi.** 2021. "Replication Data for: An Experiment in Candidate Selection." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E124501V1>.
- Colonnelli, Emanuele, Mounu Prem, and Edoardo Teso.** 2020. "Patronage and Selection in Public Sector Organizations." *American Economic Review* 110 (10): 3071–99.
- Cruz, Cesi, and Philip Keefer.** 2015. "Political Parties, Clientelism, and Bureaucratic Reform." *Comparative Political Studies* 48 (14): 1942–73.
- Cruz, Cesi, Philip Keefer, and Carlos Scartascini.** 2016. "Database of Political Institutions Codebook, 2015 Update (DPI2015)." Inter-American Development Bank.
- Dal Bó, Ernesto, and Frederico Finan.** 2018. "Progress and Perspectives in the Study of Political Selection." *Annual Review of Economics* 10: 541–75.
- Dal Bó, Ernesto, Frederico Finan, Olle Folke, Torsten Persson, and Johanna Rickne.** 2017. "Who Becomes a Politician?" *Quarterly Journal of Economics* 132 (4): 1877–914.
- De Paola, Maria, and Vincenzo Scoppa.** 2011. "Political Competition and Politician Quality: Evidence from Italian Municipalities." *Public Choice* 148 (3–4): 547–59.
- Fearon, James D.** 1999. "Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance." In *Democracy, Accountability, and Representation*, edited by Adam Przeworski, Susan C. Stokes, and Bernard Manin, 55–97. Cambridge: Cambridge University Press.
- Ferraz, Claudio, and Frederico Finan.** 2011. "Motivating Politicians: The Impacts of Monetary Incentives on Quality and Performance." NBER Working Paper 14906.
- Fisman, Raymond, Nikolaj A. Harmon, Emir Kamenica, and Inger Munk.** 2015. "Labor Supply of Politicians." *Journal of the European Economic Association* 13 (5): 871–905.
- Folke, Olle, Torsten Persson, and Johanna Rickne.** 2016. "The Primary Effect: Preference Votes and Political Promotions." *American Political Science Review* 110 (3): 559–78.
- Gagliarducci, Stefano, and Tommaso Nannicini.** 2013. "Do Better Paid Politicians Perform Better? Disentangling Incentives from Selection." *Journal of the European Economic Association* 11 (2): 369–98.
- Galasso, Vincenzo, and Tommaso Nannicini.** 2011. "Competing on Good Politicians." *American Political Science Review* 105 (1): 79–99.
- Gallagher, Michael, and Michael March, eds.** 1988. *Candidate Selection in Comparative Perspective: The Secret Garden of Politics*. London: Sage Publications.
- Gulzar, Saad, Zuhad Hai, and Binod Kumar Paudel.** Forthcoming. "Information, Candidate Selection, and the Quality of Representation: Evidence from Nepal." *Journal of Politics*.
- Hirano, Shigeo, Gabriel S. Lenz, Maksim Pinkovskiy, and James M. Snyder, Jr.** 2014. "Voter Learning in State Primary Elections." *American Journal of Political Science* 59 (1): 91–108.
- Hirano, Shigeo, and James M. Snyder, Jr.** 2014. "Primary Elections and the Quality of Elected Officials." *Quarterly Journal of Political Science* 9 (4): 473–500.
- Hirano, Shigeo, and James M. Snyder, Jr.** 2019. *Primary Elections in the United States*. Cambridge: Cambridge University Press.
- Horowitz, Joel L., and Charles F. Manski.** 2000. "Nonparametric Analysis of Randomized Experiments with Missing Covariate and Outcome Data." *Journal of the American Statistical Association* 95 (449): 77–84.
- Ichino, Nahomi, and Noah L. Nathan.** 2013. "Do Primaries Improve Electoral Performance? Clientelism and Intra-Party Conflict in Ghana." *American Journal of Political Science* 57 (2): 428–41.
- Kandeh, Jimmy D.** 1992. "Politicization of Ethnic Identities in Sierra Leone." *African Studies Review* 35 (1): 81–99.
- Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz.** 2007. "Experimental Analysis of Neighborhood Effects." *Econometrica* 75 (1): 83–119.

- Madison, James.** 1788 (2014). "Federalist Paper No. 57." In *The Federalist Papers*, edited by Jim Miller, 279–83. Mineola, NY: Dover Publications.
- Marken, Stephanie.** 2018. "Still Listening: The State of Telephone Surveys." <https://news.gallup.com/opinion/methodology/225143/listening-state-telephone-surveys.aspx?version=print>.
- Mattozzi, Andrea, and Antonio Merlo.** 2015. "Mediocracy." *Journal of Public Economics* 130: 32–44.
- Munshi, Kaivan, and Mark Rosenzweig.** 2016. "Insiders and Outsiders: Local Ethnic Politics and Public Good Provision." NBER Working Paper 21720.
- National Electoral Commission.** 2007. "Parliamentary Boundary Delimitation Data."
- Osborne, Martin J., and Al Slivinski.** 1996. "A Model of Political Competition with Citizen-Candidates." *Quarterly Journal of Economics* 111 (1): 65–96.
- Perry, James L.** 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.** 2019. "Candidate Videos and Vote Choice in Ugandan Parliamentary Elections." In *Information, Accountability, and Cumulative Learning: Lessons from Metaketa I*, edited by Thad Dunning et al., 156–87. Cambridge: Cambridge University Press.
- Serra, Gilles.** 2011. "Why Primaries? The Party's Tradeoff between Policy and Valence." *Journal of Theoretical Politics* 23 (1): 21–51.
- Statistics Sierra Leone.** 2004. "Population and Housing Census."
- Ting, Michael M., James M. Snyder, Jr., and Shigeo Hirano.** 2018. "Primary Elections and the Provision of Public Goods." *Journal of Politics* 80 (2): 647–61.
- United Nations Development Programme.** 2006. "A Handbook on Working with Political Parties," New York.
- United Nations Development Programme.** 2019. *Funding Analysis of UNDP Electoral Assistance 2004–2018*. Internal memorandum.
- Ware, Alan.** 2002. *The American Direct Primary: Party Institutionalization and Transformation in the North*. Cambridge: Cambridge University Press.
- Weaver, Jeffrey.** 2018. "Jobs for Sale: Bribery and Misallocation in Hiring." Unpublished.
- Zou, Hui, and Trevor Hastie.** 2005. "Regularization and Variable Selection via the Elastic Net." *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 67 (2): 301–20.