Going It Alone? A Structural Analysis of Coalition Formation in Elections*

Sergio Montero†

May 3, 2024

Abstract

The quality of democratic representation hinges on the alternatives available to voters. In many democracies, party leaders exercise considerable control over the electoral supply, particularly through the formation of electoral coalitions. Despite their prevalence, little is known about the tradeoffs these coalitions impose on prospective partners, how they influence other strategic choices by competing parties, or their ultimate impact on election outcomes. I develop and estimate a model of electoral competition in which: (i) parties can form coalitions to coordinate their candidate nominations, and (ii) parties invest in campaign activities in support of their candidates. The model is estimated using data from the 2012 Mexican Chamber of Deputies election, which offers district-level variation in coalition formation. A comparison of election outcomes under counterfactual coalition scenarios uncovers equilibrium campaign savings as well as significant electoral gains benefitting weaker partners disproportionately.

Keywords: electoral coalitions, campaign spending, mixed electoral systems, structural estimation, strategic complementarities

Supplementary material for this article is available in the Online Appendix.

Replication files are available in the JOP Data Archive on Dataverse (https://dataverse.harvard.edu/dataverse/jop). The empirical analysis has been successfully replicated by the JOP replication analyst.

---

*Short title: Going It Alone? Coalition Formation in Elections

†Assistant Professor, Departments of Political Science and Economics, University of Rochester, Rochester, NY 14627, email: smontero@rochester.edu.

Electoral coalitions are common in most democracies (Golder, 2006). In hopes of influencing election outcomes, like-minded political parties often coordinate their electoral strategies, typically by fielding joint candidates for office. This manipulation of the electoral supply—i.e., the alternatives available to voters—may significantly affect representation and post-election policy choices. Yet there is little evidence documenting the impact coalition formation can have on other aspects of electoral competition—e.g., campaign expenditures—or on election outcomes per se, despite scholars having long argued that democratic stability largely depends on voter satisfaction at the polls (Powell, 2004). This paper begins to fill that gap.

Early studies of electoral coalitions focused primarily on comparing their prevalence across electoral systems or on their role in shaping post-election government formation in parliamentary democracies (Cox, 1997; Ferrara and Herron, 2005; Golder, 2006; Carroll and Cox, 2007; Debus, 2009). Overwhelmingly, the two most robust predictors of coalition formation identified by such cross-country analyses have been the ideological proximity of prospective partners and whether electoral rules feature a strong majoritarian component. Although this scholarship has provided valuable insight into the extensive margin of electoral coalitions, in many democracies there is remarkable variation in the degree to which parties join forces *for a given election*. In France (Blais and Indridason, 2007), India (Bandyopadhyay, Chatterjee and Sjöström, 2011), and Japan (Catalinac and Motolinia, 2021), to give a few examples with varied institutional environments, coalition partners often nominate common candidates in only a fraction of contested constituencies, while running independently elsewhere.

Such "partial" coalition arrangements, in addition to being of considerable independent interest, offer scholars an ideal setting in which to study the incentives behind and consequences of electoral coalitions. Given a fixed party system and institutions (e.g., French parliamentary races), election outcomes with and without coalition candidates can be observed. Of course, the source of this variation is not exogenous: these are strategic nominations by parties meant to influence election results, which poses a significant methodological challenge.

Absent an experimental intervention, since a party may not be interested in joining forces in races it expects to win handily, a simple comparison of the party's performance where it ran alone versus in coordination with others might severely underestimate the electoral benefits of coalition formation. To address this, standard inferential approaches based on randomization, discontinuity designs, or parallel-trends assumptions are likely feasible only in limited settings (Hortala-Vallve, Meriläinen and Tukiainen, 2023). An alternative that is fairly underutilized in political science but widely applicable is to—in the spirit of classical approaches to sample selection (Heckman, 1979)—explicitly model and structurally estimate the determinants of parties' strategic choices.

Using data from the 2012 Mexican Chamber of Deputies election, this paper studies coalition formation in the context of legislative elections under a single-ballot mixed electoral system that combines a first-past-the-post (FPTP) tier and a proportional-representation (PR) tier. Coordination among coalition partners takes the form of joint candidate nominations for the FPTP races. Accordingly, I develop and estimate a structural model of electoral competition in which: (i) parties can make coalition formation commitments, which determine the menu of candidates competing in each FPTP contest, and (ii) parties invest in campaign activities in support of their candidates.

Mexican parties are allowed to form partial coalitions in national legislative elections. In the 2012 Chamber of Deputies election, two parties—the Institutional Revolutionary Party (PRI) and the Ecologist Green Party of Mexico (PVEM)—nominated joint candidates in only two thirds of the FPTP races, and they nominated distinct candidates and thus competed against each other elsewhere. What explains such an arrangement? Which features of the electoral system provide countervailing incentives for partners to join forces in some but not all races? How does this affect campaign competition? And what are the implications of such an alliance for the distribution of power across parties? The structural model below attempts to explicitly account for all relevant considerations—including, in particular, features of the electoral environment that are unobserved by the researcher. The model is then used to simulate election outcomes under counterfactual coalition scenarios in order to quantitatively assess the tradeoffs involved in and (equilibrium) consequences of coalition formation. To my knowledge, this paper is the first to simultaneously address these questions, doing so from a unified theoretical and empirical perspective.[1]

The estimation strategy mirrors the structure of the model and proceeds in three stages, borrowing insights from the literature in economics on entry and competition in markets with differentiated products. First, voters' preferences are estimated from FPTP district-level voting data following the aggregate discrete-choice approach to demand estimation popularized by Berry, Levinsohn and Pakes (1995). Second, the parameters of parties' payoffs driving their campaign expenditures are estimated by exploiting necessary equilibrium conditions of the campaign spending game played by parties across the FPTP districts. Lastly, the remaining parameters of parties' payoffs shaping their coalition formation decisions are identified from moment conditions analogous to market entry considerations.

With the estimated structural parameters in hand, I conduct two counterfactual experiments: I simulate the outcomes that would have ensued had PRI and PVEM decided to either not form a coalition at all or form a total coalition instead (nominating joint candidates in all districts). These experiments yield three key findings. First, I document substantial electoral gains for coalition partners. In terms of jointly held seats in the Chamber of Deputies, PRI and PVEM's partial coalition allowed them to close the gap to obtaining a legislative majority by 77% (from 206 joint seats to 241, out of a 251-seat simple majority), and they would have closed it by 89% (securing 246 joint seats) had they run together in every district. Perhaps unsurprisingly, these seat gains resulted from the FPTP component of the electoral system,

[1]Although others may disagree with key modeling choices, an advantage of the structural approach is that it makes the link between assumptions and conclusions explicit, providing a clear path toward potential improvements. Furthermore, it enables rich analyses of inherently nonlinear empirical relationships, where standard homogeneity and no-interference assumptions are unlikely to hold.

which underscores the most basic rationale for coalition formation under majoritarian voting (Cox, 1997; Ferrara and Herron, 2005): by nominating common candidates, the coalition partners avoided splitting the vote and thus raised their likelihood of victory in the corresponding FPTP races.

Second, coalition formation affected partners asymmetrically. While PRI grew its share of seats in the chamber by 7% (from 38.6% to 41.4%), PVEM almost tripled its share (from 2.6% to 6.8%). Had they formed a total coalition, only PVEM would have benefitted further. In fact, PRI would have lost one seat relative to its partial-coalition share. This offsetting pressure on the extent to which PRI and PVEM joined forces was due to the PR component of the election. Because of the way in which supporters of a coalition candidate are allowed to split their vote between the nominating parties for the PR tier (described in detail below), PRI's vote share suffered considerably under coalition candidacies, resulting in PR-seat losses for PRI despite considerable net gains for both partners.

These results add nuance to existing evidence regarding coalition formation in elections. One one hand, the 2012 PRI-PVEM alliance confirms numerous expectations from this literature. It emerged in an electoral system with a strong majoritarian component (Ferrara and Herron, 2005; Golder, 2006) and low volatility (Invernizzi, 2023). As shown in Figure 2 below, it involved two ideological neighbors (Golder, 2006; Hortala-Vallve, Meriläinen and Tukiainen, 2023). And, despite net gains, one of the partners paid a price in the PR component of the election (Hortala-Vallve, Meriläinen and Tukiainen, 2023). Yet the 2012 PRI-PVEM alliance also illustrates that, even conditional on a fixed ideological and institutional environment, there is considerable heterogeneity in strategic incentives for prospective coalition partners, with important distributional implications. Echoing Blais and Indridason (2007), who note that electoral coalition agreements in France tend to favor stronger partners, the weaker PVEM headlined only a marginal share of joint candidacies (about one fifth), concentrated in the toughest FPTP districts. However, the counterfactual experiments reveal that, after taking into account all features of the electoral environment, including strategic responses by other parties, it was actually the weaker PVEM who benefitted most from the alliance. This underscores the usefulness of the methodological approach. Moreover, while post-election legislative bargaining is beyond the scope of this paper, the results suggest strong parties may rely on electoral alliances to pre-select and foster legislative partners—even at some personal cost in the case of PRI.[2] Indeed, Catalinac and Motolinia (2021) show that, following the 2012 election, PRI and PVEM used discretionary transfers controlled by the Chamber of Deputies to reward supporters and cement their alliance.

Finally, I also examine the interplay between coalition formation and another important dimension of electoral competition: campaign expenditures. The counterfactual experiments uncover moderate equilibrium savings when parties join forces. Although average joint spending by the two coalition partners appears nearly identical if they nominate distinct candidates or if they jointly nominate a PRI candidate to represent the coalition, it decreases by about 2.5% with a joint PVEM candidate. Given PVEM's marginal share of joint candidacies, observed total spending by the two partners in the election was only 0.6% lower relative to the no-coalition scenario. Yet, while modest, these savings shine light on a typically overlooked rationale for coalition formation in elections, and they are consistent with the intuition in many formal models of endogenous valence competition that the incentives to invest in campaign advertising increase with ideological proximity (Ashworth and Bueno de Mesquita, 2009; Iaryczower and Mattozzi, 2013). In particular, the campaign benefit of a joint PVEM candidacy resulted from less intense competition with PRI's ideologically closest rival. This paper is thus the first to empirically complement theoretical work examining how electoral institutions simultaneously shape strategic candidate entry—here, with parties as gatekeepers—and the intensity of campaign competition (Iaryczower and Mattozzi, 2013).

[2] As discussed in detail below, joining forces with PVEM also resulted in a decreased share of public funding for PRI for the following three years (2013–2015) due to its lower national vote share in the 2012 Chamber of Deputies election.

Mexican Legislative Elections

To motivate the analysis, I begin by describing in some detail the institutional environment. Members of the Mexican Chamber of Deputies are elected every three years, and, prior to 2018, no incumbent could stand for consecutive re-election. The election is held under a *single-ballot* mixed electoral system. For electoral purposes, Mexico is divided into 300 districts and 5 regions (see Figure A1 in Online Appendix A). Out of 500 total deputies, 300 directly represent a district after being elected by direct ballot under first-past-the-post voting. The remaining 200 seats in the chamber, which are divided equally across regions, are allocated by proportional representation to registered national political parties as follows: the votes cast across the 300 FPTP districts are pooled by region, and each party is given a share of each region's 40 PR seats in proportion to the share of votes received by the party's candidates in the corresponding district races.[3] This allocation is subject to disproportionality restrictions that preclude any party from obtaining (i) more than 300 total seats or (ii) a share of total seats that exceeds by more than 8 percentage points the party's national vote share, in which case the excess PR seats are divided proportionally among the remaining parties.

Mexican parties are funded primarily from the federal budget. Annual funds, including for campaign purposes, are distributed as follows: 30% is divided equally among all registered parties, and the remaining 70% is divided in proportion to their national vote share in the most recent Chamber of Deputies election. Thus, Mexican parties compete in this election to secure not only seats in the legislature but also funding for their day-to-day operations and campaign activities for the following three years.

Prior to each Chamber of Deputies election, parties may form coalitions, which enable them to coordinate their candidate nominations for the FPTP district races. Coalition partners may not coordinate, however, in the PR component of the election—they are required to submit

[3]The largest remainder method with Hare quotas (Bormann and Golder, 2013) is employed to allocate the seats. Only parties that secure at least 2% of the national vote are eligible to hold seats in the chamber.

independent lists of up to 40 PR candidates per region.

Coalition agreements are negotiated by parties' national leaders and must be publicly registered before the national electoral authority, the National Electoral Institute (INE), prior to the selection of individual candidates. The agreements constitute binding commitments specifying, for each electoral district: (i) whether the coalition partners will nominate a joint candidate or independent candidates, and (ii) in the case of a joint nomination, from which party's ranks will the coalition candidate be drawn. After the election, coalition agreements imply no formal obligations for coalition victors in the legislature, who retain their original party affiliation. Thus, by supporting another party's candidate via a joint nomination, the remaining coalition partners forgo the corresponding district seat in the chamber.

When deciding whether to vote for a coalition candidate, although voters cast a single ballot that determines outcomes in both tiers of the election, they do have some control over how their vote should be counted for PR (and funding) purposes. The ballots presented to voters on election day feature one box per registered party containing the name of the party's candidate for that district.[4] If a candidate is nominated by a coalition, their name appears inside each of the coalition partners' boxes. To cast their vote in favor of a coalition candidate, voters can mark any subset of the coalition's boxes on the ballot. Regardless of the chosen subset, the vote is counted as a single vote in favor of the coalition candidate for the purpose of selecting that district's FPTP winner. However, the vote is split equally among the chosen subset for the PR allocation. In the case of a two-party coalition, this boils down to three options: giving the PR vote fully to either of the two partners or splitting it 50-50. Table A3 in Online Appendix A illustrates this for the case of PRI-PVEM coalition candidates in 2012.

Events in an election year unfold as follows. First, as described, coalitions are publicly registered. Next, candidates are selected and formally nominated. Campaigns then take place

[4]Independent candidacies or write-in campaigns are also allowed, but their vote shares are negligible. Moreover, voters supporting independent or write-in candidates effectively forgo participation in the PR component of the election.

within a fixed timeframe. And, finally, ballots are cast.

Due to one-term limits and constraints ensuring the primacy of public funding, fundraising by candidates is effectively absent from the Chamber of Deputies election.[5] Party leaders finance their candidates' campaigns directly, making a centralized decision of how much to spend in each district. In the case of coalition candidates, partners may share campaign expenditures freely, which may incentivize coalition formation. The net effect of coalition candidacies on campaign expenditures, however, depends on equilibrium responses by rival parties and is ultimately an empirical question.

The 2012 Chamber of Deputies Election

Seven parties participated in the 2012 Chamber of Deputies election: 2 parties, the National Action Party (PAN) and the New Alliance Party (NA), participated independently; 3 parties, the Party of the Democratic Revolution (PRD), the Labor Party (PT), and the Citizens' Movement (MC), joined forces in all districts in a coalition called the Progressive Movement (MP); and PRI and PVEM formed a partial coalition called Commitment for Mexico, joining forces in only 199 districts. PRI and PVEM jointly nominated a PRI candidate in 156 districts and a PVEM candidate in 43 districts—see Figure 1.

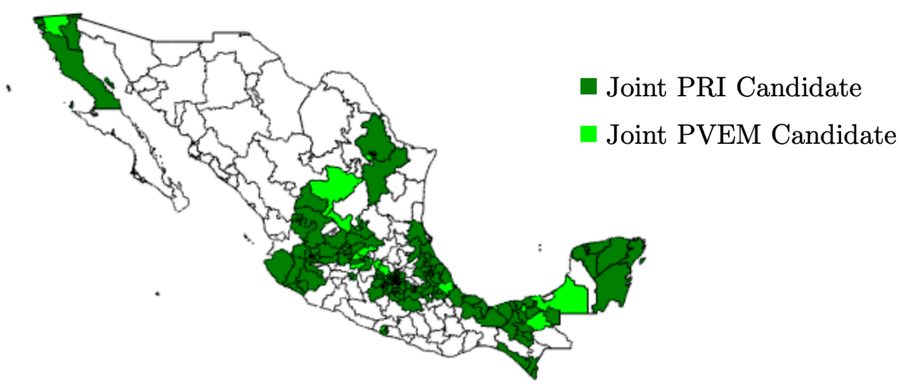

Figure 1: Districts with Joint PRI-PVEM Candidacies

As shown in Figure 2, which is based on a national poll of ideological identification conducted by a leading public opinion consultancy in 2012, the parties can be placed on a one-

[5]Private contributions, including candidates' personal funds, account for less than 1% of expenditures.

dimensional ideology spectrum as follows—from left to right: the MP parties, NA, PVEM, PRI, and PAN. Figure 2 also presents the parties' national vote shares in the 2012 election to illustrate their relative sizes. PRI, PAN, and PRD were the main political forces, in that order. Together they accounted for more than 80% of votes nationally. Of the smaller parties, the centrist PVEM was the strongest, with nearly a third of PRD's vote share. The shares in Figure 2, however, were shaped by the coalitions that formed prior to the election. One of the primary goals of this paper is to quantify this effect. Before doing so, the remainder of this section provides a descriptive look at district-level election outcomes, comparing districts with and without PRI-PVEM coalition candidates, to glean some intuition and motivate the structural model introduced below.

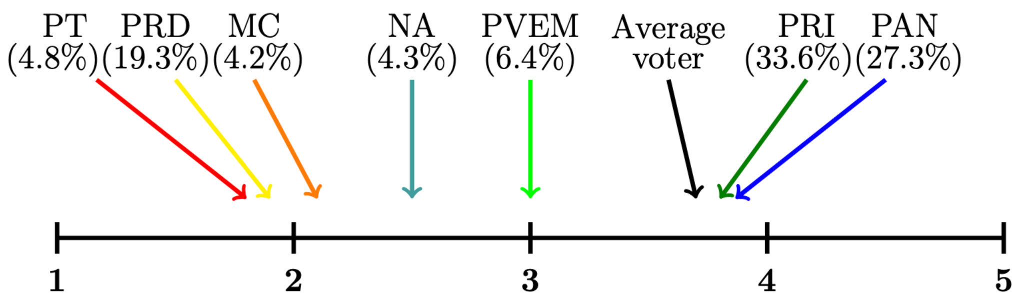

Figure 2: Left-Right Ideological Identification of Mexican Parties and Voters

Notes. Source: Consulta Mitofsky (2012). One thousand registered voters were asked in December 2012 to place the parties and themselves on a five-point, left-right ideology scale. Arrows point to national averages. Parties' national vote shares in the 2012 Chamber of Deputies election are shown in parentheses.

District-level election results are published by INE. As a coalition, PRI and PVEM were quite successful, winning 122 of the 199 districts they shared: 106 victories with a joint PRI candidate and 16 victories with a joint PVEM candidate. Independently, PRI obtained 52 additional victories, and PVEM obtained 3. The final composition of the Chamber of Deputies, including PR seats, is presented in Table A1 in Online Appendix A. Hereafter, I treat the total coalition MP as a single party.[6] Note that PRI's proportionally smaller share of PR seats was a consequence of the disproportionality restriction described above, which precludes any party's total share of seats from exceeding by more than 8 percentage points its

[6]Historically, and in terms of policy goals, the MP parties effectively acted as a single party. In 2012, they also nominated coalition candidates in the presidential and all Senate races.

national vote share.[7] Without this constraint, PRI would have obtained 67 PR seats instead of 49.

Table A2 in Online Appendix A breaks down election outcomes by type of candidate ran by PRI and PVEM. Victory rates and average vote shares are computed for each party in each subsample of districts. For PRI and PVEM, three notable comparisons emerge. First, although it doesn't account for the partners' strategic choice of where and how to run together, Table A2 suggests that, in terms of victory rates, PRI-PVEM coalition candidates outperformed their independent counterparts, underscoring the most basic rationale for coalition formation under FPTP voting. Second, despite the higher victory rates, PRI and PVEM's joint vote share appears to have suffered under coalition candidacies. The two partners commanded, on average, a per district joint vote share of 41.6% with independent candidates, but their average joint vote share was only 40.2% with a joint PRI candidate and 36.4% with a joint PVEM candidate. This suggests joint nominations may have led to a net loss of votes to rival parties (Hortala-Vallve, Meriläinen and Tukiainen, 2023).

Lastly, Table A2 also suggests that there was a transfer of votes between the coalition partners as a result of joint nominations, with PVEM's vote share benefitting significantly (increasing from 4.9% with distinct candidates to 7.0% with a joint PRI candidate or 7.7% with a joint PVEM candidate) at the expense of PRI's (which decreased from 36.7% with independent candidates to 33.2% with a joint PRI candidate or 28.7% with a joint PVEM candidate). To examine this apparent transfer more closely, Table A3 in Online Appendix A shows how PRI-PVEM coalition supporters decided to split their vote between the two partners as explained above. While most supporters gave their vote fully to one of the two parties—in proportions roughly similar to the parties' vote shares with independent candidates—a substantial fraction of coalition supporters opted for the 50-50 split. Furthermore, as shown in Figure A2 in Online Appendix A, these choices were remarkably consistent across districts, with each party's share of the no-split votes only slightly higher when headlining the coalition

[7] That is, $207 = \lfloor 500(0.335953 + 0.08) \rfloor$.

than otherwise.[8] This suggests a considerable tradeoff for PRI: joint PRI candidate nominations seem to have increased the likelihood of FPTP victory for PRI at the expense of decreasing its vote share and, thus, its share of PR seats and future funding. For PVEM, on the other hand, joint PVEM (or PRI) candidate nominations seem unambiguously beneficial: PVEM was unlikely to win FPTP races on its own (3% victory rate), and it appears to have benefitted greatly in the PR tier of the election by joining forces with PRI. In the counterfactual experiments described below, these effects are reexamined after accounting for PRI and PVEM's strategic choice of coalition configuration, influenced by differences in the competitive environment across districts.

To describe the electorate, district-level demographics from the 2010 population census are available from the National Statistics and Geography Institute (INEGI). Table A4 in Online Appendix A provides a summary of the districts by type of candidate ran by PRI and PVEM. The only noticeable difference in demographics across the three types of PRI-PVEM candidacies, though not statistically significant, concerns the percentage of rural neighborhoods in each district. Table A4 suggests that the coalition partners were more likely to nominate independent candidates or a joint PRI candidate in more rural districts, consistent with PRI's historical dominance in rural areas (Ames, 1970; Magaloni, 2006; Frey, López-Moctezuma and Montero, 2023).

Finally, Table A5 in Online Appendix A summarizes campaign spending in the district races—i.e., total expenditures in support of a candidate—by type of PRI-PVEM candidacy.

[8] Catalinac and Motolinia (2021) describe a coalition effort to encourage voters to allocate their PR vote to PRI when supporting a joint PVEM candidate, and they show that complying districts where rewarded with discretionary transfers. However, Figure A2 suggests that, overall, the effort wasn't very successful. Relatedly, Figure A2 admits negligible potential effects on PR votes of the "watermelon" candidacies analyzed by Spoon and Pulido-Gómez (2020)—i.e., formally PVEM coalition candidates (green on the outside) with prior ties to the PRI (red on the inside).

The data can be requested from INE and provide only a coarse account of monetary expenses. I focus on total spending per candidate as a broad measure of the intensity of campaign efforts. The model presented below features party leaders making strategic campaign spending decisions on a district-by-district basis, as opposed to simply dividing up resources by state or regionally. To validate this, Figure A3 in Online Appendix A maps each party's geographic distribution of campaign expenditures. As expected, there is substantial variation across neighboring districts, beyond what could be driven solely by differences in campaign costs, considering that the locations of high-spending districts are only weakly correlated across parties.

A Model of Competition in Legislative Elections

As noted, the timing of events in the Mexican Chamber of Deputies election is as follows. First, party leaders make public coalition formation commitments. In accordance with these agreements, candidates are selected and nominated. Campaigns then take place, and finally ballots are cast. The model I develop captures this timing in three stages: a coalition formation stage, a campaign stage, and a voting stage. In standard backward-induction fashion, I introduce the model starting with the voting stage.

For the coalition formation stage, the analysis that follows focuses on PRI and PVEM's choice of coalition configuration and conditions on all other parties running as observed in the data. While potentially intriguing, considering alternative coalitions would not be plausible in light of ideological incompatibilities and the electoral system's disproportionality restrictions, which discourage coalition formation among large parties.[9] Furthermore, a key objective is to quantify the equilibrium tradeoffs PRI and PVEM faced when crafting their alliance.

Before introducing the model, I develop some useful notation.

Notation. Districts are indexed by d, parties by p, and voters by i. The indicator $M_d \in \{M^{\text{PRI}}, M^{\text{PVEM}}, M^{\text{IND}}\}$ describes the menu of candidates available to voters in district d as

[9] Prior to 2018, ideologically incompatible coalitions in Mexico had only emerged in local elections (Frey, López-Moctezuma and Montero, 2023).

a result of PRI and PVEM's coalition configuration. Menu M^{PRI} corresponds to the case wherein PRI and PVEM jointly nominate a PRI candidate to represent the coalition. Menu M^{PVEM} corresponds to a joint PVEM candidate nomination. And menu M^{IND} ensues if they nominate distinct candidates and thus run independently. Each of the three menus *additionally* includes an MP candidate, an NA candidate, and a PAN candidate. Thus, menu M^{IND} features five competing candidates, while menus M^{PRI} and M^{PVEM} each feature four.

Voting stage. With a single ballot, voters simultaneously choose a candidate and a PR party list. If a candidate is nominated by a coalition, voters may split their PR vote among the nominating parties' lists. However, the selection of a candidate is the preeminent choice. The voting stage is therefore modeled as a two-tier decision: voters first select a candidate and then, if pertinent, how to split their PR vote.

When choosing a candidate, voters care about both the nominating parties' policy platforms and candidates' quality or valence. The policy platforms summarize the legislative objectives that each party or coalition hopes to achieve and that their candidates are expected to support if elected.[10] Candidate valence, on the other hand, comprises individual attributes that voters in a district may find appealing, such as charisma, intelligence, or trustworthiness, and it may be interpreted as perceived ability to represent the district's interests in legislative bargaining. Lastly, voters may also be swayed by campaign efforts in support of a candidate.

Given menu of candidates $M_d = m \in \{M^{\text{PRI}}, M^{\text{PVEM}}, M^{\text{IND}}\}$, voter i's utility from selecting candidate $j \in m$ in district d takes the form

$$\begin{aligned} u_{ijd}^{m} &= \underbrace{\alpha_{1i} c_{jd} + \alpha_{2i} c_{jd}^{2}}_{\text{effect of campaign spending}} + \underbrace{(x_{jd}^{m})'\beta}_{\text{baseline partisanship}} \\ &\quad + \underbrace{\xi_{jd}^{m}}_{\text{candidate valence}} + \underbrace{\epsilon_{ijd}^{m}}_{\text{partisanship shock}}, \end{aligned} \tag{1}$$

[10]No consecutive reelection and strict limits on campaign finance have fostered very strong parties in Mexico. Frey, López-Moctezuma and Montero (2023) document a striking degree of party unity in roll-call voting in the Chamber of Deputies.

where c_{jd} denotes campaign spending in support of candidate j, x_{jd}^m is a vector of observed menu-party-district characteristics, ξ_{jd}^m measures baseline candidate valence, and ϵ_{ijd}^m is an idiosyncratic random-utility or partisanship shock that is independent of the other components of i's payoff. As is standard in discrete-choice models, ϵ_{ijd}^m is assumed to be independently distributed according to the mean-zero Type-I Extreme Value (TIEV) distribution.

The influence of policy platforms on local voting preferences is captured through interactions between district demographics and party or coalition fixed effects. Additionally, vector x_{jd}^m includes menu fixed effects as well as party or coalition j's (logged) vote share in district d from the 2009 Chamber of Deputies election. Thus, the term $(x_{jd}^m)'\beta$ measures the baseline relative appeal—with respect to other available candidates—of j's platform for the electorate in district d.

Candidate valence, ξ_{jd}^m, is observed by voters but not by the researcher. Campaign expenditures in support of candidate j, c_{jd}, may then boost her baseline attractiveness given her platform and individual characteristics. Coefficients α_{1i} and α_{2i} determine the effectiveness of campaign efforts and are voter-specific to capture heterogeneity in impressionability across voters. Specifically, for $k = 1, 2$,

$$\alpha_{ki} = \alpha_k + \sigma_k \nu_{ki},$$

where $\nu_{ki} \overset{\text{i.i.d.}}{\sim} N(0, 1)$ is an idiosyncratic impressionability shock, and (α_k, σ_k) are parameters to be estimated. The quadratic term $\alpha_{2i} c_{jd}^2$ is introduced to allow for diminishing marginal returns to spending.

Because all idiosyncratic shocks to individual voter preferences have mean zero,

$$\delta_{jd}^m = \alpha_1 c_{jd} + \alpha_2 c_{jd}^2 + (x_{jd}^m)'\beta + \xi_{jd}^m \tag{2}$$

represents mean voter utility from selecting candidate $j \in m$ in district d. In every menu, in addition to the four or five competing candidates, voters also have available a compound outside option, $j = 0$, of either abstaining, casting a null vote, or writing in the name of an unregistered candidate. The mean utility of the outside option is normalized to zero, $\delta_{0d}^m = 0$. This normalization is without loss of generality given that only relative preferences can be identified in discrete-choice models. Furthermore, imposing a common normalization provides a shared baseline against which to interpret any menu effects. In particular, voters are assumed to behave expressively—i.e., they choose the candidate they prefer the most, disregarding strategic considerations. The overall reward structure of the electoral system—specifically, the proportional allocation of seats and future funding—arguably encourages sincere voting and warrants this assumption. Nevertheless, the menu-dependent structure of voters' preferences heuristically allows for potentially strategic responses to changes in the electoral supply.

The second tier of the voting stage—i.e., the choice of how to split the PR vote when supporting a coalition candidate—takes a similar form. If $M_d = m \neq M^{\text{IND}}$, then PRI-PVEM coalition supporters must decide whether to split their vote equally between the coalition partners or give 100% of their vote to one of them, where voter i's utility from choosing alternative j out of these three options is

$$u_{ijd}^{\text{ST},m} = (x_{jd}^m)'\beta^{\text{ST}} + \xi_{jd}^{\text{ST},m} + \epsilon_{ijd}^{\text{ST},m}. \tag{3}$$

In this case, the 50-50 split is taken as the outside option, and $\epsilon_{ijd}^{\text{ST},m}$ is again assumed to be independently distributed mean-zero TIEV. The only difference between the two voting tiers is that the second tier is unaffected by campaign spending. Campaigns are candidate-centric and, as such, are assumed to affect only the first-tier candidate choice.

Campaign stage. This stage follows the coalition formation stage and corresponding candidate nominations. The objective for all *five* parties participating in the election is to decide how much to spend in support of their registered candidates. Given impressionability shocks $\nu_i = (\nu_{1i}, \nu_{2i})$ and menu $M_d = m$, determined in the coalition formation stage, parties anticipate that voter i in district d will select candidate $j \in m$ with probability

$$P_{jd}^m(\nu_i) = \frac{\exp(\delta_{jd}^m + \sum_{k=1}^2 \sigma_k \nu_{ki} c_{jd}^k)}{1 + \sum_{j' \in m} \exp(\delta_{j'd}^m + \sum_{k=1}^2 \sigma_k \nu_{ki} c_{j'd}^k)}.$$

This conditional choice probability takes the familiar multinomial logit form due to the TIEV distribution of partisanship shocks. By a law of large numbers approximation (over 185,000 registered voters per district), candidate j's vote share in district d can then be written as

$$s^m_{jd} = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \frac{\exp(\delta^m_{jd} + \sum_{k=1}^{2} \sigma_k \nu_{ki} c^k_{jd})}{1 + \sum_{j' \in m} \exp(\delta^m_{j'd} + \sum_{k=1}^{2} \sigma_k \nu_{ki} c^k_{j'd})} \phi(\nu_{1i})\phi(\nu_{2i}) \mathrm{d}\nu_{1i} \mathrm{d}\nu_{2i}, \tag{4}$$

where ϕ denotes the standard Normal probability density function.

At this stage, candidates' valence terms, $(\xi^m_{jd})_{j \in m}$, are commonly observed by all parties (but not the researcher). Party leaders can therefore tailor their spending across districts to their candidates' relative strengths. For convenience, I assume that parties face a flexible national budget constraint. In particular, parties make independent spending decisions across districts.[11] The estimation strategy described below ensures that the spending levels predicted by the model conform to the levels observed in the data. But, rather than imposing a hard national budget constraint, which would considerably complicate the analysis, the model allows for certain flexibility with respect to parties' total spending under counterfactual scenarios. This assumption is not unreasonable, particularly for the 2012 election, which coincided with the senate and presidential contests. Indeed, parties are free to transfer resources between elections.[12] While the senate and presidential contests are outside the scope of this paper, any opportunity costs of such transfers are implicitly captured by the payoff structure described next.

Given $j \notin \{\text{PRI}, \text{PVEM}\}$ or $M_d = m = M^{\text{IND}}$ (i.e., j is not a coalition candidate), j's party's payoff in district d takes the form

$$\pi^m_{jd} = \gamma_j \log\left(s^m_{jd}\right) - c_{jd}. \tag{5}$$

Parties have a clear incentive to maximize their vote share in each district: it improves their

[11] A potential threat to this assumption is the possibility that spending in one race may have spillover effects on nearby districts. I explicitly test this below and find no evidence of significant spillovers.

[12] In fact, the (victorious) 2012 PRI-PVEM presidential candidate was accused of using expenditures in the Chamber of Deputies election as a way of skirting presidential campaign spending limits. In Online Appendix D, I show that results below are robust to this potential source of measurement error in the relevance of observed Chamber of Deputies expenditures.

chance of winning the corresponding district seat as well as their overall share of PR seats and future funding. However, campaigns are costly. Coefficient γ_j parsimoniously captures the monetary "value" of (the log of) s^m_{jd}, which should reflect both the magnitude of the party's available resources and any opportunity costs of c_{jd}. In equilibrium, j's party will equalize the marginal return of increasing its vote share in district d with an additional dollar of spending to its marginal cost (one). Since money is the numeraire, this ensures parties will equalize marginal returns across districts, optimally distributing resources.[13]

For the coalition partners, given $p \in \{\text{PRI}, \text{PVEM}\}$ and $M_d = m \neq M^{\text{IND}}$, party p's vote share can be written as

$$s^m_{pd} = \left[\frac{\exp(\delta^{\text{ST},m}_{pd}) + 0.5}{1 + \sum_{p' \in \{\text{PRI},\text{PVEM}\}} \exp(\delta^{\text{ST},m}_{p'd})} \right] s^m_{jd},$$

where s^m_{jd} denotes coalition candidate j's FPTP vote share; $\delta^{\text{ST},m}_{pd} = (x^m_{pd})'\beta^{\text{ST}} + \xi^{\text{ST},m}_{pd}$, as in the case of first-tier mean utilities; and the term in brackets represents p's share of the coalition's PR votes, adding half of the 50-50 splits to the share of coalition supporters who allocate their PR vote to party p. Letting c_{pd} denote spending by p in support of candidate j, with $c_{jd} = c_{\text{PRI},d} + c_{\text{PVEM},d}$, p's payoff is given by

$$\begin{aligned} \pi^m_{pd} &= \gamma_p \log\left(s^m_{pd}\right) - c_{pd} \\ &= \gamma_p \log\left(\frac{\exp(\delta^{\text{ST},m}_{pd}) + 0.5}{1 + \sum_{p' \in \{\text{PRI},\text{PVEM}\}} \exp(\delta^{\text{ST},m}_{p'd})} \right) + \gamma_p \log\left(s^m_{jd}\right) - c_{pd}. \end{aligned} \tag{6}$$

Note that the first term in Equation (6) doesn't depend on c_{pd}, and ultimately joint spending by the two partners, c_{jd}, determines the coalition candidate's prospects. I remain agnos-

[13] In Online Appendix B, I show that the payoff coefficients from this independent-games model coincide with the ratios of these coefficients to the Lagrange multiplier of the budget constraint from a hard-constraint model. The two models are in this sense observationally equivalent. Moreover, as previewed, spending in the counterfactuals described below is predicted to *decrease.*

tic about how PRI and PVEM divide this amount between them and simply assume that it maximizes their joint surplus $\pi^m_{\text{PRI},d} + \pi^m_{\text{PVEM},d}$, which, up to a constant (in c_{jd}), equals $(\gamma_{\text{PRI}} + \gamma_{\text{PVEM}}) \log\left(s^m_{jd}\right) - c_{jd}$. Thus, given $M_d = m \neq M^{\text{IND}}$, PRI and PVEM act as a single player in the spending game in district d against the other three parties, who chooses $c_{\text{PRI},d} + c_{\text{PVEM},d}$ with joint payoff $\pi^m_{\text{PRI},d} + \pi^m_{\text{PVEM},d}$.[14]

At the estimated parameter values reported below, and regardless of the menu of candidates, the resulting campaign spending game played in each district exhibits strict strategic complementarities—see Online Appendix B for details. A formal definition of this class of games can be found in Echenique and Edlin (2004). It suffices here to point out three key properties. First, existence of equilibrium is guaranteed. Second, mixed-strategy equilibria are not good predictions in these games, so their omission is justified (Echenique and Edlin, 2004). Third, the set of all pure-strategy equilibria can be efficiently computed (Echenique, 2007). This implies that consideration of potential multiplicity of equilibria is feasible. At the estimated parameter values, however, the campaign spending games exhibit unique equilibria. Therefore, for ease of exposition, I proceed with the description of the model and the empirical strategy under the presumption that the spending game in each district has a unique equilibrium.

Coalition formation stage. This stage completes the description of the model and focuses on PRI and PVEM's optimal choice of coalition configuration—i.e., where to run together and the party affiliation of coalition candidates. As noted, coalition formation decisions precede the candidate selection process. I assume here that, although national party leaders may have information about the identities of potential candidates, relevant electability traits are

[14] I refrain from explicitly modeling bargaining over joint spending, both to avoid overcomplicating the analysis and because only joint spending in support of coalition candidates is observed in the data. Assuming joint-surplus maximization ensures a Pareto optimal agreement for the coalition, which, given repeated interaction in diverse settings, can be plausibly supported with abundant opportunities for indirect compensation.

only fully realized in the campaign stage as competing candidates are contrasted with each other. In other words, at the coalition formation stage, party leaders don't yet know the exact candidate valence profiles, $(\xi_j^m)_{j \in m}$, that would result from each menu choice, only their distribution: $\xi_j^m \overset{\text{i.i.d.}}{\sim} N(0, \varsigma_j)$.[15]

Similarly to joint spending decisions, M_d is chosen in each district to maximize PRI and PVEM's ex-ante expected joint surplus. Given $M_d = m = M^{\text{IND}}$, party p's ex-ante payoff is simply $\bar{\pi}_{pd}^m = E[\pi_{pd}^m]$, where the expectation is taken with respect to the realized candidate valence profile in the district and corresponding campaign spending equilibrium, and π_{pd}^m is defined by Equation (5). On the other hand, given $M_d = m \neq M^{\text{IND}}$, p's ex-ante payoff is

$$\bar{\pi}_{pd}^m = (w_{pd}'\theta)\mathbf{1}_{j \neq p} + E[\pi_{pd}^m], \tag{7}$$

where j denotes the party affiliation of the coalition candidate, π_{pd}^m is defined by Equation (6), and $w_{pd}'\theta$ captures the (dis)utility for party p from not fielding a candidate in district d and thus forgoing the FPTP seat. Vector w_{pd} includes interactions between district characteristics and party fixed effects as well as, importantly, p's ex-ante probability of winning the FPTP race in district d when running independently. This probability captures the natural exit option for p from the coalition agreement and quantifies one of the key tradeoffs in this environment.

I allow the coalition partners to experience an idiosyncratic shock, η_d^m, to their joint ex-ante surplus from selecting menu m in district d. These shocks are assumed to be independently distributed mean-zero TIEV and observed by the partners but not the researcher. They allow for unobserved (by the researcher) district-specific considerations affecting bargaining that are

[15] This assumption is further supported by the relative inexperience of FPTP candidates, especially compared to PR candidates. As summarized in Table A6 in Online Appendix A, only 12% of 2012 FPTP candidates had participated in a federal legislative election in *any* of the three previous cycles (2003, 2006, or 2009).

not otherwise captured by the model with the data available.[16] PRI and PVEM then choose their coalition configuration $(M_d)_{d=1}^{300}$ to maximize $\sum_d \bar{\pi}^{M_d}_{\text{PRI},d} + \bar{\pi}^{M_d}_{\text{PVEM},d} + \eta^{M_d}_d$.

Empirical Strategy

The estimation strategy mirrors the model's three-stage structure. Step 1 recovers the voting-stage parameters in Equations (1) and (3) following the aggregate discrete-choice approach to demand estimation popularized by Berry, Levinsohn and Pakes (BLP, 1995). Step 2 obtains payoff coefficient γ_p for each party p by matching the spending levels observed in the data with the model's predictions from the campaign stage. Finally, ex-ante coalition surplus maximization is exploited in Step 3 to recover θ, which characterizes the partners' (dis)utility from not fielding a candidate.

Step 1. To emphasize the intuition, consider first the case where voters are homogeneous up to their idiosyncratic partisanship shocks—i.e., $\sigma = (\sigma_1, \sigma_2) = 0$. Taking logs of Equation (4) and replacing predicted vote shares with their observed counterparts in the data, $\hat{s}^{M_d}_{jd}$, yields the linear demand system:

$$\log(\hat{s}^{M_d}_{jd}) - \log(\hat{s}^{M_d}_{0d}) = \delta^{M_d}_{jd} = \alpha_1 c_{jd} + \alpha_2 c^2_{jd} + (x^{M_d}_{jd})'\beta + \xi^{M_d}_{jd}.$$

This is just a linear regression of the log ratio of candidate j's vote share to that of the outside option on endogenous (c_{jd}) and exogenous ($x^{M_d}_{jd}$) covariates, where the candidate's unobserved valence, $\xi^{M_d}_{jd}$, corresponds to the residual of the regression.[17] (The second-tier

[16] In Online Appendix C, I discuss an alternative formulation of the coalition formation stage without idiosyncratic bargaining shocks. That model only partially identifies an average (dis)utility from not fielding a candidate, but results are consistent across the two specifications.

[17] Assuming coalition partners don't observe candidate valence profiles when selecting M_d ensures exogeneity of $x^{M_d}_{jd}$. To see this, note that, if M_d is independent of ξ^m_{jd} for all m, then by iterating expectations: $E[x^{M_d}_{jd}\xi^{M_d}_{jd}] = \sum_m E[(\mathbf{1}_{M_d=m})x^m_{jd}\xi^m_{jd}] = \sum_m E[(\mathbf{1}_{M_d=m})x^m_{jd}E[\xi^m_{jd}|x^m_{jd}, M_d]] = 0$.

voting parameters can be recovered analogously.)

Instrumental variables are required to tackle the endogeneity of c_{jd} and identify α_1 and α_2. I exploit the prohibition on consecutive re-election and scarcity of repeat candidates (less than 4%—see Table A6 in Online Appendix A) to instrument for c_{jd} using campaign spending data from the 2009 Chamber of Deputies election. Expenditures in a district are likely to be correlated over time due to persistent features of the electoral environment as well as unobserved determinants of campaign costs (local media-market prices, ease of transportation, etc.). However, after controlling for local partisanship with observed district characteristics—in particular, vote shares in the 2009 election—lagged spending should have no direct effect on 2012 election outcomes. Thus, spending by j in district d in 2009 (and its square) should serve as a valid instrument for c_{jd} (and c_{jd}^2).[18] Coefficients $\alpha = (\alpha_1, \alpha_2)$ and β can then be estimated via two-stage least squares.

With heterogenous voter impressionability ($\sigma \neq 0$), this simple linear regression approach is no longer feasible, but BLP show that a Generalized Method of Moments (GMM) estimator of $\varphi = (\alpha, \beta, \sigma)$ can be constructed based on similar intuition. A detailed description of this estimator is relegated to Online Appendix C.

Step 2. Coefficients $\gamma = (\gamma_{\text{MP}}, \gamma_{\text{NA}}, \gamma_{\text{PVEM}}, \gamma_{\text{PRI}}, \gamma_{\text{PAN}})$ driving party leaders' campaign spending decisions are estimated by ensuring predicted and observed campaign spending levels concur. With a slight abuse of notation, let $\hat{c}_{jd}$ denote observed spending in support of candidate j in district d, let $c_{jd}(\gamma, \varphi)$ denote the corresponding equilibrium spending predicted by the model given parameters (γ, φ),[19] and let the vector z_{jd} collect the campaign spending instruments and exogenous covariates used in Step 1. These should satisfy the moment

[18]In Online Appendix D, I show results are robust (though less precise) to alternative (but weaker) choices of instruments.

[19]Predicted spending $c_{jd}(\gamma, \varphi)$ is computed as party or coalition j's best response to the observed spending of other parties in district d given (γ, φ).

condition

$$E\left[z_{jd}\left(\hat{c}_{jd} - c_{jd}(\gamma, \varphi)\right)\right] = 0 \quad \text{if and only if} \quad (\gamma, \varphi) = (\gamma_0, \varphi_0), \tag{8}$$

where (γ_0, φ_0) denotes the true value of the parameters. Letting $\hat{\varphi}$ denote the coefficient estimates from Step 1, a GMM estimator of the campaign-stage parameters can be obtained by minimizing the quadratic form $Q_N^{\text{CS}}(\gamma) = \left[\frac{1}{N} Z'\left(\hat{C} - C(\gamma, \hat{\varphi})\right)\right]' W_N \left[\frac{1}{N} Z'\left(\hat{C} - C(\gamma, \hat{\varphi})\right)\right]$, where Z, $\hat{C}$, and $C(\gamma, \hat{\varphi})$ are vertical stackings of z'_{jd}, $\hat{c}_{jd}$, and $c_{jd}(\gamma, \hat{\varphi})$ across candidates and districts, N denotes the total number of observations (candidate-districts), and $\frac{1}{N} Z'\left(\hat{C} - C(\gamma, \hat{\varphi})\right)$ is the sample analog of moment condition (8). Inference follows standard GMM theory, but standard errors must be adjusted to account for uncertainty in the first-step estimates, $\hat{\varphi}$. See Online Appendix C for details.

Step 3. Finally, θ can be estimated via Maximum Likelihood (ML) by exploiting the optimality of PRI and PVEM's observed coalition configuration. Since the coalition's idiosyncratic bargaining shocks are distributed TIEV, the likelihood of observing M_d in district d is given by

$$\mathcal{L}_d(M_d; \theta, \gamma, \varphi) = \frac{\exp\left(\bar{\pi}^{M_d}_{\text{PRI},d}(\theta, \gamma, \varphi) + \bar{\pi}^{M_d}_{\text{PVEM},d}(\theta, \gamma, \varphi)\right)}{\sum_m \exp\left(\bar{\pi}^{m}_{\text{PRI},d}(\theta, \gamma, \varphi) + \bar{\pi}^{m}_{\text{PVEM},d}(\theta, \gamma, \varphi)\right)},$$

where $\bar{\pi}^m_{pd}(\theta, \gamma, \varphi)$ denotes party p's ex-ante payoff under menu m given $(\theta, \gamma, \varphi)$.[20] Having estimated $\hat{\varphi}$ and $\hat{\gamma}$ in Step 1 and Step 2, respectively, the remaining model parameters, θ, can be recovered by maximizing the log-likelihood

$$\sum_d \log\left[\mathcal{L}_d(M_d; \theta, \hat{\gamma}, \hat{\varphi})\right].$$

Again, standard errors must be adjusted to account for estimation uncertainty in $(\hat{\gamma}, \hat{\varphi})$. See Online Appendix C for technical details.

[20]Computation of $\bar{\pi}^m_{pd}(\theta, \gamma, \varphi)$ is via simulation, drawing across 10,000 trials a candidate valence profile, $(\xi^m_{jd})_{j \in m}$, and then calculating the corresponding campaign spending equilibrium (see Online Appendix B) and ensuing election outcomes in the district. As $\bar{\pi}^m_{pd}(\theta, \gamma, \varphi)$ is linear in θ and simulation draws depend solely on (γ, φ), they need be drawn only once throughout the ML search, which considerably lowers the computational burden.

Estimation Results

Discussion of the main coefficient estimates follows the structure of the model, beginning with the voting stage. I then turn to the counterfactual experiments at the core of this paper, which quantify the tradeoffs entailed by and consequences of the PRI-PVEM electoral alliance.

Estimates of voters' preferences. For a range of model specifications, Table A7 in Online Appendix A presents estimates of the coefficients characterizing candidate choice in the voting stage. Table A8 reports analogous estimates for the second-tier choice for coalition supporters of how to allocate their PR vote. Models in columns (I)–(III) of Table A7 all include menu-party fixed effects, while those in columns (IV)–(VI) additionally include electoral region fixed effects. Columns (I) and (IV) show ordinary least squares (OLS) estimates from the homogeneous-voters version of the model with $\sigma = 0$. Columns (II), (III), (V), and (VI) report BLP estimates from the version with heterogeneous voter impressionability.

Estimates of the baseline partisanship coefficients, β, are consistent across specifications and in line with well-known historical patterns in Mexico. Districts with a higher share of voters over 60 years old tend to support the traditional establishment parties, PRI and PAN—especially the more conservative PAN. Districts with a higher share of female heads of household tend to favor the left-wing MP, which was the first to decriminalize abortion, in Mexico City. And PRI and the green party, PVEM, are advantaged in more rural districts. With regard to the analogous second-tier coefficients, β^{ST}, 50-50 PR vote splits are more likely in districts with a higher share of female heads of household and less likely in more rural districts or those with more voters over 60. Among coalition supporters who give 100% of their PR vote to one of the two partners, rural voters weakly favor PVEM, and older voters strongly favor PRI. Lastly, partisanship is estimated to be highly persistent (Frey, López-Moctezuma and Montero, 2023) given positive and extremely precise coefficient estimates for log-lagged vote share, robust across all specifications.

Campaign spending has a positive effect on candidates' vote shares, with diminishing marginal returns. Here, however, important differences between the OLS and BLP estimates emerge. As expected, the OLS estimates are an order of magnitude smaller (in absolute terms) than the BLP estimates. This reflects the nature of the endogeneity problem that the campaign-spending instruments described above are intended to address: since parties allocate resources strategically, targeting competitive races where marginal returns are higher, OLS estimates should underestimate the effectiveness of campaign expenditures given that (i) outstanding candidates are likely to receive comparatively less support, while (ii) parties nevertheless benefit from underperforming candidates' PR vote shares. This is borne out in the results. For a candidate with an average vote share (23%) and average spending (45,000 USD), the OLS estimates in column (I) indicate that a 1% increase in campaign spending would raise her vote share by about 0.16%. In contrast, the BLP estimates in column (II) imply a corresponding vote share increase of 0.9%. Figure A4 in Online Appendix A further illustrates the substantive implications of the latter, as well as considerable heterogeneity across parties, by plotting candidates' expected vote shares across the full range of observed spending by their party. On average, parties can almost triple a candidate's vote share by increasing spending from its lowest observed level to the highest, but this is heavily moderated by spending by rival parties, which underscores the importance of accounting for strategic incentives when analyzing parties' spending decisions.

Accordingly, the campaign stage of the model features party leaders playing independent campaign spending games across districts. Of particular concern for this assumption is the potential for campaign spillovers across neighboring districts. To test this, columns (III) and (VI) present estimates from alternative specifications of Equation (1) that add a term, $\alpha_3\bar{c}_{jd}$, capturing the effect of j's average spending in neighboring districts, $\bar{c}_{jd}$, on j's vote share in district d.[21] In both columns, coefficient α_3 is very precisely estimated, an order of magnitude smaller than the linear direct effect of spending (α_1), and statistically insignificant, providing no evidence of meaningful campaign spillovers.

Finally, across specifications, there is no evidence of heterogeneous voter impressionability.

[21] I use average lagged spending in neighboring districts to instrument for $\bar{c}_{jd}$.

As discussed in Online Appendix C, however, coefficients (σ_1, σ_2) are notoriously hard to identify in practice, and the estimates in Table A7 are all close to zero but extremely imprecise. Gillen et al. (2019) alternatively explore the robustness of these results to data-driven selection of demographic controls using machine learning. Regardless of model-selection approach or specification, BLP estimates are in agreement: campaign expenditures significantly improve candidates' electoral prospects, although with decreasing marginal returns, and there is little evidence in this context of heterogeneity in voters' responsiveness to campaign efforts. Due to the computational cost of estimating the remaining stages of the model and of simulating counterfactuals, I rely only on the specification in column (II) for what follows.

Estimates of parties' campaign-stage payoffs. Table 1 reports estimates of γ, the coefficients driving parties' campaign spending decisions. All coefficients are statistically significant, and their magnitudes mainly reflect the relative sizes (and thus resources) of parties. Despite the simplicity of parties' payoff specifications in Equations (5) and (6), the campaign stage of the model—together with the voting-stage estimates—is remarkably successful at fitting observed spending levels, with an adjusted R^2 of 0.67. This suggests that strategic targeting of resources to maximize electoral returns is indeed the primary consideration behind parties' campaign spending decisions. Nevertheless, there is residual variation in spending not explained by the model—due to unobserved factors such as campaign costs—which lends support to the lagged-spending instruments used in Step 1.

Table 1: Estimates of Parties' Campaign-Stage Payoffs

γ_{MP}	γ_{NA}	γ_{PVEM}	γ_{PRI}	γ_{PAN}
5.549 (1.161)	2.892 (0.616)	2.926 (0.620)	5.664 (1.167)	4.315 (0.917)

Notes. GMM estimates of coefficients driving parties' campaign spending decisions—see Equations (5) and (6)—with robust standard errors in parentheses.

Estimates of parties' coalition-stage payoffs. Table 2 shows estimates of the coefficients, θ, that characterize—as defined by Equation (7)—PRI and PVEM's (dis)utility from standing down in a district to support their coalition partner's candidate. Column (I) corresponds to a stripped-down specification wherein vector w_{pd} includes just party fixed effects as well as p's ex-ante probability of winning the FPTP race in district d running independently. The specification in column (II) features party-region fixed effects and adds interactions between district demographics and party dummies.

Notably, the coefficient on the ex-ante probability of winning is negative, statistically significant, and consistent across specifications. This reflects one of the fundamental tradeoffs in this environment: PRI and PVEM are naturally more reluctant to stand down in a district if they have a good chance of winning it on their own. Controlling for other district characteristics makes little difference. All other coefficients in column (II) are statistically insignificant, with the exception that PRI is less likely to stand down in rural districts. This may due to the higher price PRI pays in rural districts in the PR component of the election—see Table A8—or to a desire to preserve historical dominance (Ames, 1970; Magaloni, 2006; Frey, López-Moctezuma and Montero, 2023). Likewise, although not statistically significant, the only positive coefficient in Table 2 suggests PRI is relatively more willing to jointly nominate a PVEM candidate in districts with a larger share of the electorate over 60, where the PR cost for PRI is lower. These estimates thus illuminate the key tradeoffs faced by the coalition partners. Next, with the estimated model parameters in hand, I quantify these tradeoffs explicitly as well as the overall electoral impact of the PRI-PVEM alliance. I use the specification in column (II) of Table 2 to simulate the coalition formation stage, which, despite its parsimony and the multinomial nature of the data, correctly predicts close to 60% of observed coalition choices.[22]

Counterfactuals

I conduct two counterfactual experiments. First, I explore what would have happened in the 2012 Chamber of Deputies election had PRI and PVEM not formed a coalition. That is, I simulate election outcomes (averaged over 10,000 trials) imposing $M_d = M^{\text{IND}}$ in all districts where PRI and PVEM nominated a joint coalition candidate. Second, at the other extreme,

[22]The specification in column (I) correctly predicts 52% of observed coalition choices.

Table 2: Estimates of Parties' Coalition-Stage Payoffs

	(I)	(II)
Prob. of Winning FPTP Race Alone	-4.919	-4.563
	(2.734)	(2.674)
PVEM × Female		-1.370
		(4.946)
PVEM × Over 60		-6.841
		(5.484)
PVEM × Rural		-1.048
		(0.678)
PRI × Female		-1.826
		(7.495)
PRI × Over 60		0.837
		(7.300)
PRI × Rural		-3.626
		(1.140)
Party F.E.	Yes	No
Party-Region F.E.	No	Yes
Log-Likelihood	-285.8	-258.1
Observations	300	300

Notes. ML estimates of θ, which characterizes PRI and PVEM's (dis)utility from standing down in a district to support their partner's candidate as defined by Equation (7), with standard errors in parentheses.

I examine the effects of constraining PRI and PVEM to form a *total* coalition. For this scenario, in all districts where PRI and PVEM ran independently, I force PRI and PVEM to run together by restricting the choices available to them in the coalition formation stage of the model to $M_d \in \{M^{\text{PRI}}, M^{\text{PVEM}}\}$. Thus, PRI and PVEM are constrained to run together, but they optimally select the party affiliation of their coalition candidates, which results in 252 total districts with a joint PRI candidate and 48 districts with a joint PVEM candidate.

Before turning to aggregate election results, I first take a district-level look at the key tradeoffs the coalition partners faced when designing their coalition configuration. For each electoral district, I simulate PRI and PVEM's ex-ante expected equilibrium spending and vote shares across the three possible menu choices, $m \in \{M^{\text{PRI}}, M^{\text{PVEM}}, M^{\text{IND}}\}$. To summarize the main patterns that emerge, I revisit Table A2, adding model predictions of the two partners' counterfactual vote shares across menus. The first three columns of Table 3 average these vote shares over districts where the two partners decided to run independently in 2012. The next three columns do so over districts where they jointly nominated a PRI candidate. And the last three correspond to districts with joint PVEM nominations. The model's ex-ante predictions approximate observed vote shares in Table A2 remarkably well: both partners' performance is only systematically overestimated in the case of a joint PVEM nomination. This could be due to estimation error or to unfavorable valence draws for the coalition in the data, considering the relatively small number of districts (43).

Table 3: Coalition Partners' Counterfactual Ex-Ante Expected District Vote Shares

	Districts with Distinct PRI, PVEM Candidates			Districts with Joint PRI Candidate			Districts with Joint PVEM Candidate		
	Distinct Candidates (%)	Joint PRI Candidate (%)	Joint PVEM Candidate (%)	Distinct Candidates (%)	Joint PRI Candidate (%)	Joint PVEM Candidate (%)	Distinct Candidates (%)	Joint PRI Candidate (%)	Joint PVEM Candidate (%)
PVEM	4.3	9.3	10.2	4.0	8.9	9.6	4.4	9.0	9.6
PRI	37.2	35.8	38.1	34.3	33.0	34.5	29.6	28.7	30.0

Notes. Ex-ante expected vote shares are computed via simulation, drawing across 10,000 trials a candidate valence profile, $(\xi_{jd}^{m})_{j\in m}$, and then calculating the corresponding campaign spending equilibrium (see Online Appendix B) and ensuing district election outcomes. The first three columns average vote shares over districts where PRI and PVEM ran independently in 2012. The next three columns do so over districts where they jointly nominated a PRI candidate. The last three correspond to districts with joint PVEM nominations.

Table 3 corroborates many of the lessons from Table A2, yet it underscores the importance of accounting for PRI and PVEM's strategic choice of coalition configuration. Adding together the two parties' vote shares, the advantage of coalition candidates relative to independent candidates in the FPTP races remains, but the apparent net loss of votes to rival parties suggested by Table A2 is overturned. Clearly, the districts where PRI and PVEM chose to run together were more competitive, and this is what is reflected in Table A2, but counterfactual joint vote shares are systematically higher with a joint candidacy than with independent candidates. This may be thanks to pooled strengths or resources not captured explicitly by campaign expenditures or to coalition candidates more easily standing out in a less crowded field. Relatedly, joint PVEM candidates outperform counterfactual joint PRI candidates, which is consistent with clearer ideological differentiation relative to their closest strong rival, PAN (see Figure 2).

The counterfactual vote shares in Table 3 also confirm that, for PVEM, joining forces with PRI was unambiguously beneficial: the party had a marginal chance of winning FPTP races on its own, and its vote share rose considerably under joint nominations, regardless of the party affiliation of coalition candidates. For PRI, on the other hand, joining forces with PVEM entailed a clear tradeoff: either forgoing a district seat in the case of a joint PVEM nomination or paying a cost in the PR component of the election in the case of a joint PRI nomination. Accordingly, Table 3 reveals that PRI (i) ran alone where it was strongest, (ii) shored up its FPTP chances in more competitive districts with a joint PRI nomination, and (iii) only agreed to a joint PVEM nomination in the toughest races.

As previewed, with regard to campaign expenditures, there is virtually no difference in joint spending by the two partners when nominating distinct candidates versus a joint PRI nomination. Equilibrium spending, however, is on average 2.33% lower with a joint PVEM candidate. This is again consistent with increased ideological differentiation from, and thus less intense campaign competition with, their closest strong rival, PAN. Although bargaining over joint spending is not explicitly modeled or observed, these campaign savings constitute a clear potential benefit for PRI in exchange for agreeing to a joint PVEM nomination.

Turning now to aggregate election results, Figure 3 presents, relative to a benchmark counterfactual scenario with no PRI-PVEM coalition, absolute changes in all parties' vote shares (top panel), FPTP seats (middle), and PR seats (bottom) as a consequence of either the observed partial PRI-PVEM coalition configuration (in green) or a counterfactual total PRI-PVEM coalition (yellow). Results are in line with the district-level intuition from Table 3. In both cases, PVEM's vote share benefits considerably from joint nominations at the expense of PRI's (and other parties'). By not splitting the vote, joining forces allows both coalition partners to increase their share of FPTP seats. But only PVEM benefits in the PR component of the election.

Overall, by running as observed in the data, both coalition partners enjoyed net seat gains in the chamber despite PRI's PR-seat losses. In terms of jointly held seats, they managed to close the gap to obtaining a legislative majority (251 seats) by 78%—from 45 seats to 10. And they would have closed it by 89%—with five additional seats—had they run together in

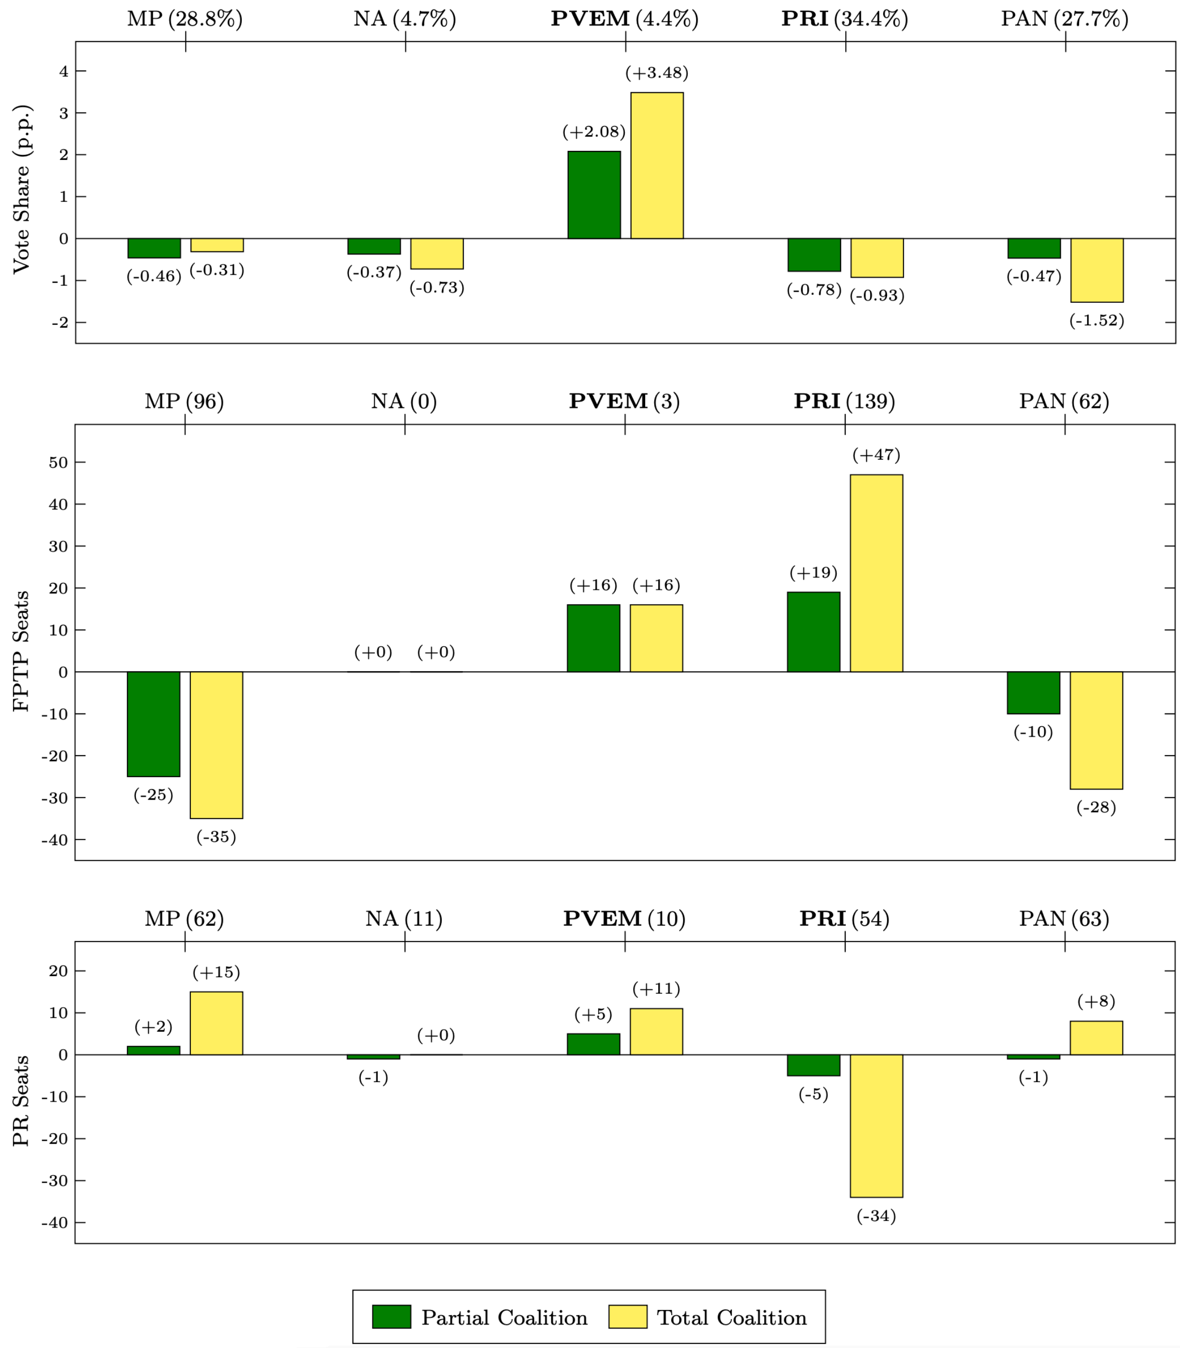

Figure 3: Electoral Impact of PRI-PVEM Alliance Relative to No-Coalition Benchmark

Notes. Bar plots of absolute changes in vote share (top), FPTP seats (middle), and PR seats (bottom), by party, relative to a benchmark counterfactual scenario with no PRI-PVEM coalition. Above each panel, benchmark values are shown in parentheses. Changes in outcomes as a result of the observed partial PRI-PVEM coalition are presented in green (change values in parentheses). Changes resulting from a counterfactual total PRI-PVEM coalition are in yellow (change values in parentheses).

all districts. However, whereas PVEM's observed net gain of 21 seats would have increased to 27 with a total coalition, PRI's observed net gain of 14 seats would have decreased by one had they joined forces everywhere. This is due to the disproportionality restriction of the electoral system described above. Reassuringly, although for tractability the restriction is not explicitly incorporated into the model,[23] the predicted breakdown of districts in the total-coalition scenario is indeed (approximately) optimal for the partners: while PRI headlined 78% of shared districts in the data, it would have done so in 95% of the remaining districts with a total coalition, which is consistent with mitigating vote share losses that constrain PRI's total share of seats. This suggests joint-surplus maximization in the coalition stage of the model provides a good first-order approximation of the partners' decision problem.

Furthermore, recall that Mexican parties compete in the Chamber of Deputies election to secure not only seats in the chamber but also their share of public funding for the following three years. Given PRI's decreased vote share, the partners' legislative gains with their partial coalition configuration came at the expense of significantly reduced funding for PRI. Based on the portion of public resources, as reported by INE, distributed in 2013, 2014, and 2015 to parties in proportion to their national vote share in the 2012 Chamber of Deputies election, PRI received close to five million USD less in funding than it would have had it not joined forces with PVEM. This amounts, given their 35-seat gain in the chamber, to about 138,000 USD in lost funding per extra seat. For comparison, the partners spent, on average, 80,000 USD per district supporting their candidates. Had they formed a total coalition, their extra five net seats in the chamber—with PRI losing one—would have cost PRI an additional 185,000 USD in lost funding per extra seat. Again, while not explicitly modeled, this suggests joint-surplus maximization in the coalition stage approximates well the key considerations behind

[23]The main challenge is computational. Although it only takes a few seconds to calculate PR seats—taking into account the disproportionality restriction—given a fixed configuration $(M_d)_{d=1}^{300}$ of the PRI-PVEM alliance, there are $3^{300} > 10^{142}$ such configurations, making it prohibitive to enumerate all possible election outcomes.

the observed PRI-PVEM coalition configuration.

With regard to overall campaign expenditures, as previewed, aggregate spending by the two partners was only 0.63% lower in the data than it would have been with no PRI-PVEM alliance. Though modest, these savings, as noted above, likely provided an ancillary incentive for joining forces, particularly in the case of joint PVEM nominations.[24] With a total coalition, however, campaign savings effectively disappear due to the increased share of PRI-headlined districts and—as Figure 3 shows—competition in the FPTP races primarily with PAN, their strongest ideological neighbor.

Discussion

Before concluding more broadly, I briefly discuss the above results in light of some closely related research on Mexican politics. In particular, Spoon and Pulido Gómez (2017) also examine joint PRI-PVEM nominations in Chamber of Deputies elections, although they do not look into the choice of which partner headlines the coalition. Their analysis yields two main findings. First, Spoon and Pulido Gómez identify that joint Chamber of Deputies nominations are more likely in states where the incumbent governor was also a PRI-PVEM candidate. This dovetails with other work on the influence of Mexican governors on federal legislators (Langston, 2010; Kerevel, 2015). In Online Appendix D, I estimate alternative model specifications that allow for potential effects of incumbent governors on both coalition choices and election outcomes. However, I do not find meaningful or robust effects. In addition to important methodological differences, a key distinction between the analysis in this paper and Spoon and Pulido Gómez's is that the latter considers three electoral cycles: 2009, 2012, and 2015. It is certainly possible that the influence of governors on coalition

[24]The potential for financial incentives in electoral coalition formation has received little scholarly attention. In settings where candidates are not publicly funded, these incentives may be even stronger, as coalition partners can share the burdens of fundraising. Moreover, potential donors may be more willing to back coalition candidates with broader support, further prompting parties to coordinate their nominations.

deliberations may rise in mid-term cycles (2009 and 2015). Relatedly, their second main finding is that past performance of PRI-PVEM coalition candidates is highly predictive of future joint nominations. While this does not explain why coalition partners join forces in the first place, it does indicate that forming an alliance is an inherently dynamic process, as further illustrated by a dramatic increase in joint PRI-PVEM nominations from 21% of districts in 2009 to 83% in 2015. This paper admittedly abstracts from dynamic considerations to take a more detailed look at the incentives at play in a given cycle. But coalitions are not mergers, and post-election disagreements among partners are not uncommon. In fact, PVEM abandoned PRI in 2021 and formed an electoral coalition instead with the incumbent president's party. Further research is needed to fully comprehend important dynamics of these alliances (Frey, López-Moctezuma and Montero, 2023; Invernizzi, 2023).

Given growing discontent in many democracies with the political establishment, understanding how these common agreements between—in many cases, unelected—party elites affect the electoral supply and political representation is increasingly pertinent. Moreover, although post-election legislative bargaining is not explicitly considered in this paper, the results are suggestive of the importance of electoral coalition formation as a preliminary stage of the legislative bargaining process. Parties may use electoral coalitions to pre-select and foster legislative partners. I hope the methodological approach in this paper provides guidance for future research on these important questions.

Acknowledgements

I thank Scott Abramson, Federico Echenique, Anderson Frey, Camilo Garcia-Jimeno, Alex Hirsch, Matias Iaryczower, Tasos Kalandrakis, Jonathan Katz, Brenton Kenkel, Nicolas Longuet Marx, Jean-Laurent Rosenthal, Matt Shum, Erik Snowberg, and multiple seminar and conference audiences for valuable comments and suggestions. Previous versions of this paper circulated under the titles "Coalition Formation, Campaign Spending, and Election Outcomes: Evidence from Mexico" and "Going It Alone? An Empirical Study of Coalition Formation in Elections."

References

Ames, Barry. 1970. "Bases of Support for Mexico's Dominant Party." *American Political Science Review* 64(1):153–167.

Ashworth, Scott and Ethan Bueno de Mesquita. 2009. "Elections with platform and valence competition." *Games and Economic Behavior* 67:191–216.

Bandyopadhyay, Siddhartha, Kalyan Chatterjee and Tomas Sjöström. 2011. "Pre-Electoral Coalitions and Post-Election Bargaining." *Quarterly Journal of Political Science* 6(1):1–53.

Berry, Steven, James Levinsohn and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63(4):841–890.

Blais, André and Indridi H. Indridason. 2007. "Making Candidates Count: The Logic of Electoral Alliances in Two-Round Legislative Elections." *Journal of Politics* 69(1):193–205.

Bormann, Nils-Christian and Matt Golder. 2013. "Democratic Electoral Systems around the world, 1946-2011." *Electoral Studies* 32:360–369.

Carroll, Royce and Gary W. Cox. 2007. "The Logic of Gamson's Law: Pre-election Coalitions and Portfolio Allocations." *American Journal of Political Science* 51(2):300–313.

Catalinac, Amy and Lucia Motolinia. 2021. "Geographically Targeted Spending in Mixed-Member Majoritarian Electoral Systems." *World Politics* 73(4):668–711.

Consulta Mitofsky. 2012. "*Geometría electoral en México* (Electoral geometry in Mexico)." Public Opinion Poll, `http://www.consulta.mx/web/images/MexicoOpina/2013/NA_GEOMETRIA_ELECTORAL.pdf` (in Spanish).

Cox, G. W. 1997. *Making Votes Count: Strategic Coordination in the World's Electoral Systems*. Cambridge University Press.

Debus, M. 2009. "Pre-electoral commitments and government formation." *Public Choice* 138:45–64.

Echenique, Federico. 2007. "Finding All Equilibria in Games of Strategic Complements." *Journal of Economic Theory* 135(1):514–532.

Echenique, Federico and Aaron Edlin. 2004. "Mixed equilibria are unstable in games of strategic complements." *Journal of Economic Theory* 118(1):61–79.

Ferrara, Federico and Erik S. Herron. 2005. "Going It Alone? Strategic Entry under Mixed Electoral Rules." *American Journal of Political Science* 49(1):16–31.

Frey, Anderson, Gabriel López-Moctezuma and Sergio Montero. 2023. "Sleeping with the Enemy: Effective Representation under Dynamic Electoral Competition." *American Journal of Political Science* 67(4):915–931.

Gillen, Benjamin J., Sergio Montero, Hyungsik Roger Moon and Matthew Shum. 2019. "BLP-2LASSO for Aggregate Discrete Choice Models with Rich Covariates." *Econometrics Journal* 22(3):262–281.

Golder, Sona N. 2006. *The Logic of Pre-Electoral Coalition Formation*. Columbus: Ohio State University Press.

Heckman, James J. 1979. "Sample Selection Bias as a Specification Error." *Econometrica* 47(1):153–161.

Hortala-Vallve, Rafael, Jaakko Meriläinen and Janne Tukiainen. 2023. "Pre-electoral coalitions and the distribution of political power." *Public Choice* (forthcoming), https://doi.org/10.1007/s11127-023-01109-1.

Iaryczower, Matias and Andrea Mattozzi. 2013. "On the Nature of Competition in Alternative Electoral Systems." *Journal of Politics* 75(3):743–756.

Invernizzi, Giovanna. 2023. "Why do Parties Merge? Electoral Volatility and Long-Term Coalitions." *Journal of Politics* (forthcoming).

Kerevel, Yann P. 2015. "(Sub)national Principals, Legislative Agents: Patronage and Political Careers in Mexico." *Comparative Political Studies* 48(8):1020–1050.

Langston, Joy. 2010. "Governors and "Their" Deputies: New Legislative Principals in Mexico." *Legislative Studies Quarterly* 35(2):235–258.

Magaloni, Beatriz. 2006. *Voting for Autocracy: Hegemonic Party Survival and its Demise in Mexico*. Cambridge: Cambridge University Press.

Powell, Jr., G. Bingham. 2004. "Political Representation in Comparative Politics." *Annual Review of Political Science* 7:273–296.

Spoon, Jae-Jae and Amalia Pulido Gómez. 2017. "Unusual Bedfellows? PRI-PVEM Electoral Alliances in Mexican Legislative Elections." *Journal of Politics in Latin America* 9(2):63–92.

Spoon, Jae-Jae and Amalia Pulido-Gómez. 2020. "Strategic alliances: Red, green, or "watermelon" candidates in Mexican legislative elections." *Party Politics* 26(6):760–769.