Representation Failure*

Matias Iaryczower Galileu Kim Sergio Montero†

March 15, 2024

Abstract

Democratic representation is constrained by the alternatives available to voters. We use rich data on thousands of candidates in three Brazilian legislative elections to (i) quantify the relative value voters place on candidates' policy positions and non-ideological attributes and (ii) assess the extent to which the "supply side" of politics hinders voter welfare. We find that the average voter suffers only a moderate loss due to policy incongruence but a large loss due to deficiencies in candidates' non-ideological characteristics. For the typical municipality, the welfare loss attributable to non-ideological factors is more than seven times larger than the policy welfare loss. To evaluate the consequences of potential institutional reforms, we develop and estimate a model of equilibrium policy determination. We show that institutional reforms aimed at improving the quality of representation may have sizable unintended welfare consequences due to equilibrium policy adjustments.

Keywords: elections, ideology, valence, voter welfare, representation

*We thank Nathan Canen, Darin Christensen, Anderson Frey, Alan Jacobs, Holger Sieg, Tara Slough, Yang-Yang Zhou, and audiences at Berkeley, the Comparative Politics and Formal Theory Conference, the Empirical Models of Political Economy Conference, EPSA, Houston, LSE, Princeton, Rochester, Tepper, Warwick, Waterloo, and Yale for helpful comments.

†Matias Iaryczower: Department of Politics, Princeton University, email: miaryc@princeton.edu; Galileu Kim: World Bank, email: galileukim@worldbank.org; Sergio Montero: Departments of Political Science and Economics, University of Rochester, email: smontero@rochester.edu.

1 Introduction

In recent years, voters in democracies around the world have expressed discontent with the entire political system. From large public demonstrations to overwhelming disapproval in opinion polls, large fractions of voters seem dissatisfied with all the alternatives available to them. Such a systemic representation failure can erode trust in democratic institutions, paving the way for authoritarian attempts.

The perceived apathy among voters raises key questions about the workings of representative democracies. How severe are actual representation failures in these political systems? What exactly is failing? In particular, do the candidates available to voters fail to advocate for their preferred policies, or do they lack competence, honesty, or other non-ideological attributes that voters value? Would institutional reforms designed to alter the pool of candidates or the strength of political parties substantially improve the value voters give to elections?

In this paper, we address these questions in the context of elections for the lower house of Brazil's National Congress (Câmara dos Deputados). We exploit rich data on thousands of candidates in three recent elections to estimate voters' preferences for candidates' ideological and non-ideological attributes. We then use voters' revealed preferences to quantify the welfare loss brought to voters by limitations in the pool of candidates they face. Finally, we combine our "demand side" estimates with estimates from a "supply side" model of policy choice to conduct counterfactual institutional experiments.

Quantifying representation failures requires that we first understand *what* is valuable to voters. If voters were purely ideological, representation failures would boil down to a lack of congruence between voters' preferences and politicians' policy positions. Indeed, this has been the most prevalent approach in the political science literature.[1] As a large body of research has shown, however, voters can and generally do have preferences over candidates' non-ideological attributes, such as their experience, competence, honesty, or gender.[2] To take both ideological and non-ideological factors into account, we rely on voters' revealed preferences over candidates' charac-

[1]See Miller and Stokes (1963), Erikson (1978), Clinton (2006), and Bafumi and Herron (2010). This approach is also prevalent in studies of *descriptive* representation; see Hero and Tolbert (1995), Cameron, Epstein, and O'halloran (1996), Griffin and Newman (2007), or Bowen and Clark (2014).

[2]See Besley and Reynal-Querol (2011), Ferraz and Finan (2011), Galasso and Nannicini (2011), Buttice and Stone (2012), Kendall, Nannicini, and Trebbi (2015), Folke, Persson, and Rickne (2016), and Beath, Christia, Egorov, and Enikolopov (2016).

teristics. Specifically, drawing on the parallels between voter choice among differentiated candidates in proportional representation (PR) electoral systems and consumer choice among differentiated products, we follow the aggregate discrete-choice demand estimation approach popularized by Berry, Levinsohn, and Pakes (1995) (BLP).[3]

The BLP approach is particularly well suited to the Brazilian electoral context. First, in Brazil's open-list PR system, voters cast their ballots overwhelmingly for individual candidates rather than political parties. This allows us to link voters' choices to individual candidate characteristics rather than those of an entire list, as would be the case in a closed-list PR system. Second, voters typically choose from among a large menu of candidates, with the typical state featuring between 92 and 114 candidates in the three elections in our sample. Such rich variation—both in menus across "markets" and in candidate characteristics within menus—gives us great purchasing power to identify voters' preferences. At the same time, the presence of a large numbers of candidates in each district makes it crucial that we can flexibly capture substitution patterns across alternatives. The typical mixed-logit BLP implementation allows us to do this in a computationally tractable manner, while overcoming the independence of irrelevant alternatives (IIA) property inherent in standard multinomial logit models.[4] Third, the BLP approach explicitly accounts for unobserved heterogeneity in candidate quality—or *valence* (Stokes 1963)—and its potential influence on candidates' policy choices. This enables us to reliably disentangle voters' preferences for policy relative to candidates' non-ideological attributes.

Our preference estimates provide several key insights concerning elections in Brazil. Consistent with previous research in political science, we find that Brazilian elections tend to be candidate-centric rather than party-centric, with voters effectively responding to candidate characteristics above party brands (Mainwaring, Scully, et al. 1995, Samuels 2003). Consistent with results in other countries, we find that Brazilian voters value candidates' education (Besley and Reynal-Querol 2011, Galasso and

[3]Other applications of the random coefficients model in electoral contexts include Rekkas (2007) (to estimate the effect of campaign spending in Canadian legislative elections), Gordon and Hartmann (2013) (political advertising in US Presidential elections), Montero (2023) (coalition formation in Mexico), and Ujhelyi, Chatterjee, and Szabó (2018) (protest voting in India). To the best of our knowledge, our paper is the first to use this approach to recover voters' preferences for policy versus non-ideological characteristics.

[4]This is particularly relevant in an electoral setting, as IIA would imply that a left-wing candidate and a right-wing candidate benefit or lose equally (in percentage terms) from a change in the policy position of, say, another right-wing candidate.

Nannicini 2011, Beath, Christia, Egorov, and Enikolopov 2016) and political experience (Buttice and Stone 2012). But we also find that voters place a significant weight on candidates' policy positions. In particular, we find that voters in more preponderantly rural districts, or with lower levels of education, tend to lean left ideologically, and we recover a significant level of heterogeneity in voters' policy preferences conditional on observed constituency characteristics. In contrast, we find no appreciable heterogeneity in voters' tastes for candidates' non-ideological attributes. This suggests that, from the perspective of our empirical application, all non-ideological candidate characteristics can be effectively considered as valence—i.e., attributes that are valued by all voters.

After recovering voters' preferences, we turn to our main objective of assessing the degree and sources of representation failures in Brazil. We begin by quantifying an overall voter welfare loss. To do this, we compute the gap between the level of welfare voters attain given the actual set of candidates they face in the data and what they would enjoy in an ideal representation benchmark. Using voter welfare as a metric allows us to weigh deficits across different dimensions *in the same way* voters resolve these tradeoffs. Comparing the actual welfare of each voter with an ideal benchmark allows us to quantify voters' losses relative to a theoretically meaningful yardstick of idealized representation. We consider three alternative representation benchmarks. In Benchmark I, we assume each voter is able to select her preferred candidate in all dimensions. In Benchmark II, we limit the number of "ideal" candidates in each state to be equal to that observed in the data, and we select these candidates to maximize average voter welfare in the state. In Benchmark III, we dispense altogether with the notion of ideal candidates and instead compare welfare in the data with what voters would obtain if they were able to choose from among all (actual) candidates running in *any* state, with valence and policy positions as observed in the data.

Our results illuminate a considerable failure of the Brazilian political system. The median welfare loss relative to our first benchmark across over five thousand municipalities is 69%. That is, in 50% of municipalities, the average voter attains a level of welfare no higher than 31% of what they would obtain in the ideal benchmark. In the comparison with Benchmark II, where we limit the number of "ideal" candidates, the median welfare loss goes down only marginally, to 66%. Thus, large estimated welfare losses are not the result of an undue inflation of the number of candidates in the first benchmark. In the comparison with Benchmark III, where the choice set is composed of all actual candidates running in any state, the welfare loss for the typical municipality goes down to 50%. This is considerably smaller than the welfare loss under the previous benchmarks but still remarkably large in magnitude. We conclude that (i) a substantial fraction of the welfare loss that emerges from the first benchmark remains when we consider alternatives that are certainly feasible for Brazil's political system, but (ii) voters in a subset of states are particularly impacted by shortages in their menu of available candidates.

To understand the sources of this representation failure, we decompose the total welfare loss in each municipality into a *policy welfare loss* (due to incongruence between voters' preferred policies and candidates' positions) and a *valence welfare loss* (due to inferior non-ideological characteristics of candidates). We find that, for the typical municipality, the valence welfare loss is more than seven times larger than the policy welfare loss. Large policy welfare losses do occur but are concentrated in a small fraction of municipalities (in a few states). In fact, the 10% worst-performing municipalities in this regard suffer a policy welfare loss of more than 54%. However, for three fourths of municipalities, the policy welfare loss is below 10%. In contrast, the welfare loss due to valence is 52% of the ideal benchmark for the median municipality and is above 69% for a quarter of municipalities. As validation, we explore whether our measures of welfare loss are related to observable measures of citizen dissatisfaction *unrelated* to voting outcomes. To do this, we use data on political protests. We find a statistically significant positive correlation between our measure of welfare loss and the likelihood of a political protest across municipalities (controlling for state fixed effects and municipal socio-economic characteristics). In particular, we find that political protests are twice as sensitive to policy welfare losses than to valence welfare losses.[5]

To evaluate institutional reforms aimed at improving voter welfare, potential strategic responses by candidates must be taken into consideration. Accordingly, we develop and estimate a model of the "supply side" of politics, where candidates' policy positions emerge explicitly as equilibrium choices. We model candidates' positions as resulting from a strategic balance between their own policy preferences and electability. Under an open-list PR electoral system, the latter has two components: candidates wish to maximize their individual vote share to further their chance of

[5] We obtain the same qualitative result using "protest votes" (blank or improperly cast ballots) instead of political protests.

obtaining a seat in the legislature, but parties may also exert some influence making candidates internalize the externalities their policy choices impose on fellow party members' vote shares. Our estimates suggest, however, that Brazilian parties have little influence over their candidates in this respect. Moreover, we find that, when trading off personal policy preferences for electability, candidates with favorable valence attributes place a larger weight on their own ideology.

We conduct two counterfactual experiments. In the first, we consider an institutional reform designed to directly alter valence in the pool of candidates (e.g., anti-corruption measures, age requirements, gender quotas). Specifically, we consider minimal education requirements. In the second experiment, we consider reforms aimed at strengthening political parties' influence over their candidates' policy choices. To reduce the computational burden, we focus our analysis on the state of Bahia, whose demographics are most representative of the nation as a whole.

Keeping candidates' policies fixed as observed in the data, a higher-education mandate leads to a 14.9% welfare increase for the typical municipality, with non-negative effects across the board. When we consider equilibrium adjustments in candidates' policy positions, the typical municipality still benefits from the reform, but the increase in welfare goes down to 5.7%. Furthermore, although the reform remains beneficial for the vast majority of voters, equilibrium adjustments lead to a downward shift in the distribution of welfare effects, including welfare *losses* for a fraction of municipalities. In the second counterfactual, we find that increasing party discipline over candidates' equilibrium policy choices benefits the average voter in 83% of all municipalities, yet average voter welfare decreases in the remaining 17%. Overall, our experiments show that the indirect equilibrium-adjustment effects of reforms aimed at improving the quality of representation can be substantial, with significant distributional implications.

To the best of our knowledge, our paper is the first to quantify welfare losses for voters due to supply-side constraints, decomposing them into ideological and non-ideological factors. We are aware of only a few papers that have estimated the relative value voters give to ideology and valence, all in the context of plurality elections between (effectively) two candidates. With a focus on evaluating informational interventions, Kendall, Nannicini, and Trebbi (2015) and Cruz, Keefer, Labonne, and Trebbi (2018) combine a structural approach with a field experiment concerning mayoral elections in Italy and the Philippines, respectively. Kendall and Matsusaka (2021) combine a structural approach with survey data on vote intentions for California ballot propositions to disentangle ideological and valence considerations. In contrast, our paper relies on aggregate candidate-choice data and is the first to examine multi-candidate PR elections.

In estimating voter preferences, we recognize the endogeneity of candidates' policy choices and instrument for them. As is standard in the industrial organization literature, however, we assume that unobserved valence (by the econometrician) is uncorrelated with other observable candidate characteristics. If candidate entry were costly and strategic, potential entrants with unfavorable observed characteristics would only enter the election if their unobserved valence compensated for that disadvantage. This correlation would bias coefficient estimates regarding voters' non-ideological preferences.[6] Patterns of entry in Brazil's legislative elections suggest that, in this context, there is little reason for concern. Pooling over the three elections in our data, 30% of candidates obtain less than 0.01% of the vote, and 60% of candidates obtain less than 0.06% of the vote. This indicates that the connection between entry and expected electoral performance is weak *in this context*.[7] As an additional check, we re-estimate the model excluding all non-ideological observable candidate characteristics. We find that the median municipality suffers a welfare loss of 66%, a figure comparable to the 69% we obtain in our main results.

2 Institutional Context and Data

We focus our analysis on elections of representatives for the lower house of the Brazilian National Congress. The *Câmara dos Deputados* is composed of 513 representatives, who are elected in 27 multi-member electoral districts, corresponding to the country's 26 states and the *Distrito Federal* of Brasilia. The magnitude of each district is determined according to population, but no state may have fewer than eight seats or more than seventy seats. Eleven states elect eight representatives, and São Paulo is the only district electing seventy representatives (see Table A1 in Appendix

[6]Kawai and Sunada (2022) address endogenous entry in US House elections and estimate candidate valence with a methodology building on the "production function" approach.

[7]The presence of a large number of "small" candidates might raise the question of whether our estimates are sensitive to the inclusion of these candidates in the analysis. In Table C2 in the Appendix, we show that our estimates are robust to excluding the top and bottom 5% and 10% of candidates, by vote share.

A). Representatives are elected for four-year terms, with no constraints on reelection (in 2014, over 74% of incumbents secured reelection).

Elections take place under an open-list proportional representation (PR) system. Each voter has one vote to cast, which can be given to a specific candidate or, rarely, to a party or coalition list (in our sample, fewer than 6% of voters do so). In each district, votes given to candidates from each list are pooled and added to the votes received by the list to form a total list vote. Seats are then distributed among lists proportionally to their total list vote according to the D'Hondt method. Within each list, seats are assigned to candidates in descending order of votes received. Note that, in the event the candidate chosen by a voter is not competitive, the vote is not wasted but gets reallocated to the member of the list closest to the threshold for obtaining a seat. Combined with large district magnitudes, this greatly diminishes incentives to vote strategically.[8]

Brazil's open-list PR system fosters a fragmented multiparty system. In the 2014 election, for instance, 28 parties placed candidates in the lower chamber (see Table A2 in Appendix A). Interestingly, vote dispersion among multiple parties is not merely driven by regional factors—it persists in vote outcomes aggregated at the municipal level (see the left panel of Figure A1 in Appendix A).

2.1 Candidates

The Brazilian electoral system puts individual candidates at the center of political choice. Indeed, the political science literature notes that (i) parties are weak, under-resourced, and often unable to constrain opportunistic behavior by individual legislators (Samuels 2003, Desposato 2006); (ii) open-list PR and a lack of formal mechanisms channeling resources to congressional party leaders promote candidate-centric legislative careers (Mainwaring, Scully, et al. 1995, Samuels 2003); and (iii) Brazilian elections tend to be candidate-centric rather than party-centric, with voters effectively responding to candidate characteristics above party labels (Mainwaring, Scully, et al. 1995, Samuels 2003).

Understanding voters' choices, therefore, requires that we analyze them at the

[8] Voting strategically in this context is a highly complex problem, which requires forming (correct) conjectures over both the threshold of seats a party attains and likely candidate ties around that threshold. In particular, for a voter whose preferred candidate is not competitive, voting for one of the top candidates on the same list is not better than voting for the voter's preferred candidate.

candidate level. To that end, we bring together data on candidates running for a seat in the Câmara dos Deputados in the 2006, 2010, and 2014 elections. In total, across these three elections and all 27 legislative districts, there are 15,698 election-specific candidates: 4,944 in 2006, 4,887 in 2010, and 5,867 in 2014. For each candidate, we obtain the number of votes received in each municipality along with a rich set of individual characteristics, including their previous professional experience, incumbency status, level of education, and gender.[9]

Figure 1 provides summary statistics of candidates' observable non-ideological characteristics. Incumbents constitute only a fraction of all candidates but are disproportionately represented among those who secure a seat in the chamber. While only about half of candidates have higher education, 75% of elected candidates do. Women compose only about a quarter of total candidates and a far lower percentage of elected legislators. Candidates with business or government (bureaucratic) experience make about 10% of the pool of candidates, and they represent a significantly lower proportion of elected candidates.

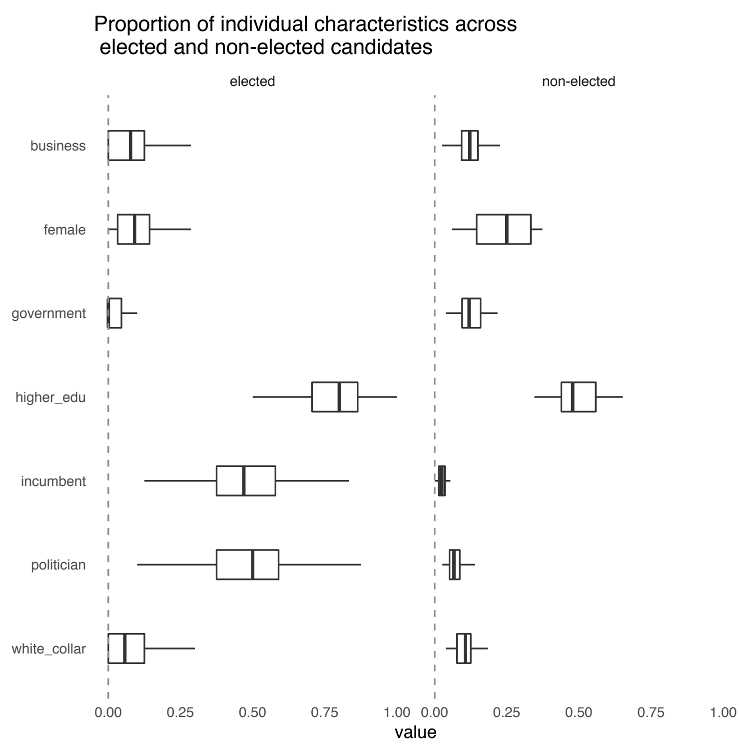

Figure 1: Candidates' Observed Non-Ideological Characteristics

[9]This information is made available by the *Tribunal Superior Eleitoral* (TSE).

Figure 1 suggests that candidates' education level, professional experience, incumbency status, and gender are relevant to voters. Whether Brazilian voters also care about candidates' policy positions—and how much weight they put on ideology relative to non-ideological considerations—is an open empirical question. Answering this question requires data for both elected and non-elected candidates. Unfortunately, there are no currently available measures of both incumbents' and challengers' policy positions for Brazilian legislative elections.[10] To address this gap, we follow the approach pioneered by Bonica (2014) in the U.S. context and produce our own estimates of candidates' ideological positions using correspondence analysis—the analog of principal component analysis for categorical data—on micro-level campaign contributions data. While Bonica interprets these estimates as politicians' preferred policies, we only view them as the positions candidates put forward, which could correspond or not to their true preferences.

In particular, we use all *individual* political contributions to federal, state, and local candidates between 2000 and 2014, encompassing 2.3 million donors and 561 thousand political candidates (see Appendix B for details).[11] Because many non-viable candidates tend to receive no contributions, we are forced to drop them from the data. Still, our final sample of candidates running for a seat in the Câmara dos Deputados includes 8,956 candidates across the three elections. In Appendix C.3, we provide a detailed comparison of candidates in and out of our sample. Moreover, we conduct a sensitivity test by imputing policy positions for excluded candidates and show that our welfare analysis remains virtually unchanged.

Figure 2 plots the distribution of our estimates of candidates' ideological positions by party in six selected states. As shown, candidates' positions vary considerably by party *and* by state within each party, which indicates that the contributions data is indeed informative about candidates' ideological positions. Furthermore, observed patterns are consistent with the typical understanding of ideological divisions in Brazil: with PCdoB and PPS on the left; PT, PDT, and PSB as center-left; PSDB,

[10]Zucco (2009) and Zucco and Lauderdale (2011) estimate *incumbents*' positions using surveys that ask them to place themselves and all the main political parties represented in the legislature on a left-right scale.

[11]Campaign contributions are published by the TSE. Since corporations and political parties may contribute to candidates strategically rather than ideologically, we exclude them from our data and focus on individual contributions by non-politicians. Under-the-table donations—*caixa dois*—are common, but previous research using the same data has shown that officially-declared donations capture the majority of contributions (Boas, Hidalgo, and Richardson 2014).

PSD, and PV at the ideological center; PMDB and PTB as center-right; and DEM and PP on the right of the policy space. (Figure A2 in Appendix A shows the overall distribution of candidates' policy positions as well as the distribution by party pooling across states.)

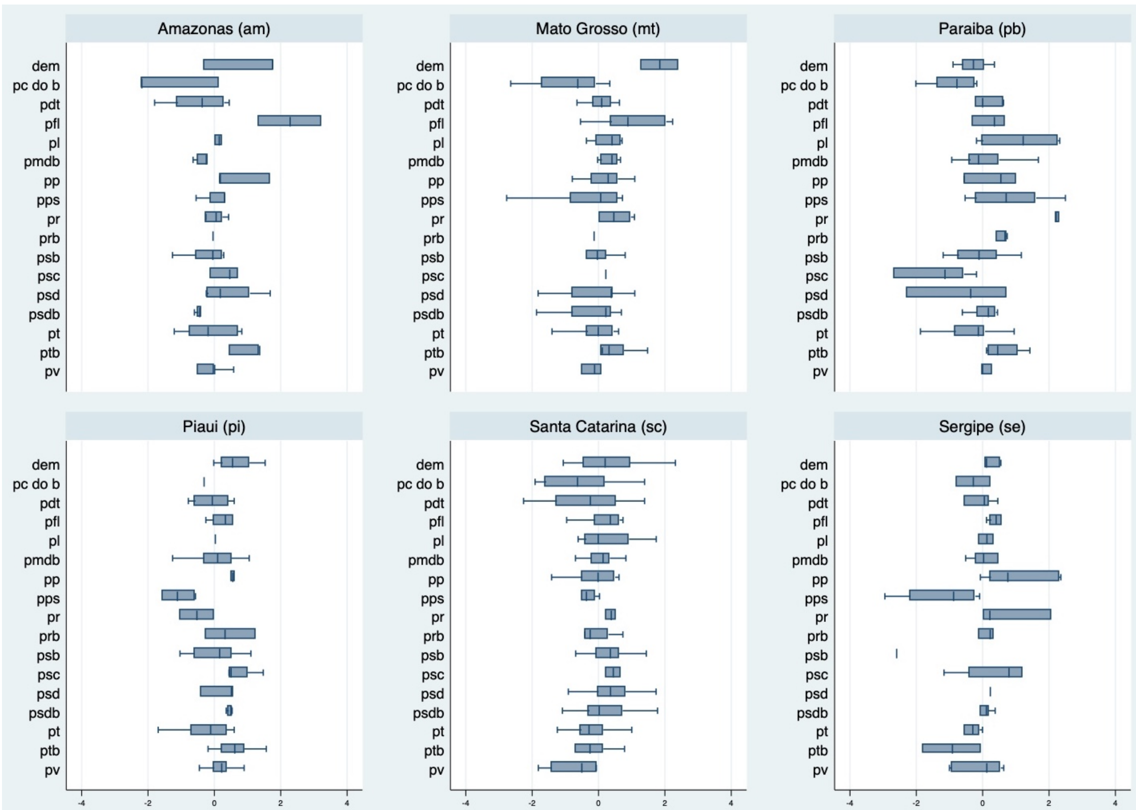

Figure 2: Candidates' Policy Positions (by Party in Six Selected States)

To further validate our policy position estimates, we conduct a battery of sanity checks. First, the left panel of Figure B1 in Appendix B presents average policy positions by party, comparing federal versus local candidates. Positions are generally consistent within party, as expected from common competitive and intra-party environments. In the right panel of Figure B1, we compare our estimates with ideology estimates for incumbents obtained from legislative surveys by Power and Zucco (2012). As the figure shows, there is general agreement between the two types of scores. Third, during the Lula presidency, there was a marked shift to the left in voters' policy preferences according to Latinobarometer survey data, depicted in the left-hand panel of Figure B2. Our estimates feature a similar leftward shift in candidates' policy positions as shown in the right-hand panel of Figure B2.

Given the prevalence of corruption in Brazil, there could be a potential concern that, even among non-corporate and non-party individual donors, campaign contributions may be motivated by public contract allocations or other forms of quid pro quo. To address this, we conduct two additional robustness checks. First, we re-estimate candidates' policy positions excluding the top 5% and 10% of donors from the sample. Since contributions seeking to buy access to politicians or to exact favors are likely to be sizable, focusing on small contributions should alleviate such concerns. As shown in Figure B3, the resulting estimates are very similar to those obtained from the full sample (correlations are 0.9 and 0.85, respectively, for estimates excluding the top 5% and 10% of donors). Second, to more directly address the possibility that campaign donors may be motivated by public contracts, we use data on public contract allocations by *deputados federais* provided by Boas, Hidalgo, and Richardson (2014). Figure B4 plots total (in logs) individual donations received in the 2006 (left) and 2010 (right) electoral cycles against the total value (in logs) of disbursed contracts for each federal deputy in the 2006-2010 legislature. We find a very weak positive association between donations and contracts, slightly more prominent for the 2006 cycle.

In the next section, we use this information on candidates' policy choices and non-ideological characteristics, along with election results, to estimate voters' preferences. An alternative that is de facto available to voters is to abstain or to cast a void vote. This "outside option" is thus effectively competing with all the candidates for votes. The data shows this outside option is a formidable alternative. Average abstention and blank-vote rates of 29% and 8.6%, respectively, in what is formally a compulsory voting system already provide suggestive evidence that voters are not enthusiastic about the candidates they face.

3 Voter Preferences

3.1 Voter Preferences: Empirical Model

To disentangle voters' preferences over ideological and non-ideological characteristics of candidates, we follow the approach of Berry, Levinsohn, and Pakes (1995) (BLP). Our data are particularly well suited to this technique, as we have rich variability—both across and within menus—from about 9,000 candidates running in multiple constituencies and electoral cycles. Combined with appropriate instruments, this ensures identification and allows us to obtain precise estimates.

We assume that voter i's utility from selecting candidate j in state (district) n is given by

$$u_{ijn} = \alpha_{0i} + \alpha_{1i} p_{jn} + \alpha_{2i} p_{jn}^2 + W_{jn}' \phi_i + X_{jn}' \beta + \xi_{jn} + \epsilon_{ijn}, \tag{3.1}$$

where $p_{jn} \in [-\overline{p}, \overline{p}]$ denotes candidate j's (endogenous) policy position, X_{jn} is a vector of (exogenous) valence characteristics of candidate j, and ϵ_{ijn} is an i.i.d. mean-zero Type-I Extreme Value (TIEV) random utility shock.[12] The vector W_{jn} includes candidate j's gender and incumbency status, which we separate from X_{jn} to allow for the possibility that voters disagree over whether male or female candidates, or incumbents or challengers, are better alternatives.[13] The term ξ_{jn} explicitly captures valence attributes of candidate j that may affect voters' preferences but are unobserved by the econometrician, such as charisma or trustworthiness. While unobserved by the analyst, ξ_{jn} is assumed to be known by voters, candidates, and parties and is therefore potentially correlated with j's policy position, p_{jn}.

Note that the coefficients determining the effect of j's policy position are voter-specific, and voter i's ideal policy can be recovered as $y_i = -\alpha_{1i}/(2\alpha_{2i})$.[14] For $k = 0, 1, 2$, we assume that

$$\alpha_{ki} = \alpha_k^0 + D_{n(i)}' \alpha_k^D + \sigma_k \nu_{ki}, \tag{3.2}$$

where D_n is a vector of demographic characteristics of state n, $n(i)$ denotes the state in which voter i resides, and $\nu_{ki} \overset{\text{i.i.d.}}{\sim} N(0, 1)$ are idiosyncratic policy preference shocks. Thus, voters' ideal points vary both with observed and unobserved voter characteristics. This enables estimation of rich substitution patterns in a computa-

[12]For brevity, we refer in this section simply to *state* n. In our empirical application, we have data spanning three electoral cycles, so n corresponds to a *state-year*. Moreover, as noted above, voters also have the option of selecting a list rather than a specific candidate. We accommodate this by treating lists as additional "candidates" whose observed attributes are averages of the member candidates' characteristics.

[13]While other scholars have disentangled the sources of incumbency advantage—see, e.g., Klašnja and Titiunik (2017)—we allow incumbency status to bundle all persistent differences between incumbent and non-incumbent candidates (name recognition, clientelistic networks, influence within parties, campaign resources, etc.), and we let ξ_{jn} capture election-specific voter tastes for unobserved candidate characteristics.

[14]Some voters may have convex policy preferences, i.e., $\alpha_{2i} \geq 0$. These voters prefer extreme policies and have ideal point $y_i = \overline{p}$ ($y_i = -\overline{p}$) if $\alpha_{1i} \geq 0$ ($\alpha_{1i} < 0$). We set $\overline{p} = 5.8$, equal to the maximum absolute policy observed in the data plus two standard deviations. Our substantive results are robust to alternative specifications of the policy space.

tionally feasible manner, while relaxing the independence of irrelevant alternatives (IIA) property of standard multinomial logit models.

In principle, the coefficients β capturing the effect of observed valence characteristics could also be allowed to vary across voters. For computational tractability, we recover only an average valence effect. Yet, recognizing that preferences for gender and incumbency effects might fundamentally differ across voters, we allow $\phi_i = (\phi_{1i}, \phi_{2i})'$ to be voter-specific, letting

$$\phi_{k-2,i} = \phi_{k-2} + \sigma_k \nu_{ki}, \quad \text{with } \nu_{ki} \overset{\text{i.i.d.}}{\sim} N(0,1) \text{ for } k = 3, 4.$$

Because random utility shocks are distributed TIEV, the probability that voter i in district n selects candidate j given shocks $\nu_i = (\nu_{0i}, \ldots, \nu_{4i})$ is

$$P^i_{jn}(\nu_i) = \frac{\exp(\delta_{jn} + \sum_{k=0}^{2} \sigma_k \nu_{ki} p^k_{jn} + \sum_{k=3}^{4} \sigma_k \nu_{ki} W_{k-2,jn})}{1 + \sum_{j' \in J_n} \exp(\delta_{j'n} + \sum_{k=0}^{2} \sigma_k \nu_{ki} p^k_{j'n} + \sum_{k=3}^{4} \sigma_k \nu_{ki} W_{k-2,j'n})},$$

where J_n denotes the set of candidates running in state n and

$$\delta_{jn} = \sum_{k=0}^{2} (\alpha^0_k + D'_n \alpha^D_k)(p_{jn})^k + W'_{jn}\phi + X'_{jn}\beta + \xi_{jn} \tag{3.3}$$

is the average voter utility from choosing candidate j.[15] We normalize the average "outside-option" utility from abstaining or casting a void vote ($j = 0$) to $\delta_{0n} = 0$. Thus, $(\alpha^0_0 + D'_n \alpha^D_0)$ captures cross-district variation in baseline abstention/blank-vote rates. Integrating over ν_i, candidate j's predicted vote share in district n can be written as

$$s_{jn} = E_{\nu_i}\big[P^i_{jn}(\nu_i)\big]. \tag{3.4}$$

3.2 Voter Preferences: Estimation

We implement the BLP estimation strategy using the Mathematical Programming with Equilibrium Constraints (MPEC) approach of Dubé, Fox, and Su (2012) for computational efficiency. Next, we summarize the main ideas, emphasizing the intu-

[15] In Appendix C.3, we explore an alternative specification of our model that allows for an interaction in voters' preferences between policy and valence considerations. (Kendall, Nannicini, and Trebbi (2015) find no evidence of a significant interaction of this sort.) Our main results are substantively unchanged.

ition. For technical details, see Appendix C.

Consider first the simpler case where voters are homogeneous up to observed covariates, which boils down to a standard multinomial-logit random utility model. Given $\sigma = 0$, we can "invert" predicted vote shares to express them in terms of average voter utilities by taking logs of (3.4): $\log(s_{jn}) - \log(s_{0n}) = \delta_{jn}$. Then, replacing predicted vote shares with their observed counterparts in the data, $\hat{s}_{jn}$, and using (3.3), we obtain

$$\log(\hat{s}_{jn}) - \log(\hat{s}_{0n}) = \sum_{k=0}^{2} (\alpha_k^0 + D_n' \alpha_k^D)(p_{jn})^k + W_{jn}' \phi + X_{jn}' \beta + \xi_{jn},$$

which is just a linear regression of the log-ratio of candidate j's vote share to that of the "outside option" on endogenous (p_{jn}) and exogenous covariates (D_n, W_{jn}, and X_{jn}). Note that candidate j's unobserved valence, ξ_{jn}, corresponds to the residual of this regression. Thus, provided we have valid instruments for the endogenous regressors, we can estimate parameters (α, ϕ, β) from this linear regression via two-stage least squares.

The multinomial logit model is computationally straightforward but imposes strong assumptions on voter preferences. In particular, since $\log(s_{jn}/s_{j'n}) = \delta_{jn} - \delta_{j'n}$, the log-ratio of the vote shares of any two candidates j and j' does not depend on the characteristics of other candidates (IIA). An important implication is that, if one candidate changes their policy position, all other candidates gain or lose votes by the same percentage. This makes little sense in a model of electoral politics, as candidates on the same side of the ideology spectrum are naturally closer substitutes than diametrically opposed candidates. The key insight of BLP is that introducing voter heterogeneity allows flexible substitution patterns to emerge. Voters with ideal points $y_i > 0$, for instance, are more likely to respond to a change in a right-wing candidate's policy than voters with $y_i < 0$, which plausibly leads to higher substitutability between right-wing candidates than between right versus left-wing candidates.

When voters are heterogeneous, the above estimation approach is no longer feasible. Yet an approach that builds on the same principles is. Given (3.4), predicted vote shares in each state n depend not only on the average utilities $\delta_n = (\delta_{1n}, \ldots, \delta_{J_n n})$ (determined by parameters α, ϕ, and β) but also on the heterogeneous preference parameters σ. We can no longer explicitly "invert" predicted vote shares $s_n = (s_{1n}, \ldots, s_{J_n n})$, but BLP show that, for any given value of σ, there exists a unique vector of average utilities $\delta_n(\sigma)$ such that predicted and observed vote shares match exactly, i.e., $\hat{s}_n = s_n(\delta_n(\sigma), \sigma)$. Then, using (3.3) and given a candidate value of $\theta = (\alpha, \phi, \beta, \sigma)$, we can compute the unobserved candidate valence consistent with $\delta_{jn}(\sigma)$:

$$\xi_{jn}(\theta) = \delta_{jn}(\sigma) - \sum_{k=0}^{2}(\alpha_k^0 + D_n'\alpha_k^D)(p_{jn})^k - W_{jn}'\phi - X_{jn}'\beta. \tag{3.5}$$

This allows us to construct a Generalized Method of Moments (GMM) estimator given a vector Z_{jn} of instruments satisfying

$$E[\xi_{jn}(\theta)|Z_{jn}] = 0 \quad \text{if and only if} \quad \theta = \theta_0, \tag{3.6}$$

where θ_0 denotes the true value of the model parameters.

The BLP estimation algorithm proceeds by iterating over two nested loops. Given a candidate value of θ, the "inner loop" inverts predicted vote shares to solve for $\xi_{jn}(\theta)$ according to (3.5). Letting Z and $\xi(\theta)$ denote vertical stackings of Z_{jn}' and $\xi_{jn}(\theta)$ across candidates and elections in the data, a sample analog of the orthogonality condition implied by (3.6) can be computed as $\frac{1}{J}Z'\xi(\theta)$, where J denotes the total number of observations. Under standard technical regularity conditions, a consistent and asymptotically normal estimator can be obtained by minimizing the (positive-definite) quadratic form $Q_J(\theta) = \left[\frac{1}{J}Z'\xi(\theta)\right]' W_J \left[\frac{1}{J}Z'\xi(\theta)\right]$. Accordingly, the "outer loop" searches over θ to minimize $Q_J(\theta)$. Inference follows standard GMM theory, including the choice of an optimal weighting matrix. We cluster standard errors at the district level, by electoral cycle, to allow for potential correlation in unobserved valence across candidates in the same race.

The BLP algorithm can be computationally inefficient—as the inner loop relies on costly fixed-point calculations—and sensitive to convergence criteria. Instead, we implement an MPEC version of the BLP estimator, which has been shown to yield better numerical performance (Dubé, Fox, and Su 2012). The idea is to impose the "equilibrium conditions" of the model, $\hat{s}_n = s_n(\delta_n(\sigma), \sigma)$, as explicit constraints on the GMM optimization. Since modern optimization algorithms satisfy constraints only at convergence, this sidesteps repeated fixed-point calculations.

Instruments. A necessary order condition for identification is that Z_{jn} must include at least as many variables as there are parameters to be estimated. As is standard in industrial organization, we assume that candidates' non-ideological characteristics (W_{jn} and X_{jn}) are fixed in the short run and uncorrelated with unobserved valence, which implies they are valid (optimal) instruments to identify ϕ and β. This assumption would not be sensible, however, if candidate entry were costly and thus strategic. As discussed above, the literature on Brazilian politics notes that parties are generally weak, which suggests they have limited gatekeeping power. Moreover, given patterns of entry in our sample, electability does not seem to be an overriding concern: 30% of candidates obtain less than 0.01% of the vote, while three fifths obtain less than 0.06%. To further address concerns about potential bias in our estimates of voters' preferences due to strategic entry, we show in Appendix C.3 that our results are virtually unchanged if we exclude outstanding candidates from our sample (i.e., those with vote shares in the top 5% or 10%). As an additional robustness check, we re-estimate the model excluding all observable non-ideological candidate characteristics, thus collapsing all valence to the unobserved term. We find that our main results are essentially unchanged.

We rely on auxiliary data and the structure of the model to obtain instruments for the remaining parameters. To identify α, notice that, given any variable that is correlated with p_{jn} but uncorrelated with ξ_{jn}, natural choices for the remaining instruments are its square and corresponding interactions with state demographics. Accordingly, to construct an instrument for p_{jn}, we exploit the policy positions of mayoral and state-level candidates in the most recent local electoral cycle. As shown in Figure B1 in Appendix B, the policy positions of local and federal legislative candidates serving the same constituency covary. This is unsurprising given that both types of candidates respond to similar electoral/party environments. However, mayoral and state-level candidates' policy positions are plausibly uncorrelated with the charisma or other unobserved non-ideological attributes of federal legislative candidates. Thus, we use a weighted average of same-party mayoral and state-level candidates' positions to instrument for p_{jn}, giving a larger weight to local candidates j' closer to j in terms of observed characteristics. Specifically, weights are inversely proportional to

$$\exp\{-(X_{jn} - X_{j'n})'\text{Cov}(X)^{-1}(X_{jn} - X_{j'n})\},$$

where $\text{Cov}(X)$ denotes the sample covariance matrix of candidates' non-ideological characteristics (including gender and incumbency status for this construction).

Finally, while the choice of instruments for (α, ϕ, β) follows standard intuition from linear regressions given (3.5), the preference variance parameters, $\sigma = (\sigma_0, \ldots, \sigma_4)$, determine the nonlinear features of the model.[16] As instruments for σ, following recent work by Gandhi and Houde (2020), we rely on a second-degree polynomial of observed differences across candidates in W_{jn}, X_{jn}, and $\hat{p}_{jn}$, the first-stage fitted value of p_{jn} using the instruments described above. These characteristics are uncorrelated with unobserved valence. Moreover, as noted, individual-level heterogeneity in voters' preferences—measured by σ—captures variability in the degree of substitutability between candidates, which in turn is determined by proximity in the attribute space. Thus, attribute differences across candidates provide the right source of variation to identify σ. A key advantage of our application is that, since elections in our data are large, we have great leverage to estimate these notoriously hard-to-identify (in practice) parameters.

3.3 Voter Preferences: Estimates

We report our parameter estimates of voters' preferences in Table 1 (observed valence, incumbency, and gender) and Table 2 (ideology) below, as well as Tables A3 (party brands) and A4 (baseline voter utility) in Appendix A. All non-dichotomous covariates are standardized, so coefficients can be compared at face value. The first column of each table presents two-stage least squares (2SLS) estimates from a multinomial logit (MNL) model that does not control for voter demographics. The second column presents 2SLS estimates from a multinomial logit model including voter demographics. The third column presents estimates from the BLP model, which allows for heterogeneity in preferences among voters conditional on covariates. As a quick examination of the tables reveals, the added complications of the BLP approach are worth pursuing, as they have considerable bite in the resulting estimates. Indeed, the three models are nested: the model in the second column is obtained by setting $\sigma = 0$, and the model in the first column additionally sets $\alpha^D = (\alpha_0^D, \alpha_1^D, \alpha_2^D) = 0$. Both restrictions are rejected by the data.

Our estimates provide several key insights regarding electoral politics in Brazil. Consistent with previous research, we find robust evidence that individual candidate characteristics (as opposed to just partisan cues) are important determinants of vot-

[16] For parsimony, we set $\sigma_0 = 0$ (the intercept of voters' utility varies only with observed demographics. We also set $\alpha_0^0 = 0$ since, as discussed below, X includes a full set of party dummies.

ers' choices. Table 1 presents the estimated effects of candidates' non-ideological attributes. As in Besley and Reynal-Querol (2011) and Beath, Christia, Egorov, and Enikolopov (2016), we find that education has a positive valence effect. Similarly, Brazilian voters have a preference for candidates with business experience, and they dislike government bureaucrats.

	MNL	MNL (w/Dem's)	BLP
Age (β_2^{age})	0.042	0.005	0.014
	(0.031)	*(0.037)*	*(0.022)*
Age Sq. ($\beta_2^{age^2}$)	-0.101	-0.005	-0.002
	(0.029)	*(0.025)*	*(0.013)*
Higher Education (β_2^{edu})	0.106	0.656	0.688
	(0.172)	*(0.086)*	*(0.044)*
Business Experience ($\beta_2^{business}$)	-0.441	-0.003	0.162
	(0.168)	*(0.156)*	*(0.059)*
Government Experience (β_2^{gov})	-0.628	-0.503	-0.494
	(0.111)	*(0.126)*	*(0.077)*
Incumbent (ϕ_1)	1.938	1.847	1.786
	(0.109)	*(0.118)*	*(0.248)*
Incumbent Preference Variance (σ_3)			0.0004
			(1160.5)
Female Candidate (ϕ_2)	-1.331	-1.173	-1.065
	(0.145)	*(0.146)*	*(0.896)*
Female Cand. Preference Variance (σ_4)			0.003
			(305.9)

Table 1: Estimates of Voters' Preferences over Non-Ideological Characteristics

In estimating the effect of candidate gender and incumbency status, we allow for heterogeneity across voters. The mean effect of being female is negative but statistically insignificant at conventional levels. On the other hand, the effect of incumbency is large, positive, and statistically significant. For interpretation of this result, recall that in our model incumbency status bundles all persistent differences between incumbent and non-incumbent candidates. Thus, the estimate indicates that the combined effect of name recognition, clientelistic networks, influence within parties, campaign resources, and other advantages incumbents might enjoy is substantial. Notably, for both gender and incumbency status, we find no evidence of heterogenous effects across voters.[17] In other words, from the perspective of this empirical application, all non-ideological attributes can be considered as valence characteristics.

[17]Our estimates of σ_3 and σ_4 are extremely imprecise. As discussed above, this is perhaps due to insufficient variation in attribute proximity in the pool of candidates across constituencies given that both gender and incumbency are dichotomous. Nevertheless, point estimates for σ_3 and σ_4 are both orders of magnitude closer to zero than those for other coefficients.

Table 2 presents estimates of voters' ideological preferences, $(\alpha, (\sigma_1, \sigma_2))$. If voters did not value candidates' policy positions, these coefficients would be zero. Our estimates reject this hypothesis and indicate that, in evaluating alternative candidates, Brazilian voters do trade off valence and ideological considerations. Since we standardize demographics in our sample and $E[\nu_{1i}] = E[\nu_{2i}] = 0$, the effect of policy on the preferences of the average voter in the country is captured entirely by the common terms α_1^0 and α_2^0. The estimate of α_1^0 is negative (-0.95) but not statistically significant, while the estimate of α_2^0 is negative, large in magnitude (-4.79), and significant at the 1% level. This implies that the average voter is centrist relative to candidates' policy offerings, and has concave policy preferences, suffering increasingly with larger deviations from their preferred policy. Policy preferences, however, effectively vary with observed demographics. In particular, a higher proportion rural, a higher median wage, or a lower proportion of educated residents in a district have a negative and statistically significant effect on α_{1i}. On the other hand, a higher proportion of educated residents has a negative and statistically significant effect on α_{2i}, and a higher proportion rural has a negative effect on α_{2i}, barely insignificant at the 10% level. Overall, this implies that voters in more preponderantly rural districts, or with lower levels of education, tend to lean left—a consistent pattern throughout most of Latin America.

In addition to the variation that can be attributed to observable voter characteristics, we find that voters' policy preferences are heterogeneous *conditional* on demographics. Indeed, while our point estimate of σ_1 is essentially zero—and imprecisely estimated—our estimate of σ_2 is positive and statistically significant. To interpret this result, note that a higher value of σ_2 implies more extreme left and right leaning ideal points, and increases the dispersion of voters' ideal points within each municipality (i.e., conditional on covariates).[18]

Evaluating our policy-preference estimates using municipality-level covariates, D_m, we can recover the distribution of voters' ideal points in each municipality. Specifically, for each municipality, we simulate a sample of registered voters, drawing for each voter i policy-preference shocks (ν_{1i}, ν_{2i}). For voters with resulting concave policy preferences ($\alpha_{2i} < 0$), we compute their ideal point as $y_i = -\alpha_{1i}/(2\alpha_{2i})$. Voters with convex ($\alpha_{2i} \geq 0$), or extreme, policy preferences are assigned an ideal point at

[18] Figure A5 in the Appendix illustrates this effect comparing the distribution of ideal points at the estimates with the distribution of ideal points consistent with alternative values of σ_2.

	MNL	MNL (w/Dem's)	BLP
Policy (α_1^0)	1.000	-2.606	-0.950
	(0.681)	*(1.556)*	*(1.009)*
Policy × Median Wage (α_1^{wage})		0.838	-2.389
		(1.564)	*(0.783)*
Policy × % Rural (α_1^{rural})		-0.679	-3.390
		(2.231)	*(1.052)*
Policy × % Higher Education (α_1^{edu})		0.239	1.120
		(0.842)	*(0.537)*
Policy × % Employed (α_1^{emp})		-1.065	1.459
		(3.183)	*(1.206)*
Policy × Average Age (α_1^{age})		-0.957	-0.763
		(1.933)	*(0.702)*
Policy × % Female (α_1^{female})		0.844	-0.598
		(1.999)	*(0.637)*
Policy Sq. (α_2^0)	0.857	-4.634	-4.879
	(0.608)	*(0.745)*	*(0.663)*
Policy Sq. × Median Wage (α_2^{wage})		2.107	1.220
		(2.087)	*(0.920)*
Policy Sq. × % Rural (α_2^{rural})		-3.638	-3.008
		(4.322)	*(1.751)*
Policy Sq. × % Higher Education (α_2^{edu})		-1.922	-1.185
		(1.252)	*(0.498)*
Policy Sq. × % Employed (α_2^{emp})		-3.052	-1.945
		(6.222)	*(2.272)*
Policy Sq. × Average Age (α_2^{age})		-0.515	-0.507
		(2.731)	*(0.952)*
Policy Sq. × % Female (α_2^{female})		0.944	1.143
		(2.780)	*(0.949)*
Policy Preference Variance (σ_1)			0.001
			(868.7)
Policy Sq. Preference Variance (σ_2)			0.384
			(0.211)

Table 2: Estimates of Voters' Policy Preferences

the boundary of the policy space (see footnote 14).

Figure 3 plots the average voter's ideal point in each municipality. The estimates show a substantial amount of ideological heterogeneity across regions, and even within states. The north, northeast, and south regions are more uniformly left-wing. On the other hand, the southeast and central-west regions (São Paulo, Goias) tend to be more conservative but highly polarized. Overall, this corroborates well-known patterns of partisanship in Brazil—see, for instance, Power and Rodrigues-Silveira (2019a).

Up to this point, we have focused on the distribution of voters' ideal policies. However, our utility specification in (3.1) also allows for heterogeneity in the weight that voters give to policy relative to ideology. In particular, note that, if we write

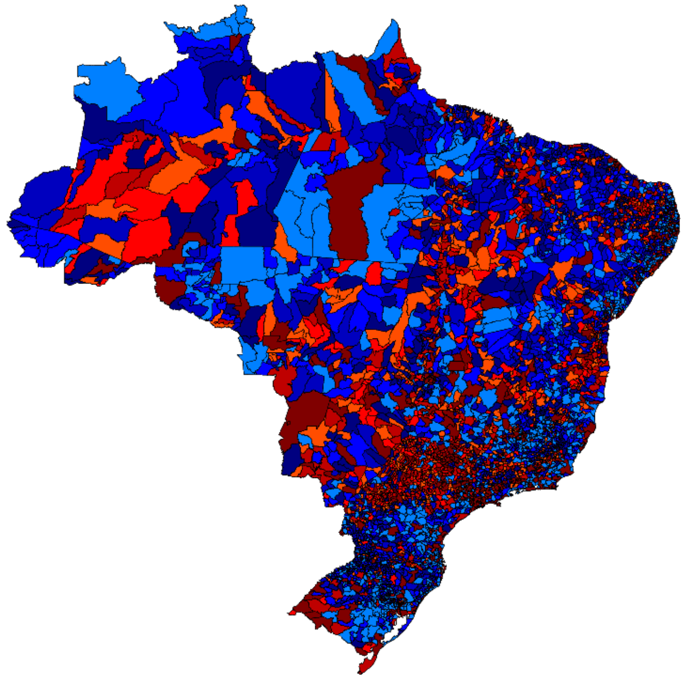

Figure 3: Voters' Ideological Preferences (by Municipality in 2014)—darker blue (red) denotes more left-leaning (right-leaning) ideal policy

voter i's utility as

$$u_{ijn} = \tilde{\alpha}_{0i} - \gamma_i (p_{jn} - y_i)^2 + W'_{jn}\phi_i + X'_{jn}\beta + \xi_{jn} + \epsilon_{ijn}, \tag{3.7}$$

we have $\alpha_{1i} = 2\gamma_i y_i$ and $\alpha_{2i} = -\gamma_i$, where voter i's ideal point is $y_i = -\alpha_{1i}/(2\alpha_{2i})$ and the relative weight of policy in voter i's utility function is

$$\gamma_i = -\alpha_{2i} = -\left[\alpha_2^0 + D'_{n(i)}\alpha_2^D + \sigma_2 \nu_{2i}\right].$$

This allows us to ascertain whether centrist or extreme leaning voters on the left or the right of the ideological spectrum care more about policy relative to non-ideological characteristics of the candidates. To evaluate this, we compute the change in utility for the average voter as we vary policies in the domain observed in the data. We find that *centrist* voters have a higher relative value for policy representation (see Figure A6 in the Appendix).

Valence vs. Ideology. A natural question is how important ideological considerations are relative to the non-ideological characteristics of candidates. Does ideology dominate differences in the education, experience, or unobserved valence of candidates? To answer this question, we compute the elasticities of candidates' vote shares with respect to their own policy position, η^p_{jj}, and valence, η^v_{jj}. That is, the ideology (valence) elasticity measures the percentage change in vote share resulting from a 1% increase in the policy position (valence) of the candidate. Thus, the ratio $r_\eta = |\eta^p_{jj}/\eta^v_{jj}|$ measures the percentage change in valence that would keep candidate j's vote share unaffected after a 1% change in her policy position (in the appropriate direction). Table 3 reports the first three quartiles of the distribution of r_η (weighted by vote share) by state.

State	Q1	Q2	Q3	State	Q1	Q2	Q3
Tocantins (to)	0.152	1.347	2.620	Distrito Federal (df)	0.116	0.436	2.301
Piauí (pi)	0.230	1.211	3.349	Goiás (go)	0.122	0.426	1.357
Paraíba (pb)	0.140	0.701	1.795	São Paulo (sp)	0.120	0.416	1.391
Rio Grande do Norte (rn)	0.134	0.606	1.574	Rio de Janeiro (rj)	0.062	0.410	2.277
Mato Grosso (mt)	0.128	0.577	1.813	Rio Grande do Sul (rs)	0.078	0.403	1.493
Acre (ac)	0.085	0.537	1.711	Amapá (ap)	0.156	0.368	1.560
Pará (pa)	0.083	0.502	1.675	Espírito Santo (es)	0.058	0.321	0.749
Bahia (ba)	0.068	0.497	1.734	Sergipe (se)	0.075	0.308	1.329
Roraima (rr)	0.069	0.493	1.775	Paraná (pr)	0.034	0.293	1.192
Maranhão (ma)	0.080	0.479	1.898	Amazonas (am)	0.083	0.227	0.871
Mato Grosso do Sul (ms)	0.162	0.469	1.127	Minas Gerais (mg)	0.029	0.200	0.618
Santa Catarina (sc)	0.105	0.441	2.237	Pernambuco (pe)	0.021	0.141	0.729
Rondônia (ro)	0.130	0.441	1.826	Alagoas (al)	0.008	0.061	0.239
Ceará (ce)	0.090	0.439	1.284	**Total**	**0.081**	**0.373**	**1.466**

Table 3: Quartiles of Ideology/Valence Candidate Vote-Share Elasticity Ratio

The median of r_η is below one in all but two states, indicating that—at the valence and policies observed in the data—voters tend to be considerably more sensitive to valence than policy. In fact, for the median candidate across the country, a 1% change in policy would require a compensating change of less than 0.4% in valence for their vote share to remain unaltered. This indicates that valence differentials in any given election weigh heavily on candidates' equilibrium policy choices. Nevertheless, there is considerable heterogeneity both across and within states: for candidates in the top quartile, a 1% change in policy would require a compensating increase in valence of more than 1.4%, whereas, for candidates in the bottom quartile, it would require an increase of less than 0.08%.

Political parties. Table A3 in Appendix A presents estimates of the value of party "brands" (β^{brands}).[19] Brazilian parties receive public funding and media time for campaign advertising in accordance with their performance in the most recent Chamber of Deputies election. Yet, if a particular party brand is not relevant for voters—carrying no information or affect—the corresponding coefficient should be zero. Indeed, consistent with the existing literature, for most parties we cannot reject the hypothesis that the party label does not affect voting behavior. We only estimate significant negative brand values for PRB, PV, and "minor" parties (those with a national vote share lower than 2.5%) and significant positive brand values for PMDB, PR, PSB, PSD, PSDB, and PT. Interestingly, PT and PMDB are the two major parties in the 2014 pro-government coalition, while PSDB and PSB are the main parties in the two opposition coalitions, *Muda Brasil* and *Unidos pelo Brasil*. Our results suggest that meaningful party-brand effects are mostly limited to major parties in the main electoral coalitions.

3.4 Quantifying Representation Failures

We now turn to our main objective of quantifying representation failures in Brazil. A prevalent approach in the political science literature has been to focus on whether there is divergence between elected candidates' policy positions and voters' stated policy preferences, as measured by surveys—see Miller and Stokes (1963), Erikson (1978), Clinton (2006), Bafumi and Herron (2010), Rogowski (2014). As a first step, we carry out a similar analysis in our data, for all the candidates competing in the election (the relevant choice set).

An informal assessment of the evidence suggests that substantive representation in Brazil is remarkably good for a large majority of voters (Saiegh 2015). Figure 4 plots the distribution of voters' preferred policies juxtaposed with the distribution of candidates' policy positions for four states, highlighting the three largest parties in each state along with the "minor" parties (according to their 2014 vote share). As the figure shows, the distribution of candidates' policy positions tracks reasonably well centrist voters' ideal points, and it is moderately shifted towards the largest mass of extreme voters in each state.

[19] Electoral coalitions among parties in Brazil are common. We account for potential coalition effects by letting the "party brand" of coalition candidates be the sum of their own party's and the mean of other parties' brands in the coalition.

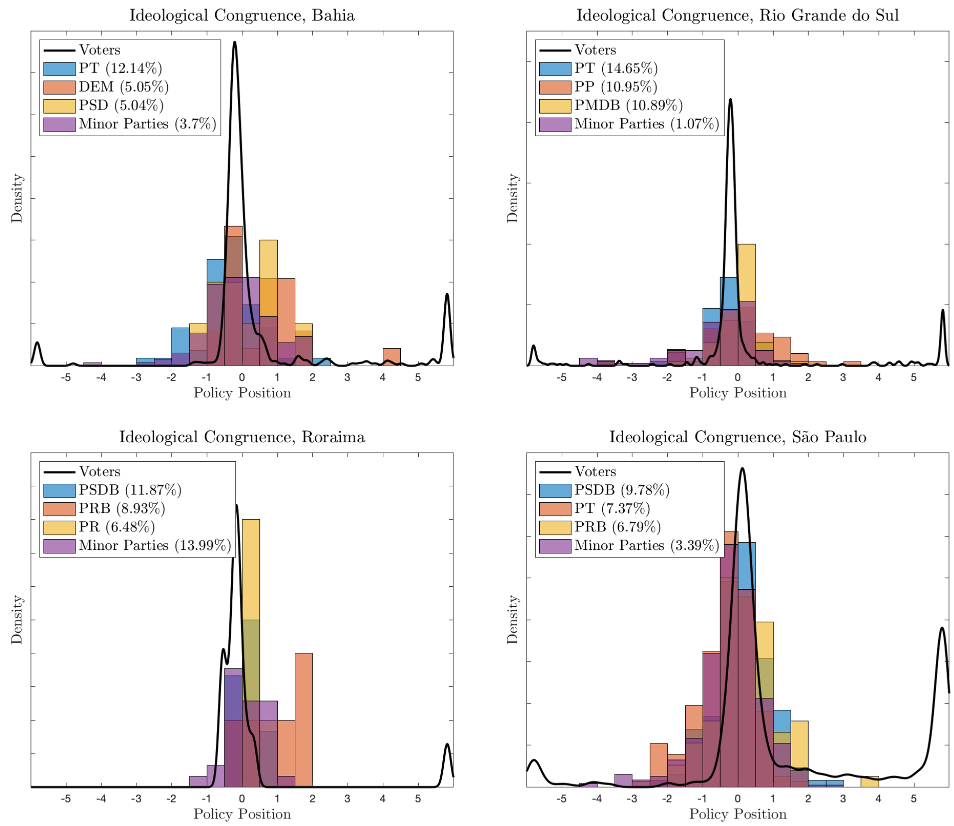

Figure 4: Distribution of Voters' Ideal Points and Candidates' Policy Positions (by Party in Four Selected States in 2014)—three largest parties' vote shares in parentheses along with "minor" ($\leq 2.5\%$ national vote share) parties'

When voters care about candidates' non-ideological attributes, though, the policy calculus misses a potentially important source of representation losses for voters. To attain a full account of representation failures, we propose a measure that accommodates both ideological and non-ideological considerations on a common scale, relying on voters' revealed preferences. Specifically, we compute the gap in voters' welfare given the actual set of candidates they face in the data relative to an *ideal* representation benchmark. Using voter welfare as a metric allows us to weigh losses in different dimensions according to the value that voters give to each attribute, thus comparing "apples to apples." And contrasting the actual welfare of each voter with an ideal benchmark provides a theoretically meaningful yardstick with which to quantify voters' losses. For our main results, we construct the ideal benchmark assuming that each voter is able to select her preferred candidate in all dimensions. We then complement these results with two alternative benchmarks that limit "ideal" candidates in different ways.

We begin by computing expected voter welfare in each municipality m given the set of candidates who participated in the 2014 election, U_m^d. To that end, we first compute the average utility voters in municipality m obtain from voting for each candidate j, δ_{jm}, evaluating (3.3) using our parameter estimates and municipality demographics, D_m. We then simulate a sample of registered voters for each municipality, drawing for each voter i preference shocks, ν_i, and random utility shocks, ϵ_{ijn}. For each simulation, we compute voter i's welfare (3.1) at her preferred candidate in the data (including abstention/blank vote) given her realized shocks. We then average over simulations to approximate the expected welfare of each voter, and we finally average over voters in each municipality.

To compute the ideal benchmark, U_m^*, we average the utility voters would derive from a hypothetical candidate with highest observed and unobserved in-sample valence and policy at their ideal point (for ease of exposition, and given the lack of heterogenous effects, we treat gender and incumbency status as "valence" characteristics). Using the realized and ideal measures of welfare, and letting U_m^{abs} denote the average utility from abstaining (which we have normalized to zero), we compute the total welfare loss in each municipality m as

$$WL_m = 1 - \frac{U_m^d - U_m^{abs}}{U_m^* - U_m^{abs}} = 1 - \frac{U_m^d}{U_m^*}.$$

Note that this measure of welfare loss is unchanged by affine transformations of utility.

Our results uncover a considerable failure of the Brazilian political system. The median welfare loss with respect to the ideal benchmark across 5,507 municipalities is 69%. That is, in 50% of municipalities, the average voter attains a level of welfare no higher than 31% of what they would enjoy in the ideal benchmark. Moreover, more than 75% of municipalities suffer a welfare loss of at least 53%, while 25% of municipalities suffer a loss of at least 84% relative to the benchmark—see Figure 5. To evaluate the robustness of these results to a potential violation of our assumptions due to endogenous entry, we re-estimate the model excluding all non-ideological observable candidate characteristics. We find a welfare loss for the median municipality of 66%, quantitatively similar to the result in our main benchmark.

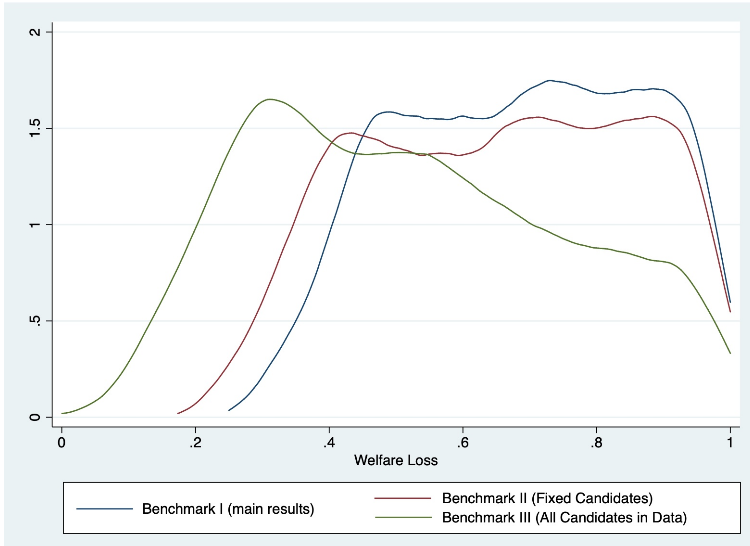

Figure 5: Distribution of Municipal Welfare Losses Relative to Alternative Benchmarks—arbitrary ideal candidates (Benchmark I), best set of candidates of size equal to number of candidates in the data (Benchmark II), and all candidates in the data (Benchmark III)

Alternative benchmarks. In our main benchmark of ideal representation, we allow each voter to select her preferred candidate with respect to all attributes. This provides a conceptually clear notion of the loss produced by limitations in the supply of candidates offered by the political system. In practice, however, the ideal benchmark may seem as imposing too heavy a burden on any political system. With this in mind, we consider two alternative benchmarks, which introduce more realistic constraints on voters' choice sets.

In Benchmark II, we restrict the number of candidates under consideration, ruling out the possibility of a personalized ideal candidate for each voter. Specifically, we limit the number of "ideal" candidates in each state to be equal to that observed in the data, and we select these candidates to maximize average voter welfare in the state. We find that the median welfare loss goes down only marginally, from 69% to 66%—see Figure 5. Thus, for the typical municipality, a large welfare loss is not the result of an undue inflation of the number of candidates in the ideal benchmark.

In Benchmark III, we dispense altogether with the notion of ideal candidates and instead compare welfare in the data with what voters would obtain if they were able to choose from among all (actual) candidates running in *any* state, with valence and policy positions as observed in the data.[20] In this case, we do observe notable changes in welfare losses. For the typical municipality, the welfare loss goes down from 69% in Benchmark I to 50% in Benchmark III. Welfare losses are still very large, to be sure, as 50% of all Brazilian municipalities' welfare is less than half of that in the benchmark. Thus, most the welfare loss captured with Benchmark I remains when we consider alternatives that are certainly feasible in Brazil's political system. However, Benchmark III suggests that voters in a subset of states are particularly impacted by deficiencies in the set of candidates they face.

Policy-valence decomposition. While the education, experience, and other valence attributes of the pool of candidates can be taken as fixed in the short run, candidates can freely choose their policy positions. Do competitive forces lead to ideological congruence between voters and candidates? To address this question, we decompose the total welfare loss in Benchmark I for the average voter in each municipality into two components: the loss due to ideological incongruence between candidates and voters (*Policy Welfare Loss*) and that due to valence (*Valence Welfare Loss*). To carry out this decomposition, we first compute an intermediate level of welfare for the average voter in each municipality from a hypothetical election in which all candidates have maximum valence, as in the ideal benchmark, but choose policies as in the 2014 election, U_m^{val}. The percentage difference between welfare at the ideal benchmark and this intermediate welfare value can be interpreted as the fraction of the welfare gap due solely to ideological incongruence. On the other hand, the difference between the intermediate welfare value and observed welfare can be attributed solely to valence:

$$WL_m = \underbrace{\frac{U_m^* - U_m^{val}}{U_m^*}}_{\text{Policy WL}} + \underbrace{\frac{U_m^{val} - U_m^d}{U_m^*}}_{\text{Valence WL}}.$$

Table 4 summarizes the results. As foreshadowed by Figure 4, the decomposition shows that, for a large fraction of municipalities, the valence welfare loss constitutes the brunt of the total welfare loss. In fact, for the median municipality, the valence welfare loss is more than seven times larger than the policy welfare loss, and, for three

[20]This benchmark should not be taken as a counterfactual of what would occur under a single national district, as candidates could adjust their policy positions to the new competitive environment.

fourths of municipalities, the valence welfare loss is more than five times larger than the policy welfare loss.[21] Large policy welfare losses do occur but are concentrated in a small fraction of municipalities, in a small set of states: Alagoas, Amapa, Distrito Federal, Pernambuco, Rio de Janeiro, and São Paulo. The 10% worst-performing municipalities in this regard suffer a policy welfare loss of more than 54%. However, for three fourths of all municipalities, the policy welfare loss is below 10%. In contrast, the picture is dramatically different for valence welfare losses: for the median municipality, the welfare loss due to valence is 52% of the ideal benchmark, and it is above 69% for 25% of municipalities.

	Mean	St.Dev.	10%	25%	50%	75%	90%
Total Welfare Loss (TWL)	0.684	0.180	0.439	0.534	0.694	0.839	0.927
Policy Welfare Loss (PWL)	0.154	0.200	0.041	0.050	0.069	0.102	0.541
Valence Welfare Loss (VWL)	0.531	0.210	0.265	0.397	0.524	0.694	0.813
VWL/PWL	8.44	6.44	0.53	5.24	7.51	10.43	16.14

Table 4: Policy and Valence Welfare Losses Across Municipalities

In Table 5, we regress our measures of welfare loss on municipality characteristics. The median wage, level of education, average age of electorate, and proportion of female voters in each municipality have countervailing associations with policy and valence welfare losses (columns III and V): the policy welfare loss is higher and the valence welfare loss is lower in municipalities that are richer, less educated, younger, and more predominantly female. Although these estimates somewhat offset each other in the total welfare loss (column I), the latter is larger in municipalities that are less educated, older, and more predominantly male, and it is statistically unresponsive to the median wage. In contrast, an indicator of how rural a municipality is associates negatively only with the policy welfare loss. As a result, the total welfare loss is larger in more urban municipalities. Finally, a municipality's employment rate has a sharp negative association with both policy and valence welfare losses. Thus, total welfare loss, policy welfare loss, and valence welfare loss are all larger in municipalities with higher levels of unemployment, and these estimates are consistently the largest in magnitude. Overall, our results suggest a strong positive association between

[21] This conclusion is qualitatively unchanged if we compute welfare losses only considering elected candidates (for the typical municipality, the valence welfare loss is more than five times larger than the policy welfare loss).

economic and political well-being, with large political welfare losses in municipalities that are older, more urban, less educated, and suffering higher levels of unemployment.

	Total WL		Policy WL		Valence WL	
	(I)	(II)	(III)	(IV)	(V)	(VI)
Median Wage	-0.003		0.080		-0.083	
	(0.002)		*(0.004)*		*(0.005)*	
% Rural	-0.071		-0.071		0.000	
	(0.002)		*(0.003)*		*(0.003)*	
% Higher Edu	-0.006		-0.014		0.008	
	(0.002)		*(0.003)*		*(0.004)*	
% Employed	-0.164		-0.073		-0.091	
	(0.002)		*(0.003)*		*(0.003)*	
Avg. Age	0.069		-0.018		0.087	
	(0.002)		*(0.004)*		*(0.004)*	
% Female	-0.036		0.018		-0.054	
	(0.001)		*(0.003)*		*(0.003)*	
Avg. Muni Idealpoint		0.002		0.003		-0.001
		(0.001)		*(0.000)*		*(0.001)*
Avg. Muni Idealpoint Sq.		0.006		0.017		-0.011
		(0.000)		*(0.000)*		*(0.000)*
Constant	0.682	0.654	0.152	0.067	0.530	0.586
	(0.001)	*(0.002)*	*(0.002)*	*(0.001)*	*(0.002)*	*(0.002)*
Obs.	5507	5507	5507	5507	5507	5507
State FE	Y	Y	Y	Y	Y	Y
R Sq.	0.634	0.118	0.453	0.931	0.251	0.390

Table 5: Welfare Losses and Municipality Characteristics

Columns II, IV, and VI of Table 5 show how our measures of welfare loss relate to the average voter ideal point in each municipality (itself a function of socioeconomic characteristics). The results indicate that the policy welfare loss increases in municipalities that have extreme policy preferences, while the valence welfare loss is larger in more ideologically moderate municipalities. To understand why this is the case, note that, in the model, candidates' vote shares reflect how attractive candidates are relative to their competitors (including abstention). All else equal, candidates who perform better (worse) at the polls must be championing a position that is attractive to voters or must have high (low) valence. Thus, if a candidate offers a policy position in a region of the policy space heavily populated by voters but performs poorly at the polls, we must infer that the candidate has low valence.

Our model estimates reflect this logic, as the relative value of each attribute is chosen to explain variation in vote shares across candidates. Recall that a large fraction of voters have moderate ideology, and this area of the policy space is covered by a large number of candidates in each state (see Figure 4), leading to a low policy welfare loss for moderate voters. For these voters to achieve a low total welfare loss, we would need *some* of these moderate candidates in each state to have high valence. However, the data reveals that, in many states, moderate candidates tend to underperform at the polls, even after accounting for the intense ideological competition they face from a multitude of close substitutes. Thus, these candidates are inferred to have relatively low valence. This results in high valence welfare losses for moderate voters, despite being well served along the policy dimension. Similarly, the data reveal that more extreme candidates, who appeal to a non-negligible minority of extreme voters, tend to outperform their policy-based advantage. As a result, valence welfare losses tend to be low when policy welfare losses are high.[22]

Welfare losses and political protests. Finally, we explore whether our measures of voter welfare loss correlate with a prominent expression of citizen dissatisfaction: political protests. The Armed Conflict Location & Event Data Project (ACLED) collects information on reported political violence and protest events around the world. For Brazil, ACLED's database records events since January 1st, 2018. In Table 6, for municipalities in our sample, we regress a binary indicator of whether a political protest occurred between January 1st and July 19th of 2018 on our measures of voter welfare loss and municipality characteristics.[23]

Across the board, point estimates in Table 6 are positive and relatively large in magnitude. Without controlling for observed municipality characteristics (only state fixed effects), column I indicates that a one percentage-point increase in total welfare loss is associated with a 0.9 percentage-point increase in the likelihood of a political protest.[24] Uncoupling policy and valence, column II suggests that political protests are twice as sensitive to policy welfare losses than to valence welfare losses. Columns III and IV show analogous results after controlling for the municipality characteristics

[22] Figure A3 in Appendix A plots candidates' estimated overall valence against their policy positions, illustrating that indeed high-valence candidates tend to adopt more extreme policy positions. A potential concern given the large number of candidates in our data is that this may be artificially driven by top-performing candidates. Figure A4 plots average voter utility from observed and unobserved valence for each candidate against their vote share. As shown, the top performing candidates are not outliers in either feature, and as such they are not the source of the estimated negative association between policy and valence welfare losses.

[23] Recall that our voter welfare estimates are based on the 2014 pool of candidates. The candidate nomination process for the next electoral cycle began on July 20th, 2018.

[24] In Table 6, we consider both peaceful protests and violent riots. Results are nearly identical if we focus on peaceful protests.

	Political Protest (mean: 0.270)			
	(I)	(II)	(III)	(IV)
TWL ($\times$100)	0.009		0.001	
	(0.001)		*(0.001)*	
PWL ($\times$100)		0.011		0.003
		(0.001)		*(0.001)*
VWL ($\times$100)		0.005		0.000
		(0.001)		*(0.001)*
Observations	5507	5507	5507	5507
Demographics	N	N	Y	Y
State FE	Y	Y	Y	Y
R Sq.	0.147	0.203	0.258	0.266

Table 6: Welfare Losses and Political Protests

in Table 5. Although coefficient estimates attenuate considerably, there remains a positive, large, and statistically significant association between our proposed measure of policy welfare loss and an important symptom of political discontent.

4 Supply-Side Politics and Institutional Reforms

Having documented that the Brazilian political system induces large welfare losses for voters—particularly *valence* welfare losses—we turn to evaluating possible institutional reforms designed to remedy representation failures: e.g, imposing education requirements. A key consideration in evaluating any change in the non-ideological characteristics of candidates, however, is that candidates may adjust their policy choices to the new environment. Thus, evaluating the full consequences of a reform requires taking into consideration both its direct and indirect (equilibrium) effects on voter welfare. In our estimation of voters' preferences, we address the endogeneity of candidates' policy positions with an instrumental-variables approach. We now take preference estimates as given and tackle the task of estimating a model of the "supply side" of politics, where candidates' positions emerge explicitly as equilibrium choices (Section 4.1). With estimates of both the "demand" and "supply" sides of politics, it is then possible to conduct counterfactual analyses of how the system would work under different conditions from those observed in the data (Section 4.2).

4.1 Policy Choice in Electoral Competition

In this section, we develop an empirical model of electoral competition among multiple candidates in an open-list PR electoral system. We model candidates' choices as emerging from a strategic balance between their own policy preferences and electability. The latter has (possibly) two components. Candidates wish to maximize their individual vote share to further their chance of obtaining a seat in the legislature. But parties may also exert some influence making candidates internalize the externalities their policy choices impose on fellow party members' vote shares.[25]

While the tradeoff between electability and ideology is at the core of many models of electoral competition, three points are noteworthy. First, differently from standard models of competition in majoritarian electoral systems, in which typically only two candidates compete for office, in our setup candidates face a large number of competitors. Thus, the key role of the median voter in standard models is replaced by more complex patterns of substitutability across candidates, which are pinned down by the cross-candidate elasticities we recover with our "demand" estimates. Second, consistent with our results in Section 3, our model is one in which candidates have valence differentials. In this setting, candidates with a valence advantage have an incentive to adopt a policy close to that of disadvantaged competitors, in order to neutralize policy differentials and make the election predominantly about valence. In majoritarian elections with two candidates, this leads to the prediction that the advantaged candidate can claim the center of the policy space, relegating the opponent to more extreme positions (Ansolabehere and Snyder 2000, Groseclose 2001, Aragones and Palfrey 2002). In our setup, this translates—all else equal—to disadvantaged candidates being displaced to positions that are ex-ante less popular with voters or that face stronger competition. Third, in our model parties may influence candidates to internalize the externalities their policy choices impose on fellow party members' vote shares.

Model. There are $L \geq 2$ parties (lists) and $J_n^\ell \geq 1$ candidates representing party $\ell = 1, \ldots, L$ in state $n = 1, \ldots, N$. Candidates are differentiated with respect to non-ideological attributes and choose their policy positions simultaneously. Thus, candidates are exogenously differentiated along a vertical dimension and endogenously

[25] Our results are unchanged if we conduct this analysis at the coalition level rather than at the party level, which suggests the key tradeoffs occur within parties—see Appendix D.3.

differentiated along a horizontal dimension.[26] Letting $\mathbf{p}_n^{-j}$ denote the vector of policy positions of all candidates in state n excluding $j \in J_n^\ell$, and letting ρ_{jn} denote candidate j's ideal policy, we assume j's payoff is

$$\Pi_{jn}(p_{jn}, \mathbf{p}_n^{-j}) = s_{jn}(p_{jn}, \mathbf{p}_n^{-j}) + \gamma_{jn} \sum_{j' \in J_n^\ell} s_{j'n}(p_{jn}, \mathbf{p}_n^{-j}) - \mu_{jn}|p_{jn} - \rho_{jn}|, \tag{4.1}$$

where $\mu_{jn} \in \mathbb{R}_+$ denotes the weight candidate j puts on her own ideology vis-à-vis electability. Note that $\gamma_{jn} \in \mathbb{R}_+$ captures the extent to which j internalizes the effect of her policy choice on the party's aggregate performance. A candidate with $\mu_{jn} = \gamma_{jn} = 0$ would solely choose policy to maximize her own vote share, whereas a candidate with $\mu_{jn} = 0$ and $\gamma_{jn} = 1$ would put equal weight on her party's aggregate vote share. Larger values of μ_{jn} would lead the candidate to place more emphasis on matching her ideal policy, ρ_{jn}, disregarding votes to her or the party. A Nash equilibrium is a profile of policies $\mathbf{p}$ such that $p_{jn} \in \arg\max_{\tilde{p}_{jn} \in [-\overline{p}, \overline{p}]} \Pi_{jn}(\tilde{p}_{jn}, \mathbf{p}_n^{-j})$ for all n and each $j \in J_n$.

To specify the empirical model, we make the following assumptions. First, we let

$$\mu_{jn} = \exp\left(\tilde{X}'_{jn}\chi + \zeta_{jn}\right), \tag{4.2}$$

where $\tilde{X}_{jn}$ is a vector of candidate characteristics that includes j's unobserved valence, ξ_{jn}, and ζ_{jn} is an idiosyncratic shock observed by candidates but not by the analyst. Second, for $j \in J_n^\ell$, we assume that

$$\gamma_{jn} = \gamma^\ell + \gamma^{\text{inc}} \tilde{I}_{jn}, \tag{4.3}$$

where γ^ℓ is a party fixed effect and $\tilde{I}_{jn}$ is a binary indicator of candidate j's incumbency status. Third, we assume that the ideal policies of party ℓ's candidates in district n are distributed $\rho_{jn} \sim N(\rho_n^\ell, (\sigma_n^\ell)^2)$, where both the mean, ρ_n^ℓ, and standard deviation, σ_n^ℓ, are functions of state demographic characteristics, which we estimate.

Equilibrium policies satisfy the following system of necessary first-order condi-

[26] See Iaryczower and Mattozzi (2013) for a model of electoral competition with multiple candidates with *endogenous* horizontal and vertical differentiation.

tions: for each candidate $j \in J_n^\ell$ in each party ℓ and state n,

$$MB_{jn}(\mathbf{p}_n, \gamma) \equiv \left| \frac{\partial s_{jn}(\mathbf{p}_n)}{\partial p_{jn}} + (\gamma^\ell + \gamma^{\text{inc}} \tilde{I}_{jn}) \sum_{j' \in J_n^\ell} \frac{\partial s_{j'n}(\mathbf{p}_n)}{\partial p_{jn}} \right| = \mu_{jn}. \tag{4.4}$$

In equilibrium, each candidate adopts a policy position such that the marginal benefit in terms of electability (of the candidate and possibly the party) equals the marginal ideological cost, μ_{jn}, of moving away from the candidate's ideal policy.

Estimation. We estimate the parameters γ and χ in (4.2) and (4.3) via GMM, exploiting the equilibrium conditions (4.4). Let $r_{jn}(\gamma) = \log\left(MB_{jn}(\mathbf{p}_n, \gamma)\right)$. Taking logs of (4.4) and substituting (4.2), we can write the equilibrium conditions as

$$\zeta_{jn} = r_{jn}(\gamma) - \tilde{X}_{jn}' \chi. \tag{4.5}$$

Note that, given γ, all components of r_{jn} are known from the data or from demand-side estimates. We can then recover coefficients (γ, χ) with a GMM approach analogous to our demand-side estimation. Given a vector of instruments $\tilde{Z}_{jn}$ such that

$$E[\zeta_{jn}(\gamma, \chi) | \tilde{Z}_{jn}] = 0 \quad \text{if and only if} \quad (\gamma, \chi) = (\gamma_0, \chi_0),$$

where (γ_0, χ_0) denotes the true value of the parameters, a GMM estimator is obtained by minimizing the (positive-definite) quadratic form

$$\tilde{Q}_J(\gamma, \chi) = \left[\frac{1}{J} \tilde{Z}' \zeta(\gamma, \chi)\right]' \tilde{W}_J \left[\frac{1}{J} \tilde{Z}' \zeta(\gamma, \chi)\right].$$

As in the demand case, we implement an (optimally-weighted) MPEC version of this estimator for computational convenience. For inference, we rely on standard results for two-step GMM estimation (Newey and McFadden 1994) to account for demand-side estimation uncertainty in $r_{jn}(\gamma)$.

The choice of instruments to identify (γ, χ) follows intuition similar to the demand case. Again, a necessary order condition is that $\tilde{Z}_{jn}$ must include at least as many variables as there are parameters to be estimated. The exogenous candidate characteristics in $\tilde{X}_{jn}$ are valid (optimal) instruments to identify χ. For γ, since the coefficients enter the moment conditions in a nonlinear fashion, instrument choice is not as straightforward. In a first iteration, we simply use party dummies and $\tilde{I}_{jn}$. We then implement an approximation of Chamberlain (1987)'s optimal instruments. See Appendix D for technical details.

Results. Table 7 presents our estimates. Two main findings emerge. First, with the exception of business experience, candidates with favorable valence attributes (unobserved valence, higher education, no government experience, male) give a larger weight to their own ideology relative to catering to the preferred policy positions of the electorate. This partially undoes the strategic centrality induced by their valence advantage. In contrast, business experience has the opposite effect, suggesting candidates with this background are more pragmatic and less ideological, although the coefficient is not statistically significant.

μ: Weight of Ideology Relative to Electability			
Constant	-6.347	Higher Education	0.845
	(75.32)		*(0.921)*
Age	0.127	Business Exp.	-0.162
	(0.295)		*(0.497)*
Age Squared	-0.013	Gov. Experience	-0.807
	(0.399)		*(0.522)*
Female Candidate	-1.073	Unobserved Valence	0.519
	(0.558)		*(0.070)*
γ: Weight of Party Vote Share Relative to Own Vote Share			
DEM	0.110	PDT	0.348
	(86.48)		*(103.2)*
MDB	0.000	PP	0.109
	(82.83)		*(87.31)*
PR	0.539	PRB	0.000
	(119.5)		*(77.30)*
PSB	0.184	PSD	0.600
	(89.78)		*(123.5)*
PSDB	0.324	PT	0.003
	(105.4)		*(83.06)*
PTB	0.264	Incumbent	-0.735
	(97.16)		*(56.46)*

Table 7: "Supply-Side" Coefficient Estimates (parties with at least three million votes). Standard errors are clustered at the party-state-year level.

Our second result concerns the extent to which candidates internalize the externalities they impose on fellow party members. We interpret this as a measure of party discipline in this electoral context. The point estimates suggest that there are non-trivial differences across parties, with PT and PMDB (the top-two parties in the pro-governing coalition in 2014) having estimates at essentially zero, while PSDB and PSB (the top parties in the two opposition coalitions) have positive effects (0.32 and 0.18). All party coefficients, however, are imprecisely estimated, so the hypothesis that discipline is similar across parties cannot be rejected. As expected, the incumbency coefficient is negative and large (-0.73), indicating incumbents are subject to weaker party discipline, although the estimate is again imprecise.[27]

4.2 Counterfactuals: Institutional Reforms

With our demand and supply estimates in hand, we now evaluate alternative institutional reforms aimed at boosting the quality of representation. First, in light of the prominence of valence in our voter welfare analysis, we consider qualification requirements designed to directly improve non-ideological characteristics in the pool of candidates. While any such change would seem obviously beneficial to voters, our supply-side results caution that these reforms might have unintended consequences through candidates' policy choices, leading to lower, or even negative, welfare effects. To account for both the direct and indirect consequences of institutional changes, we use our full equilibrium model of policy choice and voter demand. Similarly, in our second counterfactual experiment, we consider the impact of strengthening Brazilian parties' influence over their candidates' policy choices. See Appendix D.4 for technical details.

Minimal education requirements. We first quantify the change in voter welfare resulting from an institutional reform requiring candidates to have completed higher education. To reduce the computational burden, we focus our analysis on the state of Bahia, whose demographics are most representative of the nation as a whole. Moreover, in the data, the proportion of candidates running for office with a university degree is 60% for both the entire country and the state of Bahia. To illustrate the direct and equilibrium effects of the reform, we present two sets of results. First, we compute changes in welfare keeping candidates' policy positions fixed as observed in the data (the direct effect). Specifically, we draw from the empirical distribution of candidates with a university degree a menu of size equal to the observed number of

[27] Since elections in the data are very large, individual candidates' policy choices tend to have limited influence on fellow party members' vote shares, which dampens variation in $r_{jn}(\gamma)$ with which to identify γ.

candidates in Bahia in 2014, we calculate average voter welfare in each municipality as described in our welfare analysis above, and we compare the resulting level of welfare with what voters attain in the data (U_m^d). Second, we re-evaluate welfare changes after letting candidates optimally adjust their policies to the new menu of competitors (the equilibrium effect).

Figure 6 plots the distribution of the percentage change in welfare for each municipality corresponding to the direct and total effects of the reform. Keeping candidates' policies as observed in the data, the higher-education mandate leads to a 14.9% increase in welfare for the typical municipality. The effect is non-negative across the board. Three quarters of all municipalities in Bahia have an increase in average welfare above 10.8%, and a quarter have an increase above 20.3%. However, the impact of the reform is considerably different when we consider equilibrium effects. The typical municipality still benefits, but the increase in welfare goes down from 14.9% to 5.7%. Three quarters of municipalities have a welfare gain below 9.7%, and twelve percent of municipalities suffer a welfare *loss*. Overall, the reform is beneficial for the vast majority of municipalities, but the equilibrium effects are non-negligible and lead to a downward shift in the distribution of welfare changes.

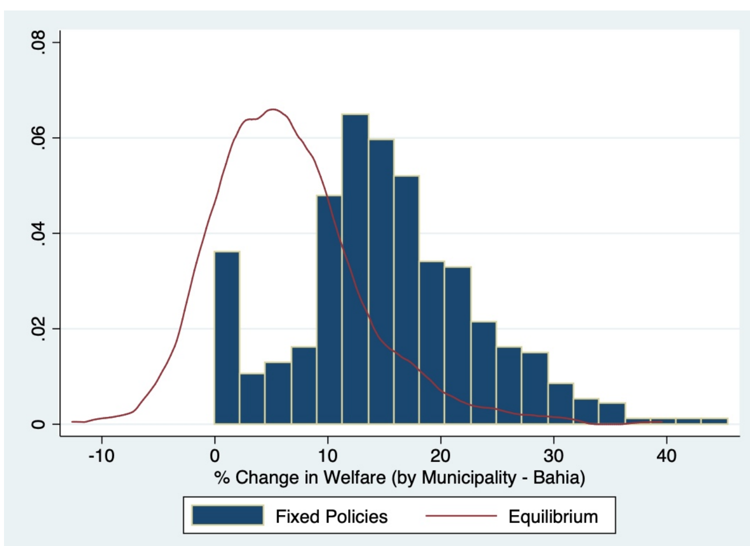

Figure 6: Distribution of Percentage Change in Municipal Welfare Due to a Higher-Education Mandate (State of Bahia, 2014)

Party discipline. Brazilian parties are generally considered to be weak (Mainwaring, Scully, et al. 1995, Samuels 2003). In our model, there are two potential channels of party influence. First, on the demand side, party brands may shape voting decisions. Second, on the supply side, parties may encourage candidates to internalize how their policy choices affect fellow party members. We refer above to the second effect as party discipline.

Our results suggest Brazilian parties indeed are weak on both counts. We now consider the consequences of strengthening party elites relative to rank-and-file candidates. Specifically, we compute the equilibrium policy choices that would result from raising party discipline to $\gamma_{jn} = 1$ for all candidates. We then compare voter welfare in this counterfactual and in the data. Note that, unlike the education counterfactual, in this instance there are no direct effects of the reform—the full change in welfare is due to equilibrium adjustments. As before, we focus on the state of Bahia.

Figure 7 plots the distribution of the percentage change in welfare for each municipality resulting from increased party discipline. We find this benefits the average voter in 83% of municipalities, but it reduces average voter welfare in the remaining 17%. The typical municipality experiences a 16.4% increase in average voter welfare, with three quarters of municipalities gaining at least 4.6% and a quarter gaining at least 21.5%. Our results reveal that strengthening political parties can be welfare improving for a majority of voters, but these gains can come at the expense of welfare losses for a minority of voters. Overall, our counterfactual experiments show that indirect or equilibrium effects can be substantial, with significant distributional implications, and should not be glossed over when evaluating the potential consequences of institutional reforms.

5 Conclusion

A well-working democracy requires that voters have access to options *they* value. In order to assess the extent to which a political system is satisfying the demands of its citizens, we first need to understand what is valuable to voters. A standard approach in the political science literature has been to focus on the level of congruence between voters' preferences and politicians' policy positions. Voters, however, generally also care about other candidate characteristics, including their education, readiness for office, gender, charisma, and trustworthiness. In this paper, we develop a methodology to gauge representation failures that accommodates ideological and non-ideological considerations, quantifying the relative importance of deficits in each dimension from

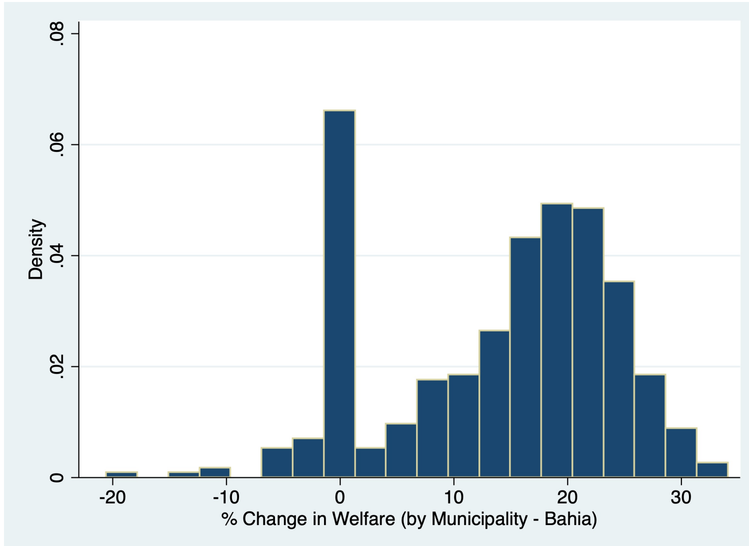

Figure 7: Distribution of Percentage Change in Municipal Welfare Due to an Increase in Party Discipline to $\gamma^{\ell} = 1$ for all Parties $\ell \in L$ (State of Bahia, 2014)

the voter's perspective.

We implement our approach in the context of elections for the lower house of Brazil's National Congress. Our results uncover a considerable failure of the Brazilian political system. To understand the sources of voters' welfare losses, we decompose the total welfare loss into a *policy welfare loss* (due to incongruence between voters' preferred policies and candidates' positions) and a *valence welfare loss* (due to inferior non-ideological characteristics of candidates). We show that, for the typical municipality, the valence welfare loss comprises the brunt of the total welfare loss.

To evaluate institutional reforms aimed at improving the quality of representation, we develop and estimate a model of the "supply side" of politics, where candidates' policy positions emerge explicitly as equilibrium choices. We conduct two counterfactual experiments. In the first, we consider a reform designed to directly alter valence in the pool of candidates (minimal education requirements). In the second, we consider reforms aimed at strengthening political parties' influence over their candidates' policy choices. Our experiments show that both types of reforms can considerably improve the welfare of a vast majority of voters. However, the reforms can have countervailing indirect effects, through candidates' equilibrium policy choices, with significant distributional implications. Thus, indirect equilibrium-adjustment effects should not be glossed over when evaluating potential institutional reforms.

References

Ansolabehere, S., and J. M. Snyder (2000): "Valence politics and equilibrium in spatial election models," *Public Choice*, 103(3), 327–336.

Aragones, E., and T. R. Palfrey (2002): "Mixed Equilibrium in a Downsian Model with a Favored Candidate," *Journal of Economic Theory*, 103, 131–161.

Bafumi, J., and M. C. Herron (2010): "Leapfrog representation and extremism: A study of American voters and their members in Congress," *American Political Science Review*, pp. 519–542.

Beath, A., F. Christia, G. Egorov, and R. Enikolopov (2016): "Electoral rules and political selection: Theory and evidence from a field experiment in Afghanistan," *The Review of Economic Studies*, 83(3), 932–968.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile prices in market equilibrium," *Econometrica: Journal of the Econometric Society*, pp. 841–890.

Besley, T., and M. Reynal-Querol (2011): "Do democracies select more educated leaders?," *American political science review*, pp. 552–566.

Boas, T. C., F. D. Hidalgo, and N. P. Richardson (2014): "The spoils of victory: campaign donations and government contracts in Brazil," *The Journal of Politics*, 76(2), 415–429.

Bonica, A. (2014): "Mapping the ideological marketplace," *American Journal of Political Science*, 58(2), 367–386.

Bowen, D. C., and C. J. Clark (2014): "Revisiting descriptive representation in Congress: Assessing the effect of race on the constituent–legislator relationship," *Political Research Quarterly*, 67(3), 695–707.

Buttice, M. K., and W. J. Stone (2012): "Candidates matter: Policy and quality differences in congressional elections," *The Journal of Politics*, 74(3), 870–887.

Cameron, C., D. Epstein, and S. O'halloran (1996): "Do majority-minority districts maximize substantive black representation in Congress?," *American Political Science Review*, 90(4), 794–812.

Chamberlain, G. (1987): "Asymptotic efficiency in estimation with conditional moment restrictions," *Journal of Econometrics*, 34(3), 305–334.

Clinton, J. D. (2006): "Representation in Congress: constituents and roll calls in the 106th House," *Journal of Politics*, 68(2), 397–409.

Cruz, C., P. Keefer, J. Labonne, and F. Trebbi (2018): "Making policies matter: Voter responses to campaign promises," Discussion paper, National Bureau of Economic Research.

Desposato, S. W. (2006): "Parties for rent? Ambition, ideology, and party switching in Brazil's chamber of deputies," *American Journal of Political Science*, 50(1), 62–80.

Dubé, J.-P., J. T. Fox, and C.-L. Su (2012): "Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation," *Econometrica*, 80(5), 2231–2267.

Erikson, R. S. (1978): "Constituency opinion and congressional behavior: A reexamination of the Miller-Stokes representation data," *American Journal of Political Science*, pp. 511–535.

Ferraz, C., and F. Finan (2011): "Electoral accountability and corruption: Evidence from the audits of local governments," *The American Economic Review*, 101(4), 1274–1311.

Folke, O., T. Persson, and J. Rickne (2016): "The primary effect: Preference votes and political promotions," *American Political Science Review*, 110(3), 559–578.

Galasso, V., and T. Nannicini (2011): "Competing on good politicians," *American political science review*, 105(1), 79–99.

Gandhi, A., and J.-F. Houde (2020): "Measuring Substitution Patterns in Differentiated-Products Industries," Working Paper.

Gordon, B. R., and W. R. Hartmann (2013): "Advertising effects in presidential elections," *Marketing Science*, 32(1), 19–35.

Griffin, J. D., and B. Newman (2007): "The unequal representation of Latinos and whites," *The Journal of Politics*, 69(4), 1032–1046.

Groseclose, T. (2001): "A model of candidate location when one candidate has a valence advantage," *American Journal of Political Science*, pp. 862–886.

Hansen, L. P. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50(4), 1029–1054.

Hero, R. E., and C. J. Tolbert (1995): "Latinos and substantive representation in the US House of Representatives: Direct, indirect, or nonexistent?," *American Journal of Political Science*, pp. 640–652.

Iaryczower, M., and A. Mattozzi (2013): "On the nature of competition in alternative electoral systems," *The Journal of Politics*, 75(3), 743–756.

Kawai, K., and T. Sunada (2022): "Estimating candidate valence," Discussion paper, National Bureau of Economic Research.

Kendall, C., and J. Matsusaka (2021): "The common good and voter polarization," Discussion paper, Mimeo, University of Southern California.

Kendall, C., T. Nannicini, and F. Trebbi (2015): "How do voters respond to information? Evidence from a randomized campaign," *American Economic Review*, 105(1), 322–53.

Klašnja, M., and R. Titiunik (2017): "The incumbency curse: Weak parties, term limits, and unfulfilled accountability," *American Political Science Review*, 111(1), 129–148.

Mainwaring, S., T. Scully, et al. (1995): *Building democratic institutions: Party systems in Latin America*. Stanford University Press Stanford.

Miller, W. E., and D. E. Stokes (1963): "Constituency influence in Congress," *American political science review*, 57(1), 45–56.

Montero, S. (2023): "Going It Alone? An Empirical Study of Coalition Formation in Elections," Conditionally accepted at *Journal of Politics*.

Newey, W. K., and D. McFadden (1994): "Large Sample Estimation and Hypothesis Testing," in *Handbook of Econometrics*, ed. by R. F. Engle, and D. L. McFadden, vol. IV, chap. 36, pp. 2111–2245. Elsevier.

Power, T. J., and R. Rodrigues-Silveira (2019a): "Mapping Ideological Preferences in Brazilian Elections, 1994-2018: A Municipal-Level Study," *Brazilian Political Science Review*, 13(1).

——— (2019b): "Mapping ideological preferences in Brazilian elections, 1994-2018: a municipal-level study," *Brazilian Political Science Review*, 13(1).

Power, T. J., and C. Zucco (2012): "Elite preferences in a consolidating democracy: the Brazilian legislative surveys, 1990–2009," *Latin American Politics and Society*, 54(4), 1–27.

Rekkas, M. (2007): "The impact of campaign spending on votes in multiparty elections," *The Review of Economics and Statistics*, 89(3), 573–585.

Rogowski, J. C. (2014): "Electoral choice, ideological conflict, and political participation," *American Journal of Political Science*, 58(2), 479–494.

Saiegh, S. M. (2015): "Using Joint Scaling Methods to Study Ideology and Representation: Evidence from Latin America," *Political Analysis*, 23, 363–384.

Samuels, D. (2003): *Ambition, federalism, and legislative politics in Brazil*. Cambridge University Press.

Stokes, D. E. (1963): "Spatial Models of Party Competition," *American Political Science Review*, 57, 368–377.

Ujhelyi, G., S. Chatterjee, and A. Szabó (2018): "None of the above," Working paper, University of Houston.

Zucco, C. (2009): "Ideology or what? Legislative behavior in multiparty presidential settings," *The Journal of Politics*, 71(3), 1076–1092.

Zucco, C., and B. E. Lauderdale (2011): "Distinguishing between influences on Brazilian legislative behavior," *Legislative Studies Quarterly*, 36(3), 363–396.

Online Appendix for *Representation Failure*

Contents

A Additional Tables and Figures

State	Representatives Number	Representatives %	Population No.	Population %	District Mag. Pop / Legs
São Paulo (sp)	70	13.6%	39,924,091	21.5%	570,344
Minas Gerais (mg)	53	10.3%	19,159,260	10.3%	361,495
Rio de Janeiro (rj)	46	9.0%	15,180,636	8.2%	330,014
Bahia (ba)	39	7.6%	13,633,969	7.3%	349,589
Rio Grande do Sul (rs)	31	6.0%	10,576,758	5.7%	341,186
Paraná (pr)	30	5.8%	10,226,737	5.5%	340,891
Pernambuco (pe)	25	4.9%	8,541,250	4.6%	341,650
Ceará (ce)	22	4.3%	8,450,527	4.4%	371,822
Maranhão (ma)	18	3.5%	6,424,340	3.5%	356,908
Goiás (go)	17	3.3%	5,849,105	3.1%	344,065
Pará (pa)	17	3.3%	7,443,904	4.0%	437,877
Santa Catarina (sc)	16	3.1%	6,178,603	3.3%	386,163
Paraíba (pb)	12	2.3%	3,753,633	2.0%	312,803
Espírito Santo (es)	10	1.9%	3,392,775	1.8%	339,278
Piauí (pi)	10	1.9%	3,086,448	1.7%	308,645
Alagoas (al)	9	1.7%	3,093,994	1.7%	343,777
Amazonas (am)	8	1.6%	3,350,773	1.8%	418,847
Rio Grande do Norte (rn)	8	1.6%	3,121,451	1.7%	390,181
Mato Grosso (mt)	8	1.6%	2,954,625	1.6%	369,328
Distrito Federal (df)	8	1.6%	2,469,489	1.3%	308,686
Mato Grosso do Sul (ms)	8	1.6%	2,404,256	1.3%	300,532
Sergipe (se)	8	1.6%	2,036,227	1.1%	254,528
Rondônia (ro)	8	1.6%	1,535,625	0.8%	191,953
Tocantins (to)	8	1.6%	1,373,551	0.7%	171,694
Acre (ac)	8	1.6%	707,125	0.4%	88,391
Amapá (ap)	8	1.6%	648,553	0.3%	81,069
Roraima (rr)	8	1.6%	425,398	0.2%	53,175
Total	513	100.0%	185,712,713	100.0%	313,514

Table A1: Number of Representatives and District Magnitude

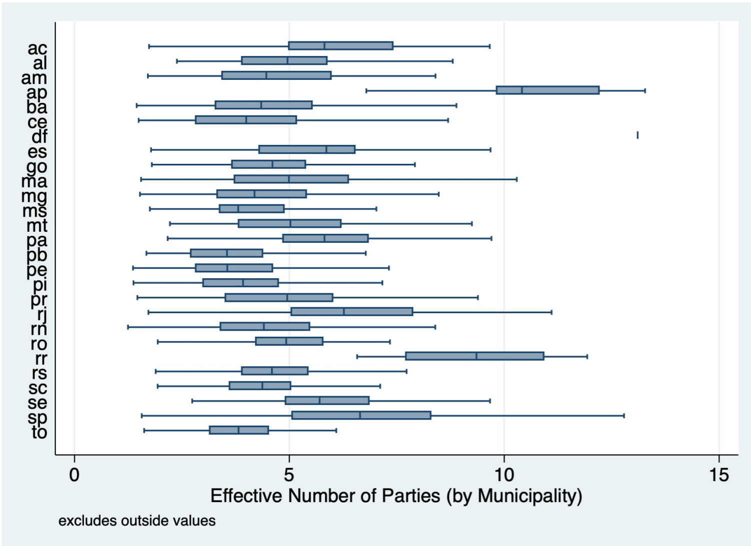

Figure A1: Effective Number of Parties (municipality-level vote shares) in 2014

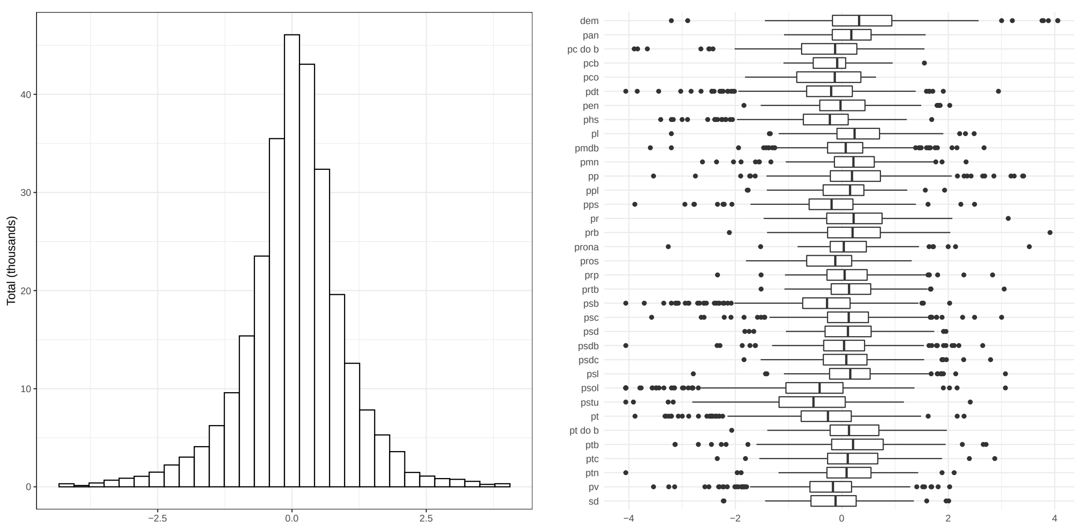

Figure A2: Candidates' Policy Positions (left: overall; right: by party)

Coalition	Parties	Votes	% of votes	Seats	% of seats
Pro-government Coalition "Com a Força do Povo"	Workers' Party (Partido dos Trabalhadores, PT)	13,554,166	13.93%	68	13.26%
	Brazilian Democratic Movement Party (Partido do Movimento Democrático Brasileiro, PMDB)	10,791,949	11.09%	66	12.87%
	Progressive Party (Partido Progressista, PP)	6,429,791	6.61%	38	7.41%
	Social Democratic Party (Partido Social Democrático, PSD)	5,967,953	6.13%	36	7.02%
	Republic Party (Partido da República, PR)	5,635,519	5.79%	34	6.63%
	Brazilian Republican Party (Partido Republicano Brasileiro, PRB)	4,424,824	4.55%	21	4.09%
	Democratic Labour Party (Partido Democrático Trabalhista, PDT)	3,472,175	3.57%	19	3.70%
	Republican Party of the Social Order (Partido Republicano da Ordem Social, PROS)	1,977,117	2.03%	11	2.14%
	Communist Party of Brazil (Partido Comunista do Brasil, PC do B)	1,913,015	1.97%	10	1,95%
	Total	**54,166,509**	**55,67%**	**303**	**59,07%**
Opposition Coalition "Muda Brasil"	Brazilian Social Democracy Party (Partido da Social Democracia Brasileira, PSDB)	11,073,631	11.38%	54	10.53%
	Democrats (Democratas, DEM)	4,085,487	4.20%	21	4.09%
	Brazilian Labour Party (Partido Trabalhista Brasileiro, PTB)	3,914,193	4.02%	25	4.88%
	Solidarity (Solidariedade, SD)	2,689,701	2.76%	15	2.92%
	Labour Party of Brazil (Partido Trabalhista do Brasil, PT do B)	828,876	0.85%	2	0,39%
	National Labor Party (Partido Trabalhista Nacional, PTN)	723,182	0.74%	4	0.78%
	National Ecologic Party (Partido Ecológico Nacional, PEN)	667,983	0.69%	2	0.39%
	Party of National Mobilization (Partido da Mobilização Nacional, PMN)	468,473	0.48%	3	0.58%
	Christian Labour Party (Partido Trabalhista Cristão, PTC)	338,117	0,35%	2	0.39%
	Total	**24,789,643**	**25,47%**	**128**	**24,95%**
Opposition Coalition "Unidos pelo Brasil"	Brazilian Socialist Party (Partido Socialista Brasileiro, PSB)	6,267,878	6.44%	34	6.63%
	Popular Socialist Party (Partido Popular Socialista, PPS)	1,955,689	2.01%	10	1.95%
	Humanist Party of Solidarity (Partido Humanista da Solidariedade, PHS)	943,068	0,97%	5	0.97%
	Social Liberal Party (Partido Social Liberal, PSL)	808,710	0.83%	1	0.20%
	Progressive Republican Party (Partido Republicano Progressista, PRP)	724,825	0.75%	3	0.58%
	Free Homeland Party (Partido Pátria Livre, PPL)	141,254	0.15%	0	0.00%
	Total	**10,841,424**	**11,15%**	**53**	**10,33%**
Out of coalition *(Fora de coligação)*	Social Christian Party (Partido Social Cristão, PSC)	2,520,421	2.59%	13	2.53%
	Green Party (Partido Verde, PV)	2,004,464	2.06%	8	1,56%
	Socialism and Liberty Party (Partido Socialismo e Liberdade, PSOL)	1,745,470	1,79%	5	0.97%
	Christian Social Democratic Party (Partido Social Democrata Cristão, PSDC)	509,936	0.52%	2	0.39%
	Brazilian Labour Renewal Party (Partido Renovador Trabalhista Brasileiro, PRTB)	454,190	0.47%	1	0.20%
	United Socialist Workers' Party (Partido Socialista dos Trabalhadores Unificado, PSTU)	188,473	0.19%	0	0.00%
	Brazilian Communist Party (Partido Comunista Brasileiro, PCB)	66,979	0.07%	0	0.00%
	Workers' Cause Party (Partido da Causa Operária, PCO)	12,969	0.01%	0	0,00%
	Total valid votes	**97,300,478**	**100,00%**	**513**	**100,00%**

Table A2: 2014 Brazilian Chamber of Deputies Election Results

	MNL	MNL (w/Dem's)	BLP
DEM	-2.210	-0.378	0.111
	(0.904)	*(0.624)*	*(0.225)*
PDT	-2.250	-0.070	-0.042
	(0.288)	*(0.265)*	*(0.160)*
MDB	-1.612	0.304	0.442
	(0.310)	*(0.253)*	*(0.130)*
PP	-2.531	-0.278	0.310
	(0.688)	*(0.472)*	*(0.204)*
PR	-1.118	0.556	0.770
	(0.617)	*(0.334)*	*(0.194)*
PRB	-3.368	-0.792	-0.323
	(0.548)	*(0.341)*	*(0.163)*
PSB	-2.129	-0.029	0.200
	(0.369)	*(0.246)*	*(0.101)*
PSC	-2.065	-0.254	0.034
	(0.473)	*(0.218)*	*(0.165)*
PSD	-0.908	0.141	0.391
	(0.766)	*(0.523)*	*(0.211)*
PSDB	-1.637	0.340	0.336
	(0.269)	*(0.242)*	*(0.163)*
PT	-1.544	0.518	0.614
	(0.353)	*(0.254)*	*(0.107)*
PTB	-2.312	-0.310	0.082
	(0.538)	*(0.280)*	*(0.154)*

Table A3: Estimates of Party-Brand Effects (we display only estimates for parties with at least three million votes)

	MNL	MNL (w/Dem's)	BLP
Median Wage (α_0^{wage})		-0.668	-0.576
		(1.385)	*(0.810)*
% Rural (α_0^{rural})		3.878	2.658
		(2.940)	*(1.295)*
% Higher Education (α_0^{edu})		1.255	0.833
		(1.007)	*(0.529)*
% Employed (α_0^{emp})		5.714	5.618
		(3.999)	*(1.625)*
Average Age (α_0^{age})		-1.638	-1.739
		(1.754)	*(0.651)*
% Female (α_0^{female})		1.023	0.659
		(1.708)	*(0.647)*

Table A4: Estimates of Baseline Voter Utility (inversely related to abstention/void-vote rates)

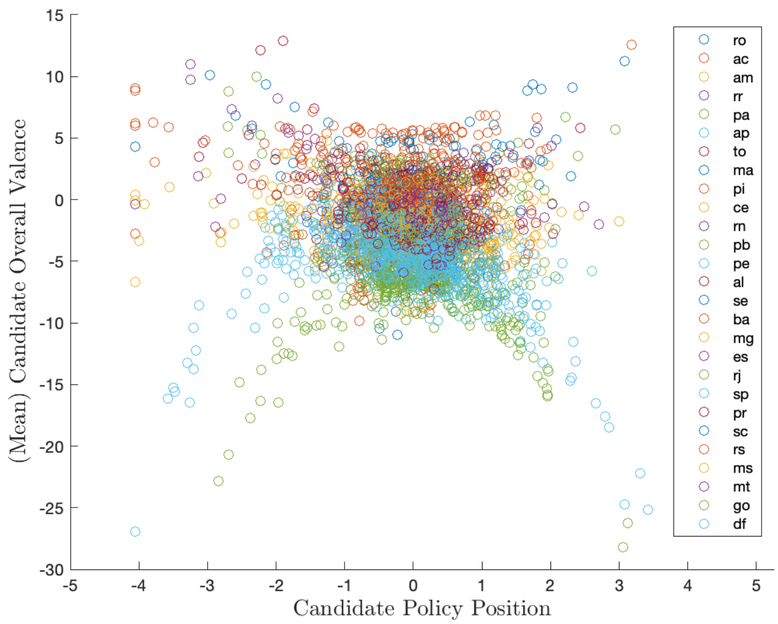

Figure A3: Candidates' Overall Valence and Policy Position

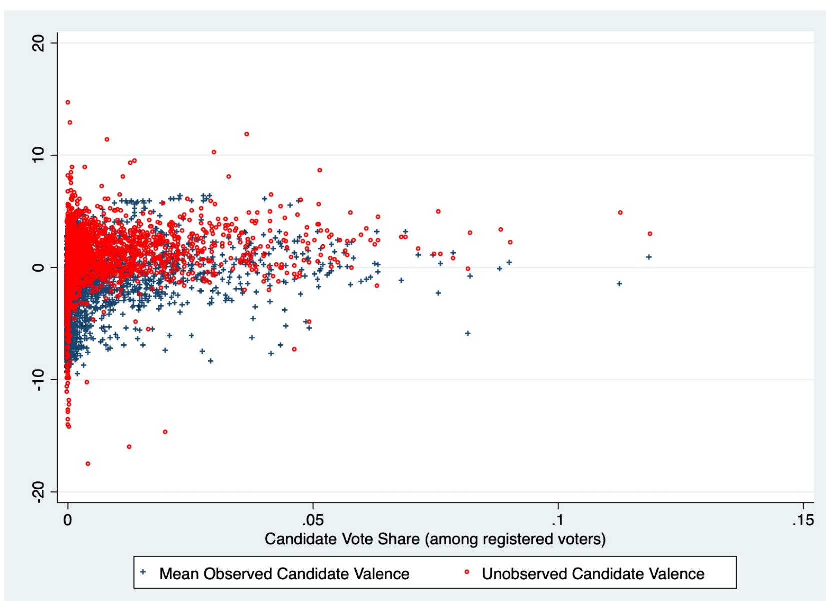

Figure A4: Candidates' (Observed and Unobserved) Valence and Vote Share

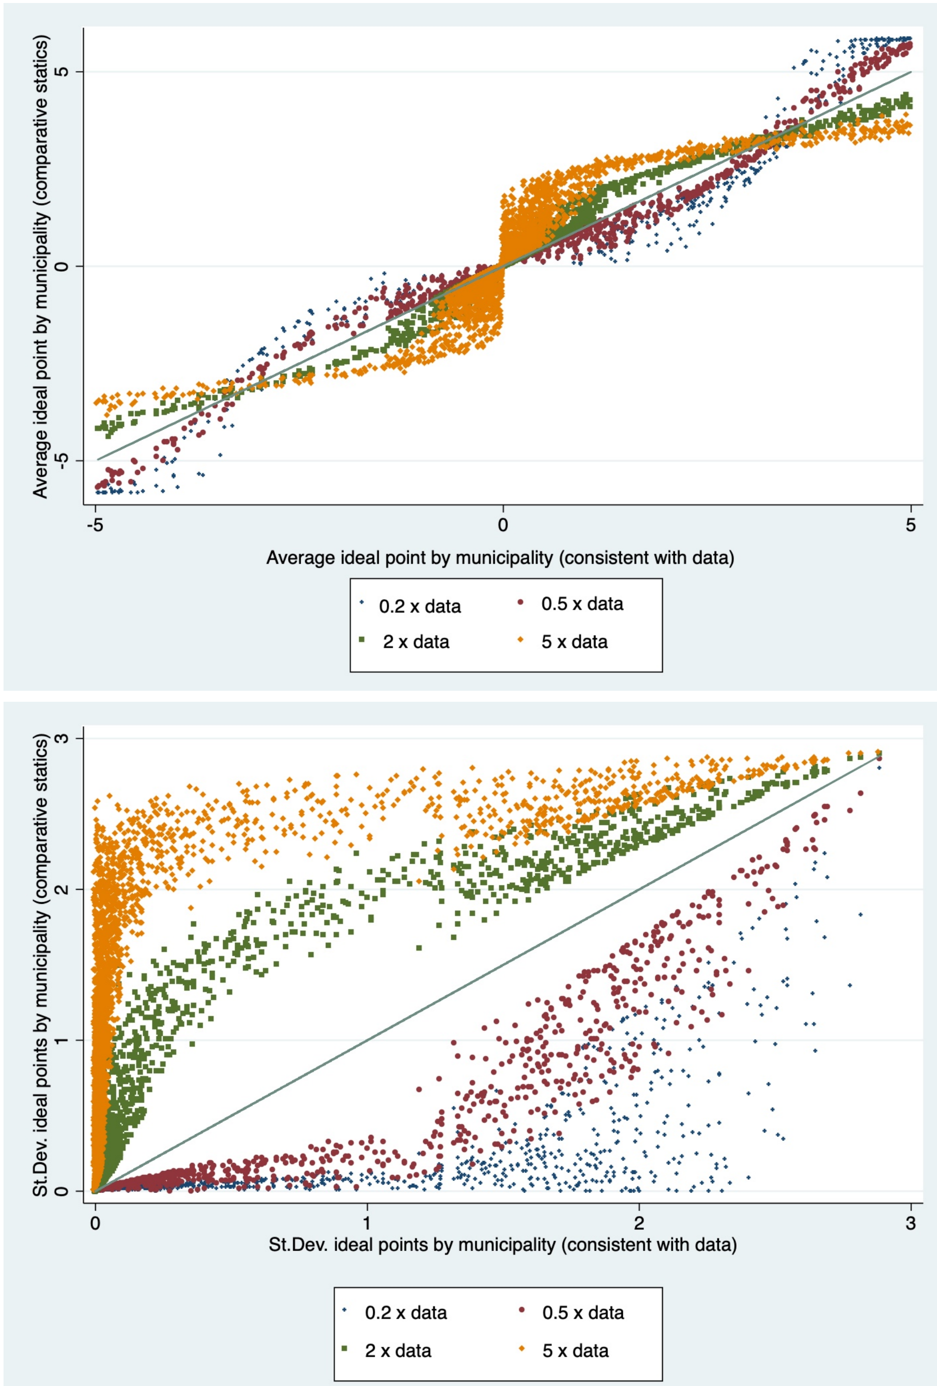

Figure A5: Sensitivity of Voter Ideal Point Estimates to Variance of Quadratic Term (σ_2). Top (bottom) panel plots the average (standard deviation) of ideal points by municipality. In each figure, the horizontal axis plots the results obtained under our estimates ("consistent with data"). The vertical axis plots the average (top) and standard deviation (bottom) of ideal points by municipality under four alternative values of σ_2, all else equal: 0.2, 0.5, 2 and 5 times the estimated value of σ_2.

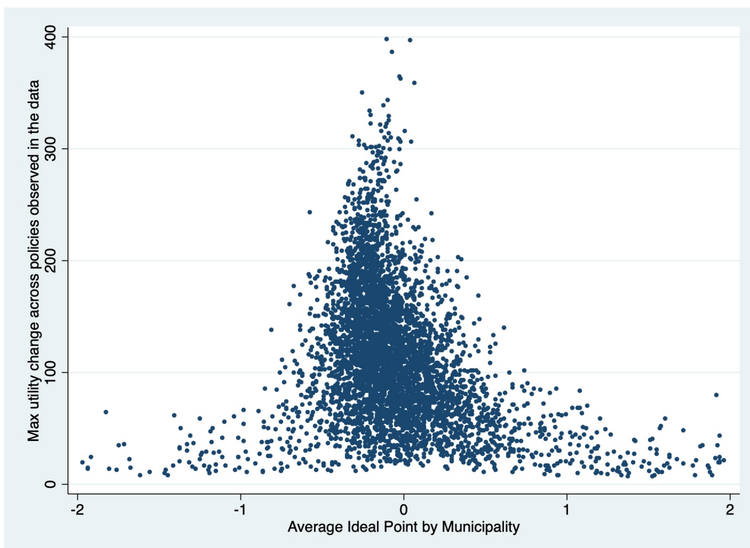

Figure A6: Average Ideal Point and Policy Sensitivity (maximum utility change across policies observed in the data), by Municipality

B Measuring Candidates' Policy Positions

As described in the text, we provide an estimate of candidates' ideological positions using correspondence analysis—the analog of principal component analysis (PCA) for categorical data—on all individual political contributions to federal, state, and local candidates between 2000 and 2014. This results in a contribution matrix R with 2.3 million rows (donors), indexed by $i = 1, \ldots, n$, and 561 thousand columns (political candidates), indexed by $j = 1, \ldots, m$. Each entry R_{ij} of the matrix R stores the total amount contributor i gives to candidate j.[28]

The first step to compute the ideology scores is to obtain the *relative* contribution matrix P, dividing each entry of R by $q = \sum_i \sum_j R_{ij}$. We then compute weights (marginals) for the rows and columns, w_r and w_c, where the i^{th} element of w_r is given by $w_r(i) = \sum_j P_{ij}$, and the j^{th} element of w_c is given by $w_c(j) = \sum_i P_{ij}$, and transform the weights into diagonal matrices $D_r = diag(1/\sqrt{w_r})$ and $D_c = diag(1/\sqrt{w_c})$. The final pre-processing step is to compute the matrix of standardized residuals $K = D_r(P - w_r w_c')D_c$, which gives weighted deviations from the "origin" under a null hypothesis of independence. Correspondence analysis then proceeds by obtaining a singular value decomposition of the matrix K, i.e., $K = U\Sigma V'$, where U and V are the left and right singular vectors of K (coordinate matrices), and Σ is a square diagonal matrix with the singular values of K on the diagonal (scaling matrix). Candidates' policy positions are obtained from the first dimension of the standard column coordinates, $p = D_c V$. As in the case of ideology scores obtained from roll-call data, an anchoring restriction is necessary. Bonica's scores for the U.S. are anchored by initializing the algorithm with all Democrats at -1 and all Republicans at 1. We initialize the algorithm with scores for candidates from each party at an ideological prior adapted from the Brazilian Legislative Surveys by Power and Rodrigues-Silveira (2019b) (see their replication package here). Scores are normalized to have mean zero and unit standard deviation.

The ideology scores obtained in this fashion have two desirable properties. First, analogous to focusing on the first principal component in PCA, these scores explain

[28] The pooled estimation of federal, state, and local candidates' ideological positions allows us to place all candidates on a common ideological scale and to leverage the greatest amount of information in the data. Since ideological proximity between candidates is identified from differences in contribution patterns by individual donors, we drop from our sample contributors who donate only to a single candidate. We also exclude corporate donors and contributions by political parties due to concerns that they may allocate their resources strategically rather than ideologically.

the largest share of variation in the data. Importantly, however, variance here is defined relative to a null hypothesis of independence between the rows and columns of R. To see this more clearly, notice that entry P_{ij} of the relative contribution matrix can be interpreted as the probability of observing donor i contributing to candidate j (or as the corresponding share of total contributions). Now, $w_r(i)$ can be viewed as the marginal probability of observing donations by i, and $w_c(j)$ corresponds to the marginal probability of donations to j. If donors assigned their contributions to candidates randomly—i.e., under a null hypothesis of independence—then the "residual" $P_{ij} - w_r(i)w_c(j)$ would be equal to zero. The first dimension of a correspondence analysis explains the largest share of variation in these residuals (appropriately normalized). Therefore, under the assumption that the primary motivation behind donors' contributions is ideology, these first-dimension scores should provide a good summary of the ideological content in the data.

Second, using this method, two candidates j and j' are assigned similar ideology scores if their donations profiles—i.e., columns P_j and $P_{j'}$ of matrix P—are similar. Candidates j and j' are assigned distant ideology scores if the set of donors who give a large fraction of their contributions to j or candidates close to j has little overlap with the set of donors who give a large fraction of their contributions to j' or candidates close to j'. Thus, assuming donors contribute primarily based on ideological considerations, these scores should reflect well the positions of candidates on the ideology spectrum.

To validate our policy position estimates, we conduct a battery of sanity checks. First, the left-hand panel of Figure B1 presents average policy positions by party, comparing federal versus local candidates. As shown, positions are generally consistent within party, as would be expected from common competitive and intra-party environments. In the right-hand panel of Figure B1, we then compare our estimates with ideology scores obtained from legislative surveys by Power and Zucco (2012). While the survey estimates are available only for elected candidates, there is general agreement between the two types of scores. Finally, under the Lula presidency, there was a marked shift to the left in voters' policy preferences, depicted in the left-hand panel of Figure B2 using Latinobarometer survey data. Our estimates feature a similar leftward shift in candidates' policy positions as shown in the right-hand panel of Figure B2.

Given the prevalence of corruption in Brazil (particularly in the wake of the largest

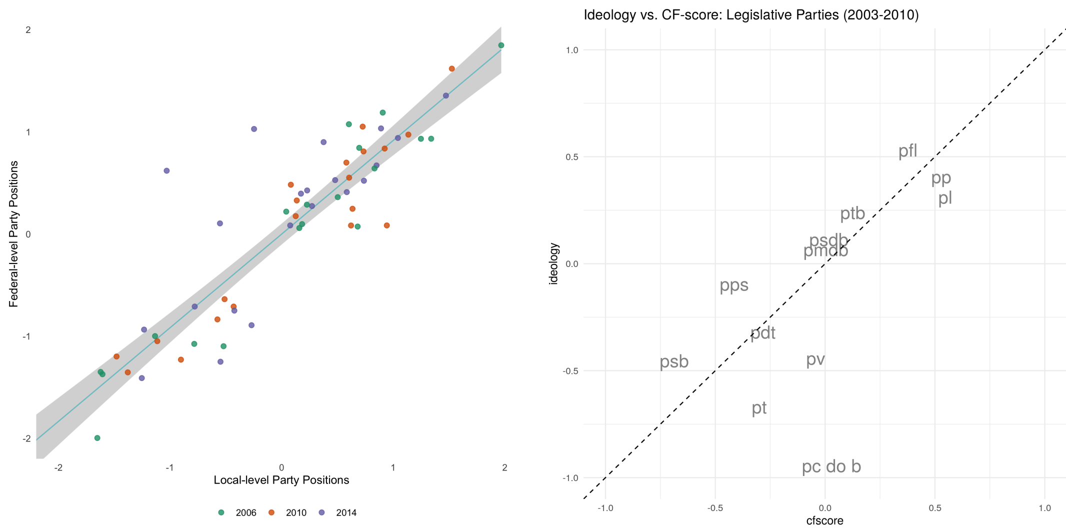

Figure B1: Average Policy Positions by Party—in federal versus local elections using campaign contributions (left) and among federal candidates using survey data (Power and Zucco, 2012) versus campaign contributions (right)

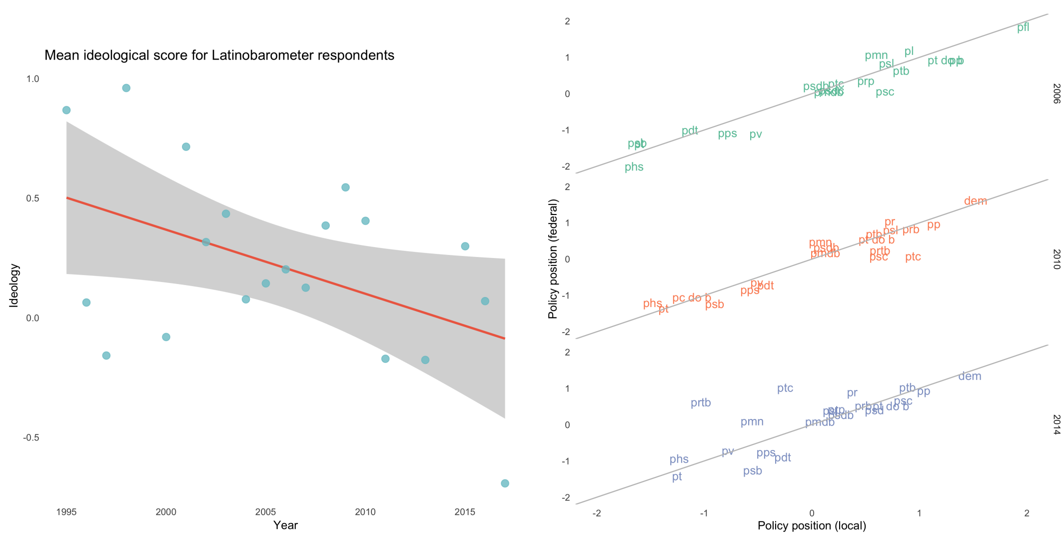

Figure B2: Leftward Shift in Policy Positions—among voters in Latinobarometer surveys (left) and among candidates in our estimates (right) scandal in Latin America with Operation Car Wash), there is a justifiable concern that, even among non-corporate and non-party individual donors, campaign contributions may be motivated by considerations other than ideology—e.g., public contract allocations or other forms of quid pro quo. To assess the sensitivity of our results to such violations of the ideological donations assumption, we conduct two tests.

First, we re-estimate candidates' policy positions excluding the top 5% and 10% of donors from the sample. Since contributions seeking to buy access to politicians or to exact favors are likely to be sizable, focusing on small contributions should alleviate such concerns. As shown in Figure B3, the resulting estimates are very similar to those obtained from the full sample. Correlations are 0.9 and 0.85, respectively, for estimates excluding the top 5% and 10% of donors.

To more directly address the possibility that campaign donors may be motivated by public contracts, we use data on public contract allocation by deputados federais provided by Boas, Hidalgo, and Richardson (2014). For each federal deputy in the 2006-2010 legislature, Figure B4 plots total (in logs) individual donations received for the 2006 (left) and 2010 (right) electoral cycles against the total value (in logs) of disbursed contracts. There is at most a very weak positive association, slightly more prominent for the 2006 cycle.

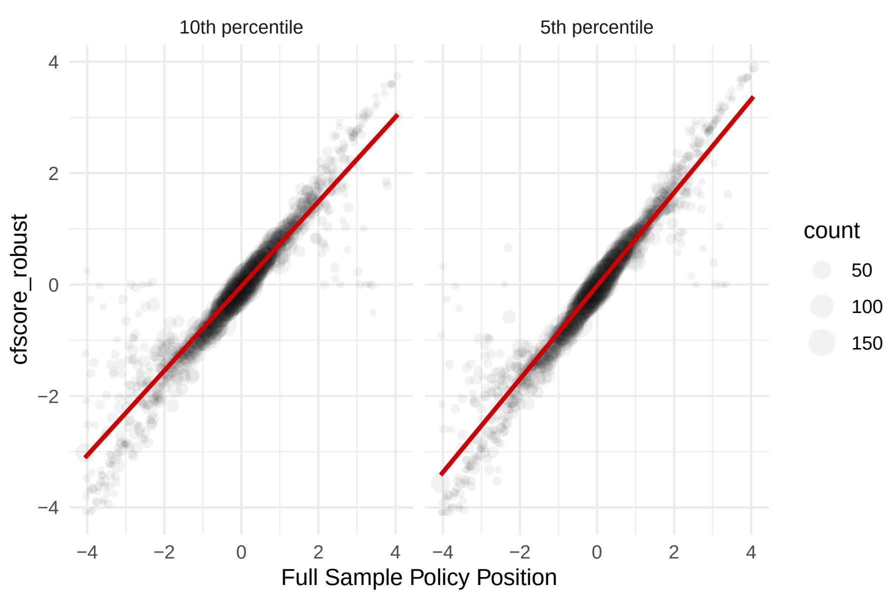

Figure B3: Estimates of Candidates' Policy Positions—excluding the top 10% (left) and 5% (right) of donors

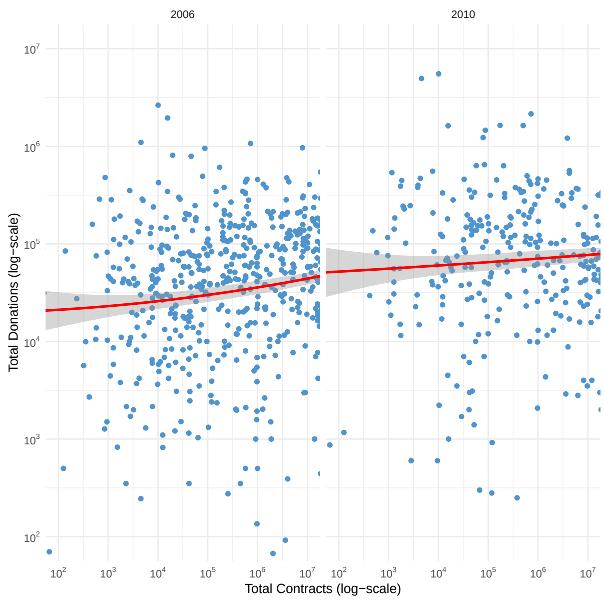

Figure B4: Individual Campaign Donations and Public Contract Disbursements (2006-2010) by Federal Deputy

C Estimation of Voters' Preferences

C.1 GMM Estimation and Inference

As discussed in Section 3.1, a GMM estimator of the demand-side parameters of our model can be obtained by minimizing the quadratic form

$$Q_J(\theta) = \xi(\theta)' Z W Z' \xi(\theta),$$

where $\xi_{jn}(\theta)$ is defined by (3.5). Under standard GMM regularity conditions (Hansen 1982, Berry, Levinsohn, and Pakes 1995), this estimator, $\hat{\theta}$, satisfies

$$\sqrt{J}(\hat{\theta} - \theta_0) \xrightarrow{d} N(0, (G'WG)^{-1} G'W\Omega W'G(G'W'G)^{-1})$$

as the sample size $J \to \infty$. Here,

$$G = E[Z_{jn} \nabla_\theta \xi_{jn}(\theta_0)] \quad \text{and} \quad \Omega = E[Z'_{jn} \xi_{jn}(\theta_0) \xi_{jn}(\theta_0)' Z'_{jn}]$$

are the gradient and variance, respectively, of the moment conditions (3.6). Notice that the optimal weighting matrix $W^* = \Omega^{-1}$ minimizes the asymptotic variance of the estimator, which then simplifies to $(G'\Omega^{-1}G)^{-1}$. This suggests a two-step estimation approach, which we follow.

In a first step, a consistent but inefficient estimate $\hat{\theta}_I$ of θ_0 can be obtained by minimizing $Q_J(\theta)$ using any positive-definite weighting matrix.[29] Then, allowing for potential correlation in unobserved valence across candidates in the same race, the optimal weighting matrix can be consistently estimated as $\hat{W}^* = \hat{\Omega}^{-1} = \left(\frac{1}{\tilde{N}} Z' V_\xi(\hat{\theta}_I)' Z\right)^{-1}$, where $(V_\xi(\hat{\theta}_I))_{jj'} = \xi_j(\hat{\theta}_I)\xi_{j'}(\hat{\theta}_I)$ if j and j' compete in the same race and $(V_\xi(\hat{\theta}_I))_{jj'} = 0$ otherwise, and $\tilde{N}$ denotes the total number of races. In a second step, reestimating the model using $\hat{W}^*$ delivers a consistent and efficient estimate $\hat{\theta}$ of θ_0. For robust inference, again allowing for potential correlation in unobserved valence across candidates in the same race, a consistent estimate of the asymptotic variance of $\hat{\theta}$ can be obtained simply as $(\hat{G}'\hat{\Omega}^{-1}\hat{G})^{-1}$, where $\hat{G} = Z'\nabla_\theta \xi(\hat{\theta})$ and $\hat{\Omega} = Z' V_\xi(\hat{\theta}) Z$.

[29] We employ an approximation of Ω^{-1} using the residuals of the homogeneous version of our model with $\sigma = 0$. Recall that estimation in this case boils down to a linear regression via two-stage least squares.

C.2 MPEC Approach

As noted in Section 3.1, the traditional BLP "nested fixed point" (NFXP) algorithm for computing $\hat{\theta}$ can be inefficient and sensitive to convergence criteria. We rely instead on the MPEC approach of Dubé, Fox, and Su (2012). The key idea is that, rather than "inverting" vote shares at each step of the optimization search, we can simply impose $s_{jn}(\delta_n, \sigma) = \hat{s}_{jn}$ as explicit constraints on the optimization program. Since state-of-the-art optimization algorithms only enforce constraints at convergence, this can considerably reduce the computational burden. Further computational gains can be obtained by exploiting sparsity. We estimate $\hat{\theta}$ by solving the following mathematical program with equilibrium constraints:

$$\min_{\theta,\xi,\psi} \psi' \tilde{W} \psi \quad \text{subject to}$$

$$\psi = Z'\xi \quad \text{and} \tag{C1}$$

$$\tilde{s}_{jn}(\delta_n, \sigma) = \hat{s}_{jn} \quad \text{for all } j, n, \text{ where} \tag{C2}$$

$$\delta_{jn} = \sum_{k=0}^{2} (\alpha_k^0 + D_n' \alpha_k^D)(p_{jn})^k + W_{jn}'\phi + X_{jn}'\beta + \xi_{jn}. \tag{C3}$$

Dubé, Fox, and Su (2012) show that this MPEC and the traditional BLP NFXP algorithm yield theoretically identical estimates of θ_0, but the MPEC approach delivers superior numerical performance. While the computational cost of estimation may seem to increase by treating ξ and the moment conditions ψ as auxiliary variables—and thus expanding the size of the optimization problem—note that (C1) and (C3) are linear constraints and (θ, ξ) no longer enter the objective function directly. This, together with the sparsity that results from ξ_{jn} having no effect on vote shares outside of j's district and electoral cycle, adds to the computational advantage over NFXP from avoiding repeated fixed point calculations.

C.3 Robustness

We evaluate the sensitivity of our main results to several key assumptions and features of our data. First, as noted in the paper, we are forced to exclude several candidates from our sample due to insufficient individual contributions with which to estimate their policy positions. Table C1 summarizes differences in observed non-ideological characteristics of candidates included and excluded from our sample, by decile of the distribution of vote shares. While some of the differences are statistically significant, it is notable that there are no systematic patterns with respect to electoral performance that would raise concerns about potential biases in our estimates. For example, female candidates are generally underrepresented in the included sample. However, they are overrepresented among the lowest performing *and* highest performing candidates, which should alleviate concerns about any systematic bias in our estimate of ϕ.

	1	2	3	4	5	6	7	8	9	10
Female	0.130	0.091	0.107	0.008	-0.021	-0.005	-0.028	-0.013	-0.051	0.017
Higher Edu.	0.070	0.026	0.061	0.020	0.085	0.089	0.056	0.054	0.079	0.024
Age	-1.402	-0.110	-0.399	0.063	0.152	-0.536	-0.153	-0.133	1.801	0.940
Gov. Exp.	0.015	0.033	-0.003	0.000	0.010	-0.018	-0.026	0.000	-0.074	-0.074
Bus. Exp.	-0.060	-0.021	-0.007	-0.013	0.002	-0.019	-0.002	0.007	-0.048	-0.148

Table C1: Differences in Means, In-Sample Vs. Out-of-Sample Candidates, by Decile of the Distribution of Vote Shares

To evaluate the robustness of our welfare analysis to the excluded sample, we conduct the following exercise. The key concern is that excluded candidates may have policy positions close to voters but not receive sufficient donations due to a perceived lack of viability. With this in mind, as a worst-case scenario, we impute policy positions for excluded candidates at the median voter's ideal point in their district, re-calculate unobserved valence so that predicted and observed vote shares match for all candidates, and then reproduce our welfare calculations. As shown in Figure C1, our results are virtually unchanged.

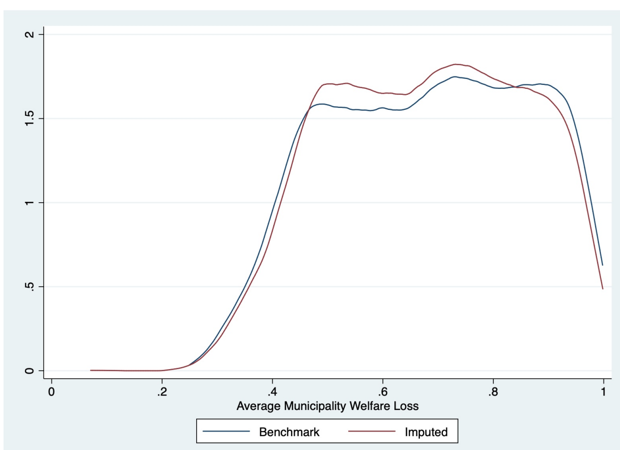

Figure C1: Welfare Analysis with Imputed Policies for Candidates Excluded from Sample

Second, our model doesn't allow for interactions between policy and non-ideological considerations. While this assumption is broadly consistent with the existing literature, it may be overly restrictive. For example, if voters care not about the policies of individual candidates per se but are trying to forecast ultimate policy outcomes in the legislature, they may use valence characteristics as a heuristic to determine how influential each candidate may be in the legislative bargaining process. Alternatively, if voters are uninformed about what the right policy for them is, they may evaluate the policies put forward by incumbents and challengers differently. To test for such possibilities, we re-estimate our model allowing for an interaction between a candidate's incumbency status and voters' evaluation of their policy position. Table C2 below shows that the estimated interaction coefficients are close to zero and statistically insignificant, and our remaining estimates of voters' preferences are nearly identical.

Third, our model assumes that voters case their ballots expressively in favor of the candidate each voter prefers the most, disregarding electability considerations. We similarly assume candidate entry is not informed by expected electoral performance. However, if voters do not care about the attributes of individual candidates, but instead are trying to forecast the eventual composition of the legislative chamber, they may evaluate the observed ideological and non-ideological characteristics of candidates perceived to be very competitive differently from those of less electable candidates. And strategic entry might induce correlation between candidates' observed and unobserved valence. To examine these possibilities, we re-estimate our model after dropping from our sample all candidates that obtain a vote share greater than 3% of registered voters in their district. These outstanding candidates constitute 5% of the total sample. We also repeat this exercise dropping candidates with a vote share greater than 2%, which constitute 10% of the total sample. Reassuringly, as shown in Table C2, our estimates of voters' preferences are substantively unchanged.

Finally, given that candidate lists in our sample are considerably large and that many candidates obtain negligible vote shares, one might question the inclusion of small candidates in the estimation. To explore the robustness of our results to such concerns, we re-estimate our model after dropping from our sample candidates with vote shares in the bottom 5% (fifth column) or 10% (sixth column) of our sample. Once again, Table C2 shows that our estimates of voters' preferences remain virtually unchanged.

	Baseline	Policy-Incumbent Interaction	Excluding Top 5%	Excluding Top 10%	Excluding Bottom 5%	Excluding Bottom 10%
Age (β^{age})	0.014	0.014	0.037	0.005	-0.019	0.001
	(0.021)	*(0.022)*	*(0.023)*	*(0.027)*	*(0.025)*	*(0.028)*
Age Sq. (β^{age^2})	-0.002	-0.001	-0.004	0.004	0.011	0.009
	(0.013)	*(0.013)*	*(0.014)*	*(0.016)*	*(0.019)*	*(0.021)*
Higher Education (β^{edu})	0.688	0.685	0.735	0.740	0.634	0.562
	(0.044)	*(0.049)*	*(0.044)*	*(0.062)*	*(0.052)*	*(0.061)*
Business Exp. ($\beta^{business}$)	0.162	0.167	0.145	0.132	0.112	0.187
	(0.059)	*(0.061)*	*(0.065)*	*(0.060)*	*(0.085)*	*(0.104)*
Government Exp. (β^{gov})	-0.494	-0.495	-0.423	-0.377	-0.491	-0.436
	(0.077)	*(0.076)*	*(0.085)*	*(0.087)*	*(0.086)*	*(0.099)*
Female Candidate (ϕ_1)	-1.065	-1.092	-1.128	-1.217	-0.675	-0.592
	(0.896)	*(0.904)*	*(1.016)*	*(1.107)*	*(1.099)*	*(1.096)*
Female Cand. Preference Variance (σ_3)	0.003	0.288	0.001	0.001	0.001	0.202
	(305.9)	*(3.796)*	*(1868.6)*	*(1600.6)*	*(913.4)*	*(6.405)*
Incumbent (ϕ_2)	1.786	1.729	1.868	1.917	1.773	1.749
	(0.248)	*(0.593)*	*(0.452)*	*(0.780)*	*(0.276)*	*(0.300)*
Incumbent Preference Variance (σ_4)	0.000	0.000	0.000	0.000	0.000	0.000
	(1160.5)	*(1907.4)*	*(3385.1)*	*(4327.4)*	*(1652.7)*	*(2321.1)*
Policy (α_1^0)	-0.950	-0.886	-1.524	-1.905	-0.654	-0.383
	(1.009)	*(1.021)*	*(0.770)*	*(1.024)*	*(1.098)*	*(1.093)*
Policy × Median Wage (α_1^{wage})	-2.389	-2.403	-1.107	-1.884	-2.089	-2.350
	(0.783)	*(0.812)*	*(0.711)*	*(0.809)*	*(0.842)*	*(0.838)*
Policy × % Rural (α_1^{rural})	-3.390	-3.385	-1.396	-0.922	-3.462	-2.885
	(1.052)	*(1.111)*	*(1.080)*	*(1.333)*	*(1.071)*	*(1.130)*
Policy × % Higher Education (α_1^{edu})	1.120	1.090	1.112	1.897	0.883	1.136
	(0.537)	*(0.537)*	*(0.536)*	*(0.633)*	*(0.545)*	*(0.610)*
Policy × % Employed (α_1^{emp})	1.459	1.341	0.499	-0.211	2.224	1.609
	(1.206)	*(1.325)*	*(1.237)*	*(1.289)*	*(1.281)*	*(1.143)*
Policy × Average Age (α_1^{age})	-0.763	-0.653	-0.766	-0.436	-1.099	-0.675
	(0.702)	*(0.833)*	*(0.666)*	*(0.715)*	*(0.821)*	*(0.798)*
Policy × % Female (α_1^{female})	-0.598	-0.708	-0.213	-0.773	-0.094	-0.630
	(0.637)	*(0.863)*	*(0.728)*	*(0.797)*	*(0.820)*	*(0.795)*
Policy × Incumbent ($\alpha_1^{incumbent}$)		0.220				
		(1.312)				
Policy Sq. (α_2^0)	-4.879	-4.855	-5.779	-6.722	-4.966	-4.924
	(0.663)	*(0.733)*	*(0.601)*	*(0.613)*	*(0.743)*	*(0.804)*
Policy Sq. × Median Wage (α_2^{wage})	1.220	1.216	0.889	0.124	1.251	1.329
	(0.920)	*(0.914)*	*(0.787)*	*(1.052)*	*(1.009)*	*(0.960)*
Policy Sq. × % Rural (α_2^{rural})	-3.008	-2.874	-1.996	-0.452	-2.938	-2.143
	(1.751)	*(1.840)*	*(1.442)*	*(1.721)*	*(1.998)*	*(1.977)*
Policy Sq. × % Higher Education (α_2^{edu})	-1.185	-1.156	-0.353	0.894	-1.185	-1.010
	(0.498)	*(0.521)*	*(0.755)*	*(0.883)*	*(0.556)*	*(0.594)*
Policy Sq. × % Employed (α_2^{emp})	-1.945	-2.077	-2.814	-3.968	-2.337	-3.356
	(2.272)	*(2.331)*	*(1.882)*	*(1.910)*	*(2.505)*	*(2.387)*
Policy Sq. × Average Age (α_2^{age})	-0.507	-0.437	-0.405	-0.123	-0.448	-0.021
	(0.952)	*(1.000)*	*(0.765)*	*(0.833)*	*(1.030)*	*(1.012)*
Policy Sq. × % Female (α_2^{female})	1.143	1.094	0.736	0.352	0.944	0.502
	(0.949)	*(0.973)*	*(0.901)*	*(0.996)*	*(1.032)*	*(1.040)*
Policy Sq. × Incumbent ($\alpha_2^{incumbent}$)		0.086				
		(0.510)				
Policy Preference Variance (σ_1)	0.001	0.000	0.000	0.000	0.001	0.001
	(868.7)	*(1051.1)*	*(1452.0)*	*(1507.7)*	*(940.8)*	*(852.5)*
Policy Sq. Preference Variance (σ_2)	0.384	0.378	0.561	0.672	0.409	0.422
	(0.211)	*(0.236)*	*(0.254)*	*(0.288)*	*(0.265)*	*(0.299)*

Table C2: Robustness Checks (party and baseline-utility effects omitted)

D Estimation of Politicians' Preferences

D.1 GMM Estimation and Inference

Given an estimate $\hat{\theta}$ of the demand-side parameters of our model, a GMM estimator of the supply-side parameters can be obtained by minimizing the quadratic form

$$\tilde{Q}_J(\gamma, \chi) = \zeta(\gamma, \chi)' \tilde{Z} \tilde{W} \tilde{Z}' \zeta(\gamma, \chi),$$

where $\zeta_{jn}(\gamma, \chi)$ is defined by (4.5). As in the demand case, we follow a two-step approach to obtain not only an estimate of the optimal weighting matrix but also to aid in the selection of appropriate instruments to identify γ, the parameters characterizing party influence over candidates' policy choices.

In a first step, we approximate $\tilde{W}^* = \tilde{\Omega}^{-1}$ by estimating a version of the model with $\gamma^\ell = 1$ and $\gamma^{\mathrm{inc}} = 0$. Note that, keeping γ fixed, estimation of χ boils down to a simple linear regression. We then use party dummies and candidates' observed incumbency status as instruments to identify γ in the first round of GMM estimation.

In the second step, we implement an approximation of Chamberlain (1987)'s optimal instruments, $Z^*_{jn} = E[\nabla_{(\gamma,\chi)} \zeta_{jn}(\gamma_0, \chi_0) | Z_{jn}]$. These correspond to the exogenous characteristics for the "linear" parameters, χ, and we use $\nabla_\gamma \zeta_{jn}(\hat{\gamma}_I, \hat{\chi}_I)$ for the "non-linear" parameters, γ, where $(\hat{\gamma}_I, \hat{\chi}_I)$ denote the first-step, inefficient estimates. Similarly to the demand case, for robust inference, we allow for arbitrary heteroskedasticity and cluster standard errors at the party-state-year level. To account for demand-side estimation uncertainty, we rely on standard results for two-step GMM estimation (Newey and McFadden 1994). We implement an MPEC version of this estimator for computational convenience, i.e., we solve

$$\begin{aligned}
&\min_{\theta,\xi,\psi} \psi' \tilde{W} \psi \quad \text{subject to} \\
&\psi = \tilde{Z}' \zeta \quad \text{and} \\
&r_{jn}(\gamma) - \tilde{X}'_{jn} \chi - \zeta_{jn} = 0 \quad \text{for all } j, n.
\end{aligned}$$

D.2 Estimation of Distribution of Politicians' Ideal Policies

Candidate j's equilibrium policy choice satisfies the following first-order condition:

$$R_{jn}(\gamma) \equiv \frac{\partial s_{jn}(\mathbf{p}_n)}{\partial p_{jn}} + (\gamma^\ell + \gamma^{\text{inc}} \tilde{I}_{jn}) \sum_{j' \in J_n^\ell} \frac{\partial s_{j'n}(\mathbf{p}_n)}{\partial p_{jn}} = \mu_{jn}(-\mathbf{1}_{p_{jn} < \rho_{jn}}). \tag{D1}$$

Having estimated $R_{jn}(\hat{\gamma})$ and $\hat{\mu}_{jn}$ as described above, note that (D1) then enables estimation of the distribution of candidates' ideal policies, ρ_{jn}, via maximum likelihood, analogous to a standard probit model. Specifically, since $\hat{\mu}_{jn} > 0$, the likelihood of observing $R_{jn}(\hat{\gamma}) < 0$ is given by $\Phi\big(\frac{p_{jn} - \rho_n^\ell}{\sigma_n^\ell}\big)$, where Φ denotes the standard normal cumulative distribution function. We specify ρ_n^ℓ as a linear index of average state demographics interacted with party dummies. Similarly, we specify ρ_n^ℓ as a linear index of within-state demographic variability along with party dummies (no interactions). These estimates are only used to simulate candidates' ideal policies for the counterfactuals as described below.

D.3 Robustness

Since coalitions are extremely common in Brazilian elections, it is possible that discipline effects over candidates' policy choices may operate at the coalition (or list) level rather than at the party level. Accordingly, we re-estimate our model of the "supply side" letting ℓ index coalitions instead of parties and letting γ^ℓ in (4.3) correspond to an average fixed effect over all parties participating in list ℓ. As shown in Table D1, the resulting parameter estimates are virtually identical, which suggests the relevant tradeoffs occur within parties.

D.4 Counterfactuals

Finally, we briefly describe implementation of our counterfactual experiments. We limit attention to the state of Bahia for computational convenience. Calculating parties' best responses according to 4.1 for all candidates in our sample, particularly in large districts like São Paulo, would be computationally prohibitive. Bahia, however, is the state most representative of the country overall in terms of demographics. As such, it provides a good testing ground for our counterfactuals.

As discussed, our first goal is to explore the effects of a minimum requirement of

μ: Weight of Ideology Relative to Electability			
Constant	-6.347	Higher Education	0.858
	(41.69)		*(0.347)*
Age	0.109	Business Exp.	-0.052
	(0.894)		*(0.087)*
Age Squared	-0.050	Gov. Experience	-0.801
	(1.714)		*(0.056)*
Female Candidate	-1.055	Unobserved Valence	0.515
	(0.045)		*(0.285)*
γ: Weight of Coalition Vote Share Relative to Own Vote Share			
DEM	0.818	PDT	0.700
	(174.2)		*(73.45)*
MDB	0.018	PP	0.435
	(56.91)		*(94.98)*
PR	0.342	PRB	0.000
	(80.35)		*(69.13)*
PSB	0.045	PSD	0.841
	(42.31)		*(109.3)*
PSDB	0.132	PT	0.054
	(109.7)		*(52.90)*
PTB	0.053	Incumbent	-1.069
	(182.1)		*(136.0)*

Table D1: “Supply-Side” Coefficient Estimates with Coalition-Level Discipline (we display only estimates for parties with at least three million votes) higher education for all candidates. To calculate welfare changes with fixed policy positions, we first draw, from the empirical distribution of candidates with higher education in Bahia, a new set of candidates of the same size as observed in the data. Drawing from the empirical distribution ensures that we account for existing correlations in the data between education and other non-ideological attributes of candidates as well as their policy choices. Average welfare in each municipality in the state is then calculated as described in our welfare analysis.

To account for equilibrium policy adjustments, we then iterate candidates' best responses according to 4.1, starting from the policies drawn above. This requires an estimate of each candidate's ideal policy, ρ_{jn}. Given (D1), since we observe $R_{jn}(\hat{\gamma})$, we draw ρ_{jn} from the distribution of candidates' ideal policies, estimated as described above, conditional on it being to the right or left of candidate j's observed policy in accordance with the sign of $R_{jn}(\hat{\gamma})$. After only a few iterations of candidates best responding to each other's policies, this procedure converges to a Nash equilibrium. We then recalculate voter welfare.

For our policy discipline counterfactual, we keep the pool of candidates as observed in the data, draw candidates' ideal policies as just described, and then iterate best responses starting from candidates' observed policies and setting $\gamma_{jn} = 1$ for all candidates. We then recalculate voter welfare given the new equilibrium policy choices.