Vote Fragmentation and Incumbent Survival: How Electoral Rules Shape Accountability

Laura Chelidonopoulos*

Word count: 10,817

Abstract

Incumbents often hold power even after voters withdraw support. I argue that whether vote loss translates into office loss depends on the extent to which the anti-incumbent vote fragments across opposition candidates under electoral rules that do not require voters to coordinate around one challenger. I test this argument with a shock to royalty transfers to oil-dependent municipalities in Brazil, where most mayoral elections are decided by single-round plurality. Using a difference-in-differences design, I show that incumbents lost votes, yet their likelihood of removal was unaffected. I argue that the shock drew more candidates into the race, splitting the opposition vote and fostering incumbent survival. Combining difference-in-differences and instrumental variables, I demonstrate that additional challengers reduced incumbents' vote shares without unseating them, and that voters fragmented their support across these alternatives. Therefore, whether withdrawn support costs incumbents office depends on the structure of competition and the rules that shape it.

Keywords: Voter coordination; Vote fragmentation; Incumbent accountability; Economic shocks; Resource dependence.

*Ph.D. Candidate in the Department of Political Science, University of Pittsburgh. laura.chelidonopoulos@pitt.edu.

I would like to thank José Antonio Cheibub, Scott Morgenstern, Jae-Jae Spoon, Mollie Cohen, Daniel de Kadt, Mark Crescenzi, Vladimir Chlouba, Lucas Novaes, Jan Vogler, Daniel Gingerich, Ikromjon Tuhtasunov, Joao Alipio-Correa, George Avelino, Rogério Arantes, and Michelle Dion for their incredibly helpful comments and suggestions. I would also like to thank the participants of the Pitt Politics Workshop, EITM 2025, DMV Conference in Political Economy & Development, MWEPS XII, MPSA 2026, MapleMeth 2026, PolMeth 2026, and the VIM Pre-Conference Workshop for their engagement and feedback.

Introduction

Elections are expected to discipline incumbents by removing those held responsible for negative outcomes, yet officeholders sometimes survive despite voters withdrawing support from them. In coalition governments, for instance, shared authority makes it harder to assign blame for adverse outcomes, weakening the electoral consequences incumbents face (Powell & Whitten, 1993). In dominant party systems, ruling parties retain power even when they do not meet voters' expectations, whether because voters lack credible alternatives, as in South Africa under the ANC (Ferree, 2011), or because the party controls resources and patronage networks that sustain its dominance, as in Mexico under the PRI (Greene, 2007; Magaloni, 2006).

Understanding when and why this pattern arises matters for how we understand electoral accountability. I argue that whether the loss of votes costs incumbents their office depends on how the anti-incumbent electorate distributes its support, concentrating it on a single alternative or dispersing it across several. This logic should operate most clearly where winning, or staying in power, requires only more votes than any other candidate. The argument builds on a coordination tradition that treats consolidation behind one challenger as an informational and behavioral achievement rather than a mechanical product of the electoral rule itself (Cox, 1997; Myerson & Weber, 1993). Previous work contends that votes cast for less viable candidates are wasted, which is expected to push competition toward two viable contenders (Duverger, 1954; Riker, 1982). Nevertheless, this requires voters who abandon the officeholder to converge on a single alternative, which is contingent rather than guaranteed (Cox, 1997). I argue that when the anti-incumbent electorate does not coordinate, its support diffuses across opposition alternatives, and the incumbent can hold office even after losing votes.

This dynamic should become more pronounced following adverse conditions that exacerbate incumbents' electoral vulnerability. Therefore, I focus on times of economic downturns. Although expectations about how accountability operates vary across systems (Powell & Whitten, 1993), downturns tend to direct blame toward incumbents through

retrospective voting (Fiorina, 1978). Research on economic voting tends to examine incumbent accountability in one of two ways: through changes in vote share or variation in reelection success. For instance, studies have documented that incumbents received fewer votes during the Great Recession in Europe (Hernández & Kriesi, 2016), following income losses in Latin America (Murillo & Visconti, 2017), and across all levels of government in the United States because of negative local economic conditions (De Benedictis-Kessner & Warshaw, 2020). From a different but seemingly complementary perspective, incumbents' likelihood of reelection is affected by both positive and negative commodity shocks in Brazil (Novaes & Schiumerini, 2022) and, across countries, by exposure to economic conditions that are exogenous to their control (Campello & Zucco, 2020; Campello & Zucco Jr., 2016).

Although both ways of assessing incumbent support indicate that officeholders are the main actors to be held accountable, the economic voting literature has largely assumed, without testing, that the two move together and can be used as equivalent measures of electoral accountability. Yet punishment can reduce an incumbent's vote share without crossing the threshold for turnover, so losing votes does not necessarily translate into losing office. This distinction is particularly important because these outcomes can lead to vastly different scenarios: the survival or removal of the officeholder. I contend that whether the former converts into the latter is structurally contingent on voter coordination around the challenger candidates. Essentially, the structure of electoral competition, alongside the rules that govern it, should influence whether the votes an incumbent loses after an economic shock translate into removal from office.

I test my theory in the context of Brazil, where mayoral elections in most municipalities are decided by plurality rule. Drawing on the 2014 Petrobras scandal and the global oil crisis, which reduced royalty inflows to oil-dependent Brazilian municipalities, I use a difference-in-differences design to estimate the electoral consequences of the shock. I show that incumbents lose vote share in royalty-losing municipalities, but their likelihood of being removed from office is not altered. Therefore, voters do punish leaders, but after economic disruptions incumbents do not necessarily face the risk of being replaced.

I contend that the reason is that the anti-incumbent vote disperses across challengers. One source of such diffusion is the entry of additional competitors: as more candidates enter the race, voters should find it harder to settle on a single alternative. In particular, I argue that the negative consequences of the shock are perceived by potential challengers as an opportunity to improve their chances of winning office, leading more of them to enter the election. In turn, the anti-incumbent vote fragments across competing candidacies rather than concentrating on one viable candidate. A difference-in-differences shows that the number of competing candidates increased in the first post-shock election. Using an instrumental variables approach to assess how much of the vote decline is due to this entry, I show that additional candidates markedly reduce the officeholder's vote share without raising turnover. A final difference-in-differences demonstrates that the shock raises the effective number of opposition candidates in affected municipalities. When faced with economic downturns that threaten their electoral support, incumbents benefit from the fact that voters face an even more fragmented field of candidates, making it easier for them to remain in office.

The findings contribute to ongoing debates on voter coordination, the electoral consequences of economic shocks (Ahlquist et al., 2020; Campello & Zucco, 2020; Campello & Zucco Jr., 2016; Healy & Malhotra, 2009; Novaes & Schiumerini, 2022; Rashkova & Su, 2024; Tertytchnaya et al., 2018), and the political consequences of resource dependence (Andersen & Ross, 2014; Brückner et al., 2012; Carreri & Dube, 2017; Caselli & Michaels, 2013; Caselli & Tesei, 2016; Goldberg et al., 2008; Haber & Menaldo, 2011; Jensen & Wantchekon, 2004; Monteiro & Ferraz, 2012; Paler et al., 2023; Robinson et al., 2006; Ross, 2001; Wantchekon, 2002). Perhaps most importantly, it speaks to the broader question of when and why elections do not translate withdrawn support into leadership change. It shows that the structure of electoral competition under certain contexts can limit the extent to which elections lead to incumbent defeat. This is especially relevant because single-round races like Brazil's mayoral elections are expected to converge around a few viable candidates and to discourage fragmentation by concentrating support behind the strongest alternatives to the incumbent (Duverger, 1954; Fujiwara, 2011).

Yet I show that, in adverse contexts such as economic shocks, instead of the anti-incumbent electorate consolidating behind a leading challenger, opposition votes disperse across multiple candidacies, and officeholders fulfill their goal of winning office. Accountability should therefore distinguish between responses to performance that are not necessarily complementary. Although previous work has argued that voters may punish officeholders randomly (Achen & Bartels, 2016, 2018) or shift blame to external actors (Alcañiz & Hellwig, 2011), I show that there are contextual and institutionally mediated reasons for incumbents to lose support, but not their hold on power.

Electoral Accountability and Opposition Coordination

Electoral systems are commonly described as realizing one of two visions of democracy: a majoritarian vision that prizes the decisive election as an instrument of citizen control over a single identifiable government, and a proportional vision that prizes representation and post-election bargaining (Powell, 2000). Within either vision, the accountability relationship is multi-channel, since voters use a single instrument, the vote, to pursue two goals at once, selecting better officeholders prospectively and sanctioning poor performers retrospectively (Fearon, 1999; Manin et al., 1999). In this sense, rulers are accountable to the extent that the probability they survive in office responds to their performance (Cheibub & Przeworski, 1999). This point is central to what follows: survival in office is one institutional translation of voter sanction, and I argue it is not the only observable expression of that sanction. Incumbents often lose votes without losing office, and I set out to explain when withdrawn support does not cost an incumbent office.

Two features of this literature are worth drawing out. First, it operates largely through the comparison of different electoral systems. I suggest that a parallel logic may operate within a single electoral rule. Even holding the rule fixed at plurality, the degree to which the anti-incumbent vote consolidates can determine which channel of sanction, vote-share loss or office removal, actually transmits the sanction voters intend. Second, an influential strand of this work operationalizes accountability as office turnover or survival

in office rather than as incumbent vote share (Cheibub & Przeworski, 1999; Powell, 2000), even as other work measures it through vote share (Hellwig & Samuels, 2008). I build on this distinction to argue that plurality rules can deliver vote-share punishment without delivering office removal, so that the two outcomes can decouple even on the majoritarian vision's own terms.

In particular, I contend that whether vote-share loss converts into office removal depends on how voters coordinate under rules that award the seat to a single winner. Under plurality, votes cast for candidates who cannot win are wasted, and this tends to push competition toward two viable candidates per district (Duverger, 1954; Riker, 1982). This tendency, however, is not automatic. Cox (1997) recasts it as a strategic-coordination problem whose resolution is contingent: single-member plurality admits both Duvergerian equilibria, in which no more than two candidates draw support, and non-Duvergerian equilibria, in which more than two remain viable, with the outcome depending on how clearly expectations converge on who can win (Cox, 1997, 1999). This contingency has a formal basis: voting equilibria under plurality rest on common perceptions of viability rather than on the rule itself (Myerson & Weber, 1993). Coordination is therefore something the electorate has to achieve, and it does not follow straightforwardly from the ballot.

The literature I draw on focuses on the political consequences of electoral rules, from the number of viable candidates a system sustains (Cox, 1997; Duverger, 1954; Lago & Martínez i Coma, 2024; Lucardi, 2019; Myerson & Weber, 1993; Riker, 1982) to the kind of accountability it can deliver across different systems (Carey & Hix, 2011; Cheibub & Przeworski, 1999; Hellwig & Samuels, 2008; Kam et al., 2020; Powell, 2000; Rudolph & Däubler, 2016). I build on this work but shift the focus to incumbent survival. Even when voters know exactly whom to punish, the structure of electoral competition should influence whether their punishment removes the incumbent from office or merely erodes the latter's support. My central argument is that electoral rules and voter coordination around the opposition jointly shape whether vote loss translates into office loss. As I discuss in more detail later, one factor that works against such coordination is the entry

of additional challengers: as more candidates compete, voters should find it harder to form the common perceptions of viability it requires.

Maeda (2010) advances a similar logic in parliamentary democracies, showing that opposition fragmentation conditions how strongly poor economic performance erodes the governing party's vote share and arguing that it also conditions the extent to which this leads the latter to lose power. Yet, despite this argument, the author does not measure removal from power, only vote-share loss. I first extend that logic to incumbent survival under rules where a single seat is at stake, which should make vote coordination especially consequential, since a single office leaves the losing side with no representation at all. Second, I show that this survival works through vote dispersion. While existing work finds that opposition fragmentation can enhance accountability, this is usually the case among ideological voters for whom the gains in ideological proximity outweigh the losses in candidate viability (Ferrer, 2023).

Overall, I contend that where the institutional environment fosters the fragmentation of the anti-incumbent vote across multiple alternatives, the consolidation needed to threaten the incumbent becomes harder to achieve, and the accountability pressure need not convert into office removal.

Incumbent Accountability Under Economic Shocks

I analyze this argument in the context of an economic shock, where, I contend, the conditions are conducive to the anti-incumbent vote fragmenting across challengers. Research extensively shows that voters punish or reward incumbents based on past economic performance, assigning credit or blame to those in office (Carlin et al., 2015; Duch & Stevenson, 2008; Lewis-Beck & Stegmaier, 2000). A long tradition holds that how well this retrospective voting produces accountability depends on whether citizens can identify who is responsible for outcomes (Powell & Whitten, 1993). In contexts with a single responsible officeholder, blame attribution should be clearer, and votes withdrawn from the incumbent are likely to be spread across multiple challengers, particularly under

electoral rules that do not foster voter coordination.

Accountability can weaken when executive and legislative elections occur at different times or when governing coalitions blur the assignment of responsibility, and it strengthens when authority is concentrated (Duch & Stevenson, 2008; Lewis-Beck & Stegmaier, 2000; Samuels, 2004). In the United States, for instance, voters often link local economic conditions to the president's party in federal and gubernatorial races, but this connection does not reliably extend to local legislators (De Benedictis-Kessner & Warshaw, 2020). Where local elections are held in a separate cycle from other higher-level contests, as in this study, responsibility for local conditions should be harder to displace onto higher levels of government.

Many economic changes stem from downturns beyond government control. In general, previous work points out that shocks complicate attribution and can distort incentives away from welfare-improving policy. For instance, commodity swings and capital reversals leave even well-informed voters unable to separate luck from competence (Campello & Zucco, 2020; Campello & Zucco Jr., 2016), and remittances move incumbent approval despite limited official control (Tertytchnaya et al., 2018). Voters also weigh visible disaster relief more heavily than preventive investment (Healy & Malhotra, 2009). In addition, retrospective voting can operate through uneven exposure that changes what issues matter, what voters prefer, and what coalitions parties can form. Previous work shows that exposure to foreign-currency mortgage risk produces strong preferences for generous bailouts that a party mobilized into a clear platform (Ahlquist et al., 2020), and that negative price shocks impose large penalties on incumbents where local economies depend on commodities (Novaes & Schiumerini, 2022).

Major crises, not only short-term shocks, are also shown to shape how voters respond to incumbents. In Western Europe, the Great Recession accelerated party change by raising electoral volatility and penalizing mainstream parties (Hernández & Kriesi, 2016), and during the euro crisis voters sanctioned incumbents while selecting among challengers along existing divides (Hobolt & Tilley, 2016). It is also the case that electoral outcomes tend to shift rightward when downturns occur under center-left governments,

a pattern linked to changes in nationalism and inequality attitudes rather than major platform shifts (Margalit & Solodoch, 2025). Finally, resource-dependent contexts, the setting I study, reveal further gaps between performance and electoral sanctions. For instance, positive oil price shocks in Colombia reduced electoral competition through coercion aimed at controlling rents (Carreri & Dube, 2017). In addition, royalty windfalls in Brazilian municipalities can increase incumbents' short-term reelection chances, suggesting patronage or visibility effects (Monteiro & Ferraz, 2012), and even before revenues arrive, the promise of future rents can boost incumbents in electorally competitive areas (Paler et al., 2023).

My argument also speaks to the literature on incumbency advantage and disadvantage. A substantial empirical literature documents an electoral advantage for incumbents in established democracies, with mechanisms such as name recognition (Mayhew, 1974a), election timing (De Benedictis-Kessner, 2018), campaign finance access and personalized communication (Dano et al., 2025), personal vote-building through constituency service (Fowler & Hall, 2014), and scare-off effects that deter quality challengers from running (Cox & Katz, 1996). A contrasting body of work documents incumbency disadvantage in developing-democracy contexts, particularly in mayoral and local races (Titiunik, 2011; Uppal, 2009), with drivers ranging from voter sanctioning of corruption (Klašnja, 2015) and poor economic performance to adverse selection of office-seekers under term-limit pressures (Klašnja & Titiunik, 2017) and exposure to exogenous economic shocks (Novaes & Schiumerini, 2022).

In the context of this study, incumbents confronting an exogenous economic shock should face a disadvantage at the ballot, since the local executive sits at the most visible institutional point of accountability and is the available target for the accountability pressure voters express. At the same time, I argue that those mayors retain an advantage in the institutional translation of that pressure into office removal, since plurality combined with the expected coordination dynamic can keep the vote-share punishment from converting into turnover. Given that voters tend to punish officeholders for economic downturns, there should be incentives for voters to withdraw their support and spread

it across multiple opposition candidates, especially as they need not coordinate around a single viable candidacy.

The Case of Brazil

To assess the theoretical expectations, I rely on the case of Brazil. In 2014, two events, which occurred only a few months apart, set the stage for empirically testing the impact of economic shocks on electoral outcomes. The first was the Petrobras corruption scandal, uncovered by the federal investigation known as Operation Car Wash (*Operação Lava Jato*), which began as a money-laundering inquiry and grew into the largest anti-corruption probe in Brazilian history. At its core was a cartel of construction and engineering firms that colluded to overcharge Petrobras for contracts, channeling a portion of the inflated payments into bribes for company executives and into illegal financing for political parties. The investigation implicated executives of the country's largest firms and politicians across the major parties, and it remained present in national news coverage for years (Gomes & de Oliveira Medeiros, 2019; Silva, 2017). The scandal drew national attention not only because of the sums involved, which ran into the tens of billions of dollars, but also because Petrobras occupies a central role in Brazil's economy.

Around mid-2014, the second event was a considerable global drop in oil prices. The decline originated on the supply side, as rising United States shale output added substantially to global production while OPEC opted in late 2014 not to cut its own output, a combination that left the market oversupplied even as global demand was softening (Baffes et al., 2015). The adjustment was pronounced: Brent crude fell from roughly 115 US dollars per barrel in mid-2014 to under 30 dollars by early 2016, and according to the U.S. Bureau of Labor Statistics prices fell by 51.7 percent between June 2014 and January 2015 (Mead & Stiger, 2015). Petrobras was especially exposed to this collapse. The company carried a heavy debt load and depended on high-cost deepwater and *pré-sal* extraction, where the economics of production are particularly sensitive to the price of crude, and it had already been weakened by the financial consequences of the

scandal. Taken together, the two shocks led Petrobras to cut investment and production.

Under this scenario, understanding the distribution of Petrobras royalties to municipalities, a structured process governed by federal agencies, is key. First, the Brazilian Institute of Geography and Statistics (IBGE) determines which are the producing municipalities that are eligible to receive royalties. For offshore production, the institute applies a projection rule from their coastal boundaries: municipalities whose projections (either orthogonal or parallel) overlap with producing offshore fields are classified as confronting municipalities and receive a share in royalties. For onshore production, the criterion is simpler: the municipality where the surface location of a well (the wellhead) physically lies is the producing municipality.

Further, the National Agency of Petroleum, Natural Gas, and Biofuels (ANP) oversees the calculation of royalties based on oil and gas production, ensuring compliance with legal frameworks. Once calculated, the ANP transfers the funds to the National Treasury Secretariat, which distributes payments to states, municipalities, and other entities. Municipalities classified as secondary and neighboring zones also receive a share of royalties, albeit a smaller fraction, as compensation for indirect social, economic, and infrastructural impacts stemming from petroleum and natural gas production. In this scenario, royalty-dependent municipalities often rely on these payments to fund public services and local development.

In the context of the political crisis and the global oil market downturn, Petrobras significantly reduced royalty transfers, which exacerbated the vulnerability of many dependent municipalities. Several municipalities suspended investments and cut public expenditures, including the postponement of road maintenance and urban infrastructure projects, reductions in municipal service contracts, and delays in payments to local suppliers (Golveia et al., 2023). In particular, the decline in royalty payments became pronounced in November 2014 and, more drastically, in 2015 and most of 2016. Figure 1 shows the total royalties allocated to Brazilian municipalities across the years, along with the number of dependent municipalities. Over the years, the number of municipalities receiving royalties increased mostly due to the discovery of new oil fields and heightened

production.

In terms of the amounts, there is a visible rapid decline from 2014 to 2016, spanning from the onset of the shock to the first municipal election after its effects had materialized, the starkest decline in royalty payments throughout the entire period. During this interval, municipal royalties fell from a little over US$1.85 billion in 2014 to almost US$1.17 billion in 2016, a decline of roughly 37 percent. In recent years, particularly between 2021 and 2024, municipalities experienced a visible resource windfall. This resulted mainly from increased production in the *pré-sal* fields, which are highly productive areas, combined with elevated crude oil prices.

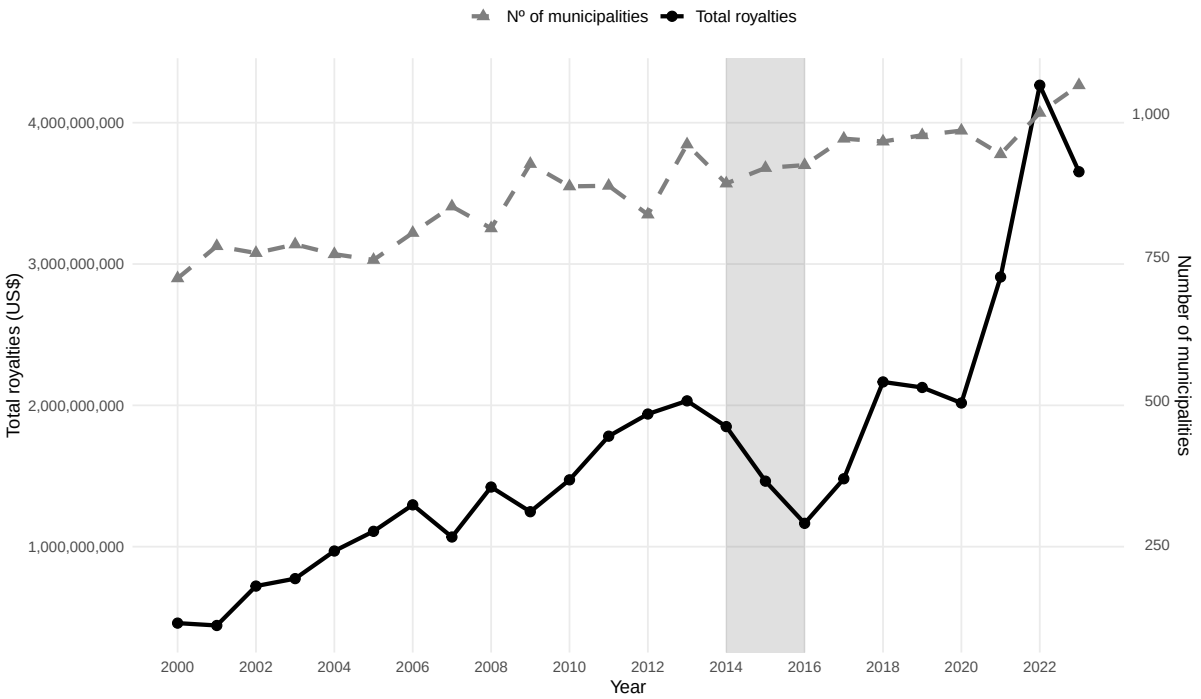

Figure 1: Royalty Allocation and Number of Dependent Municipalities (2000–2023)

Note: The solid black series plots the total royalties allocated to Brazilian municipalities by year and the dashed gray series plots the number of municipalities receiving royalties by year. The gray shaded band represents the period when royalty allocation to dependent municipalities sharply decreased. Data are sourced from the ANP. Royalty values are deflated to December 2023 prices and converted to United States dollars at the December 2023 average rate of 4.90 reais per dollar reported by the Central Bank of Brazil. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

In particular, among all royalty-dependent municipalities with a royalty allocation in 2014, about 86 percent received less two years later. ANP data also reveal that,

within the sample of royalty-losing municipalities, the average loss was 39 percent of a municipality's own allocation. More than two thirds lost over 25 percent, and about a third lost over 50 percent. Given that this decline fell within the term that preceded the 2016 election, there is strong reason to expect that any incumbent accountability effects would be short-lived and primarily concentrated in the first election following the shock (2016).

The royalty loss has also been reported in several Brazilian news sources. In the 2014–2016 period, there were diverse reports of the constraints royalty-dependent municipalities were facing in aftermath of the shock. For instance, in the municipality of Campos dos Goytacazes, transfers had been declining since October 2014, and the January 2015 installment arrived 25 percent below its level a year earlier. The municipal government put before the council bills to cut the number of secretariats and to reduce by 10 percent the salaries of secretaries and appointed staff (G1, 2015a). Moreover, São Fidélis moved public offices to a shortened morning schedule, suspended all municipality-run festivals until December 2015, and cut spending on fuel and official travel (G1, 2015b). Finally, another example is the city of Casimiro de Abreu, where transfers in the first quarter of 2015 were 48 percent lower than in the same quarter of 2014, and the local government put its investments on hold while trying to keep essential services running (Vasconcellos et al., 2015).

Overall, the royalty shock affected the municipal budget in myriad ways, ranging from administrative running costs to the suspension of local investments. These tangible consequences at the municipal level, caused by the exogenous shock, are what should motivate voters to punish incumbents at the ballot box. Previous work shows that incumbents can be blamed even for shocks external to their control (Ahlquist et al., 2020; Novaes & Schiumerini, 2022; Tertytchnaya et al., 2018), and, since they tend to be the main targets of blame in times of downturns, I expect voters to have plausible reasons to withdraw their support from them. This, nevertheless, does not necessarily mean they will be replaced, as I develop further in the paper.

Data and Empirical Strategy

I employ a difference-in-differences design to assess whether incumbents in municipalities that lost royalty revenues due to the economic shock suffer electoral disadvantage in the form of vote and office loss. The analysis focuses on mayoral elections between 2000 and 2016, specifically in municipality-year pairs where incumbents were in their first term and ran for reelection. In Brazil, mayors are directly elected every four years in elections held simultaneously across all municipalities. They are permitted to serve a maximum of two consecutive terms, after which they must step down for at least one electoral cycle before running again. This two-term limit means that, in any given election, the pool of incumbents seeking reelection consists exclusively of first-term mayors. Because municipal elections are held concurrently nationwide and separately from state and federal contests, which take place in even years between municipal cycles, local races are relatively insulated from the dynamics of higher-level elections.

In this analysis, I focus on municipalities with fewer than 200,000 registered voters in 2016. In Brazil, this population threshold determines the electoral system: municipalities above it hold two-round elections if no candidate obtains a majority of the votes, while those below it rely on a first-past-the-post rule[1]. Data are sourced from the Brazilian Superior Electoral Court (TSE), the ANP, the Public Finance Database of Municipalities (FINBRA), and the Fiscal Management Information System (Siconfi).

Treatment Definition

The treatment is defined by the loss of royalty revenue between the onset of the shock and the first election that followed it. For each royalty-dependent municipality, I measure the share of the total municipal budget composed of royalty transfers in 2014, the year the shock began, and again in 2016, the year of the first post-shock mayoral election, and I take the difference between the two. I use this measure because royalties

[1]Based on TSE data, 98.5 percent of Brazilian municipalities had fewer than 200,000 registered voters in 2016.

are spent out of the same budget as every other revenue, so a decline affects the mayor's spending capacity in proportion to the municipal budget, not only to the size of the royalties themselves. A municipality is treated when royalties shrink as a fraction of the budget over the 2014 to 2016 interval. The control group consists of municipalities that received no royalties and therefore had no exposure to the contraction.

There are 683 treated municipalities (i.e., that lost royalty revenues in the 2014–2016 period) and 4,195 control municipalities. Out of these 683 municipalities, 277 had incumbents who ran again in 2016, which then contribute to the post-shock comparison. In addition, I exclude 48 royalty-windfall municipalities from the estimation sample, since they are neither treated (because they did not lose revenues) nor true controls (because they received windfalls rather than having no dependence on royalties). Figure 2 maps Brazilian municipalities into these three groups, classified by the change in the royalty share of the municipal budget between 2014 and 2016: royalty-losing (treated) municipalities, royalty-windfall municipalities, and non-royalty-dependent (control) municipalities.

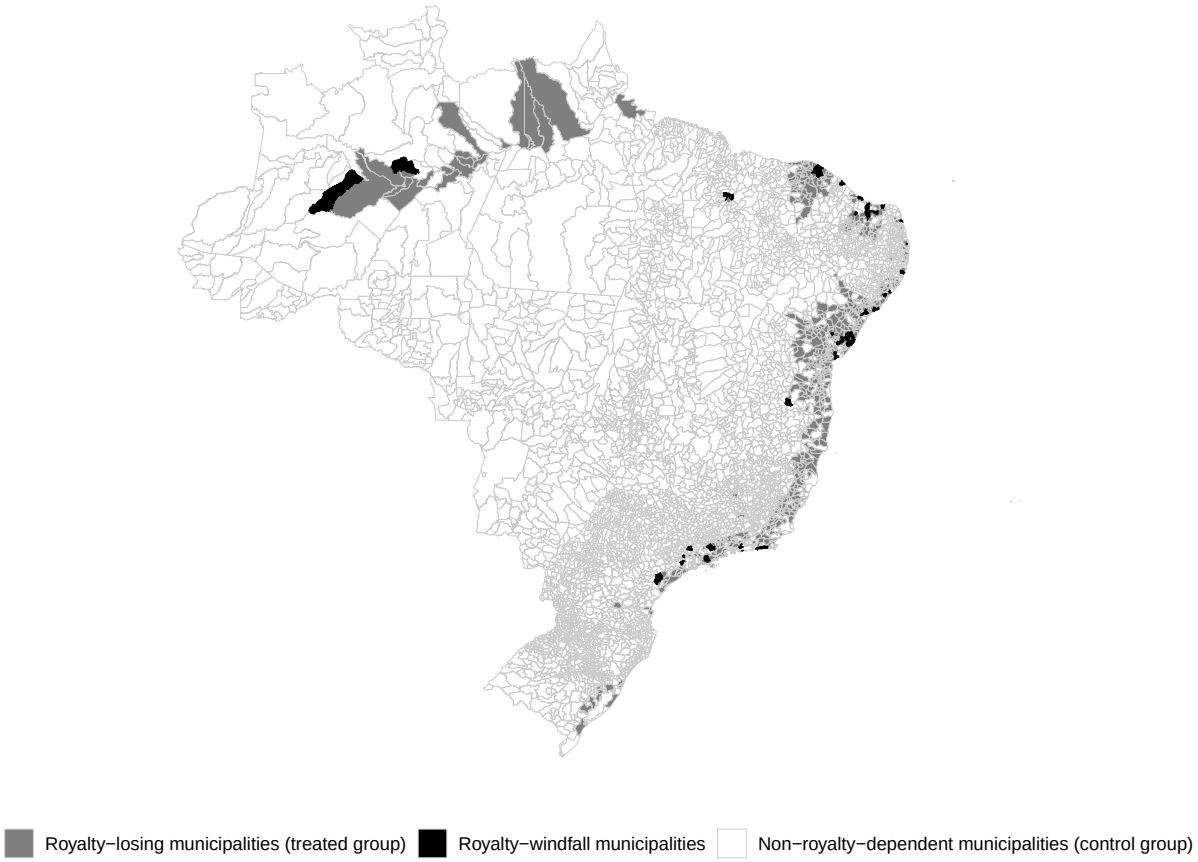

Figure 2: Treatment Status of Brazilian Municipalities in 2014

Note: The map covers the estimation sample of Brazilian municipalities. Gray denotes royalty-losing (treated) municipalities, black denotes royalty-windfall municipalities, and white denotes non-royalty-dependent (control) municipalities. Classification is by the change in the royalty share of the municipal budget between 2014 and 2016. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

Figure A.1 in the Appendix displays the cumulative distribution of royalty revenue loss among treated municipalities. The distribution shows that almost 90 percent lost up to four percentage points of their total budget. This concentration motivates the use of a binary treatment indicator rather than a continuous measure of loss intensity as in Callaway et al. (2025). With so little variation in loss intensity across treated municipalities, a continuous dose-response estimator would likely lean on the handful of municipalities with the largest losses, so I adopt the binary indicator, which keeps the affected-versus-unaffected contrast transparent without imposing a continuous functional form. To account for the existing variation in treatment intensity, the analysis also

examines heterogeneity by the size of the loss, splitting treated municipalities into terciles of the fall in royalties as a share of their total budget, measured in percentage points.

Specification

Figure 3 provides a descriptive overview of the two main electoral outcomes across treated and control municipalities over time. Panel (a) shows the average incumbent vote share. Prior to the shock, incumbents in treated and control municipalities receive comparable levels of electoral support. In 2016, however, the average vote share in treated municipalities drops visibly, while control municipalities remain relatively stable, and the confidence intervals diverge only in this year. Panel (b) displays the office loss rate (i.e., the proportion of incumbents who lost elections), and both groups exhibit an upward trend over the period. In 2016, the office loss rate is directionally higher in treated municipalities, although the confidence intervals overlap considerably.

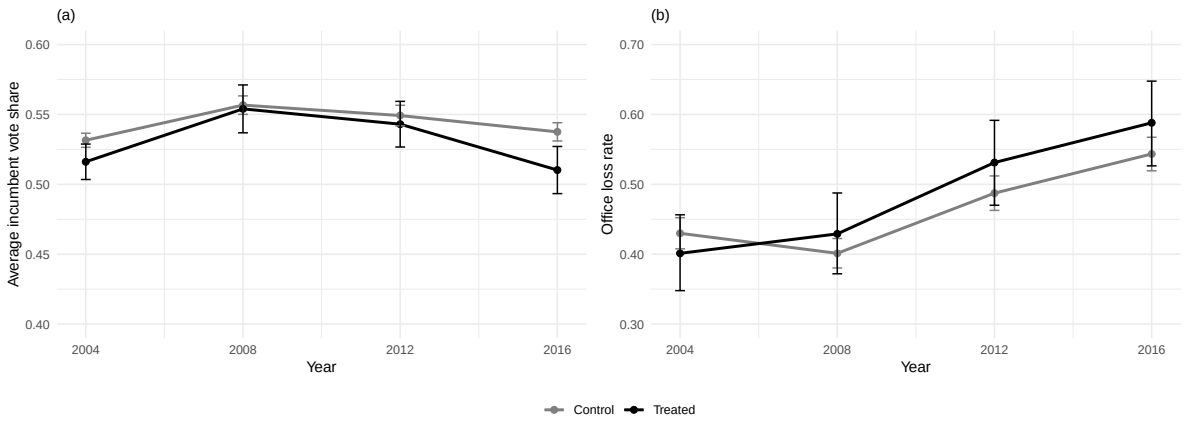

Figure 3: Descriptive Trends in Incumbent Vote Share and Office Loss by Treatment Status

Note: Panel (a) plots the mean incumbent vote share and panel (b) the mean office loss rate, the share of running incumbents who lost their election, by year for treated (royalty-losing) and control (non-royalty-dependent) municipalities. Shaded bands are 95 percent confidence intervals. Electoral data are sourced from the TSE. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

Formally, the equation that represents my empirical strategy is:

$$Y_{it} = \beta_0 + \beta_1\big(\text{Loss}_i \times \text{Post}_t\big) + \boldsymbol{\gamma}^\top\big(\mathbf{X}_i \times \text{Post}_t\big) + \alpha_i + \tau_t + \varepsilon_{it}$$

where Y_{it} denotes the outcome for municipality i in election year t (either the incumbent's vote share or a binary indicator for incumbent reelection). Loss_i is equal to 1 for royalty-dependent municipalities that experienced a negative change in the budgetary share of royalties between 2014 (the year of the shock) and 2016 (the year of the mayoral election after the shock), and 0 for non-royalty-dependent municipalities. Post_t indicates the post-shock period, the first electoral year after the shock (2016). The coefficient β_1 represents the average treatment effect on the treated (ATT), capturing the change in outcomes for treated municipalities relative to controls. $\boldsymbol{\gamma}$ is a vector of post-period adjustments that allows the post-shock shift to vary with pre-treatment municipal characteristics $\mathbf{X}_i$ (a vector of municipality-level controls)[2]. The parameters α_i and τ_t are municipality and year fixed effects, respectively.

I use all pre-treatment elections (2000–2012) and one post-treatment election (2016) in the analysis for two main reasons. First, by having more observations to estimate counterfactual outcomes, this approach allows me to estimate the ATT with higher accuracy, resulting in smaller standard errors (Egami & Yamauchi, 2023). Second, since I am interested in the immediate electoral reactions of voters to the shock, I focus only on the first post-shock mayoral election. In this case, I do not expect the shock to have long-term implications because, as seen in Figure 1, Petrobras' royalty revenues recovered to their previous levels (and above) from 2017 onward. Moreover, the next mayoral election after 2016 was held in 2020 and, due to COVID-19 complications, was likely influenced by factors other than a shock that had occurred years earlier.

Finally, three considerations are worth mentioning. First, because the scandal affected all royalty-dependent municipalities, both those that lost revenue and those that gained it, one might argue that the appropriate treatment is royalty *dependence* itself rather than royalty loss, since Petrobras, a central actor in these localities, was directly

[2]The coefficients on the baseline covariates for all models estimated in the paper are reported in Appendix A.18 in the Appendix.

involved in it. If dependence were what mattered, then royalty-windfall municipalities, which were equally dependent, should follow electoral trajectories similar to the losers on both outcomes, since what would drive the response would be exposure to Petrobras rather than the direction or magnitude of the revenue change. However, the data do not support this argument. As I show in Appendix A.14 in the Appendix, royalty-windfall municipalities exhibit null effects on both incumbent vote share and reelection. Grouping them together with the losers would therefore combine two sets of municipalities that behave differently and would dilute the loss effect by averaging it against a null. The electoral response should therefore operate through royalty loss (with magnitude heterogeneity, as I further develop), not through pre-shock dependence.

Second, my analysis tracks the individual incumbent candidate rather than the party label. In Brazilian mayoral races, the unit of accountability is often the person who occupies the office: voters tend to evaluate the incumbent's personal record (Ferraz & Finan, 2008; Pereira & Melo, 2015), challengers organize their candidacies around individual reputations and networks (Speck & Cervi, 2016), and the career and rent consequences of holding municipal office accrue to the candidate rather than to the party (Brollo et al., 2013). It is also not unusual for candidates to run in consecutive mayoral elections under different party labels, so considering only the party would mask the fact that voters may be supporting the same person across elections, even when that person's institutional affiliation changes.

Lastly, the royalty losses should not be well captured by conventional subnational economic indicators, even when measured at the municipal level. Indicators such as municipal GDP, local employment, or income per capita reflect a broad set of economic activities and therefore dilute the sudden contraction in revenues tied specifically to oil extraction and Petrobras-linked production.

Identifying Assumptions

Difference-in-differences designs rest on a set of important identifying assumptions that secure the validity of the causal estimates. First, to assess parallel trends, I estimate event study specifications that interact treatment status with election-year dummies, with 2012, the election before the economic shock, as the reference year. As shown in Figure A.2 in the Appendix, the pre-treatment coefficients are clean for incumbent vote share in 2004 and 2008, and for reelection in 2008, except in 2004, where the reelection coefficient is significant at the 10 percent level. To account for it, I apply the relative-magnitudes sensitivity analysis of Rambachan and Roth (2023). At $\bar{M} = 1$, the post-period violation is allowed to be as large as the worst observed pre-period violation. The analysis, detailed in Section A.2.2, shows that the null is robust and plausibly holds in 2016. Read together, the pre-treatment evidence is consistent with parallel trends between royalty-dependent municipalities that lost revenues and non-dependent municipalities.

Second, treatment assignment must be exogenous to factors that also determine electoral outcomes. Royalties are allocated strictly according to federal legal requirements, so local political actors cannot influence whether or how much revenue their municipality receives. Third, the composition of the sample should not change differentially across treatment and control groups as a consequence of the shock, particularly in terms of treated incumbents choosing not to rerun after the shock. Nevertheless, there is no statistically significant difference in rerunning rates between mayors in royalty-dependent municipalities and those in non-dependent municipalities (see Section A.2.3). Fourth, the no-anticipation assumption requires that treated municipalities did not adjust their behavior before the shock materialized. That assumption is plausible as both the Petrobras scandal and the global oil price collapse were sudden and largely unanticipated events.

Finally, the stable unit treatment value assumption (SUTVA) requires that the treatment status of one municipality does not affect electoral outcomes in other treated and control units. The shock was simultaneous across treated municipalities and operates

through a contraction specifically in the local budget, which substantively justifies this assumption. To assess this claim, I run placebo event studies, reported in Section A.2.4, where rings of neighboring municipalities are null for first-, second-, and third-order neighbors. Appendix A.2 discusses all the identifying assumptions in full.

Covariates and Balance

I run difference-in-differences specifications with and without the set of covariates. Table 1 reports baseline characteristics (in 2012, year of the mayoral election held before the shock) for royalty-dependent municipalities that lost revenues and non-dependent municipalities. The groups differ systematically on geographic and socioeconomic factors: royalty-dependent localities are more likely to be coastal, as shown in Figure 2, and are, on average, more populated. They further exhibit lower shares of highly educated individuals, lower average income per capita, and lower GDP rates. Urbanization rates and Workers' Party (PT) incumbency, the president's party at the time of the shock, are similar across groups.

Finally, I also consider the level of electoral volatility across dependent and non-dependent municipalities in 2012 and its average across all pre-shock electoral years (2004, 2008, and 2012). It essentially captures the stability of voter preferences over time among existing, new, and exiting political parties[3] (Mainwaring & Su, 2021; Mainwaring et al., 2017; Powell & Tucker, 2014; Roberts & Wibbels, 1999; Singer, 2013; Tavits, 2005). If one type of municipality is more volatile than the other, this would suggest that incumbents are, *a priori*, more vulnerable to losing votes because the electorate is less inclined to maintain stable preferences. According to Table 1, the differences in means are statistically significant, indicating that royalty-dependent municipalities are more volatile than non-dependent ones.

[3]Total electoral volatility is calculated as $\frac{1}{2}\sum_{i=1}^{n}|p_{i,t-1}-p_{i,t}|$, where $p_{i,t-1}$ is the percentage of votes received by party i at time $t-1$ (i.e., 2008), $p_{i,t}$ is the percentage of votes received by party i at time t (i.e., 2012), and n is the number of political parties. The factor of $\frac{1}{2}$ ensures that gains and losses are not double counted.

Table 1: Characteristics of Royalty-Dependent and Non-Dependent Municipalities

	Non-Dependent Municipalities		Royalty-Dependent Municipalities				
Variable	Mean	SE	Mean	SE	Difference in Mean	SE of Difference	N
Coastal (indicator)	0.048	0.005	0.209	0.024	0.162***	0.025	1892
Share higher ed (2012)	0.036	0.001	0.028	0.001	-0.008***	0.001	1892
Average income (2012)	1717.12	18.709	1403.14	38.570	-313.99***	42.868	1892
Log GDP (2012)	9.370	0.017	9.209	0.043	-0.161***	0.047	1892
Log population (2012)	9.284	0.025	9.837	0.058	0.552***	0.064	1892
Urban (indicator)	0.625	0.012	0.635	0.029	0.011	0.031	1892
PT incumbency (2012)	0.111	0.008	0.108	0.019	-0.003	0.020	1892
Electoral volatility (2012)	0.479	0.007	0.531	0.017	0.052***	0.018	1892
Average volatility	0.530	0.005	0.577	0.013	0.047***	0.014	1892

Note: The table reports baseline means and standard errors of each characteristic in 2012, the mayoral election before the shock, for non-dependent (control) and royalty-dependent (treated) municipalities, along with the difference in means, its standard error, and the number of observations. Electoral volatility (2012) is the value in 2012 and average volatility is the mean across the pre-shock elections of 2004, 2008, and 2012. Significance stars refer to the difference in means: $^{***}p < 0.01$. Covariates are sourced from the IBGE.

To mitigate concerns that these observable differences might confound the effects of royalty loss on electoral outcomes, the covariate-adjusted specification includes baseline controls interacted with a binary post-shock indicator. In particular, I add as controls: coastal, population (log), share of highly educated individuals, GDP (log), and electoral volatility in the party system. I exclude average income per capita as a covariate to avoid collinearity issues with education and GDP measures. Moreover, I focus on electoral volatility in 2012, the last election before the shock, rather than its average across years, since the former is more likely to capture the immediate political context shaping how voters respond to the royalty loss.

Finally, it is worth noting that the Petrobras scandal, in particular, could be expected to influence electoral patterns across all or most Brazilian municipalities due to the serious corruption allegations, especially as most citizens were exposed to news about the scandal (Silva, 2017). However, even in the presence of corruption and its potential to undermine trust in political parties in general, its effects are still heterogeneous depending on economic conditions. Several studies show that voters do not uniformly punish

corruption. Instead, their responses are often shaped by economic performance and the material benefits they associate with incumbents (Carlin et al., 2014; Choi & Woo, 2010; Ferraz & Finan, 2008; Krause & Méndez, 2009; Manzetti & Rosas, 2015; Manzetti & Wilson, 2007; Pereira & Melo, 2015). Following this logic, in royalty-dependent localities that suffered revenue losses due to the shock, economic hardship should prevail in the time of vote choice.

Results

Table 2 reports the main difference-in-differences estimates of the effect of the royalty shock on incumbent electoral outcomes[4]. The results provide clear evidence of an incumbency disadvantage in resource-dependent localities affected by the decrease in Petrobras' royalty inflows. In royalty-dependent municipalities, incumbents received significantly fewer votes after the 2014–2016 revenue loss. The point estimate in Column (1) suggests a decline of about 3.4 percentage points in incumbent vote share, consistent with the expectation that voters punished mayors for the contraction in the municipal budget.

[4]Difference-in-differences results without the set of covariates can be found in Table A.3 in the Appendix. The direction and statistical significance of the treatment coefficient are consistent with the covariate-adjusted models. I also exclude one extreme outlier municipality whose royalty loss exceeded 60 percent of its total budget to avoid a single observation exerting disproportionate leverage on the estimates. Table A.4 in the Appendix shows that the main difference-in-differences results are virtually identical when this municipality is included in the sample, so the exclusion is conservative rather than result-driving.

Table 2: The Effects of Royalty Loss on Incumbent Vote and Reelection Success

	Incumbent Vote	Reelection Success
	(1)	(2)
Loss × Post	-0.034**	0.004
	(0.016)	(0.048)
Covariates	✓	✓
Municipality FE	✓	✓
Year FE	✓	✓
Observations	8,506	8,506
Municipalities	4,873	4,873
Election years	2000–2016	2000–2016

Note: Difference-in-differences estimates of the royalty shock on the two electoral outcomes. Column (1) is the incumbent vote share and Column (2) a binary indicator for incumbent reelection. The coefficient on Loss × Post is the average treatment effect on the treated, where Loss denotes royalty-dependent municipalities that lost budgetary royalty share and Post the first post-shock election (2016). All models include municipality and year fixed effects, with covariates interacted with the post-period indicator. Standard errors (in parentheses) are clustered by municipality.

$^{**}p < 0.05$.

This disadvantage, nevertheless, is limited. Column (2) shows no corresponding effect on mayoral reelection. Although incumbents lost support in resource-dependent localities affected by the decrease in Petrobras' royalty inflows, these losses were not large enough to systematically alter their reelection success. To gauge the substantive size of the vote-share punishment, it helps to place it against baseline electoral conditions. In 2012, the election before the shock, incumbents in control and treated municipalities received very similar support, about 47.6 percent in control municipalities and 46.1 percent in treated ones. By 2016, the two groups diverged: average incumbent vote share held roughly steady in control municipalities, at about 46.6 percent, while falling to 43.1 percent in treated municipalities. Against this baseline, the estimated 3.4 percentage point decline amounts to a reduction of roughly 7 percent in the pre-shock vote share of treated incumbents.

To examine whether the electoral response varies with loss intensity, I split the royalty-loss municipalities into terciles of the fall in royalties as a share of the total municipal budget. The three terciles enter a single two-way fixed-effects (TWFE) specification

per outcome, with the non-royalty-dependent municipalities as the reference category[5]. Table A.18 in the Appendix reports these estimates. Incumbent vote share falls by 5.9 percentage points in the middle tercile and by 5.6 percentage points in the top tercile, roughly 12 percent of the estimation sample's 2012 pre-shock mean of 47.4 percent. In the bottom tercile there is no measurable change in vote share.

Read together, the pooled and tercile specifications (in particular, the middle and top bins) show that there are conditions under which electoral punishment does not translate into removal: incumbents lose vote share without losing their seats. Why this punishment does not lead to office loss is the question the next sections address.

Electoral Competition and Incumbent Survival

Voters may hold incumbents responsible at the ballot box by withdrawing support from them, yet this does not always translate into challenger victories. I contend that the reason is that the structure of electoral competition can hinder voters' ability to remove mayors from power, rather than reflecting an inability to identify whom to punish (Achen & Bartels, 2016, 2018). I argue that the mechanism through which incumbents lose votes but not office is the fact that the votes lost by the incumbent are dispersed across many candidates seeking to replace him or her, which leaves no single alternative able to gather the support needed to unseat the officeholder. The economic shock is what triggers this dynamic: by associating the incumbent with deteriorating local conditions, it signals vulnerability and draws additional challengers into the race, and the fragmentation their entry produces is what keeps the weakened incumbent in office.

The increase in challenger entry can be understood through the decision-making process of candidates who weigh the costs and benefits of running for office. Politicians' decisions to enter a race are often said to reflect a calculation of the probability of winning against the costs of running (Black, 1972; Cox, 1997; Downs, 1957). Under normal

[5]The tercile cutpoints, taken over the treated municipalities, are 0.030 and 0.399 percentage points of the budget.

circumstances, incumbents generally enjoy considerable advantages that deter potential challengers, including name recognition and the ability to claim credit for public goods (Mayhew, 1974a). These advantages produce what is known as a "scare-off" effect, in which strong incumbents discourage quality challengers from competing (Cox & Katz, 1996; Stone et al., 2004).

I argue that an economic shock should weaken these advantages. When incumbents appear vulnerable, potential opposition candidates perceive a greater chance of winning and become more willing to enter the race, even if they would have stayed out under better economic conditions. This expectation speaks to previous work on strategic entry in U.S. congressional elections, which has shown that candidates are more likely to enter when national conditions signal a vulnerable incumbent (Jacobson, 1989; Jacobson & Kernell, 1983). That work helps explain why entry rises, without focusing on how voters' support concentrates around these opponents.

Previous evidence in Brazil has stated that single-ballot plurality rule leads voters to concentrate their votes in the top two candidates (Fujiwara, 2011). Nevertheless, I argue that it depends on the context under which these elections happen. When adverse conditions expand the set of competitors, voters should find it harder to distinguish viable from non-viable candidates, since more competitors blur which opponents can realistically win. With multiple entrants and no clear focal point around which to coordinate, anti-incumbent voters end up splitting their support across several opposition candidacies, preventing any single challenger from accumulating enough support to win.

Challenger Entry as the Mechanism

To test how the patterns of competition changed after the shock, I use two complementary strategies relying on the same dataset with municipality-year observations. First, I estimate another difference-in-differences where the dependent variable is the total number of candidates competing in elections. Second, I use an instrumental variables approach, focusing on incumbents in their first term who sought reelection. I use the

decline in revenues after the shock as an instrument for the number of challengers.

The difference-in-differences specification is the same one used for the main outcomes, estimated on municipality-year observations for the election years 2000 to 2016, with municipality and election-year fixed effects, baseline covariates interacted with the post-shock indicator, and standard errors clustered at the municipality level[6].

Table 3: The Effects of Royalty Loss on the Number of Competing Candidates

	Number of Candidates
Loss × Post	0.245**
	(0.106)
Covariates	✓
Municipality FE	✓
Year FE	✓
Observations	8,506
Municipalities	4,873
Election years	2000–2016

Note: Difference-in-differences estimates of the royalty shock on the number of competing candidates, with municipality and election-year fixed effects. The coefficient on Loss × Post is the average treatment effect on the treated, where Loss denotes royalty-dependent municipalities that lost budgetary royalty share and Post the first post-shock election (2016). All models include municipality and year fixed effects, with covariates interacted with the post-period indicator. Standard errors (in parentheses) are clustered by municipality.

**$p < 0.05$.

Table 3 shows that the royalty loss increased the number of candidates competing in treated municipalities. Incumbents in these localities faced about 0.24 more challengers in 2016 than incumbents in non-dependent municipalities, roughly 8 percent of the pre-shock average of 2.90 candidates in treated municipalities. Within the sample of royalty-losing municipalities, there was an 11 percent increase in the number of competitors after the shock relative to before the shock. Disaggregating by loss intensity, the top tercile (i.e., municipalities that experienced the greatest losses) draws about 0.54 more candidates relative to non-dependent municipalities. The middle tercile shows an increase of about 0.25 candidates, although this estimate is only marginally significant. The bottom tercile

[6]Event study estimates validating the parallel trends assumption and difference-in-differences results without covariates are reported in Figure A.7 and Table A.5 in the Appendix, respectively.

shows no change (see Table A.19 in the Appendix).

Despite the difference-in-differences specification crucially identifying the causal effect of the shock on the increase in the number of competing candidates, it cannot quantify how much of the observed decline in incumbent vote share runs through that entry channel rather than through other consequences of the shock that might have reached voters directly. I therefore use $\text{Loss}_i \times \text{Post}_t$ as an instrument for competition to assess how much incumbent vote share is affected by the increased number of candidates in royalty-dependent municipalities.

Under the exclusion restriction, any effect of the shock on incumbent vote share that does not pass through challenger entry is ruled out by construction, so the two-stage least squares estimate isolates the entry channel from the other pathways through which the documented decline in incumbent support could otherwise be explained. I assess this assumption further in the paper. Moreover, that the distribution of royalty revenues is exogenous to local political control makes the independence assumption plausible. It is unlikely that the unobserved local factors driving both challenger entry and incumbent vote share also drive a federally determined royalty formula.

To assess the results, I estimate a first-stage equation for candidate entry. The sample remains the same as in the main analysis: municipality-year observations in which first-term incumbents ran for reelection. I use two-way fixed-effects (TWFE) regressions, where the interaction between belonging to the treatment group and the post-shock year is expected to positively and significantly affect the number of candidates competing in the 2016 mayoral election, using the same specification described above. Although previous work has shown that economic shocks can increase the incentives for new actors to enter the political arena (Rashkova & Su, 2024), the expectation here is not limited to new candidates: politicians who have previously competed should also perceive greater opportunities when the incumbent appears vulnerable. I therefore use the overall number of candidates as the main dependent variable.

The first stage confirms that the royalty shock raised challenger entry in treated

municipalities as reported in Table A.6 in the Appendix. As a complementary test, Table A.7 also reports results using the number of new candidates as the dependent variable, where the effect is positive but marginally significant, suggesting that the increase in competition is driven by the full pool of competitors rather than new entrants alone.

The second stage assesses whether the competitive environment shaped by the shock translated into changes in incumbent vote share by estimating a two-way fixed-effects (TWFE) model in which the number of challengers is instrumented by $\text{Loss}_i \times \text{Post}_t$. I estimate this specification for both incumbent vote share and reelection success. As shown in Table 2, the shock did not affect reelection in royalty-dependent municipalities, so the instrument should have no causal effect on this outcome.

The second-stage specification can be expressed as:

$$Y_{it} = \beta_0 + \beta_1 \hat{C}_{it} + \boldsymbol{\theta}^\top \left(\mathbf{X}_i \times \text{Post}_t \right) + \alpha_i + \tau_t + \varepsilon_{it}$$

where Y_{it} denotes incumbent vote share, and C_{it} is instrumented using $\text{Loss}_i \times \text{Post}_t$. The coefficient β_1 captures the causal effect of an additional challenger on incumbent electoral performance, identified from the exogenous variation in candidate entry induced by the royalty shock. This specification tests whether the decline in incumbent vote share is driven by the increase in the number of challengers induced by the shock, rather than reflecting a direct electoral response to deteriorating economic conditions alone.

Table 4: The Effects of Challenger Entry on Incumbent Electoral Outcomes

	Incumbent Vote	Incumbent Reelection
	(1)	(2)
Number of candidates	-0.141^{**}	0.018
	(0.064)	(0.196)
Covariates	✓	✓
Municipality FE	✓	✓
Year FE	✓	✓
Observations	8,506	8,506
Municipalities	4,873	4,873
Election years	2000–2016	2000–2016

Note: Second-stage TWFE regression estimates with municipality and election-year fixed effects. The number of candidates is instrumented by $\text{Loss}_i \times \text{Post}_t$. Standard errors (in parentheses) clustered by municipality.

$^{**}p < 0.05$.

Table 4 suggests that the negative electoral impact of the royalty shock on incumbents operated through an increase in competition. The second-stage estimates show that greater challenger entry, influenced by the shock in resource-dependent municipalities, substantially reduced incumbents' vote shares. In Column (1), each additional candidate in the race significantly decreases the incumbent's vote share by about 14.1 percentage points. Column (2), however, shows no corresponding effect on reelection, confirming that the increase in competition reduced incumbent support without altering their likelihood of retaining office.

Vote Dispersion as a Consequence of Entry

The theoretical argument above implies that the entry of additional challengers should not merely reduce the incumbent's vote share but also disperse the opposition vote across a larger set of alternatives. To assess whether the shock altered the distribution of votes in a manner consistent with this expectation, I examine changes in vote fragmentation using the effective number of candidates. In particular, I focus on the effective number of *opposition* candidates. This measure adapts the effective number of parties

index proposed by Laakso and Taagepera (1979) to the candidate level, capturing how electoral support is distributed across individual competitors rather than parties (Cox, 1997). Since my expectation is that the shock increases fragmentation among challengers specifically, not merely reduces the incumbent's dominance, I exclude the incumbent's vote share from the calculation and renormalize the remaining vote shares before applying the index.

The effective number of opposition candidates (ENC) is calculated as:

$$\text{ENC}_{it} = \frac{1}{\sum_{j=1}^{J} v_{ijt}^2}$$

where v_{ijt} is the vote share of opposition candidate j in municipality i at election t, and J is the total number of opposition candidates. The index captures the degree to which votes are concentrated or dispersed: when one candidate dominates, the effective number of opposition candidates approaches 1. When votes are more evenly distributed across many candidates, it increases. Unlike my candidate entry variable, the effective number of opposition candidates weights each candidate by their actual vote share, distinguishing between races where additional candidates are electorally relevant and those where they are marginal.

If the economic shock encouraged challenger entry and this entry dispersed the opposition vote, as the theory predicts, the effective number of opposition candidates should increase in royalty-dependent municipalities after the shock. To test this, I estimate a difference-in-differences specification identical to the one used for incumbent vote share and reelection success, using this measure as the dependent variable[7].

[7]Event study estimates validating the parallel trends assumption and difference-in-differences results without covariates are reported in Figure A.8 and Table A.8 in the Appendix, respectively.

Table 5: The Effects of Royalty Loss on Vote Fragmentation

	Effective Number of Opposition Candidates
Loss × Post	0.147***
	(0.056)
Covariates	✓
Municipality FE	✓
Year FE	✓
Observations	8,242
Municipalities	4,785
Election years	2000–2016

Note: Standard errors (in parentheses) clustered by municipality. "Loss" denotes royalty-dependent municipalities after the revenue loss. Covariates are interacted with the post-period (2016) indicator. $^{***}p < 0.01$.

Table 5 shows that the royalty shock significantly increased vote fragmentation in affected municipalities. The coefficient on the interaction term indicates that, after the shock, the effective number of opposition candidates in royalty-dependent municipalities increased by approximately 0.15 relative to control municipalities. This represents approximately an 11 percent increase relative to the pre-treatment mean of 1.34 in royalty-dependent municipalities. Figure 4 further reinforces this finding. Prior to 2016, treated and control municipalities followed comparable trajectories in vote fragmentation. After the royalty loss, however, only royalty-dependent municipalities saw a notable increase in the effective number of opposition candidates.

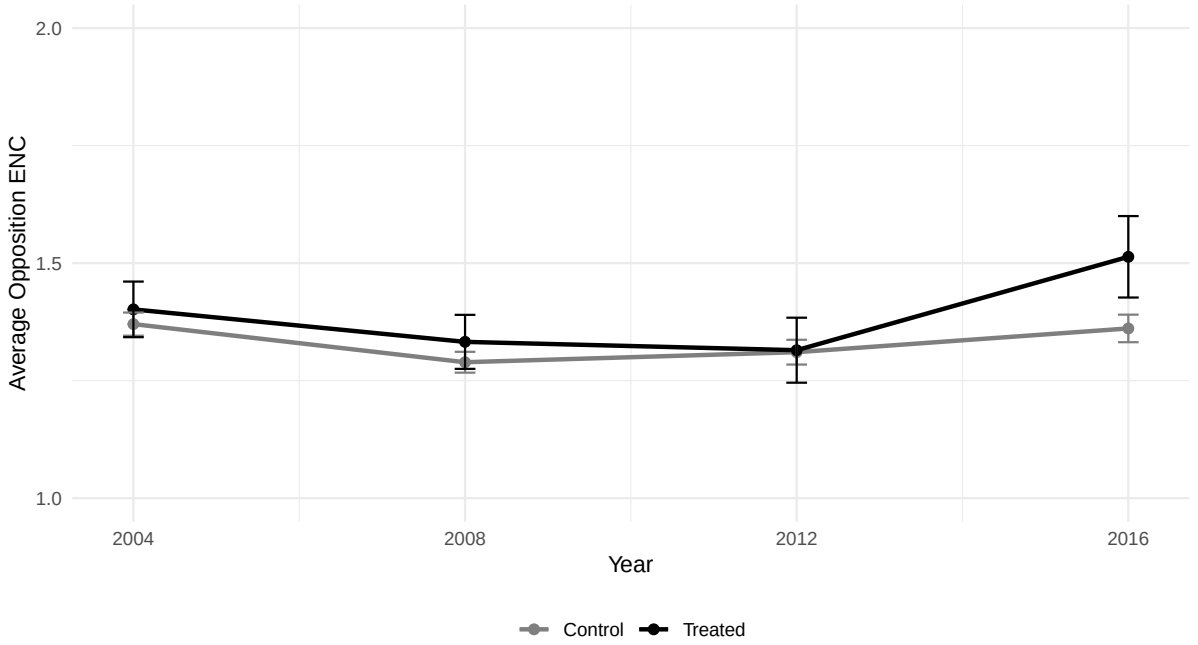

Figure 4: Average Effective Number of Opposition Candidates Over Time by Treatment Status

Note: The figure plots the average effective number of opposition candidates by election year for treated (royalty-losing) and control (non-royalty-dependent) municipalities, computed from vote shares with the incumbent excluded and the remaining shares renormalized. Vertical bars are 95 percent confidence intervals. Electoral data are sourced from the TSE. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

Finally, I split the loss group into terciles of the loss of revenues as a share of the total municipal budget. In the top tercile, the effective number of opposition candidates increased by approximately 0.28 relative to control municipalities, which represents approximately a 22 percent increase relative to the 2012 mean of 1.31. In the two lower terciles there is no significant change (see Table A.20 in the Appendix).

Together with the main results in Table 2, which show that incumbents lost vote share but not reelection, the increase in fragmentation suggests that the additional competition generated by the shock prevented voter coordination and, thus, the opposition's ability to unseat the incumbent, instead dividing its support across a wider field of challengers. Consistent with this, Table A.10 in the Appendix shows that the shock did not reduce (or increase) victory margins in treated relative to control municipalities, suggesting that the additional competitors dispersed the opposition vote broadly rather than

concentrating it behind a close second-place challenger[8].

Alternative Explanations

Having established that the royalty shock increased challenger entry and the opposition vote fragmented, I now consider alternative explanations for the observed decline in incumbent support. I focus on five potential explanations: differences in fiscal spending, variation in candidate quality, differences in victory margins, party-driven incumbent punishment, and campaign-spending advantages. This is especially important for two reasons. First, although the exclusion restriction cannot be directly tested, examining alternative channels allows us to assess whether the instrument affects the outcome through pathways other than challenger competition. Second, it allows me to assess whether incumbents' survival in office can be attributed to factors other than the fragmentation of the opposition vote.

One potential channel is that voters update their beliefs about incumbent competence after shocks that are plausibly outside officials' control, rewarding or punishing incumbents based on observed revenue management (Ashworth et al., 2018). Therefore, in the Brazilian context, it might be that voters interpreted the fiscal consequences of the economic shock as evidence of poor management by local leaders and penalized mayors accordingly. I empirically examine this with TWFE models similar to the candidate entry model, but with four different dependent variables related to public goods spending (health, education, infrastructure, and personnel), using the full dataset of municipality-year observations of incumbents that could run and ran for reelection. Table A.11 in the Appendix shows that the shock did not shift the composition of municipal spending in resource-dependent municipalities.

Another potential mechanism could operate through the quality of the candidates who choose to enter the race, since the availability of rents influences the profile of individ-

[8]Evidence on parallel trends and results from the same specification without covariates can also be found in the Appendix.

uals who pursue elected office (Brollo et al., 2013). Brollo et al. (2013) argue that when fiscal resources expand, the candidate pool deteriorates as rents attract opportunistic and less qualified actors, so under resource *contraction* qualified challengers may instead perceive greater opportunities to compete and have stronger incentives to enter. Voters might then interpret the presence of higher-quality challengers as a signal of weakened incumbent advantage and shift their support away from the officeholder. This channel implies that voters prioritize underlying candidate attributes, often proxied by education, experience, and professional background (Brollo et al., 2013; Novaes & Schiumerini, 2022), at least as much as more immediately observable electoral cues, and that the pool of qualified aspirants is sufficiently responsive to short-term changes in municipal resources. However, Table A.12 in the Appendix shows that the education and experience profiles of candidates who ran were unchanged in royalty-dependent municipalities that lost revenues after the shock.

I also examine whether the margin of victory, defined as the difference in vote share between the first- and second-place candidates, played a role in incumbent survival after the shock, given that previous work has suggested that large margins may shield officeholders from electoral defeat (Cox & Katz, 1996; Mayhew, 1974b). If electoral advantages rather than opposition fragmentation explained incumbents' survival, this should be visible in a widening gap in victory margins between incumbents who won and those who lost reelection in treated municipalities in 2016 relative to prior elections. Figure A.10 in the Appendix nevertheless shows that the post-shock gap between surviving and defeated incumbents broadly resembles the one seen in earlier elections.

Further, the Petrobras scandal, in particular, implicated the then-incumbent party of the president, the Workers' Party (PT), and revealed widespread corruption involving senior party figures and high-ranking government officials. Therefore, one might expect that municipalities with PT mayors drove the average decline in incumbent vote share in affected municipalities. To assess this possibility, I estimate a TWFE model on the full sample of incumbents that includes a triple interaction between treatment status, the post-shock period, and an indicator for PT incumbents to assess whether the treatment

effect led to differential trends in incumbent vote share for PT versus non-PT mayors. Yet, Table A.13 in the Appendix indicates that the vote decline was no steeper for PT incumbents than for non-PT incumbents in treated municipalities.

Finally, incumbents may have held office not because the anti-incumbent vote fragmented but because they spent more campaign resources than their challengers. If incumbents in royalty-losing municipalities had a spending advantage over their strongest challenger after the shock, that advantage could account for their survival. In particular, it could be the case that incumbents in treated localities increased their spending after the shock as a strategy to avoid losing office. I examine this with a TWFE model where the outcome is the incumbent's share of campaign spending relative to the strongest challenger. Table A.14 in the Appendix shows that, after the shock, treated officeholders had no spending advantage. The same holds across a set of related measures of campaign resources, including the ratio of the incumbent's spending to the top challenger's and the concentration of campaign spending across candidates.

Conclusion

This study argued that whether vote loss translates into office loss for incumbents depends on how fragmented voter support is across the pool of opposition candidates. I show that when the electorate's support disperses across challengers, the incumbent can retain office even after losing votes. I contend that this dynamic should be most visible in systems where staying in power requires obtaining more votes than any other candidate, rather than a majority. To test this argument, I turn to the adverse conditions of an economic shock.

I focus on the case of Brazil, where mayoral elections in most municipalities are decided by plurality rule. Using a difference-in-differences design, I provide causal evidence that, in oil-dependent municipalities that experienced royalty revenue losses due to the Petrobras scandal and the oil price collapse in 2014, incumbents running for reelection lost vote share compared to their counterparts elsewhere, but they were not systematically

less likely to be reelected.

I argue that the loss of revenues weakened incumbents' standing and encouraged more challenger candidates to compete in the race. The resulting increase in competitors, in turn, made it harder for voters to coordinate behind a single viable alternative, dispersing the anti-incumbent vote across multiple candidacies. Using a second difference-in-differences design and an instrumental variables approach, I first demonstrate that there was an increase in the number of competing candidates after the shock. I further show that this increased pool of competitors has a substantial negative effect on incumbent vote share, yet does not affect their reelection.

A final difference-in-differences shows that the shock increased vote fragmentation, measured as the effective number of opposition candidates, in affected municipalities. This fits my theory that, in times of adverse conditions, incumbents can become shielded from losing office when votes fragment across multiple opposition options under certain electoral rules, since no single alternative gathers sufficient votes to unseat the officeholder.

Overall, these findings contribute to a longstanding literature on how electoral rules shape coordination among voters (Cox, 1997; Myerson & Weber, 1993). That work has largely asked how many candidates voters coordinate around and when strategic voting consolidates support behind the strongest contenders (Duverger, 1954; Riker, 1982), including how plurality systems are commonly expected to discourage fragmentation and push the electorate toward a small number of viable candidates (Duverger, 1954; Fujiwara, 2011). First, I extend it by showing that the degree to which voters coordinate around opposition alternatives also shapes whether incumbents are held accountable, and in particular whether they remain in or are removed from office. Second, I identify a condition under which the coordination that plurality is expected to produce does not hold: when adverse conditions (such as an economic shock) draw additional challengers into the race, the incentives said to move voters toward a single alternative can be insufficient to prevent the opposition vote from dispersing. The structure of competition, alongside the rules that govern it, therefore conditions whether vote loss translates into removal

from office.

Finally, this paper also opens avenues for future research. One important avenue would be to expand the theoretical and analytical scope to other settings where a single office is decided by plurality with no majority requirement, such as municipal races in Mexico, mayoral and congressional races in the Philippines, and lower-house and state-assembly races in India. Another possibility is to examine whether similar dynamics operate in contexts with different electoral rules, such as proportional representation, where the fragmentation mechanism may function differently. In addition, while this paper focuses on the first election after the shock, future work could explore whether the increased entry of challengers has lasting effects on the local party system, or whether competition returns to pre-shock levels once revenues recover. Finally, the finding that incumbents may survive despite withdrawn support, even in settings where institutions are expected to promote coordination, raises normative questions about the quality of representation across democracies.

References

Achen, C. H., & Bartels, L. M. (2016). *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton University Press.

Achen, C. H., & Bartels, L. M. (2018). Statistics as If Politics Mattered: A Reply to Fowler and Hall. *The Journal of Politics*, *80*(4), 1437–1451.

Ahlquist, J., Copelovitch, M., & Walter, S. (2020). The Political Consequences of External Economic Shocks: Evidence from Poland. *American Journal of Political Science*, *64*(4), 904–920.

Alcañiz, I., & Hellwig, T. (2011). Who's to Blame? The Distribution of Responsibility in Developing Democracies. *British Journal of Political Science*, *41*(2), 389–411.

Andersen, J. J., & Ross, M. L. (2014). The Big Oil Change: A Closer Look at the Haber–Menaldo Analysis. *Comparative Political Studies*, *47*(7), 993–1021.

Ashworth, S., de Mesquita, E. B., & Friedenberg, A. (2018). Learning about Voter Rationality. *American Political Science Review*, *112*(1), 37–57.

Baffes, J., Kose, M. A., Ohnsorge, F., & Stocker, M. (2015). *The Great Plunge in Oil Prices: Causes, Consequences, and Policy Responses* (Policy Research Note No. PRN/15/01). World Bank Group. Washington, DC.

Black, G. S. (1972). A Theory of Political Ambition: Career Choices and the Role of Structural Incentives. *American Political Science Review*, *66*(4), 1260–1288.

Brollo, F., Nannicini, T., Perotti, R., & Tabellini, G. (2013). The Political Resource Curse. *American Economic Review*, *103*(5), 1759–1796.

Brückner, M., Ciccone, A., & Tesei, A. (2012). Oil Price Shocks, Income, and Democracy. *Review of Economics and Statistics*, *94*(2), 389–399.

Callaway, B., Goodman-Bacon, A., & Sant'Anna, P. H. (2025). Difference-in-differences with a Continuous Treatment. *National Bureau of Economic Research*.

Campello, D., & Zucco, C. (2020). *The Volatility Curse: Exogenous Shocks and Representation in Resource-Rich Democracies*. Cambridge University Press.

Campello, D., & Zucco Jr., C. (2016). Presidential Success and the World Economy. *The Journal of Politics*, *78*(2), 589–602.

Carey, J. M., & Hix, S. (2011). The Electoral Sweet Spot: Low-Magnitude Proportional Electoral Systems. *American Journal of Political Science*, *55*(2), 383–397.

Carlin, R. E., Love, G. J., & Martínez-Gallardo, C. (2014). Cushioning the Fall: Scandals, Economic Conditions, and Executive Approval. *Political Behavior*, *36*(4), 793–820.

Carlin, R. E., Singer, M. M., & Zechmeister, E. J. (2015). *The Latin American Voter: Pursuing Representation and Accountability in Challenging Contexts*. University of Michigan Press.

Carreri, M., & Dube, O. (2017). Do Natural Resources Influence Who Comes to Power, and How? *Journal of Politics*, *79*(2), 502–518.

Caselli, F., & Michaels, G. (2013). Do Oil Windfalls Improve Living Standards? Evidence from Brazil. *American Economic Journal: Applied Economics*, *5*(1), 208–238.

Caselli, F., & Tesei, A. (2016). Resource Windfalls, Political Regimes, and Political Stability. *Review of Economics and Statistics*, *98*(3), 573–590.

Cheibub, J. A., & Przeworski, A. (1999). Democracy, Elections, and Accountability for Economic Outcomes. In A. Przeworski, S. C. Stokes, & B. Manin (Eds.), *Democracy, Accountability, and Representation* (pp. 222–250). Cambridge University Press.

Choi, E., & Woo, J. (2010). Political Corruption, Economic Performance, and Electoral Outcomes: A Cross-National Analysis. *Contemporary Politics*, *16*(3), 249–262.

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

Cox, G. W. (1999). Electoral Rules and Electoral Coordination. *Annual Review of Political Science*, *2*, 145–161.

Cox, G. W., & Katz, J. N. (1996). Why Did the Incumbency Advantage in U.S. House Elections Grow? *American Journal of Political Science*, *40*(2), 478–497.

Dano, K., Ferlenga, F., Galasso, V., Le Pennec, C., & Pons, V. (2025). Coordination and Incumbency Advantage in Multi-Party Systems—Evidence from French Elections. *Journal of the European Economic Association*, *23*(4), 1563–1597.

De Benedictis-Kessner, J. (2018). Off-Cycle and Out of Office: Election Timing and the Incumbency Advantage. *The Journal of Politics*, *80*(1), 119–132.

De Benedictis-Kessner, J., & Warshaw, C. (2020). Accountability for the Local Economy at All Levels of Government in United States Elections. *American Political Science Review*, *114*(3), 660–676.

Downs, A. (1957). *An Economic Theory of Democracy*. Harper & Row.

Duch, R. M., & Stevenson, R. T. (2008). *The Economic Vote: How Political and Economic Institutions Condition Election Results*. Cambridge University Press.

Duverger, M. (1954). *Political Parties: Their Organization and Activity in the Modern State*. Wiley.

Egami, N., & Yamauchi, S. (2023). Using Multiple Pretreatment Periods to Improve Difference-in-Differences and Staggered Adoption Designs. *Political Analysis*, *31*(2), 195–212.

Fearon, J. D. (1999). Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance. In A. Przeworski, S. C. Stokes, & B. Manin (Eds.), *Democracy, Accountability, and Representation* (pp. 55–97). Cambridge University Press.

Ferraz, C., & Finan, F. (2008). Exposing Corrupt Politicians: The Effects of Brazil's Publicly Released Audits on Electoral Outcomes. *The Quarterly Journal of Economics*, *123*(2), 703–745.

Ferree, K. E. (2011). *Framing the Race in South Africa: The Political Origins of Racial Census Elections*. Cambridge University Press.

Ferrer, S. (2023). Opposition Fragmentation Facilitates Electoral Accountability: Evidence from a Survey Experiment. *Party Politics*, *29*(1), 112–124.

Fiorina, M. P. (1978). Economic Retrospective Voting in American National Elections: A Micro-Analysis. *American Journal of Political Science*, *22*(2), 426–443.

Fowler, A., & Hall, A. B. (2014). Disentangling the Personal and Partisan Incumbency Advantages: Evidence from Close Elections and Term Limits. *Quarterly Journal of Political Science*, *9*(4), 501–531.

Fujiwara, T. (2011). A Regression Discontinuity Test of Strategic Voting and Duverger's Law. *Quarterly Journal of Political Science*, *6*(3-4), 197–233.

G1. (2015a, January 29). *Redução do repasse dos royalties pode gerar corte salarial em Campos*. G1 Norte Fluminense. Retrieved July 29, 2026, from https://g1.globo.com/rj/norte-fluminense/noticia/2015/01/reducao-do-repasse-dos-royalties-pode-gerar-corte-salarial-em-campos.html

G1. (2015b, February 9). *Corte em repasse dos royalties gera redução de despesas em São Fidélis*. G1 Norte Fluminense. Retrieved July 29, 2026, from https://g1.globo.com/rj/norte-fluminense/noticia/2015/02/corte-em-repasse-dos-royalties-gera-reducao-de-despesas-em-sao-fidelis.html

Goldberg, E., Wibbels, E., & Mvukiyehe, E. (2008). Lessons from Strange Cases: Democracy, Development, and the Resource Curse in the U.S. States. *Comparative Political Studies*, *41*(4-5), 477–514.

Golveia, F. J. P., Matta, L. G. d., Silva Neto, R. e., & Freitas, F. (2023). Reflexos da crise do petróleo de 2014 no emprego, na renda e na educação superior em municípios de altas rendas petrolíferas do norte fluminense. *Petróleo, Royalties & Região*, 79–104.

Gomes, T. G., & de Oliveira Medeiros, C. R. (2019). Construindo e desconstruindo escândalos de corrupção: a Operação Lava-Jato nas interpretações da *Veja* e *Carta Capital*. *Revista Organizações & Sociedade*, *26*(90), 457–485.

Greene, K. F. (2007). *Why Dominant Parties Lose: Mexico's Democratization in Comparative Perspective*. Cambridge University Press.

Haber, S., & Menaldo, V. (2011). Do Natural Resources Fuel Authoritarianism? A Reappraisal of the Resource Curse. *American Political Science Review*, *105*(1), 1–26.

Healy, A., & Malhotra, N. (2009). Myopic Voters and Natural Disaster Policy. *American Political Science Review*, *103*(3), 387–406.

Hellwig, T., & Samuels, D. (2008). Electoral Accountability and the Variety of Democratic Regimes. *British Journal of Political Science*, *38*(1), 65–90.

Hernández, E., & Kriesi, H. (2016). The Electoral Consequences of the Financial and Economic Crisis in Europe. *European Journal of Political Research*, *55*(2), 203–224.

Hobolt, S. B., & Tilley, J. (2016). Fleeing the Centre: The Rise of Challenger Parties in the Aftermath of the Euro Crisis. *West European Politics*, *39*(5), 971–991.

Jacobson, G. C. (1989). Strategic Politicians and the Dynamics of U.S. House Elections, 1946–86. *American Political Science Review*, *83*(3), 773–793.

Jacobson, G. C., & Kernell, S. (1983). *Strategy and Choice in Congressional Elections* (2nd). Yale University Press.

Jensen, N., & Wantchekon, L. (2004). Resource Wealth and Political Regimes in Africa. *Comparative Political Studies*, *37*(7), 816–841.

Kam, C., Bertelli, A. M., & Held, A. (2020). The Electoral System, the Party System and Accountability in Parliamentary Government. *American Political Science Review*, *114*(3), 744–760.

Klašnja, M. (2015). Corruption and the Incumbency Disadvantage: Theory and Evidence. *The Journal of Politics*, *77*(4), 928–942.

Klašnja, M., & Titiunik, R. (2017). The Incumbency Curse: Weak Parties, Term Limits, and Unfulfilled Accountability. *American Political Science Review*, *111*(1), 129–148.

Krause, S., & Méndez, F. (2009). Corruption and Elections: An Empirical Study for a Cross-Section of Countries. *Economics and Politics*, *21*(2), 179–210.

Laakso, M., & Taagepera, R. (1979). "Effective" Number of Parties: A Measure with Application to West Europe. *Comparative Political Studies*, *12*(1), 3–27.

Lago, I., & Martínez i Coma, F. (2024). The "Odd–Even Effect": The Link Between the Number of Parties and District Magnitude. *European Journal of Political Research*, *63*(4), 1471–1490.

Lewis-Beck, M. S., & Stegmaier, M. (2000). Economic Determinants of Electoral Outcomes. *Annual Review of Political Science*, *3*, 183–219.

Lucardi, A. (2019). The Effect of District Magnitude on Electoral Outcomes: Evidence from Two Natural Experiments in Argentina. *British Journal of Political Science*, *49*(2), 557–577.

Maeda, K. (2010). Divided We Fall: Opposition Fragmentation and the Electoral Fortunes of Governing Parties. *British Journal of Political Science*, *40*(2), 419–434.

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

Mainwaring, S., Gervasoni, C., & España-Najera, A. (2017). Extra- and within-system electoral volatility. *Party Politics*, *23*(6), 623–635.

Mainwaring, S., & Su, Y.-T. P. (2021). Electoral Volatility in Latin America, 1932–2018. *Studies in Comparative International Development*, *56*(3), 271–296.

Manin, B., Przeworski, A., & Stokes, S. C. (1999). Elections and Representation. In A. Przeworski, S. C. Stokes, & B. Manin (Eds.), *Democracy, Accountability, and Representation* (pp. 29–54). Cambridge University Press.

Manzetti, L., & Rosas, G. (2015). Corruption and the Latin American Voter. In *The latin american voter: Pursuing representation and accountability in challenging contexts* (pp. 300–323). University of Michigan Press.

Manzetti, L., & Wilson, C. J. (2007). Why do corrupt governments maintain public support? *Comparative Political Studies*, *40*(8), 949–970.

Margalit, Y., & Solodoch, O. (2025). Electoral responses to economic crises [Forthcoming]. *American Journal of Political Science*.

Mayhew, D. R. (1974a). *Congress: The Electoral Connection*. Yale University Press.

Mayhew, D. R. (1974b). Congressional Elections: The Case of the Vanishing Marginals. *Polity*, *6*(3), 295–317.

Mead, D., & Stiger, P. (2015). The 2014 plunge in import petroleum prices: What happened? *Beyond the Numbers: Global Economy*, *4*(9). https://www.bls.gov/opub/btn/volume-4/pdf/the-2014-plunge-in-import-petroleum-prices-what-happened.pdf

Monteiro, J., & Ferraz, C. (2012). *Does Oil Make Leaders More Accountable? Evidence from Brazil's Offshore Oil Boom* [Working paper, Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio), Department of Economics].

Murillo, M. V., & Visconti, G. (2017). Economic Performance and Incumbents' Support in Latin America. *Electoral Studies*, *45*, 180–190.

Myerson, R. B., & Weber, R. J. (1993). A Theory of Voting Equilibria. *American Political Science Review*, *87*(1), 102–114.

Novaes, L. M., & Schiumerini, L. (2022). Commodity Shocks and Incumbency Effects. *British Journal of Political Science*, *52*(4), 1689–1708.

Paler, L., Springman, J., Grossman, G., & Pierskalla, J. (2023). Oil Discoveries and Political Windfalls: Evidence on Presidential Support in Uganda. *Political Science Research and Methods*, 1–10.

Pereira, C., & Melo, M. A. (2015). Reelecting Corrupt Incumbents in Exchange for Public Goods: Rouba Mas Faz in Brazil. *Latin American Research Review*, *50*(4), 88–115.

Powell, E. N., & Tucker, J. A. (2014). Revisiting Electoral Volatility in Post-Communist Countries: New Data, New Results and New Approaches. *British Journal of Political Science*, *44*(1), 123–147.

Powell, G. B., & Whitten, G. D. (1993). A Cross-National Analysis of Economic Voting: Taking Account of the Political Context. *American Journal of Political Science*, *37*(2), 391–414.

Powell, J., G. Bingham. (2000). *Elections as Instruments of Democracy: Majoritarian and Proportional Visions*. Yale University Press.

Rambachan, A., & Roth, J. (2023). A More Credible Approach to Parallel Trends. *Review of Economic Studies*, *90*(5), 2555–2591.

Rashkova, E. R., & Su, Y.-P. (2024). Exogenous Shocks and Party System Development: What is the Effect on New Party Entry? *Representation*, 1–13.

Riker, W. H. (1982). The Two-Party System and Duverger's Law: An Essay on the History of Political Science. *American Political Science Review*, *76*(4), 753–766.

Roberts, K. M., & Wibbels, E. (1999). Party systems and Electoral volatility in Latin America: A test of economic, institutional, and structural explanations. *American Political Science Review*, *93*(3), 575–590.

Robinson, J. A., Torvik, R., & Verdier, T. (2006). Political Foundations of the Resource Curse. *Journal of Development Economics*, *79*(2), 447–468.

Ross, M. L. (2001). Does Oil Hinder Democracy? *World Politics*, *53*(3), 325–361.

Rudolph, L., & Däubler, T. (2016). Holding Individual Representatives Accountable: The Role of Electoral Systems. *The Journal of Politics*, *78*(3), 746–762.

Samuels, D. J. (2004). Presidentialism and Accountability for the Economy in Comparative Perspective. *American Political Science Review*, *98*(3), 425–436.

Silva, E. A. B. (2017). *Corrupção e opinião pública: o escândalo da Lava Jato no governo Dilma Rousseff* [Doctoral dissertation, Universidade Federal de Minas Gerais].

Singer, M. M. (2013). Economic voting in an era of non-crisis: The changing electoral agenda in Latin America, 1982–2010. *Comparative Politics*, *45*(2), 169–185.

Speck, B. W., & Cervi, E. U. (2016). Dinheiro, Tempo e Memória Eleitoral: Os Mecanismos Que Levam ao Voto nas Eleições para Prefeito em 2012. *Dados*, *59*(1), 53–90.

Stone, W. J., Maisel, L. S., & Maestas, C. D. (2004). Quality Counts: Extending the Strategic Politician Model of Incumbent Deterrence. *American Journal of Political Science*, *48*(3), 479–495.

Tavits, M. (2005). The Development of Stable Party Support: Electoral Dynamics in Post-Communist Europe. *American Journal of Political Science*, *49*(2), 283–298.

Tertytchnaya, K., De Vries, C. E., Solaz, H., & Doyle, D. (2018). When the Money Stops: Fluctuations in Financial Remittances and Incumbent Approval in Central Eastern Europe, the Caucasus and Central Asia. *American Political Science Review*, *112*(4), 758–774.

Titiunik, R. (2011). *Incumbency Advantage in Brazil: Evidence from Municipal Mayor Elections* [Working Paper, University of Michigan].

Uppal, Y. (2009). The disadvantaged incumbents: estimating incumbency effects in Indian state legislatures. *Public Choice*, *138*(1), 9–27.

Vasconcellos, F., Schmitt, L. G., & Junior, P. R. (2015, April 9). *Dos 88 municípios com direito a royalties do petróleo, 87 têm perdas de R$ 289 milhões*. O Globo. Retrieved July 29, 2026, from https://oglobo.globo.com/rio/dos-88-municipios-com-direito-royalties-do-petroleo-87-tem-perdas-de-289-milhoes-15820037

Wantchekon, L. (2002). Why Do Resource Abundant Countries Have Authoritarian Governments? *Journal of African Development*, *5*(2), 145–176.

Appendix

Contents

Appendix **45**

A.1 Royalty Revenue Loss among Treated Municipalities **47**

A.2 DiD Identifying Assumptions **49**

A.2.1 Parallel Trends 49

A.2.2 Relative-Magnitudes Sensitivity for the Reelection Null 50

A.2.3 Exogeneity of Treatment, Sample Composition, and No Anticipation . . 51

A.2.4 Spillovers and the SUTVA Assumption 52

A.2.4.1 Relative-Magnitudes Sensitivity for the Third-Order Ring 55

A.3 Main DiD Results without Covariates **57**

A.4 Main DiD Results with Outlier Municipality **58**

A.5 Parallel Trends: Number of Candidates **59**

A.6 Number of Candidates DiD without Covariates **60**

A.7 First Stage IVs **61**

A.7.1 Candidate Entry 61

A.7.2 New Candidate Entry 62

A.8 Parallel Trends: Effective Number of Opposition Candidates **64**

A.9 ENC DiD Results without Covariates **65**

A.10 Victory Margin Analyses **66**

A.10.1 Parallel Trends: Victory Margins 67

A.10.2 Winning Margins DiD Results without Covariates 67

A.10.3 Victory Margins DiD Results 68

A.11 Alternative Explanations **70**

A.11.1 Fiscal Spending 70

A.11.2 Candidate Quality 71

A.11.3 Differences in Victory Margins 72

A.11.4 Party-driven Incumbent Punishment 73

A.11.5 Campaign Spending 74

A.12 Runoff Elections **77**

A.13 Null, Blank, and Protest Voting **79**

A.14 Windfall Municipalities **81**

A.15 Dependence-Based Treatment Specification **84**

A.16 Loss-Intensity Disaggregation **85**

A.16.1 Vote Share and Reelection by Loss Tercile 86

A.16.2 Candidate Entry by Loss Tercile 87

A.16.3 Vote Fragmentation by Loss Tercile 88

A.17 Latitude and Longitude as Baseline Covariates **89**

A.18 Table Results with Covariate Coefficients **91**

A.1 Royalty Revenue Loss among Treated Municipalities

Figure A.1 displays the cumulative distribution of royalty revenue loss among treated municipalities, measured as the decline in the share of royalties in the municipal budget between 2014 and 2016. Almost 90 percent of treated municipalities lost up to four percentage points of their total budget, which motivates the binary treatment indicator used in the main analysis and the four percentage point threshold used in the loss-intensity disaggregation.

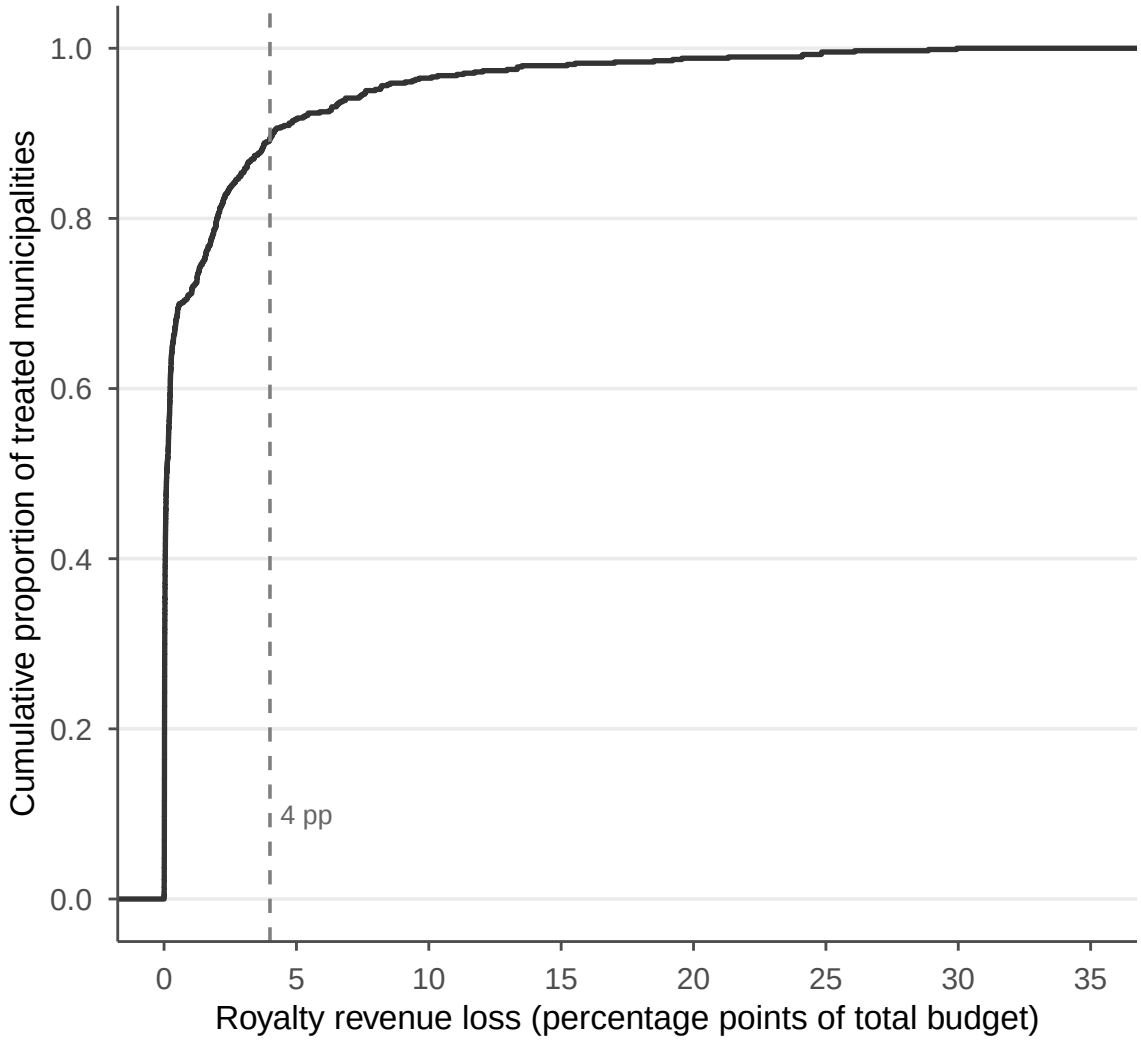

Figure A.1: Cumulative Distribution of Royalty Revenue Loss among Treated Municipalities

Note: Cumulative share of treated municipalities (vertical axis) by the size of their royalty loss between 2014 and 2016, measured as the decline in royalties as a share of the municipal budget (horizontal axis). Almost 90 percent of treated municipalities lost up to four percentage points of their budget, which motivates the binary treatment indicator and the four percentage point threshold used in the loss-intensity analysis. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

A.2 DiD Identifying Assumptions

A.2.1 Parallel Trends

To assess the validity of the parallel trends assumption, Figure A.2 presents event study specifications that interact treatment status with election-year dummies, with 2012 (the election before the economic shock) as the reference year. Pre-treatment coefficients are close to zero and statistically indistinguishable from zero, consistent with parallel trends between royalty-dependent municipalities that lost revenues and non-dependent municipalities. In 2016, the incumbent vote panel shows a clear significant decline (at the 0.05 level) for royalty-dependent municipalities. The figure also shows that the incumbency disadvantage does not extend to subsequent post-treatment years (2020 and 2024). The effect for the reelection dependent variable is insignificant in post-treatment years.

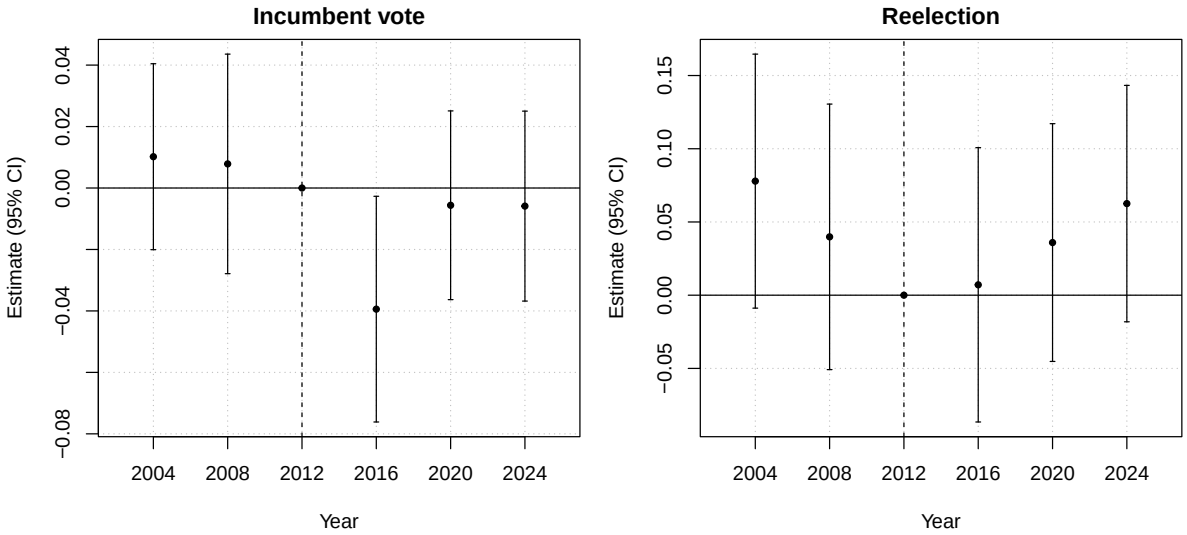

Figure A.2: Event-Study Estimates: Incumbent Vote and Reelection Outcomes

Note: Each point is the coefficient on the interaction of treatment status with an election-year dummy, with 2012 (the election before the shock) as the reference year. The left panel is the incumbent vote-share outcome and the right panel is the reelection outcome. Vertical bars are 95 percent confidence intervals from standard errors clustered by municipality. Electoral data are sourced from the TSE. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

The reelection panel of Figure A.2 shows clean pre-trends except at 2004, where the coefficient is 8.1 percentage points and significant at the 10 percent level. To account for this, I apply in the next subsection the relative-magnitudes sensitivity analysis of Rambachan and Roth (2023), which bounds the bias the pre-period deviation could generate.

A.2.2 Relative-Magnitudes Sensitivity for the Reelection Null

This subsection presents the relative-magnitudes sensitivity analysis of Rambachan and Roth (2023) for the 2004 reelection violation. The procedure treats the largest observed pre-period parallel-trends violation as a calibrated bound on the unobserved post-period violation, scaled by a factor $\bar{M}$, and reports how the confidence interval on the post-period effect widens as $\bar{M}$ grows. A "breakdown" value is the smallest $\bar{M}$ at which the effect loses significance.

The main text notes a marginal (10 percent) parallel-trends violation in 2004 for the pooled reelection event study. The 2004 pre-period coefficient is +8.1 percentage points and the 2016 reelection coefficient is +1.6 percentage points. The mechanical decomposition of the bound follows directly from these two quantities. At $\bar{M} = 1$, which allows the post-period violation to be as large as the worst observed pre-period violation, the post-period violation is bounded in $[-8.1, +8.1]$ percentage points, so the true reelection effect ranges over $[1.6 - 8.1,\ 1.6 + 8.1] = [-6.5, +9.7]$ percentage points before sampling uncertainty. Adding sampling uncertainty yields the bounded 95 percent confidence interval of $[-12.5, +20.3]$ percentage points, which contains zero. Table A.1 reports the full relative-magnitudes breakdown across restrictions. The null average reelection effect is therefore robust to a post-period violation as large as the one observed in 2004.

Table A.1: Relative-Magnitudes Sensitivity: Pooled Reelection Effect

Restriction	95% CI lower	95% CI upper
Original (FE event study)	−7.33	+10.48
Relative magnitudes, $\bar{M} = 0.5$	−9.00	+14.56
Relative magnitudes, $\bar{M} = 1.0$	−12.52	+20.31

Note: Sensitivity of the pooled 2016 reelection effect (in percentage points) to a post-period parallel-trends violation bounded at $\bar{M}$ times the largest observed pre-period violation. All intervals contain zero.

A.2.3 Exogeneity of Treatment, Sample Composition, and No Anticipation

Beyond parallel trends, the difference-in-differences design relies on additional identifying assumptions. First, treatment assignment must be exogenous to factors that also determine electoral outcomes. As discussed above, royalties are allocated strictly according to federal legal requirements and are exogenous to party dynamics: even with variation in the incumbent politician or party that holds power, there cannot be any disruptions in the distribution of royalties due to changes in political leadership. Since the criteria governing royalty transfers are determined by federal agencies, local political actors cannot influence whether or how much revenue their municipality receives, satisfying the exogeneity of treatment assignment.

Second, for the estimates to be unbiased, the composition of the sample should not change differentially across treatment and control groups as a consequence of the shock. A particular concern is that the revenue loss may have discouraged incumbents in affected municipalities from seeking reelection, which would introduce selection bias if the remaining candidates were systematically stronger or weaker. To assess this possibility, I compare the share of incumbents who ran for reelection across treated and control municipalities. Among royalty-dependent municipalities that lost revenues, 41.5 percent of incumbents sought reelection, compared to 39.4 percent in non-dependent municipalities. Nevertheless, this 2.1 percentage point difference is not statistically significant (p =

0.30). This suggests that the shock did not systematically alter which incumbents chose to run, supporting the validity of the sample composition.

Third, the no anticipation assumption requires that treated municipalities did not adjust their behavior before the shock materialized. This is plausible in the present case, as both the Petrobras scandal and the global oil price collapse were sudden and largely unanticipated events, making it unlikely that local actors could have systematically altered their electoral strategies in advance.

A.2.4 Spillovers and the SUTVA Assumption

A final identifying assumption is the stable unit treatment value assumption (SUTVA). It requires that the treatment status of one municipality does not affect electoral outcomes in other treated and control units. This assumption is supported by two features of the setting. For one, the shock was simultaneous, affecting all treated municipalities at the same time. Moreover, the treatment itself operates through a decrease in the municipal budget, which is inherently local in scope. Since royalty revenues fund services and investments within the receiving municipality, their contraction is unlikely to systematically alter economic or electoral conditions in control municipalities.

In addition to these conceptual considerations, I run placebo event studies on three rings of untreated municipal neighbors of treated municipalities, defined by queen contiguity from the GADM Brazilian shapefile. The first ring consists of untreated municipalities directly adjacent to a treated municipality, the second of untreated municipalities adjacent to a first-ring municipality, and the third of untreated municipalities one step farther out. The rings are overlapping, since a municipality can be a first-order neighbor of one treated municipality and a second-order neighbor of another, and each ring is compared separately to non-dependent municipalities at the relevant distance or beyond. If the shock genuinely spilled over, these placebo neighbors should display effects in the same direction as the treated municipalities.

Figure A.3 reports the benchmark event study, which compares the royalty-loss

treated municipalities against the "far" control of non-dependent municipalities that lie outside all three rings, the cleanest spillover-free comparison. Its 2016 vote-share coefficient is close to the headline estimate but less precisely estimated, likely because the far control is a smaller subset of the full control group. It validates the loss effect against a contamination-free control rather than reproducing the main result's precision. Figures A.4, A.5, and A.6 report the first-, second-, and third-order ring placebos. At 2016, the vote-share coefficients are essentially zero across all three rings, ranging between +0.009 and +0.016 and statistically insignificant in every case. The reelection coefficients are directionally negative. The first- and second-order tests have clean pre-trends. The third-order vote-share comparison shows a parallel-trends violation in 2004 at the 5 percent level, which I address with the relative-magnitudes sensitivity analysis of Rambachan and Roth (2023) below. Taken together, the near-zero vote-share placebos in the inner rings, where any contamination would be strongest, provide little evidence of spillover and support the SUTVA assumption underlying the main design.

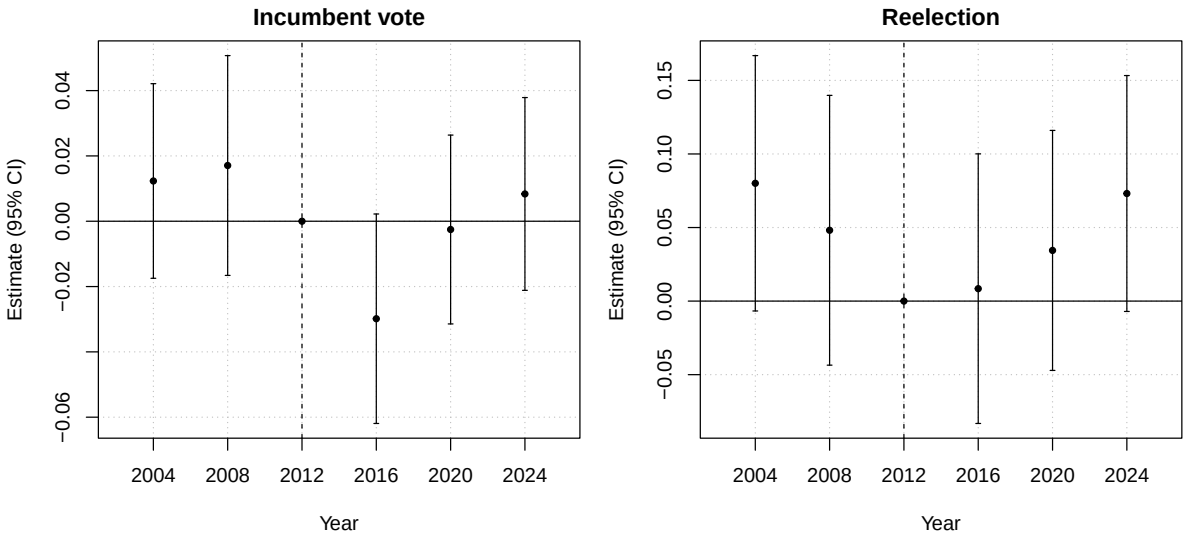

Figure A.3: Event Study: Benchmark Treated Municipalities (SUTVA Placebo Reference)

Note: Benchmark event study for treated municipalities, shown as the reference against which the neighbor-ring placebos in the following figures are compared. Points are treatment-by-year coefficients on incumbent vote share and reelection, with 2012 as the reference year and 95 percent confidence intervals clustered by municipality. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

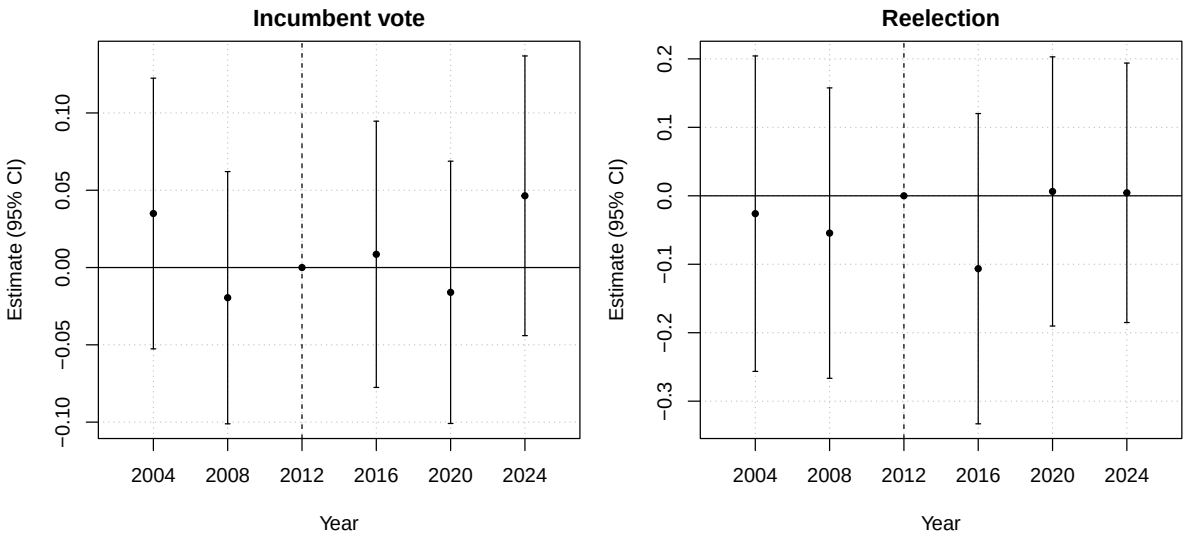

Figure A.4: Event Study: First-Order Ring of Untreated Neighbors

Note: Placebo event study for the first-order ring of untreated municipalities directly adjacent to a treated municipality, compared with non-dependent municipalities. Points are treatment-by-year coefficients with 2012 as the reference year and 95 percent confidence intervals clustered by municipality. Near-zero 2016 coefficients indicate little evidence of spillover onto immediate neighbors. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

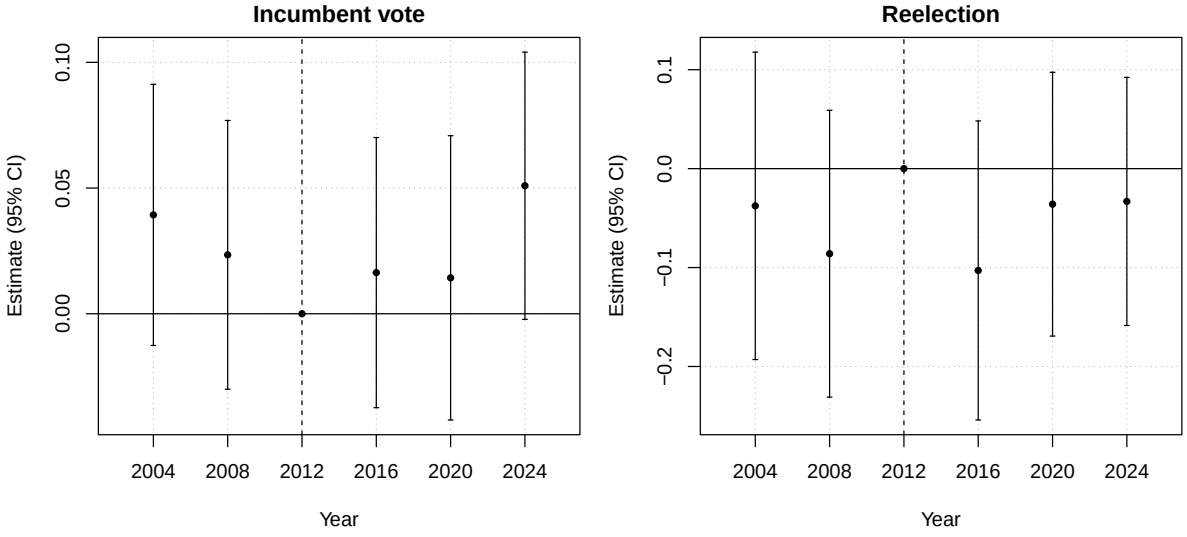

Figure A.5: Event Study: Second-Order Ring of Untreated Neighbors

Note: Placebo event study for the second-order ring of untreated neighbors, those adjacent to a first-ring municipality, compared with non-dependent municipalities. Conventions as in the first-order ring figure. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

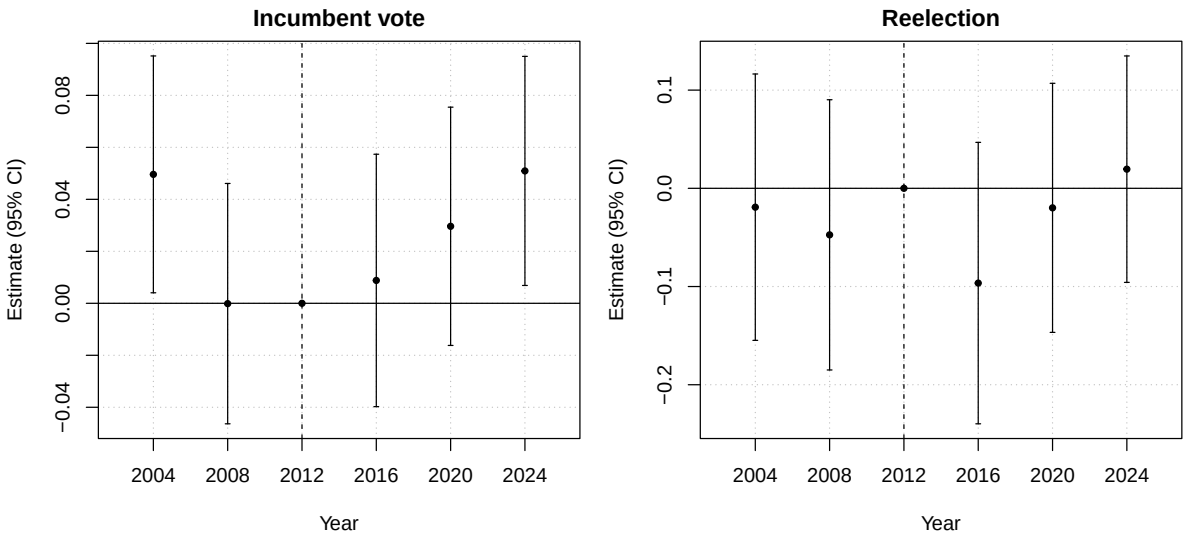

Figure A.6: Event Study: Third-Order Ring of Untreated Neighbors

Note: Placebo event study for the third-order ring of untreated neighbors, compared with non-dependent municipalities. Conventions as in the earlier ring figures. The third-order vote-share comparison shows a 2004 pre-period violation, which is bounded by the relative-magnitudes sensitivity analysis in Table A.2. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

A.2.4.1 Relative-Magnitudes Sensitivity for the Third-Order Ring

Because the third-order ring vote-share comparison violates parallel trends in 2004, I apply the Rambachan and Roth (2023) relative-magnitudes sensitivity analysis, which treats the worst observed pre-period coefficient as a calibrated bound on the unobserved post-period violation. The 2004 pre-period coefficient for the third-order ring is +5.0 percentage points. Under the relative-magnitudes restriction at $\bar{M} = 1$, the unobserved 2016 post-period violation is therefore bounded between −5.0 and +5.0 percentage points. The observed 2016 placebo coefficient of +0.9 percentage points then implies a true effect in $[-4.1, +5.9]$ percentage points absent sampling uncertainty. Adding sampling uncertainty yields the bounded 95 percent confidence interval $[-10.0, +11.4]$ percentage points, which contains zero. Table A.2 reports the full set of bounds. The null third-order spillover effect on vote share is robust to a post-period violation as large as the one observed in 2004.

Table A.2: Relative-Magnitudes Sensitivity: Third-Order Ring Vote-Share Spillover

Restriction	95% CI lower	95% CI upper
Original (FE event study)	−3.82	+5.60
Relative magnitudes, $\bar{M} = 0.5$	−6.30	+7.78
Relative magnitudes, $\bar{M} = 1.0$	−9.96	+11.35

Note: Rambachan and Roth (2023) relative-magnitudes sensitivity for the 2016 third-order ring vote-share placebo (in percentage points). The relative-magnitudes restriction at $\bar{M} = m$ bounds the post-period parallel-trends violation at m times the largest observed pre-period violation. All intervals contain zero.

A.3 Main DiD Results without Covariates

Table A.3 reports the main difference-in-differences estimates without the set of baseline covariates. The pattern is similar to the covariate-adjusted models in the main text: incumbents in royalty-loss municipalities lose vote share after the shock, while their probability of reelection is unchanged. The vote-share point estimate is somewhat larger without covariates, and statistical significance on both outcomes is consistent with the adjusted specification.

Table A.3: Baseline Effects of Royalty Loss on Incumbent Vote and Reelection Success

	Incumbent Vote	Reelection Success
	(1)	(2)
Loss × Post	-0.041***	-0.026
	(0.015)	(0.047)
Covariates		
Municipality FE	✓	✓
Year FE	✓	✓
Observations	8,511	8,511
Municipalities	4,878	4,878
Election years	2000–2016	2000–2016

Note: Municipality and election-year fixed effects. Standard errors (in parentheses) clustered by municipality.

***$p < 0.01$.

A.4 Main DiD Results with Outlier Municipality

The main sample excludes one extreme outlier whose royalty loss exceeded 60 percent of its total budget. Table A.4 re-estimates the covariate-adjusted main specification with this municipality included. The estimates are virtually identical to those in the main text, confirming that the headline results are not driven by the excluded outlier.

Table A.4: The Electoral Effects of the Royalty Shock: Outlier Included

	Incumbent Vote	Reelection Success
	(1)	(2)
Loss × Post	-0.034**	0.004
	(0.016)	(0.048)
Post × Coastal	-0.024	0.037
	(0.021)	(0.072)
Post × Log Population (2012)	-0.003	-0.037**
	(0.006)	(0.016)
Post × Log GDP (2012)	0.004	0.071**
	(0.011)	(0.030)
Post × Share Higher Ed (2012)	-0.183	-0.755
	(0.401)	(1.077)
Post × Electoral Volatility (2012)	-0.041**	-0.119**
	(0.018)	(0.054)
Municipality FE	✓	✓
Year FE	✓	✓
Observations	8,507	8,507
Municipalities	4,874	4,874
Election years	2000–2016	2000–2016

Note: Municipality and election-year fixed effects. Standard errors (in parentheses) clustered by municipality.

**$p < 0.05$.

A.5 Parallel Trends: Number of Candidates

Figure A.7 presents the event-study estimates for the total number of candidates, the entry outcome. The pre-shock coefficients are small, flat, and statistically indistinguishable from zero, which supports the parallel-trends assumption for candidate entry. The effect appears in 2016 and in 2020, although it is smaller in magnitude and significance in the latter electoral year.

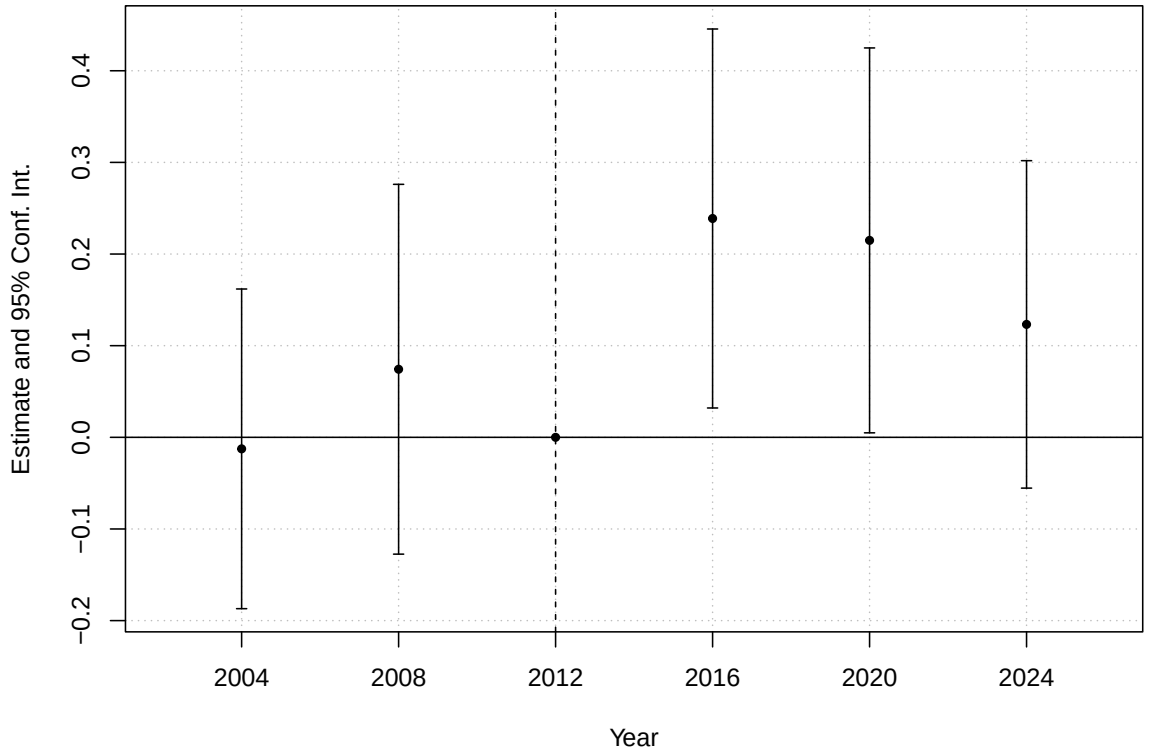

Figure A.7: Event Study: Number of Candidates

Note: Event-study estimates for the total number of candidates competing in the municipal election. Each point is the coefficient on the interaction of treatment status with an election-year dummy, with 2012 (the election before the shock) as the reference year. The specification includes municipality and election-year fixed effects. Vertical bars are 95 percent confidence intervals from standard errors clustered by municipality. The pre-shock coefficients are statistically indistinguishable from zero, supporting parallel trends for the entry outcome. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

A.6 Number of Candidates DiD without Covariates

Table A.5 reports the candidate-entry difference-in-differences result without baseline covariates. The shock increases the number of challenger candidates by approximately 0.26, very close to the covariate-adjusted estimate in the main text.

Table A.5: Baseline Effects of Royalty Loss on the Number of Challenger Candidates

	Number of Candidates
Loss × Post	0.264***
	(0.102)
Covariates	
Municipality FE	✓
Year FE	✓
Observations	8,511
Municipalities	4,878
Election years	2000–2016

Note: Difference-in-differences estimates with municipality and election-year fixed effects, estimated without baseline covariates. Standard errors (in parentheses) clustered by municipality.

***$p < 0.01$.

A.7 First Stage IVs

A.7.1 Candidate Entry

Table A.6 indicates that the economic shock significantly increased electoral competition in royalty-dependent municipalities that suffered revenue losses. In particular, it shows that incumbents in these localities faced, on average, about 0.24 more challengers in 2016 relative to comparable municipalities. The average number of candidates in 2012 was 2.66 in control and 2.90 in treated municipalities. By 2016 it rose to 2.72 in control, an increase of about 2.4 percent, and to 3.21 in treated, an increase of nearly 11 percent. These results suggest that, by encouraging both returning challengers and newcomers to enter the race, the shock expanded the supply of viable alternatives to the incumbent.

Table A.6: The Effects of Royalty Loss on the Number of Challenger Candidates

	Number of Candidates
Loss × Post	0.245**
	(0.106)
Post × Coastal	0.131
	(0.149)
Post × Log Population (2012)	0.035
	(0.032)
Post × Log GDP (2012)	0.090
	(0.055)
Post × Share Higher Ed (2012)	0.040
	(2.254)
Post × Electoral Volatility (2012)	0.016
	(0.096)
Municipality FE	✓
Year FE	✓
First-stage F	16.68
Observations	8,506
Municipalities	4,873
Election years	2000–2016

Note: First-stage TWFE regression estimates with municipality and election-year fixed effects. Standard errors (in parentheses) clustered by municipality. The first-stage F-statistic on the excluded instrument exceeds the conventional threshold of 10, indicating an adequately strong first stage.

$^{**}p < 0.05$.

A.7.2 New Candidate Entry

The first-stage instrument used above is the total number of candidates. Table A.7 reports the complementary first-stage estimate using the number of new candidates as the dependent variable. The effect is positive but marginally significant, which indicates that the rise in competition is driven by the full pool of competitors rather than by genuinely new entrants alone.

Table A.7: The Effects of Royalty Loss on the Number of New Challenger Candidates

	New Candidates
Loss × Post	0.224†
	(0.116)
Post × Coastal	0.085
	(0.163)
Post × Log Population (2012)	0.032
	(0.035)
Post × Log GDP (2012)	0.073
	(0.063)
Post × Share Higher Ed (2012)	0.675
	(2.609)
Post × Electoral Volatility (2012)	-0.072
	(0.107)
Municipality FE	✓
Year FE	✓
First-stage F	11.92
Observations	8,506
Municipalities	4,873
Election years	2000–2016

Note: First-stage TWFE regression estimates with municipality and election-year fixed effects. Standard errors (in parentheses) clustered by municipality. New candidates are defined as those who competed in the first election after the shock but not in the last election before the shock.

†$p < 0.10$.

A.8 Parallel Trends: Effective Number of Opposition Candidates

Figure A.8 presents the event-study estimates for the effective number of opposition candidates, the mechanism variable. The pre-shock coefficients are statistically indistinguishable from zero, supporting the parallel-trends assumption for the fragmentation outcome, and the effect appears only in the post-shock period, specifically 2016.

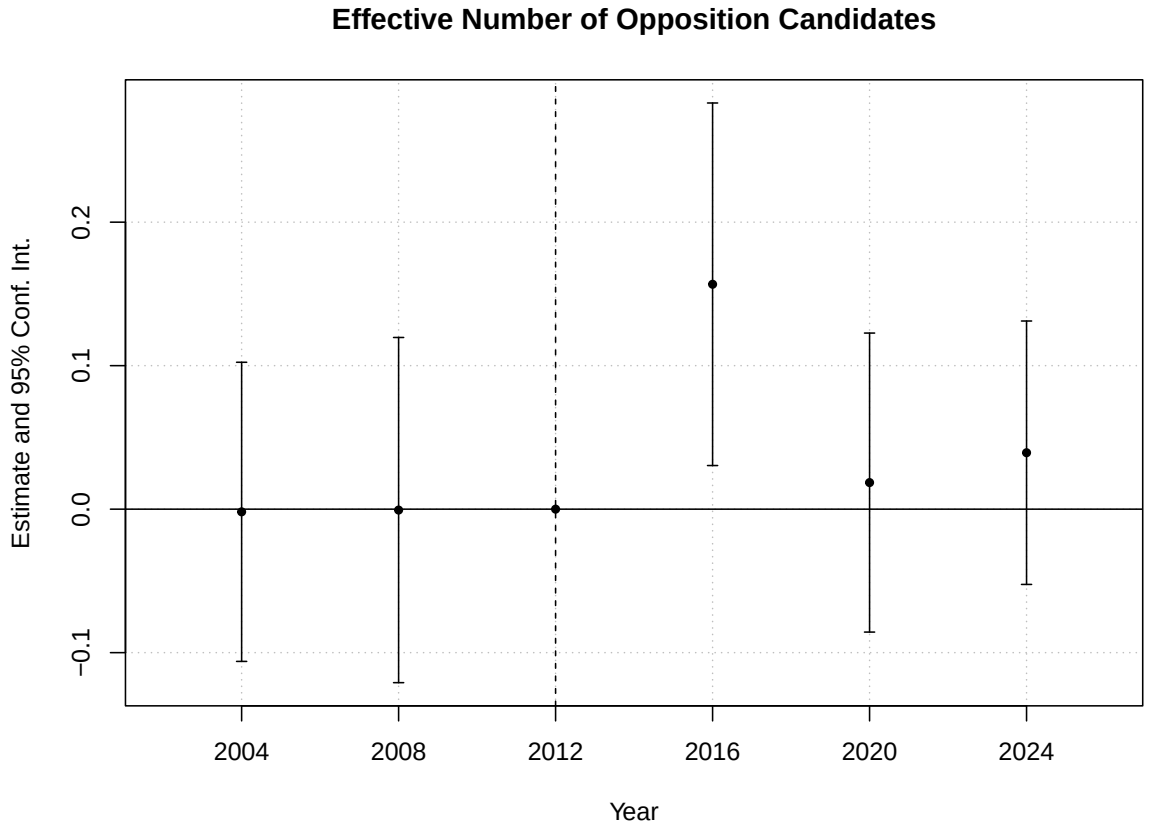

Figure A.8: Event Study: Effective Number of Opposition Candidates

Note: Event-study estimates for the effective number of opposition candidates, the mechanism outcome. Each point is the coefficient on the interaction of treatment status with an election-year dummy, with 2012 (the election before the shock) as the reference year. Vertical bars are 95 percent confidence intervals from standard errors clustered by municipality. The pre-shock coefficients are statistically indistinguishable from zero, supporting parallel trends for the fragmentation outcome. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

A.9 ENC DiD Results without Covariates

Table A.8 reports the fragmentation result without baseline covariates. The shock increases the effective number of opposition candidates by approximately 0.17, close to the covariate-adjusted estimate in the main text.

Table A.8: Baseline Effects of Royalty Loss on Vote Fragmentation

	Effective Number of Candidates
Loss × Post	0.175***
	(0.055)
Covariates	
Municipality FE	✓
Year FE	✓
Observations	8,247
Municipalities	4,790
Election years	2000–2016

Note: Standard errors (in parentheses) clustered by municipality. "Loss" denotes royalty-dependent municipalities after the revenue loss.

***$p < 0.01$.

A.10 Victory Margin Analyses

If the additional competition dispersed the opposition vote broadly rather than concentrating it behind a single close runner-up, then the shock should leave incumbent victory margins relatively unchanged in treated relative to control municipalities. Figure A.9 shows that treated and control municipalities followed comparable victory-margin trajectories before the shock, supporting parallel trends for this outcome. Tables A.9 and A.10 confirm that the shock had no detectable effect on the margin of victory: the point estimates are essentially zero and statistically insignificant.

A.10.1 Parallel Trends: Victory Margins

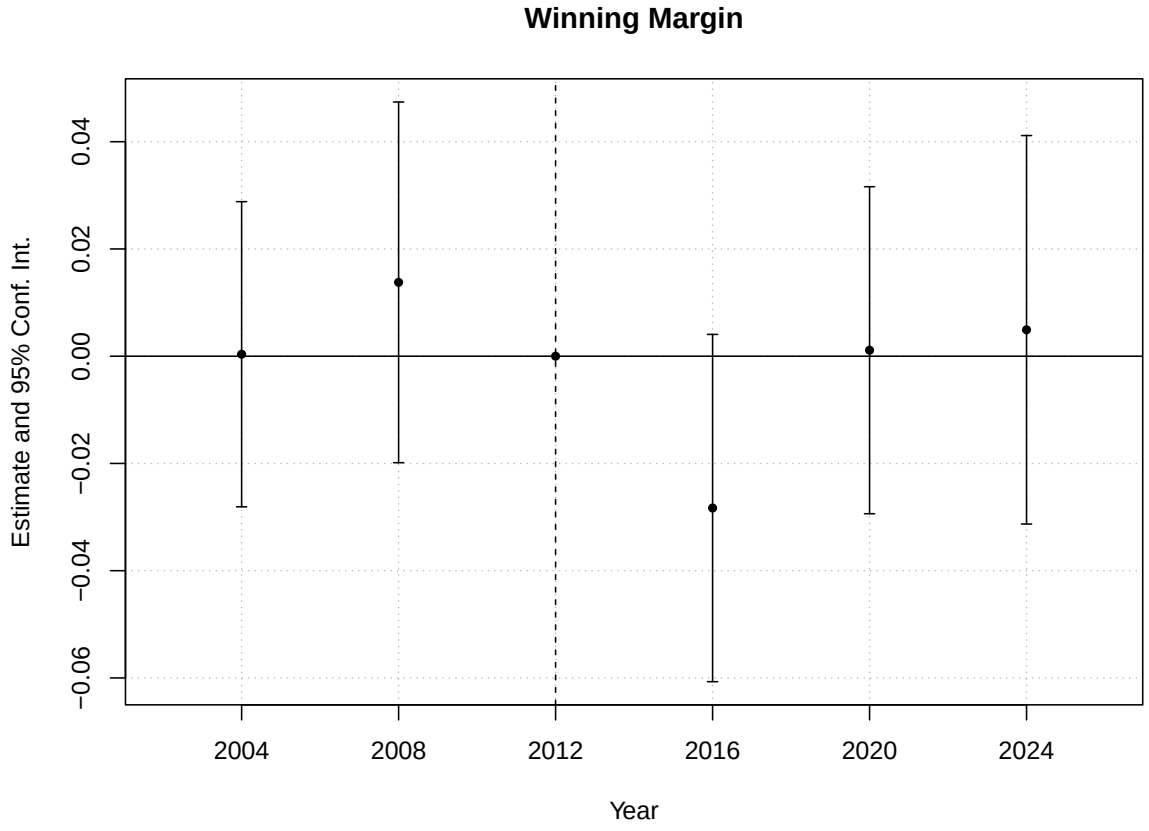

Figure A.9: Event Study: Winning Margins

Note: Event-study estimates for the incumbent's winning margin, with 2012 as the reference year. Each point is the treatment-by-year coefficient and vertical bars are 95 percent confidence intervals from standard errors clustered by municipality. Treated and control municipalities follow comparable margin trajectories before the shock. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

A.10.2 Winning Margins DiD Results without Covariates

Table A.9 reports the victory-margin difference-in-differences estimate without baseline covariates.

Table A.9: Baseline Effects of Royalty Loss on the Margin of Victory

	Winning Margin
Loss × Post	−0.009
	(0.016)
Covariates	
Municipality FE	✓
Year FE	✓
Observations	8,273
Municipalities	4,796
Election years	2000–2016

Note: Standard errors (in parentheses) clustered by municipality. "Loss" denotes royalty-dependent municipalities after the revenue loss.

A.10.3 Victory Margins DiD Results

Table A.10 reports the same specification with baseline covariates interacted with the post-shock indicator. The estimate is again indistinguishable from zero.

Table A.10: The Effects of Royalty Loss on the Margin of Victory

	Winning Margin
Loss × Post	−0.011
	(0.016)
Post × Coastal	−0.014
	(0.023)
Post × Log Population (2012)	0.017***
	(0.006)
Post × Log GDP (2012)	0.009
	(0.010)
Post × Share Higher Ed (2012)	−0.110
	(0.382)
Post × Electoral Volatility (2012)	−0.077***
	(0.018)
Municipality FE	✓
Year FE	✓
Observations	8,268
Municipalities	4,791
Election years	2000–2016

Note: Standard errors (in parentheses) clustered by municipality. "Loss" denotes royalty-dependent municipalities after the revenue loss. Covariates are interacted with the post-period (2016) indicator.
$^{***}p < 0.01$.

A.11 Alternative Explanations

A.11.1 Fiscal Spending

I consider measures of health, education, infrastructure, and personnel spending across electoral years. Data are sourced from the FINBRA and Siconfi databases, and all measures are calculated as the proportion of the budget they occupy in each municipality-year relative to total municipal expenditures. The shock to revenue inflows could have provided voters with new information about mayors' capacity to reallocate public spending. For instance, lower observed spending on health, education, or infrastructure could signal poor management, whereas higher personnel spending, often interpreted as a sign of patronage (Brollo et al., 2013; Novaes & Schiumerini, 2022), could lead voters to punish incumbents on the ballot box.

It is visible in Columns (1)–(4) in Table A.11 that the shock did not lead to systematic changes in the composition of municipal spending across the four categories examined in resource-dependent municipalities. It is worth noting that these measures capture spending as a share of total expenditures, not absolute levels. The null results therefore do not imply that public services were unaffected by the shock, but rather that mayors did not differentially reallocate resources across policy domains in response to the revenue contraction. More than that, the results suggest that the loss of royalty revenues does not influence another variable that could plausibly affect incumbent vote share. In the context of the royalty shock in Brazil, it is not surprising that this mechanism does not operate strongly after the royalty shock because mayors have limited fiscal and administrative flexibility to respond.

This limitation is driven mainly by two factors. First, part of the municipal budget is set by law. Municipalities must spend at least 25% of tax and transfer revenue on education (Art. 212 of the Federal Constitution) and at least 15% on health (Constitutional Amendment 29, regulated by Complementary Law 141/2012). These minimums keep the health and education shares stable even when revenue falls. Second, broader fiscal institu-

tions limit "compensatory" spending. Brazil's Fiscal Responsibility Law caps personnel expenditures relative to net current revenue and restricts borrowing and short-term credit operations. The law also prohibits hiring, salary increases, and unfunded obligations late in the term and in election years, with administrative and legal sanctions for violations. In practice, this leaves little room for incumbents to manage municipal budgets in the aftermath of the economic shock, which justifies voters in royalty-dependent municipalities being less likely to learn about their competence from observed policy responses.

Table A.11: The Effects of Royalty Loss on Municipal Spending

	Education (1)	Health (2)	Infrastructure (3)	Personnel (4)
Loss × Post	0.0052	-0.0024	0.0021	0.0065
	(0.0050)	(0.0040)	(0.0047)	(0.0071)
Post × Coastal	0.0178†	-0.0005	-0.0013	0.0155†
	(0.0093)	(0.0067)	(0.0067)	(0.0090)
Post × Log Population (2012)	0.0055***	0.0000	-0.0050***	-0.0037†
	(0.0015)	(0.0013)	(0.0016)	(0.0019)
Post × Log GDP (2012)	-0.0111***	0.0040†	-0.0010	-0.0082†
	(0.0026)	(0.0023)	(0.0032)	(0.0043)
Post × Share Higher Ed (2012)	-0.1626†	0.3604***	0.0357	-0.1078
	(0.0924)	(0.0781)	(0.1039)	(0.1329)
Post × Electoral Volatility (2012)	-0.0006	-0.0008	-0.0073	0.0111
	(0.0053)	(0.0045)	(0.0058)	(0.0074)
Municipality FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Observations	7,938	7,937	7,930	8,020
Municipalities	4,715	4,715	4,709	4,732
Election years	2000–2016	2000–2016	2000–2016	2000–2016

Note: TWFE regression estimates with municipality and election-year fixed effects. Standard errors (in parentheses) clustered by municipality.

$^{\dagger}p < 0.10$, $^{***}p < 0.01$.

A.11.2 Candidate Quality

To test whether candidate quality responded to the revenue contraction, I use the share of candidates with a college degree and candidates' level of political experience (i.e., the number of times that each of them has previously run for office) as proxies for quality.

The first analysis uses the full dataset of municipality-year pairs, while the second expands it to candidate-municipality-year observations. Column (1) in Table A.12 indicate no statistically discernible change in the share of candidates with higher education in royalty-dependent municipalities that lost revenues after the shock. Similarly, Column (2) shows no evidence that the royalty loss altered the experience profile of candidates who chose to run.

Table A.12: The Effects of Royalty Loss on Candidate Characteristics

	Education Level (1)	Experience (2)
Loss × Post	-0.029	-0.016
	(0.028)	(0.033)
Post × Coastal	0.008	-0.018
	(0.038)	(0.046)
Post × Log Population (2012)	-0.013	0.005
	(0.010)	(0.012)
Post × Log GDP (2012)	0.035**	0.007
	(0.016)	(0.021)
Post × Share Higher Ed (2012)	-0.896	-0.636
	(0.582)	(0.800)
Post × Electoral Volatility (2012)	0.040	-0.148***
	(0.034)	(0.042)
Municipality FE	✓	✓
Year FE	✓	✓
Observations	8,504	23,124
Municipalities	4,872	4,873
Election years	2000–2016	2000–2016

Note: TWFE regression estimates with municipality and election-year fixed effects. Standard errors (in parentheses) clustered by municipality.

$^{**}p < 0.05$, $^{***}p < 0.01$.

A.11.3 Differences in Victory Margins

Figure A.10 compares the average winning margin over time across treated municipalities, grouped by whether the incumbent won or lost reelection in 2016. Although in the post-shock election margins are directionally larger in municipalities where the incumbent survived, this gap is comparable to the one observed in prior elections.

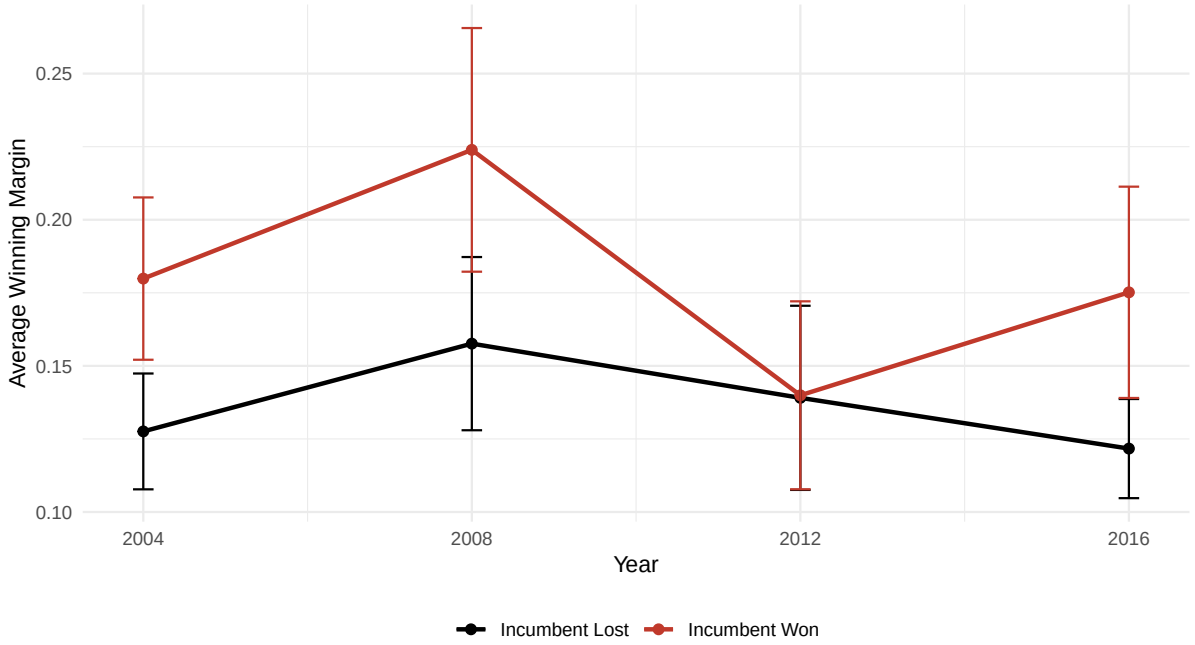

Figure A.10: Victory Margins in Treated Municipalities Over Time

Note: The figure plots the average winning margin, the difference in vote share between the first- and second-place candidates, by election year across treated (royalty-losing) municipalities, grouped by whether the incumbent won or lost reelection in 2016. Electoral data are sourced from the TSE. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

A.11.4 Party-driven Incumbent Punishment

Table A.13 indicates that there was no statistically significant differential effect on vote loss for PT incumbents in treated municipalities. Therefore, the broad nature of the punishment across parties suggests that the shock was interpreted as a local governance failure rather than a partisan one.

Table A.13: Differential Effects of Royalty Loss on PT versus Non-PT Incumbents

	Incumbent Vote
Loss × Post × PT	0.004
	(0.048)
Loss × Post	-0.036**
	(0.018)
Post × Coastal	-0.024
	(0.021)
Post × Log Population (2012)	-0.003
	(0.006)
Post × Log GDP (2012)	0.004
	(0.011)
Post × Share Higher Ed (2012)	-0.174
	(0.401)
Post × Electoral Volatility (2012)	-0.041**
	(0.018)
Municipality FE	✓
Year FE	✓
Observations	8,506
Municipalities	4,873
Election years	2000–2016

Note: TWFE regression estimates with municipality and election-year fixed effects. Standard errors (in parentheses) clustered by municipality.

**$p < 0.05$.

A.11.5 Campaign Spending

A remaining possibility is that incumbents held office because they spent more campaign resources than their challengers rather than because the anti-incumbent vote dispersed. In particular, it could have been the case that the spending patterns of incumbents in royalty-losing municipalities changed after the shock: they might have spent more as a strategy to avoid losing power. I assess this with TWFE models estimated on the sample of incumbents who ran for reelection, over the 2004 to 2016 election years (candidate campaign-finance records begin in 2004). The main outcome is the incumbent's share of campaign spending relative to the strongest challenger, defined as the incumbent's reported spending divided by the combined spending of the incumbent and

the top challenger, so that it ranges from zero to one and increases as the incumbent has a larger portion of the money spent between the two. Because it is a relative measure, it is comparable across the 2015 campaign-finance reform, which changed the level of campaign spending by banning corporate donations and imposing spending caps.[9] Column 1 of Table A.14 reports this share. The measure keeps uncontested races, where the incumbent faced no challenger, but drops municipality-years in which challengers ran yet none filed a positive record, which likely reflects missing filings rather than actual zero spending.

I complement the spending share with two related measures. Column 2 reports the ratio of the incumbent's spending to the top challenger's, which differs from the previous measure. The share is the incumbent's portion of the combined incumbent and top-challenger spending, so it always falls between zero and one. The ratio instead divides the incumbent's spending by the top challenger's, and has no upper bound: it records how many times the challenger's spending the incumbent's amounts to. Column 3 uses the same fragmentation index as the paper's effective number of candidates, applied here to campaign spending instead of votes. I compute it across all candidates, including the incumbent, so it captures whether the shock concentrated campaign money in the officeholder or a few candidates rather than spreading it out.

Across all three measures, none of the estimated effects is statistically distinguishable from zero, indicating that the shock did not hand treated incumbents a resource advantage over their challengers.

[9]The reform was enacted as Law 13.165/2015, alongside a 2015 ruling by Brazil's Supreme Federal Court that struck down corporate campaign donations.

Table A.14: Campaign Spending and the Royalty Shock

	(1) Incumbent Spending Share	(2) Incumbent-to-Challenger Spending Ratio	(3) Campaign-Spending Fragmentation
Loss × Post	−0.007	−0.465	−0.026
	(0.018)	(2.695)	(0.052)
Post × Coastal	−0.021	3.121	0.114
	(0.025)	(4.179)	(0.080)
Post × Log Population (2012)	0.003	0.315	0.065***
	(0.006)	(1.328)	(0.017)
Post × Log GDP (2012)	−0.009	−1.291	0.042
	(0.011)	(2.053)	(0.028)
Post × Share Higher Ed (2012)	0.226	11.778	−0.411
	(0.409)	(46.309)	(1.077)
Post × Electoral Volatility (2012)	−0.056***	−9.651	0.050
	(0.021)	(5.874)	(0.050)
Municipality FE	✓	✓	✓
Year FE	✓	✓	✓
Observations	9,094	8,813	9,928
Municipalities	5,063	4,968	5,181
Election years	2000–2016	2000–2016	2000–2016

Note: Two-way fixed-effects estimates of the royalty shock on measures of campaign spending, on the sample of incumbents who ran for reelection over 2004 to 2016, the years for which candidate campaign-finance records exist. Column 1 is the incumbent's share of campaign spending relative to the strongest challenger (incumbent spending over the combined spending of the incumbent and the top challenger, bounded between zero and one). Candidates without a matched spending record and contested races in which no challenger filed a positive record are dropped from the analysis. Column 2 is the ratio of the incumbent's spending to the top challenger's. Column 3 is the same fragmentation measure used for the vote (the effective number of candidates), applied to campaign spending. All models include municipality and election-year fixed effects and baseline covariates interacted with the post-shock period, with standard errors (in parentheses) clustered by municipality.

$^{***}p < 0.01$.

A.12 Runoff Elections

My argument is that plurality fosters opposition fragmentation without coordination, so incumbents lose vote share without losing office. A two-round system changes this in the second round, where voters face only two candidates and the opposition can no longer split its support. One possibility is that, if the decoupling is driven by fragmentation under plurality, the second round, by forcing a binary choice, should make royalty-losing incumbents more vulnerable to removal, that is, punished and defeated. The first round, which still presents many options, should resemble the plurality pattern of vote loss without removal.

Brazilian mayoral races in municipalities above 200,000 voters are decided by a two-round (runoff) system rather than single-round plurality. Because these municipalities are few and large, the runoff sample is small, so the estimates in this section are reported as two-way fixed-effects (TWFE) estimates. I estimate the specification separately on the first-round and the second-round samples.

The first round resembles the plurality pattern. Incumbents in royalty-losing municipalities receive about 31 percentage points less vote share, with no effect on reelection (Table A.15). The second round, where voters choose between only two candidates, shows no effect on either outcome. I therefore find no evidence that this binary choice makes royalty-losing incumbents more likely to be removed. Given how few municipalities enter this analysis, just 32 in the first round and 31 in the second, with 41 and 41 municipality-year observations respectively, the exercise is suggestive rather than conclusive.

Table A.15: Royalty Loss and Incumbent Outcomes in Runoff Municipalities

	First round		Second round	
	Incumbent Vote	Reelection Success	Incumbent Vote	Reelection Success
	(1)	(2)	(3)	(4)
Loss × Post	−0.309***	−0.316	0.070	0.217
	(0.103)	(0.255)	(0.105)	(0.446)
Municipality FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Observations	41	41	41	41
Municipalities	32	32	31	31
Election years	2000–2016	2000–2016	2000–2016	2000–2016

Note: Two-way fixed-effects estimates on the runoff (two-round) subsample, which is small (few large municipalities). Standard errors (in parentheses) clustered by municipality.

***$p < 0.01$.

A.13 Null, Blank, and Protest Voting

This section addresses expressive and protest voting as an alternative channel through which voters might register dissatisfaction with the fiscal shock. The relevant theoretical possibility is that voters who are dissatisfied but unable or unwilling to coordinate behind a challenger might instead abstain from choosing a candidate, casting blank or null ballots rather than transferring support to the opposition. If that were the dominant response, the vote-share decline documented in the main text would partly reflect a withdrawal from the candidate field rather than a reallocation across challengers.

I test this directly using three outcomes: protest votes (blank and null combined), blank votes, and null votes, each measured as a share of the electorate. The underlying municipal-detail data are drawn from the Brazilian Superior Electoral Court (TSE). Parallel trends are clean for all three outcomes. Figure A.11 reports the event study for the combined protest measure.

The 2016 effect on protest votes is null in both the event study and the covariate-adjusted difference-in-differences, and the same holds for null votes specifically. The effect on blank votes is small, negative, and only marginally significant (the 10 percent level) in the covariate-adjusted specification. In substantive terms, voters dissatisfied with the shock did not respond by withdrawing into blank or null ballots meaningfully. If anything, blank voting moved slightly downward. This pattern is consistent with the fragmentation mechanism: the anti-incumbent response operates through dispersion across a wider field of challengers, not through expressive abstention from the candidate choice.

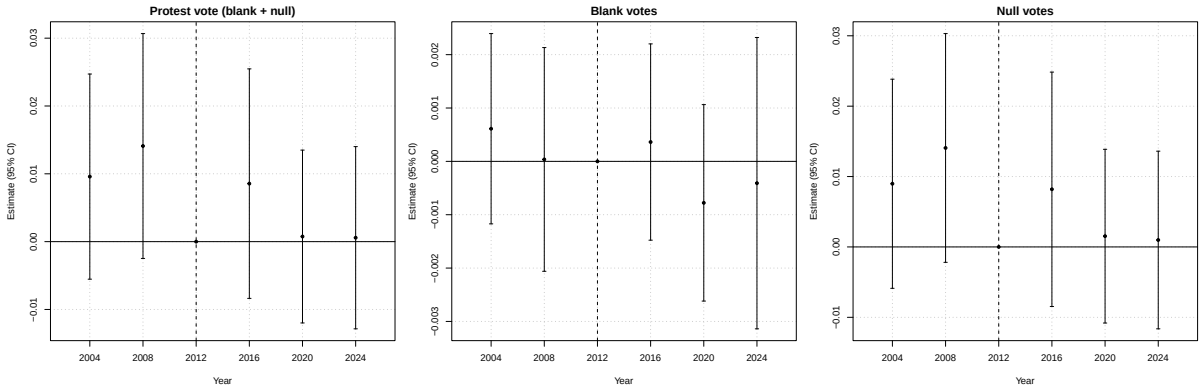

Figure A.11: Event Study: Protest Votes (Blank and Null Combined)

Note: Event study for protest voting, blank and null ballots combined as a share of the electorate, with 2012 as the reference year. Points are treatment-by-year coefficients and vertical bars are 95 percent confidence intervals clustered by municipality. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

Table A.16: Royalty Loss and Protest Voting (Covariate-Adjusted DiD)

	Protest (blank+null)	Blank	Null
	(1)	(2)	(3)
Loss × Post	−0.001	−0.002†	0.001
	(0.008)	(0.001)	(0.008)
Post × Coastal	0.003	0.002**	0.001
	(0.011)	(0.001)	(0.011)
Post × Log Population (2012)	0.006**	0.003***	0.004
	(0.003)	(0.001)	(0.003)
Post × Log GDP (2012)	0.010**	0.001	0.009**
	(0.005)	(0.001)	(0.005)
Post × Share Higher Ed (2012)	−0.037	0.045	−0.082
	(0.207)	(0.031)	(0.200)
Post × Electoral Volatility (2012)	−0.026***	−0.005***	−0.021**
	(0.009)	(0.002)	(0.008)
Municipality FE	✓	✓	✓
Year FE	✓	✓	✓
Observations	8,506	8,506	8,506
Municipalities	4,873	4,873	4,873
Election years	2000–2016	2000–2016	2000–2016

Note: Covariate-adjusted difference-in-differences with municipality and election-year fixed effects, covariates interacted with the post-shock indicator, standard errors clustered by municipality. Outcomes are vote shares. Parallel trends are clean for all three outcomes. Protest and null effects are null. The blank-vote effect is small, negative, and marginally significant.

†$p < 0.10$.

A.14 Windfall Municipalities

A second placebo, complementary to the spillover tests, examines municipalities that *gained* royalty revenue between 2014 and 2016 rather than losing it. If voters reacted to any change in royalty revenue, or if the electoral patterns documented here reflected some general feature of resource-dependent municipalities rather than the loss specifically, then windfall municipalities should show effects of their own. Nevertheless, I found no evidence for these arguments. Running the full analysis on the subset of municipalities with royalty inflow increases yields null effects on both incumbent vote share and reelection. Thus, there are no electoral rewards for incumbents where royalties rose. Figure A.12 reports the corresponding event study, in which the 2016 coefficients on both incumbent vote share and reelection are statistically indistinguishable from zero and the pre-shock coefficients are flat. Table A.17 confirms this in the difference-in-differences specification, where the windfall effect on both incumbent vote share and reelection is small and statistically indistinguishable from zero, so municipalities that gained royalties show no electoral response of either kind.

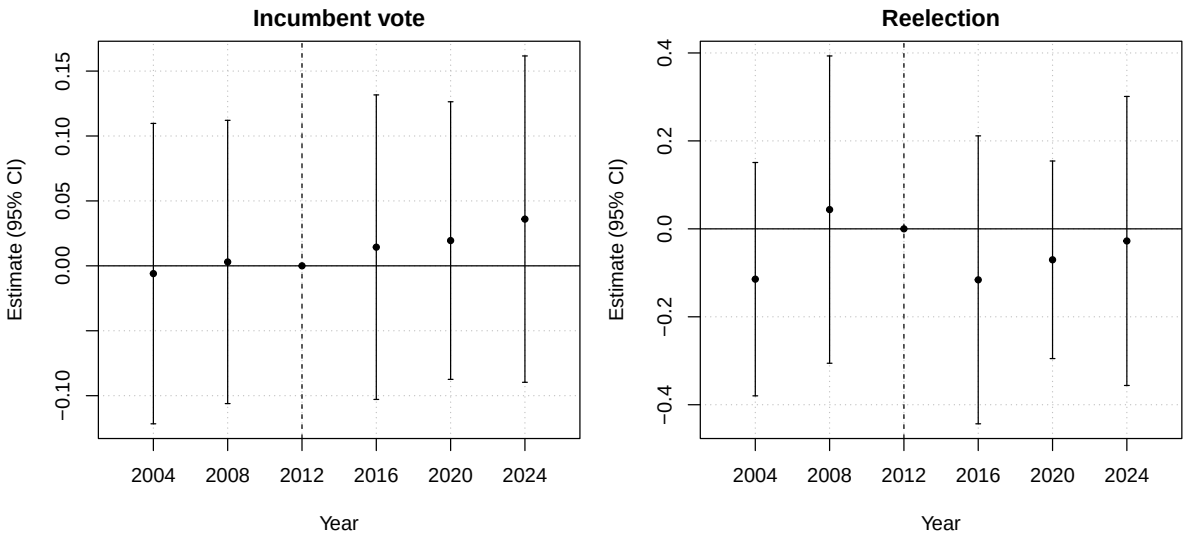

Figure A.12: Event Study: Windfall (Royalty-Gain) Municipalities

Note: Event study for royalty-windfall municipalities, those that gained royalty revenue between 2014 and 2016, run as a placebo on the direction of the shock. Points are treatment-by-year coefficients on incumbent vote share and reelection, with 2012 as the reference year and 95 percent confidence intervals clustered by municipality. The 2016 coefficients are null, indicating no electoral effect where royalties rose. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

Table A.17: The Electoral Effects of Royalty Windfalls (Placebo)

	Incumbent Vote	Reelection Success
	(1)	(2)
Windfall × Post	0.026	-0.056
	(0.065)	(0.173)
Post × Coastal	-0.042	-0.092
	(0.028)	(0.088)
Post × Log Population (2012)	-0.006	-0.032†
	(0.006)	(0.017)
Post × Log GDP (2012)	-0.005	0.033
	(0.012)	(0.033)
Post × Share Higher Ed (2012)	0.051	-0.722
	(0.416)	(1.137)
Post × Electoral Volatility (2012)	-0.042**	-0.104†
	(0.019)	(0.058)
Municipality FE	✓	✓
Year FE	✓	✓
Observations	7,406	7,406
Municipalities	4,240	4,240
Election years	2000–2016	2000–2016

Note: Difference-in-differences estimates for royalty-windfall municipalities, those that gained royalty revenue between 2014 and 2016, against the non-dependent control group. "Windfall × Post" is the interaction of windfall status with the post-shock (2016) indicator. Municipality and election-year fixed effects, with covariates interacted with the post-shock indicator. Standard errors (in parentheses) clustered by municipality. Neither treatment coefficient is statistically significant.

$^{\dagger}p < 0.10$, $^{**}p < 0.05$.

A.15 Dependence-Based Treatment Specification

The main text defines treatment by royalty loss rather than by royalty dependence. As a check on that choice, I re-estimate the design defining treatment as any royalty dependence in 2014, regardless of whether the royalty share subsequently fell (or rose). The substantive conclusions are unchanged: the vote-share effect remains negative and close to 3 percentage points, and the reelection effect remains null. What makes the loss-based definition preferable is, in particular, the behavior of the windfall municipalities. These were equally dependent on royalties but gained rather than lost revenue, and they show no electoral effect on either outcome (see Appendix A.14). If dependence itself drove the results, these equally dependent municipalities should move with the losers. That they do not indicates the response is specific to the loss rather than to royalty dependence in general. Pooling them with the losers under a dependence-based definition would average a genuine loss effect against this null and understate it. Figure A.13 reports the event study for this specification, in which the 2016 vote-share coefficient is negative but only marginally significant and the reelection coefficient is null.

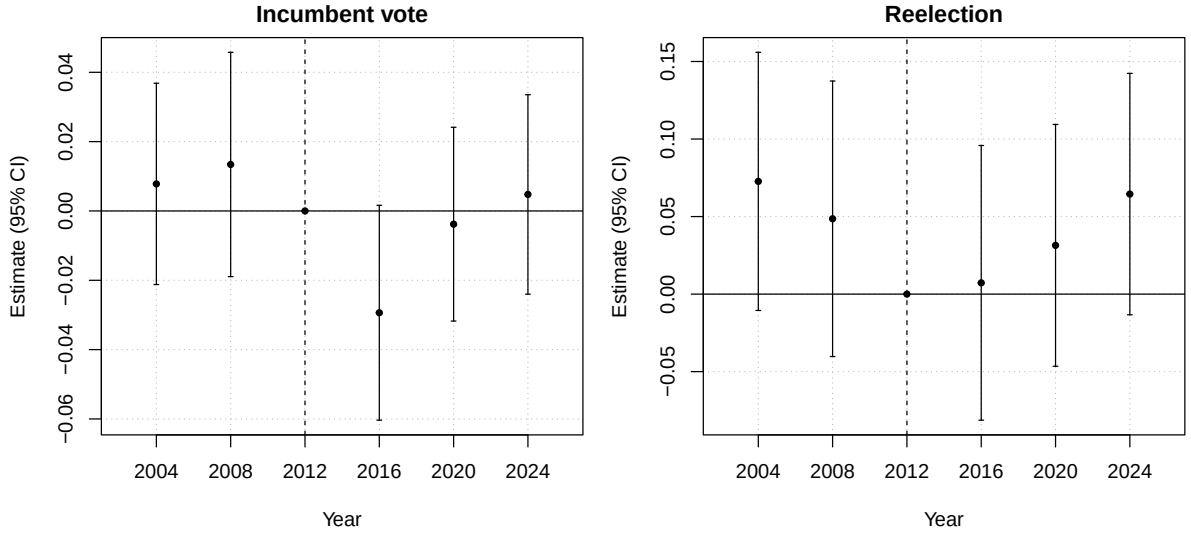

Figure A.13: Event Study: Dependence-Based Treatment Specification

Note: Event study for the dependence-based treatment definition, which codes as treated any municipality with royalty dependence in 2014 regardless of whether its royalty share later fell or rose. Points are treatment-by-year coefficients with 2012 as the reference year and 95 percent confidence intervals clustered by municipality. The R code producing this figure was drafted with assistance from Claude (Anthropic, claude-opus-5), accessed at claude.ai, in June–August 2026. All data, specifications, and output were produced and verified by the author.

A.16 Loss-Intensity Disaggregation

The main text reports the pooled effect of the royalty loss and refers here for the disaggregation by loss intensity. This appendix splits the royalty-loss municipalities into terciles of the fall in royalties as a share of the total municipal budget, and estimates the three terciles jointly against the control group of non-royalty-dependent municipalities. The question it answers is whether the electoral response varies with loss intensity. All bin estimates in this section are two-way fixed-effects (TWFE) estimates.

A.16.1 Vote Share and Reelection by Loss Tercile

Table A.18: Vote Share and Reelection by Loss Tercile

	Incumbent Vote	Reelection Success
	(1)	(2)
Q1 × Post (loss < 0.030 pp)	0.010	0.032
	(0.027)	(0.083)
Q2 × Post (loss 0.030–0.399 pp)	−0.059**	−0.088
	(0.023)	(0.069)
Q3 × Post (loss ≥ 0.399 pp)	−0.056**	0.084
	(0.023)	(0.081)
Post × Coastal	−0.019	0.025
	(0.021)	(0.072)
Post × Log Population (2012)	−0.004	−0.038**
	(0.006)	(0.016)
Post × Log GDP (2012)	0.005	0.065**
	(0.011)	(0.031)
Post × Share Higher Ed (2012)	−0.120	−0.700
	(0.401)	(1.076)
Post × Electoral Volatility (2012)	−0.039**	−0.115**
	(0.018)	(0.054)
Municipality FE	✓	✓
Year FE	✓	✓
Observations	8,506	8,506
Municipalities	4,873	4,873
Election years	2000–2016	2000–2016

Notes: TWFE estimates from a single model, with each tercile compared against the control group of non-royalty-dependent municipalities. Terciles are cut on the decline in royalties as a share of the total municipal budget. Covariates are interacted with the post-period (2016) indicator. Standard errors (in parentheses) are clustered by municipality.

**$p < 0.05$.

Table A.18 reports the tercile estimates for incumbent vote share and reelection. The two upper terciles report similar results as the main pooled vote-share result. Incumbents in the middle tercile lose about 5.9 percentage points of vote share and those in the top tercile about 5.6, which against a pre-shock mean of 47.4 percent is roughly 12 percent of their support in both.

A.16.2 Candidate Entry by Loss Tercile

Table A.19 reports the estimates for the number of mayoral candidates in the three loss terciles. Entry is flat in the bottom tercile, while the middle tercile gains about 0.25 candidates and the top tercile about 0.54 relative to the control municipalities, with only the latter being statistically significant at conventional levels. Read together, this means that the municipalities that lost a larger share of royalty revenues drew more candidates into the mayoral race than non-dependent municipalities did.

Table A.19: Number of Candidates by Loss Tercile

	Number of Candidates
Q1 × Post (loss < 0.030 pp)	−0.021
	(0.164)
Q2 × Post (loss 0.030–0.399 pp)	+0.249†
	(0.141)
Q3 × Post (loss ≥ 0.399 pp)	+0.543**
	(0.219)
Post × Coastal	+0.074
	(0.148)
Post × Log Population (2012)	+0.041
	(0.032)
Post × Log GDP (2012)	+0.075
	(0.055)
Post × Share Higher Ed (2012)	−0.310
	(2.269)
Post × Electoral Volatility (2012)	+0.016
	(0.096)
Municipality FE	✓
Year FE	✓
Observations	8,506
Municipalities	4,873
Election years	2000–2016

Notes: TWFE estimates from a single model, with each tercile compared against the control group of non-royalty-dependent municipalities. Terciles are cut on the decline in royalties as a share of the total municipal budget. Covariates are interacted with the post-period (2016) indicator. Standard errors (in parentheses) are clustered by municipality.

$^{\dagger}p < 0.10$, $^{**}p < 0.05$.

A.16.3 Vote Fragmentation by Loss Tercile

Table A.20 reports the estimates for the effective number of opposition candidates in the three loss terciles. Opposition fragmentation rises by 0.28 in the top tercile relative to the control municipalities, while the two lower terciles show no measurable change. Municipalities that lost a higher share of royalty revenues experienced more vote fragmentation relative to non-dependent municipalities.

Table A.20: Effective Number of Opposition Candidates by Loss Tercile

	Effective Number of Opposition Candidates
Q1 × Post (loss < 0.030 pp)	+0.061
	(0.072)
Q2 × Post (loss 0.030–0.399 pp)	+0.115
	(0.086)
Q3 × Post (loss ≥ 0.399 pp)	+0.283**
	(0.114)
Post × Coastal	+0.183†
	(0.095)
Post × Log Population (2012)	+0.016
	(0.019)
Post × Log GDP (2012)	+0.062†
	(0.035)
Post × Share Higher Ed (2012)	−0.083
	(1.280)
Post × Electoral Volatility (2012)	+0.005
	(0.056)
Municipality FE	✓
Year FE	✓
Observations	8,242
Municipalities	4,785
Election years	2000–2016

Notes: TWFE estimates from a single model, with each tercile compared against the control group of non-royalty-dependent municipalities. Terciles are cut on the decline in royalties as a share of the total municipal budget. Covariates are interacted with the post-period (2016) indicator. Standard errors (in parentheses) are clustered by municipality.

†$p < 0.10$, **$p < 0.05$.

A.17 Latitude and Longitude as Baseline Covariates

To address the concern that the coastal indicator alone may not capture potential geographic confounding (for example, unobserved regional shocks correlated with both oil dependence and electoral behavior), I re-estimate the main specifications adding municipality latitude and longitude as baseline covariates interacted with the post-shock indicator. Table A.21 compares the headline estimates with and without these geographic controls. All headline results are stable. The pooled vote-share effect (from the main difference-in-differences specification) remains negative and significant, moving from -3.4 to -4.1 percentage points. The pooled reelection effect remains null. The instrumental-variables second stage on vote share remains negative and significant, moving from -0.141 to -0.146. The tercile vote-share effects are likewise stable, with the middle tercile moving from -0.059 to -0.063 and the top tercile from -0.056 to -0.064, and reelection remains null in all three terciles. The effective-number-of-candidates mechanism is similarly robust, with the pooled coefficient moving from $+0.147$ to $+0.187$ and the top-tercile coefficient from $+0.283$ to $+0.339$. No coefficient loses significance once the geographic controls are added.

Table A.21: Headline Estimates with and without Latitude and Longitude Controls

Specification	Outcome	Without lat/lon	With lat/lon
Pooled main	Incumbent vote	-0.034^{**}	-0.041^{**}
Pooled main	Reelection	$+0.004$	-0.021
Pooled main	ENC (mechanism)	$+0.147^{***}$	$+0.187^{***}$
Q1 tercile	Incumbent vote	$+0.010$	$+0.003$
Q2 tercile	Incumbent vote	-0.059^{**}	-0.063^{***}
Q3 tercile	Incumbent vote	-0.056^{**}	-0.064^{***}
Q1 tercile	Reelection	$+0.032$	$+0.010$
Q2 tercile	Reelection	-0.088	-0.104
Q3 tercile	Reelection	$+0.084$	$+0.052$
Q1 tercile	ENC	$+0.061$	$+0.110$
Q2 tercile	ENC	$+0.115$	$+0.140$
Q3 tercile	ENC	$+0.283^{**}$	$+0.339^{***}$
IV second stage	Incumbent vote	-0.141^{**}	-0.146^{**}

Note: Each row reports the coefficient on the treatment-by-post interaction (or the second-stage coefficient for the IV row) without and with latitude and longitude added to the baseline covariates, all interacted with the post-shock indicator. Municipality and year fixed effects and standard errors clustered by municipality. The tercile rows are TWFE estimates in which the three terciles enter jointly, with the control group of non-royalty-dependent municipalities as the reference category. The pooled and IV rows are the headline specifications. No estimate loses significance once the geographic controls are added, and the fragmentation estimates strengthen: both the pooled effective-number-of-candidates coefficient and the top-tercile coefficient move from the 5 to the 1 percent level.

$^{\dagger}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A.18 Table Results with Covariate Coefficients

This section reports the coefficients on the baseline covariates for the four models presented in the main text: incumbent vote share and reelection, the number of candidates, the instrumental-variables second stage, and the effective number of opposition candidates. The models are the same ones estimated in the main text, with the covariate coefficients displayed alongside the treatment coefficient.

Table A.22: Baseline Covariate Coefficients: Incumbent Vote and Reelection Success

	Incumbent Vote	Reelection Success
	(1)	(2)
Loss × Post	-0.034**	0.004
	(0.016)	(0.048)
Post × Coastal	-0.024	0.037
	(0.021)	(0.072)
Post × Log Population (2012)	-0.003	-0.037**
	(0.006)	(0.016)
Post × Log GDP (2012)	0.004	0.071**
	(0.011)	(0.030)
Post × Share Higher Ed (2012)	-0.183	-0.755
	(0.401)	(1.077)
Post × Electoral Volatility (2012)	-0.041**	-0.119**
	(0.018)	(0.054)
Municipality FE	✓	✓
Year FE	✓	✓
Observations	8,506	8,506
Municipalities	4,873	4,873
Election years	2000–2016	2000–2016

Note: The models reported in Table 2 of the main text, with the baseline covariate coefficients displayed. Column (1) is the incumbent vote share and column (2) a binary indicator for incumbent reelection. Each covariate is interacted with the post-shock (2016) indicator, since municipality and election-year fixed effects absorb the levels. Standard errors (in parentheses) are clustered by municipality.

**$p < 0.05$.

Table A.23: Baseline Covariate Coefficients: Number of Competing Candidates

	Number of Candidates
Loss × Post	0.245**
	(0.106)
Post × Coastal	0.131
	(0.149)
Post × Log Population (2012)	0.035
	(0.032)
Post × Log GDP (2012)	0.090
	(0.055)
Post × Share Higher Ed (2012)	0.040
	(2.254)
Post × Electoral Volatility (2012)	0.016
	(0.096)
Municipality FE	✓
Year FE	✓
Observations	8,506
Municipalities	4,873
Election years	2000–2016

Note: The model reported in Table 3 of the main text, with the baseline covariate coefficients displayed. Each covariate is interacted with the post-shock (2016) indicator, since municipality and election-year fixed effects absorb the levels. Standard errors (in parentheses) are clustered by municipality.

**$p < 0.05$.

Table A.24: Baseline Covariate Coefficients: Challenger Entry and Incumbent Outcomes (IV Second Stage)

	Incumbent Vote	Incumbent Reelection
	(1)	(2)
Number of candidates	−0.141**	0.018
	(0.064)	(0.196)
Post × Coastal	−0.006	0.034
	(0.025)	(0.080)
Post × Log Population (2012)	0.002	−0.038**
	(0.006)	(0.019)
Post × Log GDP (2012)	0.017	0.069**
	(0.011)	(0.034)
Post × Share Higher Ed (2012)	−0.177	−0.755
	(0.355)	(1.083)
Post × Electoral Volatility (2012)	−0.038**	−0.119**
	(0.016)	(0.054)
Municipality FE	✓	✓
Year FE	✓	✓
Observations	8,506	8,506
Municipalities	4,873	4,873
Election years	2000–2016	2000–2016

Note: The second-stage models reported in Table 4 of the main text, with the baseline covariate coefficients displayed. The number of candidates is instrumented by $\text{Loss}_i \times \text{Post}_t$. The corresponding first stages are reported in Table A.6 and Table A.7. Each covariate is interacted with the post-shock (2016) indicator, since municipality and election-year fixed effects absorb the levels. Standard errors (in parentheses) are clustered by municipality.

**$p < 0.05$.

Table A.25: Baseline Covariate Coefficients: Effective Number of Opposition Candidates

	Effective Number of Opposition Candidates
Loss × Post	0.147***
	(0.056)
Post × Coastal	0.209**
	(0.096)
Post × Log Population (2012)	0.014
	(0.018)
Post × Log GDP (2012)	0.070**
	(0.034)
Post × Share Higher Ed (2012)	0.023
	(1.284)
Post × Electoral Volatility (2012)	0.003
	(0.056)
Municipality FE	✓
Year FE	✓
Observations	8,242
Municipalities	4,785
Election years	2000–2016

Note: The model reported in Table 5 of the main text, with the baseline covariate coefficients displayed. The outcome is the effective number of opposition candidates. Each covariate is interacted with the post-shock (2016) indicator, since municipality and election-year fixed effects absorb the levels. Standard errors (in parentheses) are clustered by municipality.

$^{**}p < 0.05$, $^{***}p < 0.01$.