

Democracy on a Budget

Michel Wachsmann *

September 2026

Abstract

Political parties are the gateway to elected office, and their members are the pool from which candidates are drawn. This paper studies how money shapes the competition between parties for those members. I develop a model in which parties post common offers of organizational support under budget constraints. A party's budget determines how many citizens it can reach and whether it has priority over recruits that its rivals also want. I bring the model to Brazil's party-membership registry, which records each citizen's sequence of party memberships, and combine it with party accounts around the 2015 ban on corporate donations. After the ban, recruitment at previously corporate-funded parties falls by the equivalent of 15.1 percent of its 2012 mean, relative to parties without corporate receipts in the same municipality. Recruits whose previous membership was in another party account for 42.8 percent of this decline, although they are 33.1 percent of baseline recruitment. Along directed local links between parties, recruitment shifts toward the party whose finances improve relative to its rival, and the shift operates mainly through the stronger party's intake. Campaign finance therefore affects party competition before the campaign begins, through each party's ability to recruit politically engaged citizens from its competitors.

KEYWORDS: party membership, political selection, campaign finance, party competition, Brazil.

JEL: D71, D72, P16.

1 Introduction

Institutions depend on the people who fill them. In a representative democracy, the people who fill elected office almost always arrive through a political party. Parties organize participation, assemble electoral coalitions and select candidates (Katz and Mair, 1995), and they do so with the help of their members: the citizens who staff campaigns, maintain local networks, carry demands into the organization and eventually stand for office themselves (Speck, Braga and Costa, 2015). Who joins a party, and which party they join, therefore shapes the pool from which representatives are drawn.

In Brazil the connection is direct. Independent candidacies are not permitted, and a citizen must belong to a party for a statutory period before standing for office. About 2.98 million people joined a party during the 2012 municipal cycle. Of these, 9.3 percent stood for mayor or councilor in that year's election, and together they made up 64.1 percent of all candidates for those offices. Recruitment is concentrated in the months before candidacy deadlines and before municipal elections, when several hundred thousand candidacies must be filled (Speck, 2013). The same citizens are often valuable to more than one party: someone with a local following, campaign experience or ties to a neighborhood can be recruited by several organizations, and many recruits have belonged to another party before.

*São Paulo School of Economics (FGV EESP)

This paper studies how money shapes that competition. Before parties compete for votes, they compete for members. What a party can offer a prospective member, from campaign resources to organizational support for a candidacy, depends on its budget. A change in one party's finances can therefore alter both how many citizens it attracts and whose former members it attracts. I ask how financial resources affect the competition between parties to recruit members, and whether money determines the direction of recruitment between rival parties.

Answering this question requires two things that are rarely available together: a source of variation in party finances that does not originate in the parties' own recruitment, and data on where new members come from. I use the 2015 prohibition of corporate donations in Brazil. The Supreme Federal Court banned contributions from firms in September 2015, a year before the 2016 municipal elections. Firms had supplied about three-quarters of campaign contributions (Speck, 2016), and local parties entered the reform with very different levels of dependence on them. I measure a party's exposure by its 2012 corporate receipts per registered voter in the municipality and compare its subsequent recruitment with that of parties in the same municipality that received no corporate money. I combine these accounts with Brazil's administrative membership registry, which records every membership spell with its party, municipality and start date. Because the registry follows individuals, it records whether a new member previously belonged to another party and which one. This turns counts of new members into a directed network of recruitment between rival organizations.

Figure 1 previews the margins that the analysis separates. It follows 3,054 local parties in 719 municipalities with municipal-directory accounts in both 2012 and 2016 and a fixed set of local competitors, and relates each party's loss of funding share to changes in three recruitment shares. A ten-percentage-point loss in a party's share of local funding is associated with a 0.680-percentage-point decline in its share of all recruitment and a 0.591-point decline in its share of recruitment from rivals. Its share of recorded movements toward rivals rises by 0.306 points, an estimate that is imprecise. Parties that lose financial ground attract fewer recruits, including fewer recruits from their competitors. The evidence that their own members leave for rivals at a higher rate is weaker. These are associations between endogenous shares; the rest of the paper asks whether the same pattern appears under the reform.

To interpret these margins, I develop a model of parties competing for recruits under budget constraints. Following the posted-wage formulation of Bulow and Levin (2006), each party commits to a common offer of support that applies to whoever it recruits. Competition over the most valuable citizens then takes the form of an all-pay contest with a budget cap, as in Che and Gale (1998). I add private costs of participation, so that the size of an offer also determines how many citizens are willing to join at all. A budget therefore buys two things: reach, the share of potential recruits a party can bring into politics, and priority, the right to choose first among citizens that a rival also wants. A cut that removes a party's financial advantage reduces its recruitment (Prediction 1) and, when it equalizes the two parties' offers, gives the rival first access to contested recruits (Prediction 3). The model is also explicit about when the rival does not benefit: a partial cut that leaves the exposed party ahead reduces recruitment at both parties.

The empirical results follow the order of these margins. First, recruitment falls at exposed parties. Between the 2012 and 2016 cycles, recruitment at parties that had received corporate money declines, relative to unexposed parties in the same municipality, by an amount equal to 15.1 percent of the 2012 mean. The estimate is -0.517 recruits per 1,000 voters (standard error 0.148) against a baseline of 3.426, from 14,606 local parties in 3,048 municipalities, with party and municipality fixed effects and 2008 controls. The decline increases with the amount of prior corporate funding and appears in both post-reform cycles of a five-cycle panel.

Second, the decline is concentrated among recruits who come from other parties. Recruits whose previous recorded membership was in a different party account for 33.1 percent of




Figure 1: Financial-share losses and three recruitment margins, 2012–2016. Positive horizontal values indicate a loss of the party’s share of receipts among the same financially observed local parties. Vertical values are changes in its share of all recruitment, recruitment from rivals, and recorded movements to rivals. Municipal-directory accounts are used in both years. The same 3,054 party–municipality observations in 719 municipalities and ten financial-loss bins appear in every panel. Points are raw bin means; bars are pointwise 95% municipality-clustered confidence intervals. Lines use all observations. Incoming and outgoing shares use the same local cross-party flow pool; total recruitment has its own denominator. Recorded prior-party histories do not establish active departures or solicitation. These are suggestive associations.

baseline recruitment but 42.8 percent of the estimated decline. The contraction extends to recruits from parties in the same ideological bloc and from parties in other blocs. Losing corporate money reduces a party’s ability to compete for citizens who already have political experience, not only its ability to bring new citizens into politics.

Third, money determines the direction of recruitment between rivals. I construct the network of directed recruitment flows between parties with financial records at both ends: 89,270 potential origin–destination links in 3,048 municipalities, including links with no recorded flow. Along these links, recruitment shifts toward the party whose gross receipts grow more, or whose prior corporate exposure is lower, holding the pair’s average finances fixed. The shift operates through the destination. Growth in the destination’s receipts predicts more recruitment from the origin, whereas a decline in the origin’s receipts does not predict more departures from it. Conditional allocation models point the same way. Money determines a party’s ability to attract members from its rivals. A party that loses financial ground loses much of that ability, while the evidence that it also loses more of its own members is weaker and depends on the sample.

Fourth, the reallocation is concentrated in particular parties. The Left’s recorded arrivals from other parties fall from 28,184 to 19,004, and the Workers’ Party (Partido dos Trabalhadores, PT) accounts for 88.6 percent of the decline, while the Center, the Right and unscored parties increase their shares. The rivalry between PT and the Brazilian Social Democracy Party (Partido da Social Democracia Brasileira, PSDB) shows what such a reversal looks like. In the 591 municipalities where both parties have financial records, PSDB moves from a net loss of 842 recruits to PT to a net gain of 462, and 93.6 percent of the swing comes from fewer PSDB members joining PT. Its financial gradients are imprecise, however, and the reversal coincides with the national collapse of PT after 2014. I use it to show how a change in attraction reshapes a rivalry, not as evidence of a financial cause.

Three features of the setting limit what these comparisons identify. Exposure is measured in the 2012 election, which is also the baseline for the outcome, and an earlier-window placebo shows that 2008 corporate funding predicts declining recruitment between 2008 and 2012. Part of the gradient may therefore reflect reversion from 2012 levels rather than the ban. The reform coincided with a shorter membership requirement for candidates, new spending limits for 2016 campaigns (Avis et al., 2022), and a national political crisis that affected parties unevenly. Finally, the registry records sequences of memberships, not offers or solicitations, and a previous membership need not have been active when the new one began. I report the pre-period estimates next to the main results in Section 3 and the placebo and exclusion checks in Appendix D, so that the reader can judge how much of the pattern survives each concern.

This paper contributes to two literatures. The first studies the effects of money in politics. Limits on spending change who runs and who wins: in Brazilian mayoral elections, stricter limits attract more candidates, who are less wealthy, and reduce reelection rates (Avis et al., 2022), and in British parliamentary elections higher limits raise the cost of campaigns, shrink the pool of candidates and amplify the advantages of incumbents (Fouirnaies, 2021). Spending raises vote shares in French and British elections, with lower returns for radical and extreme parties (Bekkouche, Cagé and Dewitte, 2022). Bans on corporate donations change both what politicians say and who holds office. After the French ban of 1995, candidates gave less attention to their local presence and to economic issues, and candidates from non-mainstream parties used more polarized language (Cagé, Le Pennec and Mougin, 2024). After the Brazilian ban, more challengers and new candidates ran, and incumbents who had relied on firms became less likely to be reelected (Peveri, 2024). These studies measure outcomes at or after the moment of candidacy. I study an earlier stage, the recruitment of the members from whom candidates are drawn, and show that money affects both how many people parties recruit and which rival organizations they recruit them from.

The second literature studies party membership, and the paper contributes to it in theory and in measurement. Existing models treat membership as a citizen's choice among parties with given platforms, driven by expressive benefits, policy preferences or the intensity of those preferences (Poutvaara, 2003; Gomberg, Marhuenda and Ortuño-Ortín, 2013; Kim, 2018). Parties in these models do not compete for members. Citizen-candidate models turned candidacy into an equilibrium choice by letting citizens weigh the costs and benefits of running (Osborne and Slivinski, 1996; Besley and Coate, 1997; Rivière, 1999). I write a model that does the same for membership, with competition on the side of the parties: parties with limited budgets offer support, citizens with private costs decide whether and where to join, and the allocation of recruits across parties follows from the offers. The equilibrium analysis uses standard tools from posted-offer matching and capped all-pay contests (Bulow and Levin, 2006; Che and Gale, 1998), and the model yields two predictions that the data can address: a party that loses money recruits less, and its rival gains access to contested recruits only under specific conditions. On measurement, empirical work shows that local office, descriptive representation and bureaucratic incentives shape membership in Brazil (Sells, 2020; Goyal and Sells, 2024; Frey and Santarrosa, 2026). I add budgets to these determinants and add direction to the measurement of membership: whose former members a party attracts, which rivals attract its own, and how both depend on the finances at each end of the link.

Section 2 presents the model. Section 3 describes the setting and data and estimates the effect of prior corporate funding on recruitment. Section 4 follows recruitment between parties: its prior-party component, its relation to relative finances, its persistence over time, and its distribution across ideological blocs and named rivals. Section 5 concludes. The appendices contain the proofs, extensions of the model, the construction of the data and additional results.

2 Theoretical Framework

This section models how a party's budget affects its recruitment when parties compete for the same citizens. The central object is a common offer of support. The offer plays two roles: it helps a citizen cover the cost of participating in politics, and it determines which party chooses first among citizens willing to join either one. The model delivers two predictions that the empirical analysis follows. Prediction 1 concerns the recruitment capacity of a party that loses money. Prediction 3 states when that loss gives its rival better access to contested recruits.

An offer represents the campaign resources and organizational assistance a recruit can expect from the party. It need not be a payment made when the citizen joins. As in the posted-wage model of Bulow and Levin (2006), the party commits to the same level of support for whoever it recruits, before it knows who will accept. This impersonal component links the contest for the most valuable recruit to the support available to every other member. I first characterize the allocation of citizens when everyone is willing to participate, then introduce private participation costs, so that the size of an offer also limits how many citizens a party can recruit. The comparative static that represents the funding ban combines the two margins.

The model complements theories of endogenous membership based on expressive motives (Poutvaara, 2003), policy preferences and idiosyncratic tastes (Gomberg, Marhuenda and Ortuño-Ortín, 2013), and the intensity of policy preferences (Kim, 2018). In those theories citizens choose among parties, but parties do not compete for them. The approach here is closer in spirit to citizen-candidate models (Osborne and Slivinski, 1996; Besley and Coate, 1997), which derive candidacy from citizens' costs and benefits of running, and to accounts of parties as organizations that share the costs of candidacy (Rivière, 1999) or screen applicants (Angiolillo, 2023). Here parties compete, and the object of competition is membership. I focus on the support parties can offer and on the competition it induces, and I abstract from ideology for now.

Citizens in the model have no prior party membership and can join either party. The model therefore does not derive a move from one party to another. It identifies a competitive-access margin, and the recorded membership histories provide an observable counterpart: recruitment of citizens whose previous membership was in a rival party.

2.1 Environment

There are N parties, indexed by organizational strength $Q_j \geq 0$ and ranked $Q_1 \geq \dots \geq Q_N$. Strength is an impersonal attribute of the party: its reputation, its infrastructure, its access to coalitions, and its television and radio time. It governs how the party converts a citizen's political contribution into votes. The same Q_j applies to every potential recruit and is taken as given when offers are chosen. Each party has k slots to fill and a total budget B_j , so the budget available per slot is

$$b_j = \frac{B_j}{k}. \quad (1)$$

The budget is divided equally across slots. Campaign-finance rules affect b_j through B_j , and rules on list length affect it through k . For a given total budget, the per-slot constraint is tighter in council races, where parties assemble lists, than in mayoral races, where $k = 1$.

There are $n \geq Nk$ citizens with electoral quality

$$q_i = n + 1 - i, \quad i = 1, \dots, n, \quad (2)$$

so that $q_1 > \dots > q_n$ with unit spacing. Quality measures what a citizen contributes to the party: a personal vote base, name recognition and local connections, as well as the work done as a member in organizing campaigns, mobilizing voters, carrying demands into the party and maintaining local networks. These contributions benefit the party whether or not the citizen

becomes a candidate. The party values them at $Q_j q_i$. Unit spacing makes the gain from moving up one rank the same for each party.

If party j posts w_j and recruits citizen i , its profit is $\pi_j = Q_j q_i - w_j$; if citizen i accepts w_j , the citizen's utility is $u_i = w_j - \underline{w}$, where $\underline{w}$ is the participation floor, the minimum support required to stand for office. The baseline has $N = 2$ symmetric parties with $Q_1 = Q_2 = Q$ and $b_1 = b_2 = b$, $k = 1$ slot each, and $n = 2$ citizens with $q_1 = 2 > 1 = q_2$. Ties are broken by a fair coin. Budgets satisfy $b_j \geq \underline{w}$. Section 2.4 introduces unequal budgets.

Assumption 2.1 (Impersonal offers). Each party posts a single offer w that applies to every citizen it might recruit. It is posted before the party observes who will accept, and it cannot be conditioned on the identity of the person recruited.

From the recruit's point of view, the relevant question is how much of the party's budget and organization will support the recruit's participation after joining. The offer w is a commitment to an expected level of campaign assistance and organizational support, measured in the same units as the participation cost. Assumption 2.1 isolates the common component of that support: the party sets its level before knowing who will fill the slot, and the commitment applies to whoever does.

In a mixed-strategy equilibrium, the uncertainty concerns which common offer the party posts. The support of the strategy is the set of possible levels of assistance, and the distribution gives their probabilities. Every realized offer respects the budget cap, and citizens observe it before accepting. Mixing therefore represents uncertainty about access to party resources before offers are made; the impersonal restriction concerns how an offer applies across recruits.

Assumption 2.2 (Gains from matching). $Q q_{Nk} > \underline{w}$: the lowest-ranked citizen who could be recruited is still worth more to a party than the participation floor.

Under Assumption 2.2, every match is profitable at the participation floor, including the match with the second citizen. In the deterministic two-party baseline there are two citizens and one slot per party, and both citizens accept any feasible offer. Each party therefore recruits, and the contest decides who obtains the better citizen. A party that loses first pick still recruits the second citizen and pays the support it offered. Relative to this fallback, the parties compete in an all-pay auction for the quality premium. Private participation costs, introduced below, can leave slots unfilled.

A common offer ties the resources a party spends competing for first pick to the support it gives whoever it recruits. With individual offers, a party could concentrate resources on the contested citizen and offer the floor to the other. The deterministic comparison in Appendix A shows that bidding for the better citizen then stops when the winner is indifferent between that match and recruiting the other citizen at the floor. Equilibrium profits are the same and both citizens are recruited. The quantity of recruitment becomes a question only when some citizens need more support to participate, which is the subject of Section 2.3. Private costs can also make recruitment respond to personalized offers; the restriction isolates the common component of party support without requiring all assistance to be impersonal.

The expected payoff from a tie at w can be written as

$$\Pi^{\text{tie}}(w) = Q \frac{q_1 + q_2}{2} - w = \underbrace{(Qq_2 - \underline{w})}_{\Lambda} + \underbrace{\frac{Q(q_1 - q_2)}{2}}_{\Gamma/2} - \underbrace{(w - \underline{w})}_{\Omega(w)}. \quad (3)$$

Definition 2.1 (The three quantities).

$$\Lambda \equiv Qq_2 - \underline{w}, \quad \Gamma \equiv Q(q_1 - q_2) = Q, \quad \Omega(w) \equiv w - \underline{w}. \quad (4)$$

Λ is the low-ball value: the payoff from recruiting the second citizen at the participation floor. It is the fallback of a party that gives up the contest for first pick. Because posting the floor can still win first pick in a tie, Λ is the value of the fallback match rather than every possible payoff from that offer.

Γ is the rank premium, the additional value of the first citizen over the second. It is the prize for which the parties compete, and under unit spacing it equals organizational strength. Its contribution to expected payoff depends on the probability of obtaining first pick.

$\Omega(w)$ is the offer slack, the support offered above the participation floor. It measures how far a party must go beyond minimum support to bid for the premium. The party pays it whether it obtains the first or the second citizen, and the budget limits it to at most $\Omega(b)$.

2.2 Equilibrium Under a Budget Ceiling

Lemma 2.1 (Payoff against a mixing rival). *If the rival plays a distribution F with no mass point at w , then posting w yields*

$$\Pi(w; F) = \Lambda + \Gamma F(w) - \Omega(w). \quad (5)$$

The payoff makes the all-pay structure explicit. The fallback value Λ is supplemented by the expected gain $\Gamma F(w)$ from winning first pick, at the cost $\Omega(w)$. Here $F(w)$ is the probability that the rival's offer is below w . At a tied offer, fair tie-breaking gives each party half of the rank premium in expectation, as in equation (3). A higher offer raises the chance of obtaining the better citizen, but its cost is paid even when the rival wins. At this stage every citizen participates, so budgets determine priority while both slots are filled.

An interior mass point cannot survive in a symmetric equilibrium: offering ε more replaces a tie with first pick at an arbitrarily small cost. A mass point at the budget ceiling can survive, because that deviation is infeasible. The equilibrium therefore depends on how much a party can offer relative to the value of first pick.

Proposition 2.2 (Symmetric equilibrium under a cap). *The symmetric equilibrium is unique and takes one of three forms, according to the offer slack $\Omega(b) = b - \underline{w}$:*

Offer slack	Equilibrium	Support	Π^*
$\Omega(b) \leq \Gamma/2$	pure, (b, b)	$\{b\}$	$\Lambda + \Gamma/2 - \Omega(b)$
$\Gamma/2 < \Omega(b) < \Gamma$	mixed, mass point at b	$[\underline{w}, \underline{w} + 2\Omega(b) - \Gamma] \cup \{b\}$	Λ
$\Omega(b) \geq \Gamma$	mixed, cap irrelevant	$[\underline{w}, \underline{w} + \Gamma]$	Λ

In the middle regime the density on the continuous part is $1/\Gamma$ and the mass point has size $\alpha = 2[\Gamma - \Omega(b)]/\Gamma$; in the top regime the distribution is uniform on an interval of length Γ .

Proposition 2.2 applies the capped all-pay auction of Che and Gale (1998) to recruitment. Their model of lobbying under contribution limits yields the same three regimes once offer slack replaces the bid and the rank premium replaces the valuation. Kaplan and Wettstein (2006) allow contestants to incur costs beyond a nominal cap. Here the budget is a hard constraint, because the party has no other source of resources in the model. The unconstrained contest is related to Siegel (2009), whose payoff characterization requires continuous costs and does not cover a hard ceiling.

When both parties offer b , a deviation to the floor saves $\Omega(b)$ but gives up half the rank premium. If $\Omega(b) \leq \Gamma/2$, the saving is too small to justify conceding first pick. Above this threshold the parties randomize between lower offers and the ceiling, and once $\Omega(b) \geq \Gamma$ the ceiling lies outside the unconstrained support, so further increases in the budget leave offers unchanged. Figure 2 illustrates the three cases.




Figure 2: Equilibrium offer distributions. With a tight budget, both parties offer b . As offer slack rises, the mass point at the ceiling shrinks; when the ceiling no longer binds, offers are uniform on an interval of length Γ . Filled circles mark endpoints the distribution attains; open circles mark endpoints it excludes.

2.3 Participation Costs and the Reach of an Offer

In the two-party baseline, a financial constraint reduces the number of recruits only if the available support leaves some citizens unwilling to participate. I therefore let the cost of standing for office vary across citizens. Employment and family obligations, for example, can make a candidacy costly in ways the party does not observe. The common offer now determines how much of the potential field a party can reach, in addition to its place in the recruitment order.

Assumption 2.3 (Party-blind private participation cost). Citizen i draws a private cost of standing, $c_i = \underline{w} + \epsilon_i$ with $\epsilon_i \stackrel{\text{iid}}{\sim} \text{Uniform}[0, \sigma]$, $\sigma > 0$, independent of quality and independent of the party making the offer. Citizen i accepts an available offer w if and only if $w \geq c_i$, and among acceptable offers takes the highest.

Lemma 2.3 (Reach). For $w \in [\underline{w}, \underline{w} + \sigma]$ each citizen independently accepts an available offer w with probability $p(w) = \Omega(w) / \sigma$, with $p = 0$ below and $p = 1$ above.

I call the acceptance probability $p(w)$ the reach of an offer: the expected fraction of citizens whose participation costs the offer covers. Its slope is $1/\sigma$ until $w = \underline{w} + \sigma$, where every citizen is willing to participate. Below that point, more support expands the accessible field. Beyond it, a higher offer can still improve priority but brings no additional citizens into the field.

Reach also gives competition for access its consequences. Because costs do not depend on the party, acceptance sets are nested: a citizen who accepts a lower offer also accepts a higher one. The party with the higher offer takes the best citizen in its acceptance set, and the other party takes the best remaining citizen in its own. A rival therefore changes the opportunities available at a given offer by selecting from a partly shared field. This ordering relies on party-independent costs and would not generally survive costs specific to a citizen–party pair.

Private costs do not guarantee that recruitment responds to every change in budgets. When dispersion is small and full reach is affordable, the parties return to competition over rank at a higher effective floor.

Proposition 2.4 (Displacement isomorphism). Suppose $\sigma \leq \Lambda/2$ and $b \geq \underline{w} + \sigma$. Then every feasible offer $w < \underline{w} + \sigma$ is strictly dominated by $w = \underline{w} + \sigma$, and on $[\underline{w} + \sigma, b]$ the payoffs of the stochastic game coincide with those of the deterministic game. Consequently the stochastic game and the deterministic game with participation floor $\underline{w} + \sigma$ have the same set of equilibria.

When the effective floor $\underline{w} + \sigma$ is feasible, Proposition 2.4 permits the substitution $\underline{w} \mapsto \underline{w} + \sigma$ in Proposition 2.2. Measured from the original floor, the regime thresholds become $\sigma + \Gamma/2$ and $\sigma + \Gamma$.

To study budgets that restrict participation, define the low-ball value $\varphi(u) = p(\underline{w} + u)(\Lambda - u)$, which combines the reach of an offer with the surplus from recruiting the second citizen.

Raising u expands reach and reduces that surplus, and φ is maximized at $m \equiv \min\{\sigma, \Lambda/2\}$. This yields a sufficient condition for parties to offer their full budgets.

Assumption 2.4 (Scarce budgets). $\Omega(b) \leq m = \min\{\sigma, \Lambda/2\}$.

Proposition 2.5 (Full-budget equilibrium). *Under Assumption 2.4 the profile (b, b) is the unique symmetric Nash equilibrium. All randomness in the outcome comes from participation shocks; there is no strategic mixing.*

Assumption 2.4 is sufficient for full-budget offers; Appendix B.7 characterizes the exact frontier. With $\sigma \leq \Lambda/2$, the threshold is $\bar{\Omega}(\sigma) = \sigma + \Gamma/2$: slack above the effective participation floor can reach half the rank premium before the pure profile ceases to be an equilibrium. Greater dispersion in participation costs makes low offers less attractive, because they reach fewer citizens.

Proposition 2.6 (Equilibrium composition). *At (b, b) , writing $p = p(b)$: the party recruits somebody with probability exactly p ; conditional on recruiting it obtains q_1 or q_2 with probability $\frac{1}{2}$ each; and $\Pi^* = p(b)[\Lambda + \Gamma/2 - \Omega(b)]$.*

In the symmetric full-budget equilibrium, a common budget cut lowers the probability of recruitment and leaves the conditional quality mix unchanged. Unequal budgets add the competitive-access margin: the party with the higher offer chooses first from the field it reaches. The funding ban raises the question of what happens when the exposed party loses this advantage, which connects its own recruitment capacity to its rival's opportunity to recruit.

2.4 The Ban as a Comparative Static

I consider a specific loss of financial advantage. Budgets satisfy $0 < \Omega(b_2) < \Omega(b_1) \leq m$, so both parties offer their full budgets, with $p_1 = p(b_1) > p_2 = p(b_2)$. Party 1 reaches more citizens and chooses first; party 2 selects from the remaining citizens it can reach. The ban reduces party 1's budget exactly to b_2 , where the common budget satisfies Assumption 2.4, and both parties then play the symmetric full-budget equilibrium. The cut lowers the exposed party's offer and replaces its first-pick advantage with a tie. The expressions below compare the two allocations, assigning zero quality to an unfilled slot so that expected quality combines the number and the composition of recruits.

Proposition 2.7 (Exact redistribution). *The equalizing cut reduces both margins of the exposed party, raises the unexposed party's expected quality, and weakly raises its recruitment. The changes are for the exposed party,*

$$\Delta_1^{\text{qual}} = \frac{q_1 + q_2}{2} p_2 - [q_1 p_1 + q_2 p_1 (1 - p_1)] < 0, \quad \Delta_1^{\text{quant}} = p_2 - p_1 (2 - p_1) < 0, \quad (6)$$

and for the unexposed party,

$$\Delta_2^{\text{qual}} = p_2 \left[\frac{q_1 + q_2}{2} - q_2 p_1 \right] > 0, \quad \Delta_2^{\text{quant}} = p_2 (1 - p_1) \geq 0. \quad (7)$$

The rival's recruitment gain is strict when $p_1 < 1$; it is zero at $p_1 = 1$.

Proposition 2.7 separates the two consequences of the cut. The initially richer party loses reach and priority, so its chance of filling a slot falls. Its rival keeps the same offer and reach but gains access, because it now chooses first in half of the ties. The rival's opportunities improve without any increase in its own resources. This is the competitive-access mechanism behind Prediction 3, alongside the exposed party's loss of capacity in Prediction 1.

The mechanism motivates the study of recruitment from prior parties. A membership record naming another party shows where a party's recruits come from as financial advantages

change. It does not show which party would have recruited that citizen under the model's counterfactual offers. The proposition concerns a change between two full-budget equilibria; Appendix B shows how additional parties, additional slots and other budget orderings change the rival's response.

Exact equalization is essential to the rival's gain. A partial cut that leaves party 1 strictly richer reduces its reach without giving the rival priority in a tie. With both parties still in the scarce region, party 2 recruits with probability $p_1 p_2$, which falls with p_1 , so both parties recruit less. Appendix B.12 derives this comparison. Prediction 3 therefore does not say that every loss at an exposed party benefits its rival. Outside the stated region, equilibrium offers and their ordering must be established before the sign of a recruitment response can be determined.

Under additional conditions, the capacity margin stops responding to resources. Fix a rival with $0 < \Omega_2 < m$ and suppose $\sigma > \Gamma$. The first-picker objective is then unimodal over feasible reach. Writing its target as τ_1 (derived in Appendix A), party 1 offers $w_1 = \underline{w} + \min\{\Omega_1, \tau_1\}$ for $\Omega_1 > \Omega_2$, so

$$\frac{\partial p_1}{\partial b_1} = \frac{1}{\sigma} \mathbb{1}\{\Omega_1 < \tau_1\}, \quad (8)$$

away from the kink. Recruitment $p_1(2 - p_1)$ has a declining marginal response below this target and is constant above it. With general dispersion, the constrained optimum compares an interior candidate with the budget endpoint, and a single truncated target need not describe it. Heterogeneity in thresholds does not by itself imply a particular shape for the aggregate dose-response.

2.5 Predictions

Prediction 1 (Recruitment falls in the exposed party). Its probability of filling a slot falls: $\Delta_1^{\text{quant}} < 0$.

Prediction 2 (Expected electoral quality falls). Expected electoral quality, including zero for an unfilled slot, falls: $\Delta_1^{\text{qual}} < 0$.

Prediction 3 (The unexposed rival gains). Under the equalizing cut, the rival fills its slot with weakly higher probability and draws higher expected quality: $\Delta_2^{\text{quant}} \geq 0$ and $\Delta_2^{\text{qual}} > 0$. The quantity gain is strict for $p_1 < 1$. Both statements are conditional on the two-party experiment. With three parties, the quality gain reverses when the pre-ban budget gap is large, with threshold $\tilde{b} \geq 2b - \underline{w}$: a party that recruited mostly in a segment only its money could reach was not crowding its rivals, and the ban adds a competitor to their pool. In the multiple-slot approximation, the quantity gain reverses unless $k > np_1/2$: when many citizens remain available, choosing second is not costly, and the rival loses only its exclusive access to the residual pool. Both calculations are in Appendix B; the multiple-slot cutoff is an approximation, not an exact finite-field result.

Prediction 4 (The municipality total falls). Both parties end up posting the lower budget, so fewer citizens accept at all and aggregate recruitment in the municipality falls even though one party gained.

Prediction 5 (The effect is capped). Under $\sigma > \Gamma$, a fixed rival with $0 < \Omega_2 < m$, and the first-picker ordering above, the exposed party's offer stops responding once $\Omega_1 \geq \tau_1$, so its loss under equalization stops increasing with its initial budget. This is a conditional saturation result. It does not imply a one-for-one response below the target or a prediction for all parties.

Predictions 1 and 3 distinguish recruitment capacity from competitive access, and the empirical analysis follows them in turn. The party-level comparison asks whether exposed parties record lower recruitment. The directed network then asks where recruitment comes

from and how it depends on the finances of both parties. For PT and PSDB, I report arrivals in each direction and each party's balance against its other rivals; the full network extends the same accounting to every supported party. Following flows in both directions separates lower intake from a larger outflow of former members. It does not separately identify the model's two causal responses or recover the aggregate effect in Prediction 4.

The mapping from the model to the data has three limits. First, citizens in the model have no prior membership, so a recorded prior-party flow is an empirical sourcing outcome rather than a derived switch; the previous membership need not have been active, and neither offers nor solicitation are observed. Second, gross receipts and prior corporate exposure do not measure the model's usable recruitment budget, common offer or participation cost. A positive association between relative finances and recruitment therefore supports the relevance of competition for recruits without reproducing the equalizing-cut experiment of Proposition 2.7. Third, the model has no ideological dimension, so the bloc ledger describes where flows go across the party spectrum rather than testing a prediction about ideological commitment.

Candidate attributes are observed only for the members who stand for office and bear only indirectly on electoral quality, which in the model also assigns zero to an unfilled slot. Prediction 5 motivates allowing the recruitment response to vary with exposure, although flexible dose profiles and directed flows do not identify its structural threshold. The per-slot budget in (1) suggests why list length matters: a fixed total budget provides less support per recruit when it is spread across more slots. The next section describes the data, the measures and the identifying assumptions.

3 Recruitment and Financial Resources

This section asks whether parties that relied more on corporate donations recruit fewer members after the ban. I estimate four specifications: binary and continuous exposure in 2012–2016 first differences, a dynamic panel over five electoral cycles, and a pooled panel. A continuous-treatment estimator then lets the recruitment response vary with the amount of prior corporate funding.

The comparison between exposed and unexposed parties combines the two margins of the model. An exposed party can lose recruits because its own reach contracts (Prediction 1), and an unexposed rival can gain access to contested citizens when relative offers change (Prediction 3). The estimates below therefore measure the difference between the two, not the exposed party's response in isolation. Section 4 examines the flows between parties that underlie this difference.

3.1 Setting, Data and Financial Exposure

The Supreme Federal Court prohibited corporate donations on September 17, 2015, and the 2016 municipal elections were the first held under the ban. Parties that had received corporate money lost that source of funding. I measure their exposure by their 2012 corporate receipts. Because parties could replace the lost funds or receive transfers from other levels of the party, exposure need not equal the eventual loss of resources available for recruitment.

Other rules changed at the same time. An electoral reform enacted later that month (Law 13.165/2015) reduced the minimum period of party membership required for a candidacy from twelve to six months and introduced spending limits for the 2016 campaigns (Avis et al., 2022). A constitutional amendment in February 2016 opened a thirty-day window in which elected officials could change party without losing their mandates. The first change applies to all parties and moves the timing of recruitment within the cycle; the others could affect exposed and unexposed parties differently, and I return to them when interpreting the results.

The data combine the membership registry of the Superior Electoral Court (TSE) with party accounts, election results and registered voters. The outcome is the recruitment flow into a party in a municipality during an electoral cycle, defined as the number of recorded membership spells that begin in the cycle. A cycle starts the day after a municipal election and ends on the next election day, so the 2016 cycle contains substantial recruitment before the ban. I divide recruitment by the number of registered voters in 2012, per 1,000 voters, and hold this denominator fixed across cycles, so that changes in the rate reflect changes in recruitment rather than in the electorate. The outcome measures entries into membership, not the stock of members.

The historical registry records membership start dates through April 2021. To follow the same parties over time, I apply one party classification, based on that snapshot, to the membership, financial and electoral records. Renamed parties keep a single identifier. Where the registry no longer separates parties that were incorporated or merged by 2021, their predecessors are grouped with the successor; mergers after 2021 are not applied retroactively. I refer to these groups as parties. Appendix C describes the grouping and its implications for measuring movements between parties.

Financial exposure comes from the accounts of the local party organization: its electoral committee and its municipal directory. Candidate accounts are analyzed separately and never added to party receipts, because transfers from parties to candidates would be counted twice. Observed accounts with no corporate receipts are distinct from missing accounts. Let p denote a party, m a municipality, and V_m the number of registered voters in m in 2012. Exposure is corporate receipts per registered voter, in nominal 2012 Brazilian reais (BRL):

$$D_{pm} = \frac{\text{2012 corporate receipts of party } p \text{ in } m \text{ (BRL)}}{V_m}. \quad (9)$$

The binary specification uses $\mathbb{1}\{D_{pm} > 0\}$ and the continuous specification uses D_{pm} . Both are determined before the ban. Appendix C describes the accounts and the coverage of the sample.

3.2 Recruitment Before and After the Reform

The first comparison follows the same local parties from the 2012 to the 2016 cycle and asks whether recruitment changed differently at parties that had received corporate money, conditional on their earlier electoral and membership record. With two periods, the difference-in-differences specification can be written in first differences:

$$\Delta y_{pm} = \beta \mathbb{1}\{D_{pm} > 0\} + X'_{pm} \gamma + \delta_p + \mu_m + \varepsilon_{pm}, \quad \Delta y_{pm} = y_{pm,2016} - y_{pm,2012}, \quad (10)$$

where y_{pmt} is recruitment per 1,000 registered voters in 2012, and δ_p and μ_m are party and municipality effects in first differences. The municipality effects restrict the comparison to parties competing in the same local electorate. Standard errors are clustered by municipality.

The vector X_{pm} contains the party's 2008 membership stock, its recruitment over the preceding cycle, and its 2008 electoral record: whether it contested the election, its vote share, its seat share, its candidates per seat, whether it held the mayoralty and its number of incumbents. These controls allow recruitment to follow different trajectories for parties with different histories. Ratios that are undefined for parties absent in 2008 follow the rules in Appendix C; missing records are not treated as zero activity.

Binary exposure. Recruitment at parties that received corporate contributions in 2012 declines, relative to parties that did not, by an amount equal to 15.1 percent of the 2012 mean of the estimation sample (Table 1). The coefficient is -0.517 recruits per 1,000 voters, with a standard error of 0.148, against a baseline of 3.426; the sample contains 14,606 local parties in 3,048

municipalities. The percentage scales the coefficient by baseline recruitment and is not a change in the share of the electorate that belongs to a party. It is a relative difference between exposed and unexposed parties, not an estimate of the change in national recruitment.

Table 1: Recruitment and prior corporate funding

Coding	Estimate	SE	2012 mean	N
Any corporate receipts	−0.517	0.148	3.426	14,606
BRL per voter	−0.120	0.060	3.426	14,606

The binary estimate is equivalent to 15.1% of the retained sample’s 2012 mean. The outcome is recorded recruitment flows per 1,000 voters. Party by municipality, using the common historical party groupings. First differences, 2016 minus 2012; party and municipality fixed effects and the 2008 controls. Standard errors cluster by municipality. Baseline means use retained observations; exposure is nominal 2012 BRL per voter.

Funding intensity. The second specification replaces the indicator with the amount of prior corporate funding:

$$\Delta y_{pm} = \beta D_{pm} + X'_{pm} \gamma + \delta_p + \mu_m + \varepsilon_{pm}. \quad (11)$$

It uses the same sample and controls as equation (10). The coefficient is −0.120 (0.060) per real of 2012 corporate receipts per voter (Table 1). This coding imposes a linear relationship on a distribution with a large mass at zero and a long upper tail, so the binary and continuous estimates answer different questions: the first compares parties with and without corporate money, the second compares parties with more and less of it.

Both estimates are negative. Recruitment falls more at parties that had corporate money, and more at parties that had more of it. In the model, these differences combine the exposed party’s loss of capacity with any change in its rivals’ access. Parties without corporate receipts are the comparison group, but their recruitment need not be unaffected by the reform.

Allowing the response to vary with exposure. The continuous-treatment difference-in-differences estimator of Callaway, Goodman-Bacon and Sant’Anna (forthcoming), hereafter CGBS, relaxes the linear restriction. I use 2012 as the base period and average the 2016 and 2020 post periods, comparing parties with positive exposure to parties with none. The sample contains 15,906 local parties in 4,347 municipalities, without the baseline controls or municipality–cycle effects, and a cubic spline with no interior knots allows a nonlinear relationship between recruitment and exposure. The average level difference (ATT) is −0.454 (0.102) recruits per 1,000 voters, and the average marginal response (ACRT) is −0.194 (0.019) per additional real of prior corporate receipts per voter. Standard errors are clustered by municipality. Table 17 reports these estimates next to two-way fixed-effects comparisons with their respective base periods.

At the reported percentiles of positive exposure, the estimated loss grows with exposure while its marginal magnitude declines (Table 18). The profile is concave over this range, which is consistent with the diminishing sensitivity to resources described in Prediction 5, but it does not identify a budget threshold. With two interior knots, the average marginal response becomes −3.187 (0.348) while the level estimate is unchanged (Table 31). The level difference is stable across specifications; the shape of the marginal response is not.

A causal reading of the CGBS level comparison requires parallel trends for untreated outcomes, and a causal reading of the marginal response requires stronger restrictions across exposure levels. Allowing a nonlinear response does not deliver either restriction, and the estimator does not remove spillovers onto the comparison parties.

3.3 When Does Recruitment Diverge?

The third specification follows recruitment over five electoral cycles. The earlier cycles show whether differences associated with 2012 exposure were already present before the ban, and the later cycles trace the post-reform path. The dynamic specification is

$$y_{pmt} = \sum_{\tau \neq 2012} \beta_{\tau}(D_{pm} \times \mathbb{1}\{t = \tau\}) + \alpha_{pm} + \lambda_{pt} + \theta_{mt} + \varepsilon_{pmt}, \quad (12)$$

for the 2004–2020 cycles, with the binary version replacing D_{pm} by $\mathbb{1}\{D_{pm} > 0\}$. The denominator is again the 2012 electorate. The 2004 and 2008 coefficients measure differences associated with later exposure before the reform; the 2016 and 2020 coefficients are post-reform.

Party–municipality effects absorb persistent differences between local organizations, party–cycle effects absorb each party’s national trajectory, and municipality–cycle effects absorb changes common to all parties in a municipality. These fixed effects replace the 2008 controls of the first-difference model. Table 2 reports both codings, and Figure 3 plots the continuous one. The dose coefficients are -0.124 (0.060) in 2016 and -0.137 (0.042) in 2020, so the negative gradient appears in both post-reform cycles.

Table 2: Recruitment dynamics relative to 2012

Cycle	Binary	SE	p	Dose	SE
2004	−0.423	0.166	0.011	−0.038	0.050
2008	−0.113	0.163	0.487	−0.032	0.055
2016	−0.527	0.137	<0.001	−0.124	0.060
2020	−0.853	0.134	<0.001	−0.137	0.042

Party–municipality, party–cycle and municipality–cycle fixed effects. Fixed 2012 electorate; municipality-clustered standard errors. Joint pre-period tests: binary $p = 0.026$; dose $p = 0.753$. $N = 73,030$.




Figure 3: Recruitment dynamics under dose exposure, relative to 2012. Fixed 2012 electorate; municipality-clustered intervals. The saturated panel uses $N = 73,030$ observations.

The binary coefficients are -0.527 (0.137) in 2016 and -0.853 (0.134) in 2020. The 2012–2016 comparison precedes the public campaign fund (FEFC), created in October 2017 and first distributed in 2018. The 2020 coefficient also reflects that fund, later changes in electoral rules

and the wider political environment, and it cannot isolate a persistent effect of the ban without further assumptions.

The pre-period evidence differs by coding. The joint test of the 2004 and 2008 coefficients gives $p = 0.026$ under binary exposure and $p = 0.753$ under dose exposure. Under binary exposure, the 2004 coefficient is negative and significant, so exposed parties recruited less in 2004 than in 2012 relative to unexposed parties. The earlier-window placebo in Appendix D points in the same direction: 2008 corporate funding predicts a decline in recruitment between 2008 and 2012, when there was no ban. That placebo uses committee accounts only, nominal amounts from a different year and a different sample, so it does not replicate the main design. Both results are consistent with corporate money flowing to local parties that were recruiting unusually well in the year it was measured, in which case part of the subsequent decline reflects reversion from that level. The dose coefficients before 2012 are small and imprecise, and the dose coefficients after 2012 are several times larger, which is the pattern the ban would produce.

3.4 Average Recruitment Differences in the Panel

The fourth specification summarizes the post-reform difference with a single coefficient on exposure interacted with an indicator for the 2016 and 2020 cycles:

$$y_{pmt} = \beta^{\text{pool}}(D_{pm} \times \mathbb{1}\{t \in \{2016, 2020\}\}) + \alpha_{pm} + \lambda_{pt} + \theta_{mt} + \varepsilon_{pmt}. \quad (13)$$

The binary version again replaces D_{pm} by $\mathbb{1}\{D_{pm} > 0\}$. This comparison averages over the three pre-reform cycles, whereas the dynamic coefficients and the CGBS estimator use 2012 as the reference; the CGBS comparison also omits the party–cycle and municipality–cycle effects, so the estimates need not coincide. Table 3 reports both codings under four sets of fixed effects. With all effects included, the dose estimate is negative: parties with more prior corporate funding recruit less after the reform relative to their own earlier record and to the other parties in their municipality.

Table 3: Pooled post-ban comparisons

Fixed effects	Binary	SE	Dose	SE	N
Unit + cycle	−0.349	0.080	−0.093	0.023	79,530
+ party–cycle	−0.253	0.078	−0.085	0.022	79,525
+ municipality–cycle	−0.546	0.091	−0.108	0.032	73,035
Both cycle interactions	−0.511	0.089	−0.107	0.031	73,030

Post includes both 2016 and 2020. Every specification includes party–municipality effects; standard errors cluster by municipality.

The corresponding binary estimate is −0.511 (0.089). Its size changes across fixed-effect sets because both the adjustment and the comparison change. None of these specifications identifies the sign or size of spillovers onto unexposed rivals.

The four specifications give a consistent result. Parties that had corporate money recruit less after 2012, the decline increases with the amount of money, and it persists in the panel. A causal interpretation requires that, after adjustment, recruitment would have followed the same path at exposed and unexposed parties in the absence of the reform, and that no other change related to exposure explains the post-reform pattern. The binary pre-period estimates, the placebo and the national shocks that hit particular parties make that attribution uncertain. Appendices D and E report the functional-form and timing comparisons.

If money governs competition for people who already belong to other parties, the decline should also appear among recruits with a previous party membership. The next section

examines that margin, which connects the capacity result to the competitive-access mechanism of Prediction 3.

4 Poaching and Local Party Competition

This section follows recruitment from its origin to its destination. I first show that the decline in recruitment at exposed parties is concentrated among recruits who previously belonged to another party. I then relate recruitment along each directed link between two local parties to the finances at both ends, follow these links over five electoral cycles, and describe how the flows are distributed across ideological blocs and between PT and PSDB. I refer to recruitment of citizens whose previous recorded membership was in another party as poaching. The term describes the recorded sequence of memberships; the registry does not show whether the earlier membership was still active or whether the new party approached the recruit.

4.1 Recruitment from Rival Parties

Table 4 repeats the specification in equation (10), with the same sample, controls and fixed effects, and replaces total recruitment with recruitment of citizens whose previous membership record names another party. At exposed parties, this flow declines by an amount equal to 19.5 percent of its 2012 mean. The estimate is -0.221 (0.050) per 1,000 voters. Prior-party recruitment is 33.1 percent of baseline recruitment but accounts for 42.8 percent of the estimated decline in total recruitment. This share is a ratio of two estimated coefficients rather than a mediation estimate, but it shows that the loss of recruitment capacity falls disproportionately on citizens with previous party experience.

Table 4: Recruitment with a recorded different prior party

Recorded origin	Estimate	SE	2012 mean	<i>p</i>
All prior-party recruitment flows	-0.221	0.050	1.135	<0.001
From the same bloc	-0.094	0.027	0.298	<0.001
From a different bloc	-0.075	0.033	0.575	0.021
Distance not scorable	-0.052	0.017	0.263	0.003
From a Left party	-0.031	0.015	0.194	0.040
From a Center party	-0.066	0.020	0.354	<0.001
From a Right party	-0.097	0.029	0.446	<0.001
From an unscored party	-0.027	0.011	0.142	0.014

Binary exposure and the headline specification. Origin is the preceding observed spell, with both parties assigned to the common historical groupings. It does not establish an immediate transfer, event-time party identity, or event-time municipality. Unscored includes merged party groups without a defensible common 2012 ideology. Each partition includes its unscored residual.

The decline includes recruits from ideologically close and distant parties. Recruitment from a party in the same bloc falls by 0.094 (0.027) per 1,000 voters, or 31.5 percent of that flow's 2012 mean. Recruitment from a party in another bloc falls by 0.075 (0.033), or 13.0 percent of its mean. The proportional decline is larger within blocs, although the comparison of the two point estimates is not a test that they differ. By origin, recruitment from Left, Center and Right parties all declines at exposed destinations. Right-wing origins account for the largest absolute decline and also supply the largest baseline flow.

These estimates describe whom an exposed party stops recruiting. They do not show whether a financially weakened party loses more of its own members to its rivals. That question requires the finances of both parties, which the next comparison introduces. The prior-

party comparison above requires financial records only at the recruiting party; the comparisons below require them at both ends of each link.

4.2 Relative Funding and the Direction of Poaching

Figure 1 measures each party's position within its municipality. Among the same local parties, I compute each party's share of local funding and its share of three recruitment flows: all recruitment, recruitment from rivals, and recorded movements to rivals. The share of incoming poaching, for example, is

$$\text{Party's share of incoming poaching} = \frac{\text{Recruits it attracts from eligible rivals}}{\text{All recruitment between eligible local parties}}.$$

The outgoing share uses the same denominator and counts the recruits that other parties attract from the focal party. The share of all recruitment has its own denominator, which includes recruits with no previous membership in the local set.

Losing financial share and losing recruitment share. The comparison follows 3,054 local parties in 719 municipalities between 2012 and 2016, using municipal-directory accounts in both years. A ten-percentage-point loss of local funding share is associated with a 0.680-percentage-point decline in the party's share of all recruitment and a 0.591-point decline in its share of recruitment from rivals; the second estimate is less precise. Its share of movements to rivals rises by 0.306 points, an imprecise estimate. Adding party and municipality fixed effects leaves the pattern unchanged (Table 15).

The clearest relationship concerns attraction. Parties that lose financial share account for a smaller share of recruitment, including recruitment from their competitors, and there is less evidence that their rivals draw more heavily on their members. The shares describe positions within the local recruitment pool, not individual probabilities of leaving a party. Appendix D.1 gives the exact denominators and reports the comparison without municipalities that have only two eligible parties, where the incoming and outgoing shares mirror each other by construction.

Following each origin–destination pair. The next comparison follows poaching between particular rivals. Let Poaching_{odmt} count recruits who join destination party d in municipality m during cycle t and whose preceding membership record names origin party o . The network contains 89,270 directed local pairs in 3,048 municipalities, including pairs with no recorded flow, and both parties in each pair have financial records in 2012. Because a single directed link carries few recruits, I express pair-level recruitment per 100,000 registered voters in 2012, whereas party-level recruitment in Section 3 is per 1,000. The network contains 112,023 flows in 2012 and 102,637 in 2016, about 11.2 percent of all transitions between parties recorded in the same municipality. Appendix C documents its construction.

I measure relative finances in two ways. The first uses realized funding: the destination has an advantage when its gross receipts grow more, or decline less, than the origin's between 2012 and 2016. I relate this advantage to the change in poaching per 100,000 registered voters, controlling for the pair's average growth, directed national party-pair effects and municipality effects. The specification asks how resources at each end relate to the flow between them, in the spirit of a gravity equation, without imposing a structural gravity model. Appendix D.1 gives the formula for the symmetric growth measure, which accommodates small baseline amounts.

The second measure uses the reform, through the exposure D_{pm} of Section 3 measured at both ends of the link. A destination is relatively less exposed when its origin received more corporate money. Let $r_{odmt} = 100,000 \times \text{Poaching}_{odmt} / \text{Voters}_{m,2012}$ denote the pair-level recruitment rate. The specification includes both parties' exposures:

$$\Delta r_{odm} = \beta_o D_{om} + \beta_d D_{dm} + \alpha_{od} + \mu_m + u_{odm}, \quad (14)$$

where Δ is the change from 2012 to 2016. The outcome is a recruitment rate, not the local shares of Figure 1. The two coefficients separate the origin’s exposure from the destination’s. Equivalently, I estimate the association with origin-minus-destination exposure holding the pair’s average exposure fixed. Directed pairs receive equal weight, and standard errors are clustered by municipality.

Table 5: Relative resources and changes in directional recruitment, 2012–2016

Measure and account sample	Slope (SE)	Pairs	Municipalities
Gross-receipt growth advantage	0.535 (0.219)	29,730	1,732
Prior corporate-exposure advantage	0.373 (0.170)	89,222	3,048
Directory-only growth advantage	0.591 (0.297)	14,826	978
Gross growth on directory overlap	0.820 (0.269)	14,826	978

The outcome is the change in recorded recruitment flows per 100,000 fixed 2012 voters. Each regression controls the pair mean of its financial regressor, directed national party-pair effects and municipality effects; standard errors cluster by municipality. Row 2 uses origin minus destination corporate receipts, in nominal 2012 BRL per voter. Rows 1, 3 and 4 use destination minus origin symmetric growth in real gross receipts, $2(B_{2016} - B_{2012}) / (B_{2016} + B_{2012})$. They require observed accounts in both years. Rows 3 and 4 share the same retained pairs; row 3 uses municipal directories in both years, while row 4 retains the original 2012 committee-plus-directory measure. Coefficients across the two units and different samples are not comparable elasticities.

Table 5 shows the same direction under both measures: poaching shifts toward destinations whose receipts grow relatively more and toward destinations that were relatively less exposed to corporate money. The slope on the destination’s receipt-growth advantage is 0.535 (0.219) per unit of symmetric growth, and the slope on its exposure advantage is 0.373 (0.170). Figure 11 in Appendix D.1 plots both relationships. A party’s ability to attract recruits from a rival depends on its resources relative to that rival’s, which is the competitive-access margin of Prediction 3.

The endpoint coefficients show which party drives the relationship. The slope on the destination’s receipt growth is 1.248 (0.308), and the slope on the origin’s receipt growth is 0.178 (0.325). Recruitment increases when the destination’s finances improve, but a decline in the origin’s finances does not increase departures from it. The specification does not impose equal and opposite responses at the two ends, and the data do not show them.

The growth relationship remains positive when municipal directories are used in both years, so that the same type of account is compared over time. That comparison and the original measure on the same overlap retain 14,826 pairs; the full growth and exposure regressions retain 29,730 and 89,222 pairs. Gross receipts respond to recruitment and to local political conditions, so these estimates are associations with realized receipts and prior exposure rather than elasticities of recruitment with respect to usable budgets.

4.3 Relative Exposure over Five Cycles

I follow the same directed local pairs through the 2004, 2008, 2012, 2016 and 2020 cycles. The panel allows each pair its own permanent level of recruitment and accounts for municipality-wide shocks and for the national trajectory of each party as origin and as destination:

$$r_{odmt} = \sum_{s \neq 2012} \mathbb{1}\{t = s\} (\beta_{o,s} D_{om} + \beta_{d,s} D_{dm}) + \alpha_{odm} + \mu_{mt} + \lambda_{ot} + \kappa_{dt} + u_{odmt}. \quad (15)$$

As in the first-difference comparison, the relative coefficient is $(\beta_{o,t} - \beta_{d,t})/2$, holding the pair’s average exposure fixed.

The relative coefficient is 0.369 (0.170) in 2016 and 0.307 (0.153) in 2020 (Table 6). Destinations that were less exposed than their origins receive more recruits from them after the reform, after accounting for these fixed effects. The pre-reform relative coefficients are also positive, 0.204 in 2004 and 0.171 in 2008, about half the post-reform values.




Figure 4: Relative and mean exposure in the five-cycle recruitment panel. The 2012 cycle is the reference. Relative exposure is origin minus destination nominal 2012 corporate receipts per voter; mean exposure is their average. The outcome is recruitment per 100,000 fixed 2012 voters. Coefficients transform the unrestricted endpoint estimates using their full covariance; intervals cluster by municipality. The two components impose different pre-period restrictions. The 2016 cycle includes pre-ban months.

Table 6: Relative-exposure dynamics and their pre-period diagnostics

Cycle relative to 2012	Relative-dose slope	SE	<i>p</i>
2004	0.204	0.136	0.134
2008	0.171	0.142	0.229
2016	0.369	0.170	0.030
2020	0.307	0.153	0.045
Joint pre-period test, 2004 and 2008			<i>p</i>
Continuous relative exposure			0.220
Continuous exposure: both endpoints			0.003
Binary relative exposure			0.004
Continuous relative exposure, narrower lineages			0.020

The outcome is recorded recruitment flows per 100,000 fixed 2012 voters. The full sample contains 446,350 pair-cycle observations in 3,048 municipalities over 2004–2020. Fixed effects are local directed pair, municipality–cycle, origin–cycle and destination–cycle. The relative coefficient is $(\beta_{o,t} - \beta_{d,t})/2$, conditional on pair mean exposure; its standard error uses the full municipality-clustered covariance matrix. All tests use that covariance. The full-endpoint test has four restrictions; relative tests have two. The narrower-lineage comparison excludes PSD, REP, PSOL and the mixed PCDOB, PODE and PATRIOTA/PRP families, retaining 57,898 pairs in 2,809 municipalities. Nonrejection is not evidence of parallel counterfactual trends.

The pre-period tests limit the causal interpretation. The test of the two pre-reform relative coefficients gives $p = 0.220$, but the test for mean exposure gives $p = 0.001$ and the joint test of both endpoints gives $p = 0.003$. The binary relative test rejects ($p = 0.004$), as does the continuous relative test after excluding newer and merged party groups ($p = 0.020$; 57,898 pairs in 2,809 municipalities). All tests use the full municipality-clustered covariance matrix. Because they impose different restrictions, failing to reject one of them does not establish parallel trends (Roth, 2022). Most of the pre-period difference is in the destination’s exposure: the joint test of the destination’s pre-reform coefficients gives $p < 0.001$, and the corresponding test for the origin gives $p = 0.342$. This is consistent with the reversion from 2012 levels discussed in Section 3.3.

Which destinations attract an origin’s recruits? Two conditional Poisson models examine allocation within observed recruitment pools (Santos Silva and Tenreiro, 2006). The first holds fixed each origin’s outgoing pool in a municipality and cycle and asks which destinations attract those recruits. The second holds fixed each destination’s intake and asks which origins supply it. Both use raw counts, keep zero flows, and interact corporate exposure with an indicator for the 2016 cycle, with local-pair effects and the other endpoint’s national cycle effects.

Table 7: Financial exposure and conditional recruitment allocation

Specification	Coefficient (SE)	p	Municipalities
Destination allocation	−0.022 (0.012)	0.061	1,849
Origin allocation	0.003 (0.006)	0.615	1,820

Conditional Poisson log-mean slopes for counts per nominal 2012 corporate BRL per voter, with municipality-clustered standard errors. Destination allocation holds origin–municipality–cycle effects fixed; origin allocation holds destination–municipality–cycle effects fixed. Both include local directed-pair effects and opposite national endpoint–cycle effects. Retained samples are 82,338 observations in 1,849 municipalities and 77,566 in 1,820, respectively. These are different observed recruitment pools and supports; the coefficients do not identify total induced flows.

The allocation models point to the destination again. Holding the origin’s outgoing pool fixed, a more exposed destination receives a smaller share of that pool: the slope is −0.022 (0.012), with $p = 0.061$. Holding the destination’s intake fixed, the slope on origin exposure is 0.003 (0.006), with $p = 0.615$. The two models use different pools and retained samples, so the difference in precision is not a test that the two responses differ. Both are consistent with exposure reducing a party’s ability to attract recruits more than it increases the rate at which its own members leave.

4.4 Ideological Blocs

Table 8 distributes the network’s flows across ideological blocs. Arrivals into Left parties fall from 28,184 to 19,004, and the Left’s share of arrivals falls by 6.6 percentage points. PT alone accounts for 8,133 of the bloc’s 9,180 fewer arrivals, or 88.6 percent. PSB’s arrivals, by comparison, are nearly unchanged at 6,605 and 6,692. The contraction of the Left’s intake is therefore the contraction of PT’s.

The other blocs gain shares: 2.4 percentage points for the Center, 1.0 for the Right and 3.3 for unscored parties. These gains do not require more arrivals. Center arrivals are almost unchanged, and the Center’s net balance improves mainly because fewer of its members join other parties. Among unscored parties, arrivals into the PODE group grow from 3,277 to 5,750. The party ledger in Section 4.5 shows which organizations account for these changes.

The blocs use the fixed 2013 scores of Zucco and Power (2024): Left below −0.3, Center between −0.3 and 0.3, and Right above 0.3. PT scores −0.532 and PSDB 0.219, which places

Table 8: Recruitment and redistribution within the supported network

Bloc	Recorded arrivals		Net recorded flow	
	2012	2016	2012	2016
Left	28,184	19,004	3,261	−4,872
Center	28,623	28,680	−9,730	−1,255
Right	29,478	27,990	−2,823	−810
Unscored	25,738	26,963	9,292	6,937

The same 89,270 directed party-pair-municipality cells in 3,048 municipalities are retained in both cycles, including zero flows. Both endpoints have observed 2012 finance support. Arrivals have a different previous observed lineage; net flow subtracts recorded outgoing transitions within this graph. Net flows sum to zero across the four rows and are not membership-stock changes. Fixed 2013 party positions define blocs; unscored lineages remain separate.

PSDB in the Center, so the rivalry between the two parties is distinct from a comparison between Left and Right. Unscored parties remain in the totals and in the share denominators. The ledger describes the distribution of recruitment across parties. Appendix D.4 tests whether ideological distance or extremity changes the relationship between relative exposure and poaching; the interactions do not show that ideological separation protects parties from financial competition.

4.5 PT, PSDB and the Anatomy of a Reallocation

The rivalry between PT and PSDB shows how a large reversal in the balance of recruitment can arise from a change in one party's intake. In the 591 municipalities where both parties have financial records, PT-to-PSDB flows rise from 883 to 966 between the 2012 and 2016 cycles, while PSDB-to-PT flows fall from 1,725 to 504. PSDB's net balance with PT moves from −842 to 462. Of this reversal of 1,304 recorded transitions, 1,221, or 93.6 percent, come from fewer PSDB members joining PT.




Figure 5: PT–PSDB and the wider supported local network in 591 common municipalities. Arrivals and origin counts use the same eligible counterparties in both cycles. Origin counts are later recruitment records naming the focal party as the previous observed party, not verified exits from active membership. Wider comparisons include the named rival.

The pattern extends beyond the bilateral relationship. Figure 5 holds the 591 municipalities fixed and compares each party with all its eligible rivals. PT's arrivals fall from 9,084 to 2,805, and its net balance moves from 2,606 to $-4,167$. PSDB's arrivals rise from 6,859 to 7,180, and its balance moves from $-2,295$ to 791. PT attracts far fewer recruits from the rest of the party system, while PSDB's intake grows only modestly. Both comparisons include the named rival, so the flows overlap.

Table 9: Party-by-party recruitment within the fixed supported network

Party	Bloc	Municipalities	Recorded arrivals		Net recorded flow	
			2012	2016	2012	2016
PT	Left	1,391	12,682	4,549	2,901	$-5,814$
MDB	Center	1,475	10,690	11,613	$-6,599$	$-1,101$
PSDB	Center	1,238	10,207	11,107	$-3,317$	1,178
PSD	Right	923	9,111	8,165	8,867	4,034
PDT	Left	842	7,683	6,251	$-1,023$	$-1,095$
PP	Right	1,186	6,796	6,730	$-3,293$	$-1,423$
PSB	Left	819	6,605	6,692	644	1,056
PTB	Right	862	5,816	5,449	$-3,048$	$-1,104$
PCDOB	Unscored	448	5,685	4,228	2,766	800
PL	Right	745	4,406	4,415	$-1,090$	-342
PV	Center	564	4,385	3,449	1,077	-388
DEM	Right	725	3,349	3,231	$-4,259$	$-1,975$
CID	Center	560	3,341	2,511	-891	-944
PODE	Unscored	355	3,277	5,750	1,100	2,506
PSC	Unscored	479	3,224	2,908	715	411
REP	Unscored	404	2,983	3,667	1,288	1,609
AVANTE	Unscored	216	2,076	942	1,120	-732
PSDC	Unscored	169	1,805	985	663	-142
PSL	Unscored	260	1,701	1,489	477	123
Patriota/PRP	Unscored	209	1,258	3,703	270	2,178
PMN	Unscored	215	1,253	961	-2	-340
PSOL	Left	120	1,214	1,512	739	981
PRN	Unscored	166	1,168	1,260	436	413
PRTB	Unscored	181	1,081	971	299	126
PSTU	Unscored	43	217	85	161	-9
PCB	Unscored	11	9	8	-1	-6
PCO	Unscored	1	1	6	0	0
Full network		3,048	112,023	102,637	0	0

All 27 parties, harmonized to common 2021 groupings, in the fixed 89,270-directed-pair graph are shown, ordered by 2012 arrivals. Each party uses the same eligible municipalities and rivals in both cycles; these sets differ across parties and overlap. Both endpoints require 2012 finance support. Arrivals are recruitment flows with a different previous recorded party in the same recorded municipality. Net subtracts origin counts (subsequent recruitment events elsewhere with this party recorded as the prior party), not verified active departures. These are event counts, not unique recruits or membership-stock changes. Blocs use fixed 2013 positions; incompatible merged families PCDOB, PODE and Patriota/PRP remain Unscored. Common party groupings and party entry limit comparisons of levels. The PT/PSDB rows cover each party's full eligible support, rather than the 591 municipalities common to PT and PSDB in the named-rivalry figure.

Table 9 extends the comparison to each party's full set of eligible municipalities. PT again loses arrivals and becomes a net source of recruits, and PSDB gains arrivals and becomes a net recipient. MDB's arrivals rise from 10,690 to 11,613, and its net deficit narrows from 6,599 to 1,101, so several parties share in the redistribution. Each party's set of municipalities is fixed across cycles but differs across parties and from the 591 common municipalities. Parties with short histories have fewer members who can later appear in another party, so net balances should be read together with arrivals rather than as a ranking of organizational strength.

The rivalry illustrates the allocation mechanism, but it does not identify its financial cause. Across municipalities, the change in the PSDB-favoring net rate has a slope of 0.215 (0.461) on relative exposure and 0.991 (1.834) on relative receipt growth in the 195 municipalities with accounts in both years. With one observation per municipality, these comparisons cannot include municipality effects. The history of the rivalry also shows that 2012 was an unusual cycle: PSDB's net balance with PT in the same municipalities is -484 in 2004, 60 in 2008, -842 in 2012, 462 in 2016 and 292 in 2020. The reversal coincides with the national collapse of PT after 2014, which the party-cycle effects of the network regressions absorb. The directed network, not the named rivalry, carries the evidence that relative finances move recruitment; the rivalry shows how a fall in one party's attraction can reshape the balance between two major competitors.

Taken together, the results in this section locate the recruitment decline in the competition for people with previous party experience. The decline includes recruits from within and across ideological blocs, relative financial improvement is associated with stronger recruitment by the destination, and the largest reallocation of the period, the reversal between PT and PSDB, operates through reduced recruitment into PT. The pre-period evidence and the contemporaneous changes in rules and national politics limit how much of this pattern can be attributed to the ban. Appendix D reports annual and election-window comparisons, and Appendix E reports the sensitivity and inference details.

5 Conclusion

Institutions depend on the people who fill them, and in a representative democracy most of those people arrive through political parties. Parties recruit the citizens who staff campaigns, carry local demands into politics and eventually stand for office. In Brazil, where every candidate must be a party member, the competition among parties for these citizens is the first stage of political selection. This paper studies how money shapes that competition.

The paper makes two contributions. The first is a model of competition for members. Citizen-candidate models made candidacy an equilibrium choice; this model does the same for membership, with parties on the competing side. Each party offers a common level of support under a budget constraint, and citizens with private costs decide whether and where to join. A budget buys reach, the number of citizens a party can bring into politics, and priority over the citizens that its rivals also want. The model predicts that a party that loses money recruits less, and that its rival gains access to contested recruits when the loss removes the exposed party's advantage. The second contribution is to measure membership as a flow between parties, using administrative histories that record where each recruit came from.

This competition is local. Brazilian parties recruit in more than five thousand municipalities, and before each municipal election the same organizers, community leaders and former candidates are sought by several parties at once. Following recruitment within municipalities turns membership totals into a map of local rivalry: which party draws on which, and how the balance moves when resources change.

The evidence around the 2015 ban on corporate donations is consistent with the model. Recruitment at parties that had received corporate money falls, relative to parties without corporate receipts in the same municipality, by an amount equal to 15.1 percent of its 2012 mean. The decline is concentrated among citizens with previous party experience: recruits from other parties are 33.1 percent of baseline recruitment but account for 42.8 percent of the decline. Along local links between parties, recruitment moves toward the party whose finances improve relative to its rival, and the adjustment operates through that party's intake. Relative money governs a party's ability to attract its rivals' members; losing financial ground reduces that attraction more than it increases departures. The largest reallocation of the period took

this form: the reversal between PT and PSDB came almost entirely from fewer PSDB members joining PT.

These results concern the pool from which parties select their candidates. If money decides which parties can attract politically engaged citizens, and whether some citizens enter party politics at all, campaign finance shapes representation before any candidate reaches the ballot. The question that remains is whether the ban changed who enters that pool, and not only which party they join. A reform that narrows the pool of engaged citizens, or sorts them toward the parties that can still afford to support them, could be costly for representative democracy even if it makes campaigns cleaner.

References

- Angiolillo, Fabio. 2023. "Authoritarian Ruling Parties' Recruitment Dilemma: Evidence from China." *Journal of East Asian Studies* 23 (3): 491–515.
- Avis, Eric, Claudio Ferraz, Frederico Finan, and Carlos Varjão. 2022. "Money and Politics: The Effects of Campaign Spending Limits on Political Entry and Competition." *American Economic Journal: Applied Economics* 14 (4): 167–199.
- Barber, Michael J. 2016. "Ideological Donors, Contribution Limits, and the Polarization of American Legislatures." *The Journal of Politics* 78 (1): 296–310.
- Bekkouche, Yasmine, Julia Cagé, and Edgard Dewitte. 2022. "The Heterogeneous Price of a Vote: Evidence from Multiparty Systems, 1993–2017." *Journal of Public Economics* 206: 104559.
- Besley, Timothy, and Stephen Coate. 1997. "An Economic Model of Representative Democracy." *Quarterly Journal of Economics* 112 (1): 85–114.
- Bulow, Jeremy, and Jonathan Levin. 2006. "Matching and Price Competition." *American Economic Review* 96 (3): 652–668.
- Cagé, Julia, Caroline Le Pennec, and Elisa Mougin. 2024. "Firm Donations and Political Rhetoric: Evidence from a National Ban." *American Economic Journal: Economic Policy* 16 (3): 217–256.
- Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. Forthcoming. "Difference-in-Differences with a Continuous Treatment." *American Economic Review*.
- Che, Yeon-Koo, and Ian L. Gale. 1998. "Caps on Political Lobbying." *American Economic Review* 88 (3): 643–651.
- Cinelli, Carlos, and Chad Hazlett. 2020. "Making Sense of Sensitivity: Extending Omitted Variable Bias." *Journal of the Royal Statistical Society, Series B (Statistical Methodology)* 82 (1): 39–67.
- Dasgupta, Partha, and Eric Maskin. 1986. "The Existence of Equilibrium in Discontinuous Economic Games, I: Theory." *The Review of Economic Studies* 53 (1): 1–26.
- Fourinaies, Alexander. 2021. "How Do Campaign Spending Limits Affect Elections? Evidence from the United Kingdom 1885–2019." *American Political Science Review* 115 (2): 395–411.
- Frey, Anderson, and Rogerio Santarrosa. 2026. "The Politicization of Bureaucrats: Evidence from Brazil." *The Journal of Politics* 88 (2): 554–571.
- Gomberg, Andrei M., Francisco Marhuenda, and Ignacio Ortuño-Ortín. 2013. "Endogenous Party Platforms: 'Stochastic' Membership." Centro de Investigación Económica, ITAM, Working Paper 13-07.
- Goyal, Tanushree, and Cameron Sells. 2024. "Descriptive Representation and Party Building: Evidence from Municipal Governments in Brazil." *American Political Science Review* 118 (4): 1840–1855.
- Kaplan, Todd R., and David Wettstein. 2006. "Caps on Political Lobbying: Comment." *American Economic Review* 96 (4): 1351–1354.
- Katz, Richard S., and Peter Mair. 1995. "Changing Models of Party Organization and Party Democracy: The Emergence of the Cartel Party." *Party Politics* 1 (1): 5–28.
- Kim, Kwang-ho. 2018. "A Model of Endogenous Party Membership and Platforms with Opportunistic Politicians." *The Korean Economic Review* 34 (1): 51–73.

- Mair, Peter, and Ingrid van Biezen. 2001. "Party Membership in Twenty European Democracies, 1980–2000." *Party Politics* 7 (1): 5–21.
- Osborne, Martin J., and Al Slivinski. 1996. "A Model of Political Competition with Citizen-Candidates." *Quarterly Journal of Economics* 111 (1): 65–96.
- Peveri, Julieta. 2024. "Corporate Contributions and Political Selection: Evidence from Brazil." Unpublished manuscript.
- Poutvaara, Panu. 2003. "Party Platforms with Endogenous Party Membership." *Public Choice* 117 (1–2): 79–98.
- Reny, Philip J. 1999. "On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games." *Econometrica* 67 (5): 1029–1056.
- Rivière, Anouk. 1999. "Citizen Candidacy, Party Formation and Duverger's Law." ECARE, Université Libre de Bruxelles. Unpublished manuscript.
- Roth, Jonathan. 2022. "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends." *American Economic Review: Insights* 4 (3): 305–322.
- Santos Silva, J. M. C., and Silvana Tenreyro. 2006. "The Log of Gravity." *The Review of Economics and Statistics* 88 (4): 641–658.
- Sells, Cameron J. 2020. "Building Parties from City Hall: Party Membership and Municipal Government in Brazil." *The Journal of Politics* 82 (4): 1576–1589.
- Siegel, Ron. 2009. "All-Pay Contests." *Econometrica* 77 (1): 71–92.
- Speck, Bruno Wilhelm. 2013. "Nem ideológica, nem oportunista: a filiação partidária no contexto pré-eleitoral no Brasil." *Cadernos Adenauer* XIV (2): 37–60.
- Speck, Bruno Wilhelm, Maria do Socorro Sousa Braga, and Valeriano Costa. 2015. "Estudo exploratório sobre filiação e identificação partidária no Brasil." *Revista de Sociologia e Política* 23 (56): 125–148.
- Speck, Bruno Wilhelm. 2016. "Game over: duas décadas de financiamento de campanhas com doações de empresas no Brasil." *Revista de Estudos Brasileños* 3 (4): 125–135.
- Zucco, Cesar, and Timothy J. Power. 2024. "The Ideology of Brazilian Parties and Presidents: A Research Note on Coalitional Presidentialism Under Stress." *Latin American Politics and Society* 66 (1): 178–188.

Appendix

Appendix A contains the proofs of the results in Section 2, and Appendix B extends the model. Appendix C describes the administrative records and the construction of the samples. Appendices D and E report additional results and robustness checks.

A Proofs

Throughout, $\Lambda = Qq_2 - \underline{w}$, $\Gamma = Q(q_1 - q_2)$ and $\Omega(w) = w - \underline{w}$ as in Definition 2.1, and a strategy is a CDF on $[\underline{w}, b]$. An equilibrium requires constant expected payoff almost surely under each player's strategy (CEU) and no profitable feasible deviation (NPD). Equality of payoffs also holds at support points where the expected payoff is continuous, but it need not hold at a zero-probability endpoint of the support that faces an atom of the rival. Existence in the symmetric baseline follows from the explicit construction below. I make no general existence claim for asymmetric hard caps.

Equilibrium under a Budget Ceiling

Proof of Lemma 2.1. Step 1. Acceptance is deterministic and there are two citizens and one slot each, so both citizens are matched and each party fills its slot with certainty. Whatever the rival does, the party pays w .

Step 2. The party is first picker exactly when its offer strictly exceeds the rival's, an event of probability $F(w^-)$; since F has no mass at w , $F(w^-) = F(w)$. On that event it takes q_1 , otherwise q_2 , so expected quality is $Q[q_1F(w) + q_2(1 - F(w))]$.

Step 3. Subtracting the payment and adding and subtracting $\underline{w}$,

$$\Pi(w; F) = Qq_2 + Q(q_1 - q_2)F(w) - w = (Qq_2 - \underline{w}) + \Gamma F(w) - (w - \underline{w}) = \Lambda + \Gamma F(w) - \Omega(w).$$

□

Lemma A.1 (No mass point below the cap). *In any symmetric equilibrium the common distribution F^* has no mass point in $[\underline{w}, b)$. A mass point can sit only at the cap.*

Proof. Suppose F^* places a mass point of size $a > 0$ at $w_0 \in [\underline{w}, b)$. Posting w_0 wins first pick outright with probability $F^*(w_0^-)$ and ties with probability a , winning half of those, so $\Pi(w_0) = \Lambda + \Gamma[F^*(w_0^-) + a/2] - \Omega(w_0)$. For $\varepsilon > 0$ small enough that $w_0 + \varepsilon < b$ (possible precisely because $w_0 < b$), posting $w_0 + \varepsilon$ beats the entire mass point, so $\Pi(w_0 + \varepsilon) \geq \Lambda + \Gamma[F^*(w_0^-) + a] - \Omega(w_0) - \varepsilon$. Subtracting, $\Pi(w_0 + \varepsilon) - \Pi(w_0) \geq \Gamma a/2 - \varepsilon > 0$ for $\varepsilon < \Gamma a/2$. The gain is discrete (half the mass point times the rank premium), while the cost is arbitrarily small, contradicting (CEU) at the common positive-mass action w_0 .

At $w_0 = b$ the deviation is infeasible, so the argument does not apply and a mass point can survive. This is why a mass point can arise only at the ceiling. □

Lemma A.2 (Indifference pins the distribution). *In any symmetric equilibrium with slack budgets, $\Pi^* = \Lambda$, $F^*(w) = \Omega(w)/\Gamma$ and $f^*(w) = 1/\Gamma$ on $[\underline{w}, \underline{w} + \Gamma]$.*

Proof. Step 1. Let $w_0 = \inf \text{supp } F^*$ and suppose $w_0 > \underline{w}$. Any offer in $[\underline{w}, w_0)$ loses first pick with certainty, and so does w_0 itself, since by Lemma A.1 there is no mass point there. Both therefore deliver the second citizen, but w_0 costs strictly more, so w_0 is dominated by $\underline{w}$ and cannot lie in the support. Hence $w_0 = \underline{w}$ and $F^*(\underline{w}) = 0$.

Step 2. Evaluating Lemma 2.1 at $\underline{w}$ gives $\Pi(\underline{w}) = \Lambda$, and by Step 1 $\underline{w}$ is in the support, so $\Pi^* = \Lambda$. The equilibrium payoff is known before the equilibrium strategy is: competition dissipates the entire rank premium and each party keeps exactly what it could have guaranteed by not competing.

Step 3. For w in the support, (CEU) and Lemma 2.1 give $\Lambda + \Gamma F^*(w) - \Omega(w) = \Lambda$, so $F^*(w) = \Omega(w)/\Gamma$. Differentiating, $f^* = 1/\Gamma$. The support ends where the mass is exhausted, $\Omega(\bar{w}) = \Gamma$, so $\bar{w} = \underline{w} + \Gamma$, which is feasible iff $\Omega(b) \geq \Gamma$. $\square$

Proof of Proposition 2.2. Three regimes in turn.

Regime 1: $\Omega(b) \leq \Gamma/2$. Existence. Suppose the rival posts b . Raising above b is infeasible. Any deviation to $w \in (\underline{w}, b)$ loses first pick with certainty and so delivers the second citizen above the floor, dominated by deviating to $\underline{w}$. So only $w = \underline{w}$ need be checked, worth Λ . Staying at b delivers the tie payoff (3), and staying is optimal iff $\Lambda + \Gamma/2 - \Omega(b) \geq \Lambda$, that is $\Omega(b) \leq \Gamma/2$. Uniqueness. A symmetric mixed equilibrium would have density $1/\Gamma$ and a mass point at b of size $\alpha = 2[\Gamma - \Omega(b)]/\Gamma$ by Lemma A.2 and Lemma A.1. But $\Omega(b) < \Gamma/2$ gives $\alpha > 1$, impossible. Only the degenerate distribution at b survives.

Regime 2: $\Gamma/2 < \Omega(b) < \Gamma$. Conjecture support $[\underline{w}, \tilde{w}] \cup \{b\}$ with density $1/\Gamma$ below $\tilde{w}$ and a mass point α at b .

The gap is empty. For $w \in (\tilde{w}, b)$ the rival's remaining mass sits at the cap, so $F^*(w) = \Omega(\tilde{w})/\Gamma$ is constant there and, by Lemma 2.1, $\Pi(w) = \Pi^* - (w - \tilde{w}) < \Pi^*$. Bidding just under the ceiling buys no extra chance of first pick and costs strictly more.

Pinning $\tilde{w}$. Write $w' = \Omega(\tilde{w})$ and $\alpha = 1 - w'/\Gamma$. Posting b wins against the continuous part with probability w'/Γ and ties against the mass point, so

$$\Pi(b) = Q \frac{q_1 + q_2}{2} + \frac{w'}{2} - b.$$

Setting $\Pi(b) = \Pi^* = \Lambda$ gives $\Gamma/2 + w'/2 = \Omega(b)$, hence $w' = 2\Omega(b) - \Gamma$, so $\tilde{w} = \underline{w} + 2\Omega(b) - \Gamma$ and $\alpha = 2[\Gamma - \Omega(b)]/\Gamma$. Both lie in the required ranges exactly when $\Gamma/2 < \Omega(b) < \Gamma$, and the regimes glue: $\alpha \rightarrow 1, \tilde{w} \rightarrow \underline{w}$ as $\Omega(b) \downarrow \Gamma/2$, and $\alpha \rightarrow 0, \tilde{w} \rightarrow \underline{w} + \Gamma$ as $\Omega(b) \uparrow \Gamma$.

Regime 3: $\Omega(b) \geq \Gamma$. The whole support of Lemma A.2 is affordable and the constraint never binds, so the equilibrium is uniform on $[\underline{w}, \underline{w} + \Gamma]$ with $\Pi^* = \Lambda$. That no pure profile survives here is immediate: if $w_1^* > w_2^*$ party 1 can shave and still win; if $w_1^* = w_2^* = w^* < b$ deviating up by ε wins first pick for a gain of $\Gamma/2 - \varepsilon > 0$; and if $w^* = b$ deviating down to $\underline{w}$ gains $\Omega(b) - \Gamma/2 \geq \Gamma/2 > 0$. $\square$

Remark A.1 (Where slackness is used). The budget enters only in the last sub-case, and it needs $\Omega(b) \geq \Gamma/2$ for the downward deviation to pay. Below that threshold the argument fails and the pure profile at the ceiling returns, which is regime 1. Nothing in the proof touches Assumption 2.1 or 2.2: mixing is not a feature of impersonality as such, but of unconstrained competition under it.

Proposition A.3 (Personalized offers). *In the deterministic baseline with $k = 1$, $n = 2$ and $Q_1 = Q_2 = Q$, personalized offers give $w^{(1)} = \min\{\underline{w} + \Gamma, b\}$ to the better citizen and $w^{(2)} = \underline{w}$ to the other, in pure strategies, with equilibrium profit $\Pi^* = \Lambda$ whenever $\Omega(b) \geq \Gamma$.*

Proof. Citizen 2 faces no competition once the rival is targeting citizen 1, so any offer above $\underline{w}$ is dominated and $w^{(2)} = \underline{w}$. Competition for citizen 1 is then Bertrand in the surplus $Q(q_1 - q_2) = \Gamma$, since a party that loses pays nothing on the offer it made. Bidding stops where the winner is indifferent to taking citizen 2 at the floor, $Qq_1 - w^{(1)} = Qq_2 - \underline{w}$, giving $w^{(1)} = \underline{w} + \Gamma$, truncated at b . When $\Omega(b) \geq \Gamma$, both parties earn $Qq_2 - \underline{w} = \Lambda$; with a tighter cap their expected payoff is $\Lambda + [\Gamma - \Omega(b)]/2$. $\square$

Participation Costs

Proof of Lemma 2.3. $c_i \leq w \iff \epsilon_i \leq w - \underline{w}$, and $\epsilon_i \sim U[0, \sigma]$ gives $\mathbb{P}(\epsilon_i \leq x) = x/\sigma$ on $[0, \sigma]$, zero below and one above. Independence across citizens is inherited from the iid shocks. $\square$

Lemma A.4 (The low-ball optimum). $\varphi(u) = p(\underline{w} + u)(\Lambda - u)$ is strictly increasing on $[0, m]$ and strictly decreasing on $[m, \infty)$, where $m = \min\{\sigma, \Lambda/2\}$, with $\varphi(m) = \Lambda - \sigma$ if $\sigma \leq \Lambda/2$ and $\varphi(m) = \Lambda^2/(4\sigma)$ otherwise.

Proof. On $[0, \sigma]$, $\varphi(u) = u(\Lambda - u)/\sigma$, a downward parabola with vertex at $u = \Lambda/2$; on $[\sigma, \infty)$, $\varphi(u) = \Lambda - u$, strictly decreasing. If $\sigma \leq \Lambda/2$ the vertex lies right of the kink, so the parabola is still rising at σ and the maximum is the kink, with $\varphi(\sigma) = \Lambda - \sigma$. If $\sigma \geq \Lambda/2$ the vertex is interior and $\varphi(\Lambda/2) = \Lambda^2/(4\sigma)$. $\square$

Proof of Proposition 2.4. Coincidence above $\underline{w} + \sigma$. For $w \geq \underline{w} + \sigma$, $p(w) = 1$: the first picker takes q_1 and pays w , the second takes q_2 and pays w , which are the deterministic payoffs.

Dominance below. Fix a rival offer x and compare $w < \underline{w} + \sigma$ with $\underline{w} + \sigma$.

Case $x > \underline{w} + \sigma$. Both lose, and the payoffs are $p(x)\varphi(w - \underline{w}) < p(x)\varphi(\sigma)$ because φ is strictly increasing on $[0, \sigma]$ when $\sigma \leq \Lambda/2$ (Lemma A.4).

Case $x < w$. Both win, so it suffices that Π_{first} is increasing on $[\underline{w}, \underline{w} + \sigma]$. Writing $w = \underline{w} + \sigma p$,

$$\frac{d\Pi_{\text{first}}}{dp} = (2\Lambda + \Gamma) - 2(\Lambda + 2\sigma)p + 3\sigma p^2, \quad (16)$$

a convex quadratic, positive at $p = 0$ and equal to $\Gamma - \sigma > 0$ at $p = 1$, where the inequality uses $\sigma \leq \Lambda/2 \leq \Gamma/2$ under the normalization of Section 2.1. Its minimum is at $p = (\Lambda + 2\sigma)/(3\sigma)$, which exceeds one whenever $\Lambda > \sigma$, implied by $\sigma \leq \Lambda/2 < \Lambda$. The minimum therefore lies to the right of the interval and the derivative is positive throughout.

Case $w < x < \underline{w} + \sigma$. The higher offer wins, whereas w loses. Its payoff exceeds $\varphi(\sigma)$ by Γ , and the lower offer earns at most $\varphi(\sigma)$. If $x = w$, the old tie payoff is below the first-picker payoff at w , which increases strictly up to $\underline{w} + \sigma$. If $x = \underline{w} + \sigma$, the new tie earns $\varphi(\sigma) + \Gamma/2$, again strictly more than the old losing offer. The comparison uses the feasibility condition $b \geq \underline{w} + \sigma$.

In every case the inequality is strict, so the two games have the same equilibria under $\underline{w} \mapsto \underline{w} + \sigma$. $\square$

Proof of Proposition 2.5. If $b = \underline{w}$, the feasible interval has one action and the conclusion is immediate. Otherwise, write $u = w - \underline{w}$ and $p = u/\sigma$ over the scarce range. Losing to a fixed higher offer x yields $p(x)\varphi(u)$, which increases strictly with $u \leq m$. Winning yields

$$\Pi_{\text{first}}(\underline{w} + u) = (2 - p)\varphi(u) + \Gamma p, \quad \frac{d\Pi_{\text{first}}}{du} = (2 - p)\varphi'(u) + \frac{\Gamma - \varphi(u)}{\sigma} > 0.$$

Here $\varphi'(u) \geq 0$ and $\varphi(u) < \Lambda \leq \Gamma$. At a change of ordering, first pick exceeds second pick at the same offer:

$$\Pi_{\text{first}}(w) - p(w)\varphi(u) = 2(1 - p)\varphi(u) + \Gamma p > 0 \quad \text{for } w > \underline{w}.$$

The tie is their average. Thus increasing one's offer improves payoff within either ordering and at the upward crossing of a rival's offer. Against any feasible rival strategy, the cap strictly dominates every lower offer. This establishes the full-budget equilibrium and its uniqueness in the symmetric class. At (b, b) , averaging first and second pick gives

$$\Pi^{\text{tie}}(b) = p(b)[\Lambda + \Gamma/2 - \Omega(b)]. \quad (17)$$

$\square$

Proof of Proposition 2.6. At a tie the coin decides. As first picker the party takes q_1 with probability p and q_2 with probability $(1 - p)p$; as second picker it takes q_2 with probability p^2 and nothing otherwise. Averaging, $\mathbb{P}(q_1) = \frac{1}{2}p$ and $\mathbb{P}(q_2) = \frac{1}{2}(1 - p)p + \frac{1}{2}p^2 = \frac{p}{2}$, so the two are equal, and $\mathbb{P}(\text{nobody}) = \frac{1}{2}(1 - p)^2 + \frac{1}{2}(1 - p^2) = 1 - p$. The three sum to one. Expected quality is $\frac{q_1 + q_2}{2}p$ and the probability of recruiting is p , which gives (17). $\square$

The Funding-Ban Comparison

Proof of Proposition 2.7. After the ban both parties are symmetric at b_2 , with expected quality $\frac{q_1+q_2}{2}p_2$ and recruitment probability p_2 by Proposition 2.6. Subtract the pre-ban values.

Party 1. $\Delta_1^{\text{quant}} = p_2 - p_1(2 - p_1) < 0$ because $p_1(2 - p_1) \geq p_1 > p_2$. For quality, $q_1p_1 + q_2p_1(1 - p_1) = p_1[q_1 + q_2(1 - p_1)] \geq p_1\frac{q_1+q_2}{2}$, using $q_1 \geq \frac{q_1+q_2}{2}$ and $q_2(1 - p_1) \geq 0$, so $\Delta_1^{\text{qual}} \leq \frac{q_1+q_2}{2}(p_2 - p_1) < 0$.

Party 2. $\Delta_2^{\text{quant}} = p_2 - p_1p_2 = p_2(1 - p_1) \geq 0$, strictly if $p_1 < 1$, and $\Delta_2^{\text{qual}} = p_2[\frac{q_1+q_2}{2} - q_2p_1]$ with $\frac{q_1+q_2}{2} - q_2p_1 \geq \frac{q_1-q_2}{2} > 0$. $\square$

Remark A.2 (Scope of the comparison). The allocation calculation uses $p_1 > p_2 > 0$ and a post-cut tie with reach p_2 . It does not establish the pre- or post-cut equilibrium outside the stated budget conditions. In particular, when a rival has enough budget to outbid a proposed first-picker offer, separate one-player optima need not be mutual best responses. No general sign is asserted in that region.

The first-picker objective and its constrained optimum. Writing $w = \underline{w} + \sigma p$, payment occurs only upon recruitment, so

$$H(p) \equiv \Pi_{\text{first}}(\underline{w} + \sigma p) = (2\Lambda + \Gamma)p - (\Lambda + 2\sigma)p^2 + \sigma p^3.$$

For a budget Ω_1 , set $c = \min\{\Omega_1/\sigma, 1\}$ and $D = (\Lambda + 2\sigma)^2 - 3\sigma(2\Lambda + \Gamma)$. If $D \geq 0$, define $p_- = (\Lambda + 2\sigma - \sqrt{D})/(3\sigma)$. The constrained maximizer is obtained by comparing $H(c)$ with $H(p_-)$ when $0 < p_- < c$; if that candidate is absent, choose c . This follows because $H'(0) > 0$, the smaller stationary root is the only possible interior local maximum, and offers beyond full reach only add cost. A global target cannot in general be truncated at an arbitrary cap: H can fall after a local maximum and then rise again before $p = 1$.

For example, $Q = 1$, $\underline{w} = 0.9$, $\sigma = 1$ gives $H(p) = 1.2p - 2.1p^2 + p^3$: $H(0.4) = 0.208 > H(1) = 0.100$. The condition $\sigma \leq \Gamma$ is therefore not sufficient for full reach to be optimal: the participation floor also matters, through Λ .

Proof of the restricted saturation result, equation (8). Assume $\sigma > \Gamma$ and $0 < \Omega_2 < m$, with $\Omega_1 > \Omega_2$. Then $H'(0) > 0$ and $H'(1) = \Gamma - \sigma < 0$. Its smaller root p_- lies in $(0, 1)$ and its larger root exceeds one. Thus H is unimodal on $[0, 1]$, with target $\tau_1 = \sigma p_-$, and the constrained optimum is $u_1 = \min\{\Omega_1, \tau_1\}$.

The ordering is a best-response restriction, not an assumption about budgets alone. On $0 \leq u \leq m$, write $\Pi_{\text{first}}(\underline{w} + u) = (2 - p)\varphi(u) + \Gamma p$. Its derivative is $(2 - p)\varphi'(u) + (\Gamma - \varphi(u))/\sigma > 0$, using $\Gamma \geq \Lambda > \varphi(u)$ and $\varphi'(u) \geq 0$. Hence $\tau_1 \geq m > \Omega_2$ and $u_1 > \Omega_2$. Party 2 cannot outbid and maximizes its second-picker payoff at its cap. Party 1's chosen first-picker payoff exceeds any losing deviation (whose payoff is at most $p_2\varphi(\Omega_2)$) and the tie at b_2 . The proposed offers are therefore mutual best responses.

Reach is $p_1 = u_1/\sigma$, proving (8) away from the kink. Applying the chain rule to $p_1(2 - p_1)$ and $q_1p_1 + q_2p_1(1 - p_1)$ gives marginal effects $2(1 - p_1)/\sigma$ and $[q_1 + q_2(1 - 2p_1)]/\sigma$ below the target, and zero above it. The quality measure here is q , as in the main text; multiplying by Q gives electoral value to the party. $\square$

B Model Extensions

This appendix extends the baseline to differences in organizational strength and budgets, multiple slots and alternative participation costs. Sections B.8 and B.9 show how the unexposed rival's response can change when there are more parties or more positions to fill.

Throughout, $\Gamma_j \equiv Q_j(q_1 - q_2)$ is party j 's rank premium and $\Lambda_j \equiv Q_j q_2 - \underline{w}$ its low-ball value.

B.1 Asymmetric Quality

Let $Q_1 > Q_2$ with equal and slack budgets. Each party's density is pinned by its opponent's rank premium: $f_1^* = 1/\Gamma_2$ and $f_2^* = 1/\Gamma_1$, both on $(\underline{w}, \underline{w} + \Gamma_2]$. The weaker party bounds the support, runs a mass deficit $\alpha = 1 - \Gamma_2/\Gamma_1$ and places it at the floor; payoffs are $\Pi_1^* = Q_1 q_1 - \underline{w} - \Gamma_2$ and $\Pi_2^* = \Lambda_2$.

Argument. Constancy of Π_j on the support requires $\Gamma_j f_{-j}^* = 1$, so each party randomizes at the rate that makes its rival indifferent, and the stronger party's density is higher when $\Gamma_1 > \Gamma_2$. The supports must share an upper endpoint, since otherwise the party with the higher endpoint could lower its offer without losing. Party 1's continuous mass over $[\underline{w}, \bar{w}]$ is $\Omega(\bar{w})/\Gamma_2 \leq 1$, which binds first, giving $\bar{w} = \underline{w} + \Gamma_2$. Party 1's distribution then reaches one exactly and has no mass point, while party 2's continuous part carries only $\Gamma_2/\Gamma_1 < 1$. An interior atom can be overbid, but an atom at the lower boundary is compatible with the asymmetric profile, so the deficit sits at the floor. Party 1 places no mass at $\underline{w}$ itself, since tying with party 2's mass point is dominated by posting $\underline{w} + \varepsilon$.

This is the result of Bulow and Levin (2006): each density is pinned by the rival's valuation, offers are compressed, and the strongest party gains most. Matching is assortative in equilibrium without being assumed. The payoff gap comes from a difference in valuations, not from a better sorting technology: party 1 values first pick more and is willing to spend more to obtain it.

B.2 Asymmetric Hard Caps

With unequal hard caps, a natural candidate profile places the richer party's atom just above the poorer party's cap b_2 . This profile is not an equilibrium of the continuous-action game with fair ties: any offer $b_2 + \varepsilon$ can be replaced by $b_2 + \varepsilon/2$, which still wins and costs less. Moving the atom to b_2 does not help, because the constrained party then gains by deviating to its cap and sharing the atom.

Remark B.1 (Scope). I do not characterize the equilibrium distribution, payoffs or first-pick probabilities with asymmetric hard caps, and I do not claim that the constrained party cannot win. The argument above rules out the candidate profile; it does not show that no equilibrium exists. A discrete offer grid or a different tie rule would change the game and requires a separate analysis. With equal budgets, the payoff is $\Lambda + \Gamma/2 - \Omega(b)$ in the tight regime and Λ above it.

B.3 Quality and Budget Together

Combining differences in strength and in budgets would require the asymmetric hard-cap equilibrium, which is not characterized here. Organizational strength changes the value of first pick, and resources restrict feasible offers, but these observations do not imply that a rich weak party always defeats a poor strong one. The stochastic comparison with scarce budgets in the main text is proved separately and does not rely on the deterministic construction.

B.4 Multiple Slots

Let each party post one impersonal per-member offer and recruit k members for a total outlay of $k\underline{w}$, the higher offer taking the top- k block. The quality-sum gap between blocks is exactly k^2 : position by position the difference is k , and there are k positions. Raising the offer by

dw costs $k dw$ and buys $f^*(w) dw$ extra probability of taking a block worth Γk^2 , so indifference gives

$$f^* = \frac{1}{k\Gamma} \text{ on } [\underline{w}, \underline{w} + k\Gamma]. \quad (18)$$

Remark B.2 (The budget tightens quadratically). Lengthening the list has two effects that compound. It divides the per-slot budget $b = B/k$ by k , and it widens the support the party must cover to $k\Gamma$, also by a factor of k . The regime is determined by the ratio $\Omega(b)/(k\Gamma) = (B/k - \underline{w})/(k\Gamma)$, whose leading term falls at rate k^2 . This statement concerns the deterministic comparison with a fixed pool. It does not establish exclusion or assign offices to stochastic recruitment regimes, which also depend on the size of the pool, participation and capacity.

B.5 Capacity as the Shocked Argument

Holding the total budget B fixed, a reduction of capacity from k to $k' < k$ raises the per-slot ceiling from B/k to B/k' . A binding limit on list length also limits the number of candidates a party can nominate. These are accounting statements. If a realized ordered pool is held fixed and its lowest-ranked members are removed, its mean quality rises. Once parties can change offers and the allocation changes, neither that composition result nor higher spending per candidate follows from the ceiling alone. Re-entry at rival parties requires acceptable offers and spare capacity, and seat shares require an electoral aggregation rule that the model does not specify. I do not derive these outcomes as comparative statics.

B.6 The Participation Floor as a Policy Instrument

The participation floor may depend in part on statutory eligibility requirements. In reduced form, write $\underline{w} = \kappa T$, where T is the minimum period of party membership required before standing for office and κ converts waiting time into forgone value. Lowering T admits citizens who were previously screened out, and since $\Omega(b) = b - \underline{w}$ it raises offer slack for every party at unchanged budgets, moving the whole system up the ladder of Proposition 2.2.

Brazil reduced the minimum period of party membership from twelve months to six in 2015. The change helps interpret the timing of recruitment around the eligibility deadline. Because it applied nationally, cycle fixed effects absorb its common effect; the empirical design does not identify its effect separately (Appendix C).

B.7 The Exact Frontier and the Support Floor

Assumption 2.4 is sufficient but conservative. The full-budget profile (b, b) is a Nash equilibrium if and only if $\Lambda + \Gamma/2 - \Omega(b) \geq \varphi(m)$, equivalently $\Omega(b) \leq \bar{\Omega}(\sigma)$ with

$$\bar{\Omega}(\sigma) = \sigma + \Gamma/2 \text{ if } \sigma \leq \Lambda/2, \quad \bar{\Omega}(\sigma) = \Lambda + \Gamma/2 - \Lambda^2/(4\sigma) \text{ if } \sigma \geq \Lambda/2. \quad (19)$$

$\bar{\Omega}$ is continuous and strictly increasing, with $\bar{\Omega}(0^+) = \Gamma/2$ and $\bar{\Omega} \uparrow \Lambda + \Gamma/2$. The common factor $p(b)$ cancels in the comparison, so the frontier depends only on where the low-ball optimum lies, not on the reach achieved at the ceiling.

For $\sigma \leq \Lambda/2$ the condition reads $b - (\underline{w} + \sigma) \leq \Gamma/2$: the pure equilibrium survives as long as the budget slack above the effective floor is at most half the rank premium. A deviation becomes attractive when its cost saving exceeds the expected gain from keeping priority in the matching order.

The support floor. In any symmetric mixed equilibrium, the lowest offer in the support is $\underline{w} + m$. In the deterministic model parties sometimes post the bare reservation wage; here such an offer reaches almost nobody and is never made in the mixed regime. This restriction does not apply to pure profiles with scarce budgets below the mixed support floor.

Two forces break the pure equilibrium, and small σ is not one of them. The first is saturation of reach: once $\Omega(b) > \sigma$, the offer already reaches the whole field and additional money buys only position, so undercutting pays as soon as the money saved exceeds $\Gamma/2$. The second is exhaustion of the margin: beyond $u = \Lambda/2$, the reach gained no longer covers the money spent. The first applies when σ is small relative to the budget, the second when σ is large relative to the surplus. The pure region lies between them and grows with σ : dispersion discourages undercutting, because a low offer reaches too few citizens to be worth making.

B.8 Three Parties: The Conditional Quality Spillover

This section derives the first of the two conditions in Prediction 3. I work in the large-field limit of Section B.10, where the expected quality of the best acceptor of a field reached with probability p is $g(p) = n + 1 - 1/p$ and successive order statistics are spaced by $1/p$. Take one party with corporate money (budget $\tilde{b}$) and two small parties (budget b each), all posting their full budgets. Write $p_b = p(b)$, $\delta \equiv 1/p_b$, $a \equiv g(p_b)$, and

$$\theta \equiv \mathbb{P}(M(\tilde{b}) \in A(b)) = \frac{p_b}{p_{\tilde{b}}} \in (0, 1), \quad (20)$$

the probability that the rich party's recruit is someone the small parties could also have afforded.

Before the ban a small party expects $a - \frac{\delta}{2}(1 + 2\theta)$; after it, with three-way symmetry, $a - \delta$. Hence

$$\Delta\bar{q}_2 = \frac{1}{p_b} \left(\theta - \frac{1}{2} \right), \quad \text{so} \quad \Delta\bar{q}_2 > 0 \iff p_{\tilde{b}} < 2p_b \iff \tilde{b} < 2b - \underline{w}. \quad (21)$$

The exposed party's own quality change is unambiguous, $\Delta\bar{q}_1 = 1/p_{\tilde{b}} - 2/p_b < 0$.

Remark B.3 (Why the rival's gain is the fragile prediction). Before the ban, the number of parties effectively competing for the top of the accessible field is $2 + \theta$; after it, the number is exactly three. When $\theta > \frac{1}{2}$, the rich party often took the best citizens that the small parties could reach, and equalizing budgets helps them. When $\theta < \frac{1}{2}$, the rich party recruited mostly in a segment that only its money could reach, and the ban adds it to everyone else's pool as one more competitor, which crowds the small parties and lowers their quality. The threshold is $\tilde{b} = 2b - \underline{w}$.

The deterministic limit gives $\Delta\bar{q}_2 > 0$ in all cases, but only because $\theta \rightarrow 1$ as $\sigma \rightarrow 0$. With dispersion in participation costs and a large budget advantage, the typical situation of the largest pre-ban recipients of corporate money, the same offers give the opposite comparison. These are allocation calculations in a large field, conditional on offers; they do not establish a three-party equilibrium or a general quantity response.

B.9 Multiple Slots with Asymmetric Budgets

This section derives the second condition in Prediction 3. Two parties have k slots each and unequal offers with $p_1 > p_2$. In a finite field, the first party's expected number of recruits is $\Phi_1 = \mathbb{E}[\min\{k, X_1\}]$, where $X_1 \sim \text{Binomial}(n, p_1)$. Because quality and costs are independent, the expected size of the residual pool is

$$\mathbb{E}[\text{pool}_2] = np_2 - \frac{p_2}{p_1} \Phi_1. \quad (22)$$

Party 2's expected fill is the expectation of its residual pool capped at k , which in general differs from the cap of the expected pool. With equal offers its expected fill is $\frac{1}{2} \mathbb{E}[\min\{2k, X_2\}]$ for $X_2 \sim \text{Binomial}(n, p_2)$.

A deterministic-count approximation replaces random counts by their means: $\Phi_1 \approx \min\{k, np_1\}$, $\Phi_2 \approx \min\{k, p_2(n - \Phi_1/p_1)\}$, and $\bar{\Phi} \approx \min\{k, np_2/2\}$. Within this approximation, counts are unchanged in an abundant field; in the intermediate region $np_1 \geq k > np_2/2$, the rival gains only if $k > np_1/2$; and in the thin region $np_1 < k$, the richer party exhausts the approximate pool before the cut. These calculations are conditional on offers and do not establish a multi-slot equilibrium or exact finite-field cutoffs. A law-of-large-numbers interpretation requires a joint scaling of n , k and p_j that stays away from capacity boundaries; in particular, $n \rightarrow \infty$ with np_j fixed retains Poisson uncertainty in counts. For $n = 100$, $k = 1$ and $p_2 = 0.01$, the exact symmetric fill is 0.449, compared with 0.500 under the approximation.

B.10 The Large-Field Limit

For $M(w)$ the best acceptor of $A(w)$, $\mathbb{P}(M(w) \leq m) = (1 - p)^{n-m}$ and

$$\mathbb{E}[M(w)] = n - \frac{1-p}{p} + O((1-p)^n) \equiv g(p) + O((1-p)^n), \quad (23)$$

with $g'(p) = 1/p^2 > 0$ and $g''(p) = -2/p^3 < 0$: the expected best acceptor is increasing and concave in reach. Sweeping qualities from the top, the number of absent citizens before the first present one is geometric with mean $(1-p)/p$, which gives (23) directly and also $\mathbb{E}[M_{(r)}] = g(p) - (r-1)/p + o(1)$. The premium of first pick over second is therefore $1/p$ positions, the large-field counterpart of Γ .

A single party maximizing $\pi(w) = Qg(p(w)) - w$ has an interior optimum $w^* = \underline{w} + \sqrt{Q\sigma}$ when $Q \leq \sigma$, which equates the marginal gain in field quality to the unit marginal cost; the party then stops short of its budget. A second picker's payoff is $\pi(w) - Q/p(b)$: losing the contest shifts the payoff function down by a constant and does not change its maximizer, which is the large-field counterpart of the factorization used in Appendix A.

With J parties sharing a budget, the full-budget profile is an equilibrium if and only if $\pi(\hat{w}) - \pi(b) \leq Q(J-1)/(2p(b))$. The right-hand side increases with J : a tie is easier to sustain when there are more parties, because leaving a crowded tie costs more positions. The closed-form frontier for $J = 2$ has three branches; it is not needed for the results in this paper and is omitted.

Fill rates. With $k > 1$ and $\Phi(w) = \mathbb{E}[\min\{k, |A(w)|\}]$, Φ is nondecreasing in w and strictly increasing whenever the cap does not bind. When the field is abundant relative to the list, $\Phi \approx k$ and the number of recruits is insensitive to the offer; when the field is thin, $\Phi \approx np(w)$ and the number of recruits rises approximately with slope n/σ . These approximations require concentration away from capacity boundaries. A race with $k = 1$ is insensitive in quantity only if np is large, and a longer list is sensitive only if the accessible field is thin. The type of office alone does not determine which case applies; the identity $b = B/k$ motivates the comparison.

B.11 Correlated Participation Costs

Assumption 2.3 makes costs independent of quality. Correlation between the two can change both the acceptance set and the share of the richer party's recruits that the rival could afford. The sign of the covariance alone does not determine the conditional cost distributions needed to compute these changes. Signing the comparative statics under correlated costs would require a distributional restriction and a new equilibrium argument, and I leave it for future work.

B.12 Dropping the Budget Ordering

Section 2.4 lowers party 1's budget exactly to b_2 . For a partial cut with $b_1 > b'_1 > b_2$ and all offers in the scarce region, let $p'_1 = p(b'_1)$. The order of offers remains strict, and direct

subtraction gives the changes in recruitment,

$$\Delta_1^{\text{quant}} = (p'_1 - p_1)(2 - p_1 - p'_1) < 0, \quad \Delta_2^{\text{quant}} = p_2(p'_1 - p_1) < 0,$$

and the changes in expected quality,

$$\Delta_1^{\text{qual}} = (p'_1 - p_1)[q_1 + q_2(1 - p_1 - p'_1)] < 0, \quad \Delta_2^{\text{qual}} = q_2 p_2(p'_1 - p_1) < 0.$$

The total changes by $(p'_1 - p_1)(2 + p_2 - p_1 - p'_1) < 0$. The rival loses because it recruits the second citizen only when the richer party recruits the better one; its gain under exact equalization comes from choosing first in half of the ties. The exposed party's fill is quadratic in reach, so a partial cut is not an affine rescaling of the equalization experiment.

For example, $Q = 1$, $\underline{w} = 0$, $\sigma = 1$, $b_2 = 0.2$ and $b_1 : 0.4 \rightarrow 0.3$ satisfy scarcity. The exposed party's fill falls from 0.640 to 0.510, and the rival's fill falls from 0.080 to 0.060.

Remark B.4 (Levels, ordering and boundaries). Both reach levels enter the responses even within a regime, so a budget gap alone does not determine them. At exact equality the tie allocation applies; if the ordering reverses, the first- and second-picker allocations must be recomputed. Crossing the thresholds of the symmetric deterministic game does not establish a global comparative static for asymmetric budgets, and I do not claim a general sign for the spillover across these cases.

B.13 The Role of the Scarcity Conditions

Assumption 2.4 contains two restrictions, $\Omega(b) \leq \Lambda/2$ and $\Omega(b) \leq \sigma$. The second is needed for the linear expression $p(w) = (w - \underline{w})/\sigma$ to apply throughout $[\underline{w}, b]$, and it does not follow from the first.

The distinction matters when budgets differ. If $\sigma < \Lambda/2$ and $\Omega(b_2) \in (\sigma, \Lambda/2]$, party 2 already reaches every citizen at $\underline{w} + \sigma$. At a full-budget profile with a higher-offering rival, lowering its offer to that level saves resources without reducing reach.

Example. Set $Q = 1$, $q_1 = 2$, $q_2 = 1$, $\underline{w} = 0$, so $\Lambda = \Gamma = 1$; let $\sigma = 0.3$ and $(\Omega(b_1), \Omega(b_2)) = (0.45, 0.40)$, both satisfying $\Omega(b_j) \leq \Lambda/2$. At the full-budget profile party 2 earns $p_1 \varphi(0.40) = 0.600$; deviating to $w = 0.30$ it earns $\varphi(0.30) = 0.700$. The deviation gains 0.100, so this full-budget profile is not an equilibrium. The restriction $\Omega(b) \leq \sigma$ excludes the example.

C Data Construction

C.1 Registries

The analysis combines four registries of the Superior Electoral Court (TSE). The party membership registry (*Filiação partidária*) records individual membership histories. Its historical snapshot contains 24,666,625 source rows, which yield 24,448,824 cleaned membership spells, and covers start dates through April 30, 2021. Campaign and party accounts (*Prestação de contas*) cover the 2004–2024 elections and supply the financial measures. Election results and candidate records provide candidacies, votes, seats and candidate attributes. Electorate records supply the number of registered voters, the denominator of every per-voter variable.

The same party groupings are applied to the membership, financial and electoral records. The groupings follow the historical registry as of 2021: they combine renamed parties and the predecessors that the registry no longer separates after an incorporation or merger. For example, PPL and PCdoB are combined, as are PHS and PTN/Podemos, and PEN/Patriota and PRP. I call these groups party lineages in the technical descriptions below. A lineage can pool distinct 2012 organizations, and mergers after 2021 are not incorporated. Financial

and electoral controls are aggregated to lineages before matching, and a movement within a lineage is not counted as a transition between parties. Electoral summaries by the party labels in use at each election are kept separate from these analysis units. The 2013 wave of the Power–Zucco party ideology scores (Zucco and Power, 2024) classifies the origins of recruits and the distribution of funding across the ideological spectrum.

Membership and candidacy. In Brazil, recruitment is closely tied to candidacy. Independent candidacies are not permitted, and a citizen must belong to a party for a statutory minimum period before standing for office. A 2015 reform reduced that period from twelve to six months, which can shift the timing of recruitment; the annual and two-year comparisons in Appendix D do not separate this change from the funding reform. About fifteen million people were party members on election day in 2012. Separately, 2.98 million distinct people have a recorded membership start during the 2012 electoral cycle, some of whom may have left before election day. Of these, 9.3 percent stood for mayor or councilor that year, and they accounted for 64.1 percent of candidates for those offices. Party members were 10.7 percent of the electorate in 2012 (Speck, 2013), compared with an average of about 5 percent across twenty Western European democracies in the late 1990s (Mair and van Biezen, 2001). New memberships cluster before the deadline for candidacy eligibility and are especially numerous before municipal elections. Speck (2013) relates this pattern to the number of positions at stake: municipal elections field several hundred thousand candidates, compared with about twenty thousand in national elections. Local parties therefore recruit most of the people who enter electoral politics.

C.2 Reconstructing the Membership Stock

The registry records spells, so the stock of members must be accumulated from entries and netted of terminations. Terminations are recorded imperfectly, because periodic administrative purges are not always recorded as dated events. The reconstruction therefore produces two scenarios: a lower value that treats ambiguous cases as terminated and an upper value that treats them as continuing. The two values are not guaranteed bounds on the true stock.

Table 10: Reconstructed membership interval against TSE’s published totals

Year	TSE official (mn)	Lower (mn)	Upper (mn)	Inside
2010	13.779	13.331	14.429	Yes
2012	15.075	14.632	15.785	Yes
2014	14.728	14.826	15.989	No
2016	16.527	16.447	17.643	Yes
2018	15.965	16.786	17.986	No
2020	16.511	16.347	17.557	Yes

Frozen-2021 lineages; distinct voters within each lineage and municipality. The two reconstructions differ only in handling undated terminations. Unrecorded administrative purges can put official totals outside the interval. Four of the six published totals lie inside it.

The interval contains the TSE’s published total in four of six years, including both years used in the design. It exceeds the official total in 2014 and 2018, when published membership declined; unrecorded terminations are one possible reason. Recruitment flows use start dates and do not depend on assumptions about later terminations.

C.3 Campaign Resources

The treatment is the corporate receipts reported in the local party’s 2012 campaign accounts, divided by registered voters in that year. Candidate receipts are a separate measure and are not added to party receipts, because transfers from parties to candidates would otherwise be




Figure 6: Gross committee-plus-directory receipts by organizational level, 2004–2024, in nominal BRL millions. Electoral committees enter the municipal level; accounts of party directories, including state and national organizations, are available from 2012. Percentages from 2012 onward show the municipal share of these gross receipts. Transfers between accounts are not consolidated.




Figure 7: Monthly recruitment flows and candidacy among cycle recruits, 1998–2024. Bars show recorded recruitment flows (left axis). Points show the share of distinct people with a recorded membership entry between consecutive municipal elections who stood for mayor or councilor at the closing election (right axis). Dashed lines mark municipal elections; dotted lines mark statutory membership deadlines. The shaded period after April 2021 uses the modern registry, which omits canceled memberships; the 2024 share therefore has a survivor-selected denominator and is not directly comparable with earlier cycles. Recruitment estimates use complete historical cycles through 2020.

counted twice. The main sample requires positive observed total receipts in 2012 with complete nominal amounts, after excluding quarantined records; accounts of other recipients support the descriptive comparisons.

The comparisons based on realized money express receipts in December 2016 BRL and require observed, complete amounts in both years. The initial measure combines committees and directories in 2012 but uses only directories in 2016, because committees were abolished. A second measure uses municipal directories in both years, and I re-estimate the initial measure on exactly that overlap. Neither measure consolidates transfers or measures the model's net recruitment budget. The denominator of the symmetric growth measure is the sum of receipts in both years, so growth is undefined when both values are zero.

C.3.1 Local Party Resources

The financial measure is the campaign resources recorded by the local party organization; the model's usable recruitment budget is not observed directly. Candidate receipts are analyzed separately because they include transfers from parties. Figure 8 shows total receipts and their corporate component for each type of recipient. Recorded corporate contributions fall to zero in 2016.




Figure 8: Total funding of candidates and of the local party, with the corporate component, 2004–2024, in millions of December 2016 BRL. Receipts are deflated by the month of the donation using the IPCA. Candidate and party receipts are shown separately. Corporate contributions are prohibited from 2016 onward.

Figure 6 reports gross committee-plus-directory receipts by organizational level, assigning electoral committees to the municipal level. These totals describe where receipts were reported, not how the money was finally used in each municipality.

C.3.2 The First Stage

Corporate receipts disappear from the records after the ban by construction. The association between prior exposure and the change in total recorded local-party funding is -1.048 (0.021) per real of prior corporate receipts per voter, and the corresponding coefficient for candidate funding is -0.552 (0.091). These are changes in gross nominal receipts, not in consolidated real budgets. The abolition of committees changes the set of accounts, and internal transfers can be recorded on both sides. The estimates therefore do not show a one-for-one loss of net resources and do not rule out substitution from other sources.

Table 11: Recorded funding changes and prior corporate exposure

Outcome	Estimate	SE	N
Corporate money to the local party, per voter	−1.000	0.000	14,606
Gross local party receipts, nominal BRL per voter	−1.048	0.021	14,606
Campaign funding of the party’s candidates, per voter	−0.552	0.091	14,606

Dose coding; first differences in nominal receipts per 2012 voter. Corporate receipts disappear mechanically. Gross committee-plus-directory receipts are not consolidated net resources: internal transfers and committee abolition limit comparison. Candidate receipts are separate.

Table 12: The estimation universe by election year

Year	Municipalities	Parties	Cells	Candidates	Mayoral	Council	Parties per municipality
2004	5,561	27	57,236	354,485	15,790	338,695	10.3
2008	5,562	27	61,409	336,516	15,016	321,500	11.0
2012	5,562	30	70,967	435,562	15,417	420,145	12.8
2016	5,561	35	76,669	449,284	16,018	433,266	13.8
2020	5,562	34	41,603	505,679	18,352	487,327	7.5
2024	5,561	29	42,805	428,365	15,070	413,295	7.7

All candidates in the balanced 5,562-municipality panel. A cell is a party lineage by municipality. Party counts use the rebrand-harmonized lineage, so a party that renamed is not counted twice.

Table 13: Sample construction: from contesting cells to the estimation sample

Step	Cells	Percent of universe
Party-lineage x municipality cells contesting in 2012	69,400	100.0
less: no observed campaign receipts (treatment undefined)	3,795	5.5
Estimation sample (treatment defined)	65,605	94.5
of which: corporate receipts > 0 (treated)	17,159	24.7
of which: corporate receipts = 0 (untreated)	48,446	69.8

Councilor scope, 2012. Cells with no observed campaign receipts are excluded rather than coded untreated: a party that raised nothing is not a clean zero-corporate control.

C.4 Sample Flow and the Universe

These tables distinguish the universe of parties contesting each election and the candidate-finance sample from the party-account sample used for the main estimates. The units differ, and sample sizes or counts of local competitors from one should not be compared with the other.

In the directed network, $Poaching_{Godmt}$ is the number of membership starts during cycle t at destination party d whose immediately preceding recorded membership names a different origin party o in the same recorded municipality m . Parties are the lineages described above. The start date of the new membership must be strictly later than that of the previous one. Financial support requires positive observed total receipts in 2012 and complete nominal amounts at both endpoints, after excluding quarantined records. Supported parties with zero corporate receipts are distinct from missing accounts.

The network starts from 15,906 supported party–municipality cells in 4,347 municipalities and requires support at both endpoints. This yields 89,270 directed pairs in 3,048 municipalities, which are kept even when their flow is zero. The five-cycle panel contains 446,350 rows and the 2008–2020 annual panel 1,160,510. The 2012 and 2016 cycles contain 112,023 and 102,637 cross-party flows in this network, together 11.2 percent of the 1,917,329 transitions between parties recorded in the same municipality. The network is a selected set of financially observed parties, not a census of movements between parties. In the party-account universe, 10,950

cells have zero corporate receipts and 4,956 have positive receipts. Mean positive exposure is R\$1.622 per voter, its median is R\$0.404, and the standard deviation of exposure in the full universe is 2.361. There are 371 supported cells with no recorded recruitment in 2012 and 1,853 in 2016. The party-level comparison retains 14,606 cells in 3,048 municipalities after controls and the removal of singletons.

Each person's complete membership history is ordered before any period or network restriction is applied, so intermediate memberships outside the network are not skipped. Counts can include several entries by the same person. A previous membership record does not show that the person was an active member when the new membership began, or that the new party made an offer, and a matching recorded municipality does not verify residence at the date of the event. Recruits without a recorded previous party are not necessarily entering politics for the first time.

Incoming, outgoing and net recorded counts for party d are

$$I_{dmt} = \sum_{o \neq d} \text{Poaching}_{odmt}, \quad O_{dmt} = \sum_{q \neq d} \text{Poaching}_{dqmt}, \quad N_{dmt} = I_{dmt} - O_{dmt}.$$

The sums run over eligible rivals in the fixed local network. Outgoing counts are dated by the later membership start rather than by a recorded termination. Net counts sum to zero within the network, whereas total recruitment and membership stocks need not balance. The cycles have the same boundaries as in the party-level comparison. Rates divide these counts by $\text{Voters}_{m,2012}$, the number of registered voters in municipality m in 2012, and multiply by 100,000.

The PT-PSDB comparison uses the same 591 municipalities for both directions and for each party's recruitment from all eligible rivals. The PSDB-favoring bilateral net rate is

$$r_{\text{PT,PSDB},m,t} - r_{\text{PSDB,PT},m,t},$$

in recruits per 100,000 registered voters in 2012. The change in this rate is related to PT-minus-PSDB corporate exposure or to PSDB-minus-PT receipt growth. The receipt-growth comparison retains the 195 municipalities with accounts in both years. With one observation per municipality, these cross-sections cannot include municipality effects. The full party ledger instead uses each party's own fixed set of eligible municipalities, so its totals differ from those of the common-municipality comparison.

The network uses the same party lineages as the party-level analysis, so the early cycles condition on financial support observed in 2012 and can include parties with different histories. A narrower comparison excludes PSD, REP, PSOL and the mixed PCDOB, PODE and PATRIOTA/PRP lineages, leaving 57,898 directed pairs in 2,809 municipalities. This changes the population compared rather than recovering a uniquely correct historical network.

C.4.1 Selection into Exposure

Corporate funding is associated with geography and with the earlier organizational and electoral record of the party (Figure 9). These patterns motivate the municipality fixed effects and the 2008 controls, but they do not rule out selection on unobserved trends. The binary comparison within a municipality requires funded and unfunded parties in the same municipality, although municipalities without that variation still contribute to other parameters. The continuous comparison also uses differences in intensity among funded parties. Funding need not be a persistent characteristic of a party: resources and recruitment conditions specific to 2012 remain a potential confounder.




Figure 9: Predictors of 2012 corporate receipts per voter among supported local party organizations. Continuous regressors are standardized over the sample; intervals cluster by municipality. These are selection diagnostics.

Table 14: Coverage of candidate characteristics by election year (percent non-missing)

Year	Sex	Race	Education	Occupation	Civil status	Age	Wealth	Prior-start age
2004	100.0	0.0	98.7	96.5	98.7	99.9	0.0	85.3
2008	99.6	0.0	99.6	99.6	99.6	99.6	97.1	95.3
2012	99.9	0.0	99.9	99.9	99.9	99.9	62.1	98.5
2016	100.0	100.0	100.0	100.0	100.0	100.0	64.6	99.3
2020	100.0	98.9	100.0	100.0	100.0	100.0	61.3	99.6
2024	100.0	99.4	100.0	100.0	100.0	100.0	63.9	69.9

Candidate race is zero before 2016 because TSE did not collect it: the field is a literal placeholder in earlier years and is set to missing here. Declared wealth coverage falls from 97 percent in 2008 to about 62 percent thereafter, a change in the source file rather than in candidate behavior. Months since the latest recorded prior membership start are not verified active-party tenure. The 2024 values use a history truncated in April 2021 and can miss later recruitment records.

C.5 Coverage of Candidate Attributes

The availability of candidate attributes differs by year and variable. Race is not recorded in the candidate files before 2016, and wealth coverage is incomplete. The table reports coverage, and each regression reports its own sample. The historical membership registry does not cover recruitment in the 2024 cycle.

C.6 Corporate Funding and Ideology

Corporate funding is distributed across the ideological spectrum (Figure 10). This comparison provides context for the results on the origins of recruits; it does not rule out differences in exposure within ideological groups or other changes associated with party ideology.

The competition analysis holds the 2013 scores fixed: Left is below -0.3 , Center is between -0.3 and 0.3 , and Right is above 0.3 . PT scores -0.532 and PSDB 0.219 , which places PSDB in the Center. Lineages that combine parties without a defensible common score and parties without scores remain unscored in the ledger and in the denominators of network shares. Both endpoints have scores for 66.9 percent of 2012 flows and 61.3 percent of 2016 flows. The continuous distance and extremity models use the same 46,760 scored pairs in 2,837 municipalities.




Figure 10: Corporate BRL per voter against party ideology, 2012. Each point is a party; area is the number of party-municipality observations. Ideology is the Power–Zucco score, 2013 wave; the fitted line is quadratic. Unsourced parties and merged 2021 lineages without a defensible common 2012 score are excluded from this plot and are not coded as centrist.

The municipal Center comparison requires financial support in the Left, the Center and the Right, which leaves 1,227 municipalities. Its four groups by prior corporate funding contain 524, 198, 224 and 281 municipalities with corporate receipts on neither side, the Left only, the Right only and both sides. The 21 municipalities with no network flows in either cycle are omitted only from the share changes. The broader registry outcomes count all recorded entries into the named bloc in these municipalities, including parties outside the financially supported set, and are therefore distinct from recruitment within the network.

D Additional Results

This appendix reports the functional-form and identification diagnostics for the main recruitment comparisons, detailed counts of recruits’ origins, ideological and municipal extensions of the network analysis, and downstream outcomes. Table 1 in Section 3 reports the benchmark for total recruitment and Table 4 in Section 4 its prior-party component. These party-level comparisons require financial records only at the recruiting party, whereas the network comparisons require them at both ends of each link; the results below keep the two samples separate.

D.1 Local Financial Shares and Recruitment Margins

Figure 1 uses local shares of funding and recruitment. The directed comparisons in Section 4 instead use pair-level recruitment per 100,000 registered voters in 2012. The definitions below record their different denominators, accounts and samples.

For the comparison of local shares, let $\mathcal{P}_m$ be the set of parties in the supported network with positive municipal-directory receipts in 2012 and complete directory accounts in both 2012 and 2016. The set is fixed across cycles and contains at least two parties in each municipality. Let B_{pmt}^D denote real gross directory receipts and H_{pmt} the total recruitment flow into party p , including recruits from outside $\mathcal{P}_m$. With the directed counts $\text{Poaching}_{odmt}^{\mathcal{P}}$ defined in Appendix C, the flows within the set are

$$I_{pmt}^{\mathcal{P}} = \sum_{o \in \mathcal{P}_m \setminus \{p\}} \text{Poaching}_{opmt}^{\mathcal{P}}, \quad O_{pmt}^{\mathcal{P}} = \sum_{q \in \mathcal{P}_m \setminus \{p\}} \text{Poaching}_{qpmt}^{\mathcal{P}}.$$

Write $K^{\text{total}} = H$, $K^{\text{in}} = I^{\mathcal{P}}$ and $K^{\text{out}} = O^{\mathcal{P}}$. The local shares are

$$s_{pmt}^B = \frac{B_{pmt}^D}{\sum_{j \in \mathcal{P}_m} B_{jmt}^D}, \quad s_{pmt}^k = \frac{K_{pmt}^k}{\sum_{j \in \mathcal{P}_m} K_{jmt}^k}, \quad k \in \{\text{total, in, out}\}. \quad (24)$$

The share of total recruitment uses total flows into the eligible parties, while the incoming and outgoing shares use the same pool of flows between those parties. The three coefficients therefore do not add up to a decomposition.

The loss of financial share and the changes in recruitment shares are measured in percentage points:

$$x_{pm} = 100(s_{pm,2012}^B - s_{pm,2016}^B), \quad y_{pm}^k = 100(s_{pm,2016}^k - s_{pm,2012}^k). \quad (25)$$

All three outcomes use the same 3,054 party–municipality cells in 719 municipalities. Individual zero outcomes are kept; a municipality is excluded if any required local denominator is zero in either cycle. Figure 1 reports ten common bins of financial loss and the unadjusted OLS relationship $y_{pm}^k = \alpha^k + \theta^k x_{pm} + u_{pm}^k$, and companion regressions add party and municipality fixed effects. Cells receive equal weight, standard errors are clustered by municipality, and the intervals around the bin means are pointwise.

These associations describe relative positions within the selected local parties. A loss of financial share need not mean that receipts fell, and a larger outgoing share need not mean more recorded departures or a higher probability that a member leaves. Observed receipts are endogenous gross amounts, so the comparison does not identify a causal effect of funding. In municipalities with two parties, $y_{pm}^{\text{in}} = -y_{pm}^{\text{out}}$ by construction.

Table 15 reports the unadjusted slopes behind Figure 1 and the corresponding estimates with party and municipality fixed effects.

Table 15: Financial-share losses and three recruitment margins

Margin	Raw (SE)	p	With FE (SE)	FE p
All recruitment	−0.680 (0.201)	<0.001	−0.674 (0.199)	<0.001
From rivals	−0.591 (0.323)	0.068	−0.631 (0.310)	0.042
To rivals	0.306 (0.310)	0.324	0.339 (0.296)	0.252

Associations with a 10-percentage-point loss of municipal funding share, 2012 to 2016. All outcomes are changes in local shares, in percentage points, using the same 3,054 party–municipality observations in 719 municipalities. Municipal-directory receipts are observed in both cycles. Adjusted models include party and municipality fixed effects. Standard errors cluster by municipality. Total recruitment includes all recorded recruitment flows at eligible parties; incoming and outgoing shares use the common pool of recorded transitions between those parties. The coefficients are not an additive decomposition. Incoming and outgoing shares mirror each other in two-party networks; Table 16 reports the comparison restricted to municipalities with at least three eligible parties. These are descriptive associations.

Table 16 restricts the sample to municipalities with at least three eligible parties, which removes the mechanical mirror between incoming and outgoing shares, and reports the original measure that mixes committee and directory accounts. In the restricted sample, the association with outgoing shares becomes larger and statistically significant, while the associations with total and incoming shares also grow. Each comparison keeps its own sample, and none replaces the main sample of Figure 1.

Relative resources along directed links. The directed comparison asks how a destination’s financial position relative to a particular origin relates to recruitment along that link. Figure 11 reports it using changes in the pair-level recruitment rate per 100,000 voters rather than local shares, with different accounts and samples from Figure 1.

For gross receipts $B_{pm,t}$ in December 2016 BRL, define

$$g_{pm} = \frac{2(B_{pm,2016} - B_{pm,2012})}{B_{pm,2016} + B_{pm,2012}}, \quad a_{odm} = g_{dm} - g_{om}. \quad (26)$$

Table 16: Account definitions and local competitive support

Accounts	Support	Margin	Change (SE)	p	N	Municipalities
Directories	3+ parties	All recruitment	−0.877 (0.208)	<0.001	2,688	536
Directories	3+ parties	From rivals	−1.006 (0.300)	<0.001	2,688	536
Directories	3+ parties	To rivals	0.615 (0.277)	0.027	2,688	536
Mixed	All	All recruitment	−0.738 (0.141)	<0.001	5,684	1,318
Mixed	All	From rivals	−0.602 (0.219)	0.006	5,684	1,318
Mixed	All	To rivals	0.171 (0.205)	0.405	5,684	1,318
Mixed	3+ parties	All recruitment	−0.828 (0.137)	<0.001	4,968	960
Mixed	3+ parties	From rivals	−0.701 (0.201)	<0.001	4,968	960
Mixed	3+ parties	To rivals	0.122 (0.173)	0.479	4,968	960

Unadjusted changes in local outcome share per 10-percentage-point loss in local funding share. All three outcomes use the same observations within each specification; standard errors cluster by municipality. Directories uses municipal-directory accounts in both years. Mixed combines committees and directories in 2012 but uses directories in 2016. The 3+ restriction excludes exactly two-party networks, where incoming and outgoing shares mirror each other. Account definitions and support change the samples; these are descriptive comparisons.

Positive a means that the destination’s resources grew more, or declined less, than the origin’s. The symmetric growth measure accommodates small baseline amounts without dividing by baseline receipts alone; it is not log growth or an elasticity. The first directed comparison relates the change in the pair-level rate, Δr_{odm} , to a_{odm} , controlling for the pair’s mean growth, directed national party-pair effects and municipality effects. Here Δr_{odm} is the 2016 rate minus the 2012 rate, both computed with the 2012 electorate.

The accounts must be observed with complete amounts in both years. The initial measure combines committees and municipal directories in 2012 but uses only directories in 2016, when the set of accounts changed. I also use directories in both years and re-estimate the initial measure on exactly that overlap, which separates the definition of the accounts from the selection of the sample. Gross receipts include transfers and in-kind resources, do not measure the model’s usable budget, and can respond to recruitment and local political conditions, so this comparison is descriptive.

For the exposure comparison in equation (14), D_{pm} is nominal 2012 corporate receipts per voter, as in equation (9). The equivalent relative and mean regressors are

$$A_{odm} = D_{om} - D_{dm}, \quad M_{odm} = (D_{om} + D_{dm})/2, \quad (27)$$

with coefficients $\gamma = (\beta_o - \beta_d)/2$ and $\delta = \beta_o + \beta_d$. Positive A means that the destination was less exposed than its origin. Holding mean exposure fixed avoids imposing equal and opposite effects at the two endpoints. All pairs receive equal weight and standard errors are clustered by municipality, so municipalities with more supported parties contribute more observations.

D.2 Functional Form and Continuous Treatment

The CGBS estimates discussed in Section 3.2 give an average level difference of -0.454 (0.102) and an average marginal response of -0.194 (0.019). They use 2012 as the base period and 2016 and 2020 as post periods, without the main controls or municipality-cycle effects. Table 17 compares them with five-cycle two-way fixed-effects estimates, which average three pre-reform cycles, and with a binary two-way fixed-effects estimate that uses the same base period.

The profile uses a cubic spline with no interior knots, and standard errors aggregate the estimator’s unit influence functions by municipality. A causal interpretation of the marginal responses requires strong parallel trends across exposure levels, beyond the parallel-trends condition for untreated outcomes that the level effects require.

The marginal response is sensitive to the specification. Adding two interior knots changes the average marginal estimate from -0.194 to -3.187 (0.348), while the average level estimate



Figure 11: Relative financial advantage and changes in recruitment, 2012–2016. The outcome is the change in recorded recruitment flows per 100,000 fixed 2012 voters. The left panel uses destination-minus-origin symmetric real gross-receipt growth, combining committees and directories in 2012 and directories in 2016. The right panel uses origin-minus-destination nominal 2012 corporate BRL per voter. Twenty bins summarize partial residuals after removing pair mean finances, national directed-pair effects and municipality effects; lines use all retained pairs. Samples are 29,730 pairs in 1,732 municipalities and 89,222 pairs in 3,048 municipalities, respectively. These are conditional associations.

Table 17: Continuous-treatment estimates and TWFE comparisons

Estimand	Estimate	SE	Municipalities
Average level difference	−0.454	0.102	4,347
Average marginal response	−0.194	0.019	4,347
TWFE dose, five cycles	−0.093	0.023	4,347
TWFE binary, five cycles	−0.349	0.080	4,347
TWFE dose, positive only	−0.079	0.023	2,564
TWFE binary, 2012 base	−0.454	0.102	4,347

The continuous-dose estimator uses 2012 as its base and 2016 and 2020 as post periods, with no covariates. Five-cycle TWFE averages 2004, 2008 and 2012 before the ban. The same-base binary TWFE checks timing and weights. Cubic spline, zero interior knots; municipality inference aggregates unit influence functions. A causal marginal response requires strong parallel trends across doses.

Table 18: Unadjusted spline profile of the response

Exposure point	BRL/voter	Level	SE	Marginal	SE
Median funded	0.404	−0.567	0.127	−0.898	0.283
75th percentile	1.520	−1.384	0.315	−0.577	0.164
90th percentile	4.170	−2.151	0.565	−0.060	0.238

Nearest evaluated grid point to each percentile of positive exposure. Municipality-clustered pointwise standard errors. Cubic spline with zero interior knots. Shape depends on specification; it does not identify a structural budget threshold.

does not change. The signs agree, but the magnitude of the marginal response depends on the spline.

The model delivers a conditional saturation prediction, not a mapping from corporate receipts to a structural budget cap. A small or insignificant derivative at high exposure does not establish saturation, an unconstrained budget or a discontinuity at zero. The spline comparison and the binned estimates in Figure 12 show how much the shape depends on the specification.

D.3 Prior-Party Sources and Ideological Partitions

The detailed counts below complement the regression decomposition in Table 4. They use the sample of recruiting parties with financial records and do not impose the network's requirements of financial records at both endpoints and a matching recorded municipality.

The partitions by origin and by ideological distance keep an explicit unscored group. Power-Zucco scores come from the 2013 wave, with bloc cutoffs at ± 0.3 . Lineages that combine distinct 2012 parties are left unscored when a successor's score cannot represent all predecessors. Differences between proportional point estimates across these groups are not tests of equality.

Table 19: Recruitment flows in the supported sample by the different previous observed frozen-2021 lineage


Prior party	Bloc	2012 cycle	2016 cycle	Change	Percent
MDB	Center	32,302	21,922	-10,380	-32.1
DEM	Right	27,284	15,099	-12,185	-44.7
PSDB	Center	26,164	17,882	-8,282	-31.7
PP	Right	20,831	15,244	-5,587	-26.8
PTB	Right	20,716	14,001	-6,715	-32.4
PDT	Left	19,299	14,708	-4,591	-23.8
PT	Left	17,107	17,522	415	2.4
PL	Right	14,738	11,021	-3,717	-25.2
CID	Center	12,913	8,894	-4,019	-31.1
PSB	Left	12,774	11,207	-1,567	-12.3
PV	Center	8,287	7,820	-467	-5.6
PSC	-	6,813	6,557	-256	-3.8

A recorded cross-party recruitment entry has a different lineage on the person's previous observed membership record. This does not establish continuous prior membership, an immediate transfer, or competing recruitment offers. Counts cover destinations with an observed 2012 party account. Bloc uses the Power and Zucco (2023) 2013 wave, cut at plus or minus 0.3. A dash marks an unavailable score or a frozen group pooling distinct 2012 parties.


The counts identify the immediately preceding recorded lineage, not verified flows between organizations or municipalities at the time of the event. They describe the sample and are distinct from the conditional exposure coefficients. Regressions for a single lineage compare municipalities, because municipality fixed effects would absorb a cross-section with one observation per municipality, and their interpretation therefore differs from the main comparison.

D.3.1 Interpretation

The prior-party measure is relevant to the competition for experienced recruits, but a general reduction in organizational activity would also lower it. Neither its coefficient nor the ideological partition distinguishes competition over offers from such alternative channels. The mechanism is an interpretation consistent with several observed margins rather than a separately identified causal pathway.



(a) Positive-dose bins



(b) Descriptive local-linear fit

Figure 12: Functional-form diagnostics. Zero exposure is the binned model's reference category. The local-linear plot uses a residualized outcome and raw positive exposure; it has no inferential band and is not a partial-regression causal estimator.

D.4 Ideological Distance, Extremity and Financial Responsiveness

The reversal between PT and PSDB could reflect the fortunes of particular parties or a broader ideological structure in recruitment. If ideological attachment protects parties from financial competition, the relationship between relative exposure and poaching should vary with the positions of the origin and the destination. Section 4.4 reports the distribution of flows across blocs. Here I use the scored subset of the same network to estimate how the exposure slopes vary with ideology. These are regression slopes, not structural budget elasticities.

The distance interactions test whether the relative-exposure relationship weakens when origin and destination are further apart. Table 20 separates distance from the extremity of each endpoint on the same 46,760 scored pairs in 2,837 municipalities. The scored endpoints cover 66.9 percent of the network's 2012 flows and 61.3 percent of its 2016 flows. Lineages without a defensible common score remain unscored rather than being assigned to the Center, and holding the scored sample fixed separates changes in the specification from changes in the flows included. Distance is the absolute difference between party scores, and extremity is the absolute score of each endpoint. The distance interactions are imprecise, alone and together with extremity, so the data do not show that ideological separation weakens the financial gradient.

Table 20: Ideological distance, endpoint extremity and the relative-exposure contrast

Quantity or joint test	Estimate (SE)	<i>p</i>
Distance interactions: distance-only model	—	0.271
Distance interactions: joint model	—	0.183
Extremity interactions: joint model	—	< 0.001
Origin exposure $\times$ origin extremity	3.040 (0.812)	< 0.001
Destination exposure $\times$ destination extremity	2.493 (0.951)	0.009
Common-extremity gradient of relative slope	0.273 (0.536)	0.610

The continuous models retain the same 46,760 scored pairs in 2,837 municipalities. The outcome is the change in recruitment flows per 100,000 fixed 2012 voters. Distance is the absolute difference between fixed 2013 party scores; extremity is the absolute score of an endpoint. Regressions include national directed-pair and municipality effects. Standard errors and joint tests use the full municipality-clustered covariance. The last row is half the difference between the two extremity interactions, holding distance fixed; it is not the joint test that both interactions vanish. No unscored lineage is coded as centrist.

The extremity interactions are jointly significant, and both are positive. More extreme origins have a more positive association between their exposure and later recorded outflows, and more extreme destinations have a less negative association between their own exposure and arrivals. Whether common extremity changes the relative response depends on the difference between the two interactions, which is 0.273 (0.536), with $p = 0.610$. The significant endpoint interactions therefore do not show that ideology insulates parties from relative money. Party positions do not measure members' ideological commitment, and the baseline model has no ideological dimension, so this exercise is an empirical extension rather than a test of a structural prediction.

D.5 Municipal Exposure: Does the Center Expand When Both Flanks Lose Money?

When both ideological flanks are financially exposed, the difference between a larger share and more recruitment becomes a question about the whole local party system. If competitive access shifts toward the Center, does the Center recruit more, or does it only account for more of a smaller pool? This comparison relates earlier funding patterns to later recruitment at the municipal level. It is a municipal counterpart to the party-level evidence rather than an implication of the two-party model.

Requiring financial records in all three blocs leaves 1,227 municipalities. Table 21 forms four groups by whether the supported Left and Right parties had any corporate receipts. Missing accounts are never classified as zero funding, and the groups measure prior exposure rather than realized losses.

Table 21: Does the Center gain when both sides had corporate funding?

Corporate-positive Left/Right sides	Municipalities	Network inflow change /1,000	Inflow-share change, pp	Total recruitment change /1,000
Neither	524	0.230	7.5	−0.503
Left only	198	−0.045	4.5	−2.226
Right only	224	0.182	3.9	−1.399
Both	281	0.121	4.1	−1.051

The sample requires a supported 2012 account in each of Left, Center and Right. Groups indicate any positive corporate receipts among supported Left/Right parties, not realized losses. Entries are unweighted means of municipal changes from 2012 to 2016. Rates use fixed 2012 electorate. Inflow shares include unscored parties in the denominator; municipalities with zero network flows in either cycle are excluded only from share changes (511, 195, 221 and 279 municipalities, respectively). Total recruitment rates include all recorded Center destinations in these selected municipalities, including Center lineages outside finance support. Those broader outcomes are distinct from the closed transition network.

Center network inflow rises by 0.121 per 1,000 voters in the group where both sides had corporate funding and by 0.230 in the group where neither did. The difference is -0.109 , with a 95 percent descriptive interval of $[-0.364, 0.146]$. Center inflow shares rise in every group, and the largest increase is where neither side had corporate receipts; the difference between the two groups is -3.5 percentage points $([-7.5, 0.6])$. The Center does not gain more where both flanks were exposed.

The full supported network gives a separate accounting comparison. Center arrivals are nearly unchanged, and the Center's net deficit narrows mainly because fewer of its members join other parties. In 2016 the Center receives a net flow of 976 from the Left but sends a net 178 to the Right and 2,053 to unscored parties. Its bilateral balance improves with every group, but it does not become a net recipient from both the Left and the Right, or a net recipient overall. Net balances sum to zero within the network and are not changes in membership stocks.

Relative shares and absolute recruitment move differently here. Broader registry recruitment into Center and Right parties declines in every group, including parties outside the financially supported network. Continuous regressions that control for all four bloc exposures and municipal characteristics also do not show a precise absolute gain for the Center. The groups differ in size and in Center funding, so the contrast is descriptive. Together with the bloc ledger, it shows that a rising Center share or a narrowing deficit does not imply that the Center recruited more at the expense of both flanks.

D.6 The Control Ladder

Figure 13 shows the recruitment coefficient as controls are added. The ladder also includes binary and standardized exposure for each outcome. Changes along the ladder reflect both adjustment and, where the sample changes, differences in support.

D.7 Alternative Fixed Effects

Figure 14 reports the event study under six sets of fixed effects. Section 3.3 reports the saturated specification and both pre-period tests. Failing to reject parallel pre-trends in one specification does not offset a rejection or an adverse placebo in another; the comparisons show how sensitive the estimates are to the assumed comparison structure.




Figure 13: The recruitment coefficient across the control ladder, in nominal 2012 corporate BRL per voter. Each specification uses its own retained sample.




Figure 14: Event studies under six fixed-effect sets.

D.8 Placebos, Confounding and Competing Shocks

D.8.1 Earlier-Window Placebo

Table 22: Earlier-window placebo, committee scope

Window	Estimate	SE	p	N
2008 to 2012	−0.281	0.077	<0.001	18,367
2012 to 2016	−0.071	0.072	0.328	7,870

The earlier-window exercise uses 2008 corporate exposure and 2004 controls; the later one uses 2012 exposure and 2008 controls. Both use committee scope, unlike the headline committee-plus-directory scope. Samples and nominal currency years differ. A pre-reform association is evidence against a broad no-differential-trends claim, not an identical-design placebo for the headline coefficient.

The committee-scope placebo for 2008–2012 gives −0.281 (0.077), with $p < 0.001$, in a window with no ban. It uses earlier exposure and controls, committee accounts rather than the combined accounts of the main specification, a different sample, and amounts in nominal reais of a different year, so it is not an exact replication of the main design. It shows, however, that corporate funding predicted declining recruitment before the ban. The same-scope estimate for 2012–2016 is smaller than the placebo. Temporary funding and recruitment conditions, selection and other differential trends remain material concerns for the dose specification.

D.8.2 Sensitivity to Unobserved Confounding

Table 23: Classical omitted-variable sensitivity diagnostics

Specification	RV, estimate	RV, interval	Municipalities
Two-period, dose	4.368%	2.606%	3,048
Two-period, binary	3.391%	1.610%	3,048
Two-period, dose (prior party)	4.067%	2.300%	3,048
Panel, dose	2.865%	1.975%	3,048
Panel, binary	2.466%	1.572%	3,048

Cinelli–Hazlett robustness values use classical OLS standard errors and residual degrees of freedom including fixed effects. They are not thresholds for the municipality-clustered confidence intervals reported for the regressions.

The Cinelli–Hazlett robustness values in Table 23 are classical OLS diagnostics. Replacing the classical statistic with a clustered t and the residual degrees of freedom with $G - 1$ does not produce a valid clustered robustness threshold, so the regression coefficients continue to use municipality-clustered inference. The observed benchmarks are computed separately for binary and dose exposure on their fitted samples (Appendix E); they describe observed associations and do not bound unobserved confounding.

D.8.3 Competing Shocks

The Lava Jato investigation and other political events could have affected exposed parties differently. Table 24 excludes, cumulatively, PT, MDB, PP and PSDB, the parties at the center of the investigation. The binary estimate does not shrink as these parties are removed: it rises from −0.517 to −0.750 (0.173) for total recruitment, and from −0.221 to −0.333 (0.072) for prior-party recruitment. Removing parties also changes the comparison and the sample, and the persistence of the estimate outside these parties does not rule out a wider corruption shock or other omitted trends.

Table 24: Dropping the parties at the center of Operation Lava Jato

Specification	New members	Prior party	N
All parties	−0.517 (0.148)	−0.221 (0.050)	14,606
Excluding PT	−0.587 (0.148)	−0.254 (0.057)	12,969
Excluding PT, MDB	−0.599 (0.164)	−0.297 (0.065)	11,258
Excluding PT, MDB, PP	−0.764 (0.161)	−0.321 (0.068)	9,931
Excluding PT, MDB, PP, PSDB	−0.750 (0.173)	−0.333 (0.072)	8,569

Two-period first difference on binary treatment with party and municipality fixed effects and the 2008 controls. Cumulative exclusions change the comparison and retained sample. Persistence outside these parties does not rule out wider corruption shocks or omitted trends. Prior party means a different preceding recorded lineage, not verified direct switching. Standard errors cluster by municipality.

D.9 Calendar Timing and Network Measurement

Table 25: Alternative recruitment-timing comparisons

Specification and endpoint	Coefficient (SE)	<i>p</i>	Municipalities
Annual, phase adjusted: origin	0.023 (0.019)	0.239	3,048
Annual, phase adjusted: destination	−0.092 (0.050)	0.067	3,048
Two-year window 2016: origin	0.260 (0.185)	0.160	3,048
Two-year window 2016: destination	−0.446 (0.281)	0.112	3,048
Two-year window 2020: origin	0.019 (0.093)	0.839	3,048
Two-year window 2020: destination	−0.548 (0.279)	0.050	3,048

The outcome is recorded recruitment flows per 100,000 fixed 2012 voters; coefficients use nominal 2012 corporate BRL per voter and errors cluster by municipality. Annual rows omit 2015 and control exposure by election phase (1,071,240 observations). Two-year rows sum calendar years 2007–08, 2011–12, 2015–16 and 2019–20 (357,080 observations), with 2011–12 as reference. They include months after election day. The single earlier-window joint test has $p = 0.491$; the annual pretest is $p = 0.169$ and remains seasonality-sensitive. These timing comparisons do not isolate the funding reform from concurrent changes.

Annual recruitment is concentrated around candidacy deadlines. The annual pre-period test, relative to 2014, gives $p = 0.169$; it mixes phases of the electoral calendar and cannot replace the cycle diagnostics. A pooled annual comparison omits 2015 and adds interactions between each endpoint’s exposure and the phase of the electoral calendar; its origin and destination coefficients are not significant at the five percent level. Two-calendar-year windows combine 2007–08, 2011–12, 2015–16 and 2019–20, which partly addresses shifts between the year before an election and the election year. They include months after election day and allow only one earlier-window comparison, so the pre-period test’s $p = 0.491$ is not a general validation of the design.

Additional specifications winsorize exposure, count flows without the geographic restriction, and restrict the sample to transitions whose previous membership has a dated termination no more than a year before the new start. The endpoint signs are similar, and the precision varies. These checks address particular concerns about measurement and the upper tail without turning recorded histories into verified poaching. Municipalities with only two eligible parties identify a relative contrast after municipality effects, and separate endpoint slopes require variation from larger local networks. On the overlap with realized accounts, the slope on destination exposure is close to zero ($p = 0.684$), so comparisons between regressors must account for the selection of accounts. None of these checks provides an exogenous change in net recruitment resources.

Table 26: Electoral outcomes, dose coding

Outcome	Estimate	SE	<i>p</i>	<i>N</i>
Candidates per 1,000 voters	−0.005	0.004	0.198	14,606
Exited (no candidate in 2016), pp	−0.366	0.098	<0.001	14,606
Votes per registered voter, pp	−1.139	0.229	<0.001	14,606
Won any seat, pp	−0.292	0.225	0.194	14,606
Vote share, pp	−0.588	0.119	<0.001	11,888
Seat share, pp	−0.138	0.075	0.066	11,888
Candidates per seat	−0.004	0.002	0.063	11,888

Headline fixed effects and controls; municipality-clustered standard errors. Each row uses its own retained sample. Conditional shares describe observations with defined denominators. Rows marked pp report shares and indicators in percentage points.

Table 27: Candidate composition, dose coding

Outcome	Estimate	SE	<i>p</i>	<i>N</i>
Mean declared wealth of the slate, asinh	−0.016	0.006	0.010	10,564
Different prior candidate party (returners), pp	−0.541	0.186	0.004	9,308
Mean education of the slate, 1–8 scale	0.007	0.004	0.078	11,885
Previously observed candidate share, pp	0.019	0.123	0.878	11,885
Share incumbent, pp	0.258	0.177	0.146	11,888

Headline fixed effects and controls; municipality-clustered standard errors. Each row uses its own retained sample. Conditional shares describe observations with defined denominators. Rows marked pp report shares and indicators in percentage points. Candidate party changes use the common historical groupings and the previous observed candidacy, conditional on returning. Previously observed candidacy counts all history since 2000, giving 2016 longer coverage than 2012; this is not a fixed-window experience measure.

D.10 Downstream Outcomes

Count outcomes treat electoral nonparticipation as zero where the source supports that interpretation. Candidate attributes and shares have their own missing values and conditions on participation and should be read together with the probability of exit. The candidate composition table uses each candidate’s cumulative history of candidacies since 2000, which covers more years in 2016 than in 2012. The share of returning candidates with a different previous lineage is defined among returning candidates and uses the 2021 lineages. The separate analysis of prior recruits uses a common retrospective window, but that choice does not turn the cumulative candidacy history into a fixed-window measure. The attributes of candidates who stand do not identify the model’s latent quality for the whole pool of recruits.

D.11 Recruitment Composition

The historical registry does not record sex, so I infer it from first names and keep unclassified names as a separate category. Share outcomes require a defined denominator and therefore condition on recruitment and on classification. A null estimate for the share does not show that demographic composition is unchanged. Comparisons with the composition of the electorate weight the electorates of the represented municipalities accordingly. These descriptive comparisons do not identify representativeness in the population when recruitment or name coverage is selective.

Table 28: Recruitment composition, binary coding

Outcome	Estimate	SE	<i>p</i>	<i>N</i>
Female recruitment per 1,000 voters	−0.149	0.063	0.018	14,606
Male recruitment per 1,000 voters	−0.330	0.079	<0.001	14,606
Unclassified recruitment per 1,000 voters	−0.038	0.014	0.008	14,606
Female share of classified recruitment, pp	−0.084	0.541	0.876	12,413

Headline fixed effects and controls; municipality-clustered standard errors. Each row uses its own retained sample. Conditional shares describe observations with defined denominators. Rows marked pp report shares and indicators in percentage points. Historical recruit sex is inferred from names, with unclassified recruitment retained as a separate outcome.

D.12 The Modern Membership Registry

The modern membership snapshot supports the descriptive series in Figure 7 but is not used to test the historical recruitment estimates. Mergers after 2021 and retrospective party labels prevent a validated mapping to the 2021 lineages, and a snapshot of surviving records cannot by itself identify differences in historical retention.

E Robustness Detail

E.1 Sensitivity Analysis

Table 23 reports the classical robustness values. The benchmarks below compare regressions with and without each observed covariate, conditioning on the other covariates and the fixed effects on the same sample. Binary exposure must be compared with its own benchmarks and dose exposure with its own. Municipality size is absorbed by the municipality fixed effects. These diagnostics do not bound unobserved confounding or the municipality-clustered confidence intervals.

Table 29: Observed benchmarks by treatment coding

Covariate	Treatment partial R^2	Outcome partial R^2
binary: Candidates per seat, 2008	0.191%	0.000%
binary: Held the mayoralty, 2008	0.188%	0.016%
binary: Contested in 2008	0.152%	0.608%
binary: Vote share, 2008	0.097%	0.025%
binary: Incumbents, 2008	0.030%	0.104%
binary: Members per 1,000 voters, 2008	0.030%	0.253%
binary: Seat share, 2008	0.004%	0.004%
binary: New recruitment flows per 1,000, 2004–08	0.002%	0.068%
dose: Held the mayoralty, 2008	0.225%	0.014%
dose: Vote share, 2008	0.076%	0.025%
dose: Contested in 2008	0.035%	0.617%
dose: New recruitment flows per 1,000, 2004–08	0.009%	0.066%
dose: Incumbents, 2008	0.006%	0.103%
dose: Seat share, 2008	0.003%	0.003%
dose: Candidates per seat, 2008	0.002%	0.000%
dose: Members per 1,000 voters, 2008	0.000%	0.248%

Each row compares models with and without the covariate, conditioning on remaining controls and fixed effects on the same fitted sample. Compare binary with binary and dose with dose. These associations are not upper bounds on omitted confounding.

E.2 Functional Form

Table 30: Is the dose–response linear?

Specification	Term	Estimate	SE	N
Linear (reported)	Corporate BRL per voter	−0.120	0.060	14,606
Quadratic	Corporate BRL per voter	−0.161	0.069	14,606
Quadratic	Squared term	0.001	0.003	14,606
Linear, top 1% of doses dropped	Corporate BRL per voter	−0.188	0.068	14,445

All specifications carry the party and municipality fixed effects and the 2008 control set of the reported design, and cluster by municipality. The squared term is reported with a caution: squared exposure can be sensitive to the upper tail. The trimmed row is the same linear specification after dropping the top one percent of doses. Exposure is nominal 2012 BRL per voter. The binned specification estimates a separate contrast for each chosen positive-dose bin and is in Figure 12.

Zero exposure is the omitted category in the binned comparison. The kernel plot is descriptive, because it residualizes the outcome without jointly partialling out exposure. The unadjusted continuous estimator also reports a cubic specification with two interior knots next to the baseline with none. Differences between these profiles diagnose sensitivity to the specification and are not a basis for choosing the most favorable shape. Dropping the top one percent of exposure makes the linear coefficient larger in magnitude, so the dose estimate is not driven by the upper tail.

Table 31: Sensitivity to the declared spline specification

Specification	Average level	SE	Average marginal	SE
cubic, zero interior knots	−0.454	0.102	−0.194	0.019
cubic, two interior knots	−0.454	0.102	−3.187	0.348

Identical unadjusted 2012-base comparison, post 2016 and 2020. Both profiles use cubic splines and municipality influence-function inference. The two specifications use zero and two interior knots, respectively. These summaries do not establish invariance of the entire dose profile.

E.3 Inference

Municipality counts for fitted regressions refer to retained observations, after missing values and singletons are removed. Inference for the continuous estimator aggregates unit influence functions into municipality scores and applies a finite-cluster correction. Spline degree and knots are stated with each specification.

The network comparisons use different denominators, samples and pre-period restrictions from the party-level comparisons. In the directed panel, the test of the relative contrast uses the full covariance matrix across both endpoints and all pre-reform cycles, and the joint endpoint test adds restrictions on mean exposure. Section 4 reports the continuous, binary and narrower-lineage tests together rather than selecting one. Because the network is defined by financial support in 2012, the early outcomes are conditioned on that network, which limits their interpretation as predetermined placebos.

Annual and two-calendar-year specifications address timing but remain exposed to different responses to the change in the membership deadline. The conditional Poisson models remove different groups of all-zero fixed effects, so differences in significance between attraction and sourcing do not test whether the effects differ. Directed comparisons use municipality-clustered standard errors, municipal group means use descriptive t intervals, contrasts between groups use Welch intervals, and municipal cross-sectional regressions use HC1 standard errors. The set of exercises is exploratory, and I apply no adjustment for multiple testing. None of these checks turns recorded membership histories into observed offers or exits from active membership.

E.4 Additional Notes on Party-Level Comparisons

Section 3.2 specifies the first-difference design and Section 3.3 the saturated panel. Their outcomes require financial records only at the recruiting party, unlike the network, which requires them at both ends of each link. The tables for total and prior-party recruitment appear in Sections 3 and 4, the dynamic and pooled results in Section 3, and the functional-form comparisons and the earlier-window placebo in Appendix D. The construction and inference rules above apply to all of them.

E.5 Ideological Financial-Response Interactions

The interaction exercise in Appendix D.4 is separate from the descriptive bloc ledger in the main text. It uses the fixed 2013 party scores and bloc definitions of Section 4.4.

The comparison asks whether distance and extremity change the association between financial exposure and recruitment. Distance measures the separation between two parties, and endpoint extremity asks a different question about their individual positions. Distance is the absolute difference between the party scores, and extremity is the absolute score. Regressions on a common sample interact each endpoint's exposure with these measures and use contrasts based on the full covariance matrix to test whether ideological position changes relative responsiveness. These tests concern heterogeneity in slopes; the available outcomes and financial measures do not identify structural funding elasticities, and neither measure observes a member's ideological commitment.

E.6 Municipal Exposure and Center Outcomes

A relative gain in a bilateral contest need not increase recruitment in the municipality. The municipal comparison therefore asks whether exposure on both ideological flanks is associated with higher Center inflow, a larger Center share or a more favorable net balance.

Municipalities with financial records in the Left, the Center and the Right are grouped by whether corporate receipts were present on neither side, the Left only, the Right only, or both sides. The groups measure prior exposure, not realized losses. Municipalities are equally weighted, and broader registry recruitment in the same municipalities is reported separately from flows inside the supported network. The continuous specifications include all bloc exposures jointly and then add controls for state, municipality size and the number of supported parties.