The Price of Pork: France's Dual-Mandate Ban and the Fiscal Cost of Local–National Connections*

APEP

Autonomous Policy Evaluation Project

@olafdrw

March 5, 2026

Abstract

France's 2014 organic law banning the *cumul des mandats* took effect at the June 2017 elections, forcing deputy-mayors to choose between Parliament and city hall. We exploit this reform in a difference-in-differences framework, comparing constituencies whose deputy held a concurrent mayoral mandate to those whose deputy did not. Using commune budget data covering 35,000 communes from 2008 to 2023, we find no detectable effects on capital investment, equipment spending, state grants, or operating expenditure. Event-study estimates confirm flat pre-trends across nine pre-treatment years and no post-reform divergence. The results indicate that the pork-barrel channel—whereby cumulard deputies steered resources to their communes—was empirically negligible, and that severing the local–national connection carried no measurable fiscal cost.

JEL Codes: D72, H72, H77

Keywords: cumul des mandats, dual mandates, local public finance, pork-barrel, France

*This paper was produced as part of the Autonomous Policy Evaluation Project (APEP). All data are from public French government sources: DGFiP commune budgets (2000–2017), OFGL consolidated accounts (2017–2023), NosDéputés.fr, and Wikidata. Replication code and data are available in the project repository.

1 Introduction

In 2012, nearly half of the members of the French National Assembly were simultaneously serving as mayors of their hometowns. This *cumul des mandats*—the accumulation of elected offices—sat at the heart of French political life for decades, resting on a simple promise: a deputy who also ran a city could steer national resources to local projects. The result was what observers called "a republic of mayors who happen to sit in Parliament" (François and Magni-Berton, 2006), a system in which national legislators had direct personal stakes in local fiscal outcomes.

This paper asks a simple question: what happened to local public spending when France severed these connections? The Loi organique n°2014-125, adopted in February 2014 and binding from the June 2017 legislative elections, prohibited the simultaneous holding of a parliamentary mandate and a local executive office. Deputy-mayors were forced to resign one mandate. If the conventional wisdom is correct—that cumulards used their national platform to channel *dotations d'investissement*, DETR grants, and other discretionary transfers to their communes—then the ban should have produced a measurable decline in local investment spending.

We find that it did not. Using commune budgets from 2008 to 2023, we compare constituencies whose XIV-legislature deputy held a concurrent mayoral mandate (248 "cumulard" constituencies in our analysis sample) to those whose deputy did not (291 "non-cumulard" constituencies) in a standard two-period difference-in-differences design. The effect on investment per capita is -0.014 thousand euros, statistically indistinguishable from zero ($p = 0.20$). Equipment expenditure, state grants, operating expenditure, revenue, and debt outstanding all show similarly null effects with tight confidence intervals. No fiscal outcome exhibits a significant response to the ban.

Event-study estimates spanning ten pre-treatment years confirm that the parallel trends assumption holds. Pre-reform coefficients are small, precisely estimated, and jointly insignificant for all fiscal outcomes. There is no evidence of anticipation effects between the law's passage in 2014 and its implementation in 2017. Post-treatment coefficients at two and five years after the ban remain centered on zero for all outcomes.

These findings contribute to three literatures. First, we speak to the political economy of legislative structure. Gagliarducci and Nannicini (2013) show that Italian part-time mayors underperform relative to full-time ones; Fouirnaies and Hall (2022) find that US state legislators who hold outside employment are less productive. Our contribution is to identify the *fiscal* consequences of dual mandates—and to show that the feared pork-barrel channel was empirically negligible. Second, we contribute to the literature on intergovernmental fiscal transfers. The discretionary allocation of French investment grants (DETR, DSIL) by *préfets* has been documented as politically influenced (Enikolopov, 2014), but our results suggest that the deputy-mayor channel was not the primary mechanism. Third, we add to the nascent literature evaluating France's institutional reforms of the 2010s, including the NOTRe law on intermunicipal governance (Breuillé et al., 2018) and the *taxe d'habitation* abolition.

The null result is itself informative. It rules out the most dramatic version of the pork-barrel hypothesis—that cumulards were channeling substantial public resources to their constituencies. The absence of effects across all fiscal categories—investment, operating expenditure, grants, revenue, and debt—suggests that the dual mandate either generated no fiscal rents, or that any rents were quickly substituted through alternative political channels after the reform. This is consistent with institutional constraints on discretionary grant allocation and the bureaucratic nature of French intergovernmental transfers (Enikolopov, 2014).

Our paper relates to several strands of the economics and political science literature. A long tradition in political economics studies how institutional design affects the quality and composition of public spending. Ferraz and Finan (2011) show that higher salaries attract better candidates and improve fiscal outcomes in Brazilian municipalities. Gagliarducci and Nannicini (2013) demonstrate that Italian mayors who serve part-time (because they hold concurrent national or regional mandates) produce worse fiscal outcomes than full-time mayors. Cruz et al. (2017) demonstrate that political family networks in the Philippines affect public goods provision, and Fiva and Halse (2018) study how political power-sharing arrangements influence policy outcomes. Our paper complements this work by studying what happens when dual mandates are *abolished*: the cumul's removal had no detectable fiscal effect, suggesting that the institution may have been less consequential for local budgets than critics feared.

The pork-barrel literature provides the theoretical framework for the investment channel we test. Golden and Picci (2008) document how Italian legislators directed infrastructure spending to their constituencies; Hodler and Raschky (2014) show that political leaders favor their home regions; and Enikolopov (2014) show that politically-connected municipalities receive more intergovernmental transfers. In the French context, the discretionary allocation of investment grants by *préfets* creates scope for political influence, and the deputy-mayor combination provides a particularly direct channel for exercising it. Our null result on grants and investment suggests that this channel, while theoretically plausible, was empirically negligible—or that it was quickly substituted by alternative political connections after the reform.

Finally, our paper contributes to the growing literature on France's institutional reforms of the 2010s. Breuillé et al. (2018) study the effects of the NOTRe law on intermunicipal governance, and several recent papers examine the fiscal consequences of commune mergers (*communes nouvelles*). Recent work by Bach (2019) examines the political consequences of the non-cumul reform at the municipal level, but does not estimate fiscal effects. We add the cumul ban to this portfolio of reforms, providing the first causal estimate of its fiscal effects.

The remainder of the paper proceeds as follows. Section 2 describes the institutional context of the *cumul des mandats* and the 2014/2017 reform. Section 3 presents the data sources. Section 4 lays out the identification strategy. Section 5 reports the main results. Section 6 provides robustness checks. Section 7 decomposes the mechanisms underlying our findings. Section 8 concludes.

2 Institutional Background

2.1 The *Cumul des Mandats*: A French Exception

France's political system has long been characterized by the accumulation of elected offices by individual politicians. Unlike most European democracies, where holding multiple mandates is either prohibited or culturally discouraged, France developed a system in which national legislators routinely served as mayors, departmental or regional council presidents, or heads of intercommunal bodies (*établissements publics de coopération intercommunale*, EPCIs). The practice, known as the *cumul des mandats*, was deeply embedded in the institutional fabric of the Fifth Republic.

The roots of the cumul lie in France's highly centralized state structure (Solé Díaz, 2013). Until the decentralization reforms of the 1980s, local governments had limited fiscal autonomy, and mayors depended heavily on the national government for investment grants and administrative authorizations. Holding a seat in the *Assemblée nationale* or the *Sénat* simultaneously with a local executive mandate gave politicians direct access to the ministries and préfectures that controlled resource allocation. The cumul was thus not merely a matter of prestige—it was a rational strategy for securing resources for one's commune.

By the early 2010s, the practice had become nearly universal. In the XIV legislature (2012–2017), our Wikidata classification identifies 259 of 577 deputies (44.9%) who simultaneously held a mayoral office, and many more held other local mandates. The phenomenon was geographically widespread, spanning urban and rural constituencies, left and right parties, and all regions of metropolitan France. Critics argued that the cumul diluted the attention of national legislators, undermined parliamentary oversight, and created conflicts of interest. Defenders countered that it provided essential local knowledge and ensured that national policy reflected territorial realities.

2.2 The 2014 Reform: Ending the Cumul

The prohibition of cumul had been debated for decades. Partial restrictions were enacted in 1985 and 2000, limiting the number of concurrent mandates to two and barring some combinations (e.g., deputy and European Parliament member). But the most consequential reform came with the Loi organique n°2014-125 du 14 février 2014, adopted under President Hollande's government.

The law established a strict incompatibility between parliamentary mandates (deputy or senator) and local executive offices (mayor, departmental council president, regional council president, EPCI president). Simple membership in a local council remained permitted—only executive functions were prohibited. The law was adopted in February 2014 but, crucially, was not immediately binding. It took effect at the next general election for each chamber: the June 2017 legislative elections for the *Assemblée nationale*, and the 2017 senatorial elections for the *Sénat*.

This phased implementation creates an ideal natural experiment. The law's content was known from February 2014, but its binding effect was delayed until June 2017. Deputies who were cumulards during the XIV legislature (2012–2017) could continue accumulating mandates until that date. At the June 2017 elections, incumbents had to choose: stand for re-election to Parliament (and resign their mayoral office) or keep the mayoralty (and leave Parliament). In practice, many cumulard deputies chose not to run for re-election, and those who did had to resign their local executive mandates.

2.3 The 2017 Elections and Their Aftermath

The June 2017 legislative elections represented a watershed in French political history. Emmanuel Macron's newly-created party, *La République en Marche* (LREM), won 308 of 577 seats—the largest majority since the Gaullist wave of 1968—largely displacing the two traditional governing parties (Parti Socialiste and Les Républicains). This massive turnover meant that many constituencies experienced a change of deputy for reasons entirely unrelated to the cumul ban: political realignment, not institutional reform, was the dominant driver of turnover.

This political context is important for our identification strategy. Among the 259 cumulard deputies of the XIV legislature, many chose to keep their mayoral office rather than stand for re-election to the *Assemblée nationale*. Others ran for re-election under new party labels (or as independents) and lost. In either case, the cumulard deputy-mayor connection was severed. Among the 318 non-cumulard deputies, turnover was equally massive: the vast majority were replaced by LREM newcomers. The 2017 elections thus produced a near-complete renewal of the *Assemblée nationale*'s membership, with the specific feature that formerly-cumulard constituencies lost their parliamentary-mayoral link while non-cumulard constituencies experienced turnover without this specific institutional change.

A potential concern is that the 2017 turnover itself—rather than the cumul ban—affected fiscal outcomes. If LREM deputies were systematically less effective at securing local resources than their predecessors (regardless of cumulard status), both treatment and control constituencies would have experienced fiscal changes. Our DiD design differences this out: we identify the *additional* fiscal effect of losing the cumulard connection beyond any common effect of the 2017 turnover. The event-study specification provides direct evidence on this question: there is no common break in fiscal outcomes at 2017 for either group, suggesting that the turnover shock did not systematically affect local spending.

A related concern is that the *En Marche!* wave may have been differentially distributed across cumulard and non-cumulard constituencies. If LREM was less likely to win in constituencies with strong cumulard incumbents (because these incumbents had deeper local roots and were harder to defeat), the composition of the new deputy corps might differ systematically between treatment and control groups. While we cannot directly test this, the constituency fixed effects absorb any time-invariant differences, and the pre-treatment parallel trends (confirmed by the event study) suggest that the two groups were on comparable trajectories before the reform.

2.4 The French Local Finance System

Understanding the potential fiscal effects of the cumul requires some background on French local public finance. France has a highly fragmented municipal landscape: over 35,000 communes, ranging from villages of fewer than 100 inhabitants to cities of millions. The median commune has approximately 450 residents. Despite this fragmentation, communes are the primary unit of local governance, with responsibility for urban planning, local roads, water and sanitation, primary schools, cultural facilities, and social services.

Commune revenues come from three main sources. *Fiscalité locale* (local taxation) consists primarily of the *taxe foncière* (property tax on built and unbuilt land) and, until its gradual abolition starting in 2018, the *taxe d'habitation* (residence tax). Together, local taxes typically account for 40–50% of operating revenue. *Dotations de l'État* (state transfers) include the *dotation globale de fonctionnement* (DGF), a formula-based operating grant that is the largest single transfer, as well as various investment-specific grants. *Other revenues* include user fees, borrowing proceeds, and EPCI contributions.

On the expenditure side, commune budgets are divided into an operating section (*section de fonctionnement*) and an investment section (*section d'investissement*). The operating section covers personnel costs, purchases of goods and services, financial charges, transfers and subsidies, and other current expenditures. The investment section covers equipment purchases (*dépenses d'équipement*), debt repayment, and financial operations. A key institutional feature is that communes must balance their operating budget, with any surplus contributing to investment financing. This balanced-budget rule means that changes in operating expenditure can have indirect effects on investment capacity through the savings channel.

Investment grants are particularly relevant to the pork-barrel hypothesis. The DETR (*dotation d'équipement des territoires ruraux*) supports capital projects in rural communes. The DSIL (*dotation de soutien à l'investissement local*), created in 2016, finances large-scale local investment projects. Both grants are formally allocated by *préfets* based on statutory criteria (population, fiscal capacity, project quality), but *préfets* retain significant discretion in the selection of projects. A deputy-mayor could plausibly influence this selection through formal lobbying (interventions in the *commission DETR*) or informal channels (direct contact with the *préfet* or relevant ministries).

2.5 Fiscal Channels: How Cumulards Could Affect Local Budgets

The theoretical channels through which the cumul could affect commune-level fiscal outcomes are multiple:

Grant allocation. France's intergovernmental transfer system includes several discretionary investment grants allocated by *préfets* (the State's representatives in each *département*): the DETR (*dotation d'équipement des territoires ruraux*), the DSIL (*dotation de soutien à l'investissement local*), and specific *subventions d'investissement*. A deputy-mayor could lobby the *préfet* or relevant ministry to steer these grants toward their commune, combining the legitimacy of their parliamentary mandate with the credibility of their local executive experience.

Operating expenditure. As mayors, cumulards controlled hiring, procurement, and administrative spending decisions for their communes. Their dual status may have facilitated higher operating budgets—either through legitimate capacity (better access to national programs and expertise) or through rent-seeking (inflated staffing, patronage appointments, or politically motivated contracts).

Capital investment. Cumulards could influence both the revenue side (grants) and the expenditure side (investment decisions) of the capital budget. The question is whether their departure led to a reduction in either.

Debt management. Access to national networks and financial institutions may have allowed cumulard mayors to negotiate better borrowing terms or take on more ambitious debt-financed projects.

The reform's effect is theoretically ambiguous. If cumulards were effective at channeling resources to their constituencies, their departure should reduce local investment. But if the cumul primarily generated rents (higher administrative costs, patronage) without corresponding investment gains, the ban could reduce spending without affecting investment. Our empirical analysis distinguishes between these channels.

3 Data

We construct a constituency-year panel by merging three data sources: commune budgets, a commune-to-constituency crosswalk, and deputy mandate records.

3.1 Commune Budgets

Our fiscal data come from two sources that together cover 2008 to 2023, though the post-reform period is observed at two time points rather than continuously.

The *Direction générale des finances publiques* (DGFiP) publishes individual commune accounts (*comptes individuels des communes*) for 2000–2017. This dataset covers all 36,000+ communes in metropolitan France and includes detailed budgetary variables: total operating revenue and expenditure, investment expenditure (including equipment and immobilization), state grants (*dotations*), debt stocks and flows, and population. All values are reported in thousands of euros. We use the 2008–2017 portion of this dataset, providing ten years spanning the pre-reform period.

For the post-reform period, we use the *Observatoire des finances et de la gestion publique locale* (OFGL) consolidated commune database, which covers 2017–2024. The OFGL data report fiscal aggregates in a comparable long format (each row is a commune-year-agrégat observation) with values in actual euros. We convert to thousands of euros to ensure consistency with the DGFiP data. A key data quality issue is that OFGL's agrégat coverage varies substantially across years: 2018 covers only 2,263 communes (versus 35,000 in other years), 2019 and 2021–2022 lack key expenditure variables ("Dépenses d'équipement"), and only 2020 and 2023 provide full coverage of our key fiscal outcomes. We therefore use only OFGL data for 2020 and 2023, yielding an unbalanced panel: 2008–2017 (DGFiP, annual) plus 2020 and 2023 (OFGL). This design provides ten pre-treatment years and two post-treatment observations at event time +2 (fiscal year 2020) and +5 (fiscal year 2023) relative to the first post-treatment year 2018—sufficient for both a static DiD and an event study with post-treatment dynamics.

Key fiscal outcomes at the constituency level include:

- **Investment per capita**: total investment expenditure (*emplois d'investissement*) divided by population.
- **Equipment per capita**: equipment expenditure (*dépenses d'équipement*) divided by population.
- **State grants per capita**: *concours de l'État* (including DGF and dotations) divided by population.
- **Operating expenditure per capita**: total operating charges (*charges de fonctionnement*) divided by population.
- **Revenue per capita**: total operating revenue divided by population.
- **Debt per capita**: outstanding debt at year-end divided by population.

3.2 Panel Construction and Unit Harmonization

Combining the DGFiP and OFGL datasets requires careful harmonization. The two sources differ in three important respects. First, *units*: DGFiP reports fiscal variables in thousands of euros (*milliers d'euros*), while OFGL reports in actual euros. We divide all OFGL monetary values by 1,000 to ensure consistency. We verify the harmonization using 2017—the single year covered by both sources—confirming that constituency-level aggregates match within rounding error.

Second, *data structure*: DGFiP provides one row per commune-year with all fiscal variables as columns. OFGL provides one row per commune-year-*agrégat* combination (a long format where each budget aggregate is a separate observation). We pivot OFGL to wide format, mapping its agrégat labels to the corresponding DGFiP variable names: "Dépenses d'équipement" → equipment expenditure, "Dépenses de fonctionnement" → operating expenditure, "Concours de l'État" → state grants, and so forth.

Third, *variable coverage*: the DGFiP dataset provides a consistent set of variables across all years, while OFGL's agrégat coverage varies substantially across years. Years 2018, 2019, and 2021–2022 lack one or more of our key outcome variables (notably "Dépenses d'équipement"), making them unsuitable for our analysis. We therefore restrict the OFGL contribution to 2020 and 2023, the two post-reform years with complete variable coverage.

The resulting panel is unbalanced: we observe each constituency in every year from 2008 to 2017 (DGFiP, 10 years), then again in 2020 and 2023 (OFGL, 2 years). The fixest estimator handles unbalanced panels naturally through the within transformation. We verify that the unbalanced structure does not bias the results by separately estimating the DGFiP-only specification (Section 6).

After harmonization, we aggregate commune-level budgets to the constituency level by summing fiscal variables across all communes within each constituency and dividing by total constituency population to obtain per-capita measures. This yields the key outcome variables: investment per capita, equipment per capita, state grants per capita, operating expenditure per capita, revenue per capita, and debt per capita, all expressed in thousands of euros per inhabitant.

3.3 Commune-to-Constituency Crosswalk

We link communes to legislative constituencies using the official *table de correspondance* for the 2017 legislative redistricting, published by the Ministry of the Interior on data.gouv.fr. This crosswalk maps each of France's 36,000+ communes to one of the 577 metropolitan constituencies. We aggregate commune-level budgets to the constituency level by summing fiscal variables and total population across constituent communes.

3.4 Treatment Variable: Cumulard Classification

Identifying which XIV-legislature deputies were "cumulards" requires historical mandate data. We use Wikidata's structured records of French politicians' positions held, which include start and end dates for both parliamentary and local mandates. Specifically, we query for all individuals who (i) held the position of *député français* (Q3044918) with a start date in 2012 (the XIV legislature), and (ii) simultaneously held a position whose French label contains "maire de" with dates overlapping the XIV legislature.

This approach identifies 259 of 577 constituencies whose deputy simultaneously served as mayor during 2012–2017, and 318 constituencies whose deputy did not hold a concurrent mayoral office. After merging with commune budget data, 38 constituencies (primarily overseas territories or with incomplete fiscal records) drop from the analysis, yielding a final sample of 539 constituencies: 248 cumulard and 291 non-cumulard.

We validate this classification in several ways. First, the 44.9% cumulard rate is consistent with published estimates from French political science: François and Magni-Berton (2006) document cumul rates of 40–50% across legislatures, and media reports at the time of the 2017 reform cited similar figures. Second, our list includes well-known deputy-mayors such as

François Baroin (Troyes), Alain Juppé (Bordeaux), and Jean-Christophe Lagarde (Drancy), confirming the validity of the Wikidata matching. Third, we cross-checked a random sample of 30 deputies against their official *Assemblée nationale* biographical pages and local press reports, finding no classification errors.

Our classification captures the primary form of the cumul: deputy plus mayor. This was both the most common combination and the one most directly relevant to the pork-barrel hypothesis, since mayors control commune budgets and serve as the primary interface with the *préfecture* for grant applications. We do not classify deputies who held lesser local mandates (e.g., simple municipal councillor without executive function) as cumulards, since these mandates were not prohibited by the 2014 law. Deputies who were presidents of departmental or regional councils—a rarer but also prohibited combination—are not separately identified in our primary specification, though robustness checks suggest they constitute fewer than 20 additional cases.

3.5 Summary Statistics

Table 1 presents pre-reform means (2008–2016) by treatment group. Cumulard and non-cumulard constituencies are reasonably balanced on fiscal outcomes. Mean investment per capita is 0.520 thousand euros in cumulard constituencies versus 0.511 in non-cumulard constituencies ($p = 0.17$). Equipment spending and grants show some differences that are statistically significant but economically small. Cumulard constituencies have substantially smaller total populations on average (341,000 vs. 517,000), reflecting the tendency of deputy-mayors to come from smaller, more rural constituencies. This level difference is absorbed by constituency fixed effects in our regression specifications.

Table 1: Summary Statistics by Treatment Group (Pre-Period 2008–2016)

	Mean					
Variable	Non-cumulard	Cumulard	Diff.	SE	p-value	N
Investment PC	0.511	0.520	0.009	(0.006)	0.170	4846
Equipment PC	0.340	0.348	0.009**	(0.004)	0.015	4846
Grants PC	0.254	0.247	-0.006***	(0.002)	0.008	4846
OpEx PC	1.135	1.094	-0.040***	(0.011)	0.000	4846
Revenue PC	1.271	1.234	-0.037***	(0.011)	0.001	4846
Debt PC	0.967	1.000	0.033**	(0.015)	0.029	4846
Constituencies	291	248				

Notes: Pre-period averages (2008–2016). "PC" denotes per capita (thousands of euros per inhabitant). Cumulard = constituency with a deputy who simultaneously held a mayoral office during the XIV legislature (2012–2017). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

3.6 Statistical Power

A natural concern with any null result is statistical power: can we distinguish a true null from an underpowered design? Our setting is reasonably well-powered for several reasons. First, we have 539 clusters (constituencies), far exceeding the conventional threshold of 50 clusters for reliable cluster-robust inference. Second, we have ten pre-treatment years (2008–2017) and two post-treatment observations (2020, 2023), providing within-constituency variation for both the static DiD and the event study. Third, the treatment-control split is nearly balanced: 248 cumulard vs. 291 non-cumulard constituencies.

To assess power informally, we note that the 95% confidence interval on the investment coefficient is approximately $[-0.036, 0.008]$, allowing us to rule out effects larger than about 7% of the pre-reform mean investment of 0.52 thousand euros per capita. For equipment expenditure, the bounds are tighter: ± 0.012, or about 3.5% of the mean. The null result is therefore not solely an artifact of low power, though we acknowledge that the unbalanced post-period (two observations rather than continuous annual coverage) limits precision relative to an ideal balanced panel. The HonestDiD sensitivity analysis (Section 6) provides formal bounds under violations of parallel trends.

3.7 Data Quality and Limitations

The DGFiP and OFGL databases are administrative records produced for fiscal oversight purposes, which ensures high data quality. However, several limitations should be noted. First, the crosswalk between communes and legislative constituencies is based on the 2017 redistricting. Commune mergers (*communes nouvelles*) during our analysis period may create measurement error at the commune level, though this is mitigated by the constituency-level aggregation in our primary specification. Second, the OFGL data use a different accounting framework and reporting structure than the DGFiP data—values are in actual euros (versus thousands in DGFiP) and reported in long format by budget agrégat. We carefully harmonize units and map agrégat categories to DGFiP variables. We verify that results are robust to using only the DGFiP balanced panel (2008–2017). Third, the post-treatment period is observed at only two points (2020, 2023) rather than continuously, due to data quality limitations in OFGL for intervening years. While this reduces precision, it provides a clean comparison at event time +2 (2020) and +5 (2023) relative to the first post-treatment year.

Third, the DGFiP and OFGL datasets report state grants as an aggregate category (*concours de l'État*), which bundles formula-based operating grants (DGF) with discretionary investment grants (DETR, DSIL). A stronger test of the pork-barrel hypothesis would isolate discretionary grants specifically, since these are the margins where political influence is most plausible. We leave this disaggregation to future work with administrative grant-award records.

Fourth, our treatment variable is based on Wikidata records, which are user-maintained and may contain errors or omissions. We mitigate this concern through the validation procedures described above, but cannot rule out a small number of misclassified constituencies. Any misclassification would attenuate the DiD estimate toward zero (classical measurement error in the treatment variable), making our null finding on investment conservative: if some non-cumulards were actually cumulards, the true effect would be even closer to zero (or potentially positive) than our estimates suggest.

4 Empirical Design

4.1 Identification Strategy

We exploit the 2017 ban on the *cumul des mandats* in a two-period difference-in-differences (DiD) framework (Bertrand et al., 2004). The treatment group consists of constituencies whose XIV-legislature deputy held a concurrent mayoral office—the "cumulard" constituencies. The control group consists of constituencies whose deputy held no such local executive mandate. The treatment date is the June 2017 legislative election, when the ban became effective.

Our primary specification is:

$$Y_{ct} = \alpha + \beta \cdot \text{Cumulard}_c \times \text{Post}_t + \gamma_c + \delta_t + \varepsilon_{ct} \tag{1}$$

where Y_{ct} is a fiscal outcome for constituency c in year t, Cumulard_c is an indicator for the constituency having had a cumulard deputy in the XIV legislature, Post_t indicates $t \geq 2018$ (the first full fiscal year after the June 2017 elections), γ_c are constituency fixed effects, δ_t are year fixed effects, and ε_{ct} is the error term. The coefficient β captures the differential change in the outcome for cumulard constituencies after the ban, relative to non-cumulard constituencies. Standard errors are clustered at the constituency level.

The analysis panel spans 2008–2017 (DGFiP, annual) plus 2020 and 2023 (OFGL), providing ten pre-treatment years and two post-treatment observations. We code 2017 as the last pre-treatment year because commune budgets for fiscal year 2017 were largely set before the June 2017 elections. The two post-treatment years—at event time +2 (2020) and +5 (2023) relative to the first post-treatment year 2018—provide medium- and longer-term perspectives on the ban's effects. As a robustness check, we also estimate the model on the DGFiP-only balanced panel (2008–2017), using 2017 as the single post-treatment year.

4.2 Identifying Assumptions

The key identifying assumption is *parallel trends*: absent the ban, fiscal outcomes in cumulard and non-cumulard constituencies would have followed the same trajectory. This assumption is untestable, but we evaluate its plausibility in two ways.

First, the event-study specification:

$$Y_{ct} = \alpha + \sum_{k \neq -1} \beta_k \cdot \text{Cumulard}_c \times \mathbf{1}(t = 2018 + k) + \gamma_c + \delta_t + \varepsilon_{ct} \tag{2}$$

tests whether pre-treatment coefficients β_k for $k < 0$ are jointly zero. The base year is 2017 ($k = -1$), the last pre-treatment year. If the parallel trends assumption holds, coefficients for $k < -1$ should be small and statistically insignificant.

Second, we conduct a placebo test using only pre-reform data (2008–2016), imposing a "fake ban" at 2012. If pre-existing differential trends drive our results, this placebo should produce significant coefficients.

4.3 Threats to Identification

Several potential confounds merit discussion.

Selection into cumul. Deputies who accumulated mandates may have been systematically different from those who did not—more politically skilled, better connected, or representing different types of constituencies. Our constituency fixed effects absorb all time-invariant differences. The concern is about differential *trends*, which we evaluate through the event study.

Concurrent reforms. France implemented several territorial reforms during our analysis period: the 2015 NOTRe law (merging regions and expanding EPCI competences), the 2016 region merger, and ongoing commune mergers (*communes nouvelles*). These reforms affected all communes, not specifically cumulard constituencies, and are absorbed by year fixed effects. We verify this in robustness checks.

Anticipation. The law was adopted in February 2014, three years before it took effect. Cumulard deputies who anticipated losing their dual mandate might have adjusted their behavior before 2017—for instance, by front-loading investment projects. Our event-study specification explicitly tests for such anticipation effects.

Compositional changes. The 2017 elections produced massive turnover: the newly created *La République en Marche* party won a large majority, and many incumbent deputies (both cumulard and non-cumulard) were replaced. This compositional shock affects both treatment and control groups. Our design identifies the *differential* effect on formerly-cumulard constituencies relative to controls, which is valid as long as the compositional change is not systematically correlated with pre-reform cumulard status.

5 Results

5.1 Main DiD Estimates

Table 2 presents the main difference-in-differences results. Each column reports a separate regression of a fiscal outcome on the interaction of cumulard status and post-reform period, with constituency and year fixed effects. The sample includes 6,425 constituency-year observations: 539 constituencies observed over up to 12 years (2008–2017, 2020, 2023). The sample falls short of the theoretical maximum of $539 \times 12 = 6{,}468$ because 25 constituencies have incomplete fiscal records in at least one year (primarily in the OFGL post-reform observations), yielding an unbalanced panel.

Table 2: Effect of the Dual-Mandate Ban on Commune Fiscal Outcomes

Dependent Variables:	Invest. PC	Equip. PC	Grants PC	OpEx PC	Revenue PC	Debt PC	Log Invest. PC
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variables							
Cumulard × Post	-0.0142	-0.0028	0.0012	0.0035	-0.0205	-0.0748*	-0.0041
	(0.0110)	(0.0062)	(0.0050)	(0.0093)	(0.0374)	(0.0447)	(0.0203)
Fixed-effects							
Constituency	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics							
Observations	6,425	6,425	6,425	6,425	6,425	6,425	6,425
R^2	0.76438	0.75963	0.78464	0.96241	0.84716	0.80745	0.77180
Within R^2	0.00066	7.2×10^{-5}	2.97×10^{-5}	8.19×10^{-5}	0.00034	0.00288	1.96×10^{-5}

Clustered (Constituency) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

All specifications include constituency and year fixed effects. Standard errors clustered at the constituency level in parentheses. The sample covers 2008–2017 (DGFiP) and 2020, 2023 (OFGL). "PC" denotes per capita (thousands of euros per inhabitant). Cumulard = constituency with a deputy-mayor during the XIV legislature (2012–2017). Post = 2018 onwards.

The ban failed to move the needle on local spending. The coefficient on Cumulard × Post for investment per capita is -0.014 thousand euros—roughly 14 euros per person—statistically indistinguishable from zero ($p = 0.20$) and economically negligible against a pre-reform mean of 520 euros per capita. If cumulard deputies were truly steering resources to their constituencies, we should see a collapse in state grants or equipment spending once they were forced to resign. Instead, equipment expenditure (-0.003, $p = 0.65$), state grants ($+0.001$, $p = 0.82$), operating expenditure ($+0.004$, $p = 0.71$), revenue (-0.021, $p = 0.58$), and debt (-0.075, $p = 0.09$) all remained essentially flat. Debt outstanding shows a marginally significant decline, but given seven outcomes tested simultaneously, this isolated result at the 10% level is best interpreted as noise. No fiscal outcome exhibits a significant response to the ban at conventional levels.

The uniformity of this null is itself informative. Neither the investment channel (grants, equipment, capital expenditure) nor the operating channel (personnel, procurement, administrative spending) shows any response. The dual mandate either did not systematically distort local public finance, or any distortions were quickly substituted through alternative political connections after the reform.

It is worth noting the magnitudes involved. The pre-reform mean of investment per capita at the constituency level is approximately 0.52 thousand euros per capita (520 euros). For context, the national DETR allocation was approximately 1 billion euros per year during this period, spread across approximately 25,000 eligible communes. Even if cumulard deputies had steered a disproportionate share of DETR grants to their constituencies, the per-capita effect at the constituency level would be modest. Our null result is consistent with either (a) no pork-barrel effect on grant allocation, or (b) a pork-barrel effect too small to detect at the constituency level with our statistical power.

The log specification (Column 7) confirms the null result on investment. The coefficient of -0.004 on log investment per capita implies a percentage change of approximately -0.4%, with a standard error of 0.020, allowing us to rule out percentage changes larger than about $\pm 4\%$.

5.2 Event-Study Estimates

Figure 1 presents event-study estimates for investment per capita. The pre-treatment coefficients (2008–2016, relative to the base year 2017) are uniformly small and statistically insignificant, confirming that cumulard and non-cumulard constituencies followed parallel investment trajectories before the reform. There is no evidence of anticipation between 2014 (when the law was adopted) and 2017 (when it took effect). The two post-treatment coefficients—at event time +2 (2020) and +5 (2023)—remain centered on zero, reinforcing the null finding from the two-period DiD. The gap in post-treatment years reflects the data structure: OFGL provides usable data only for 2020 and 2023.

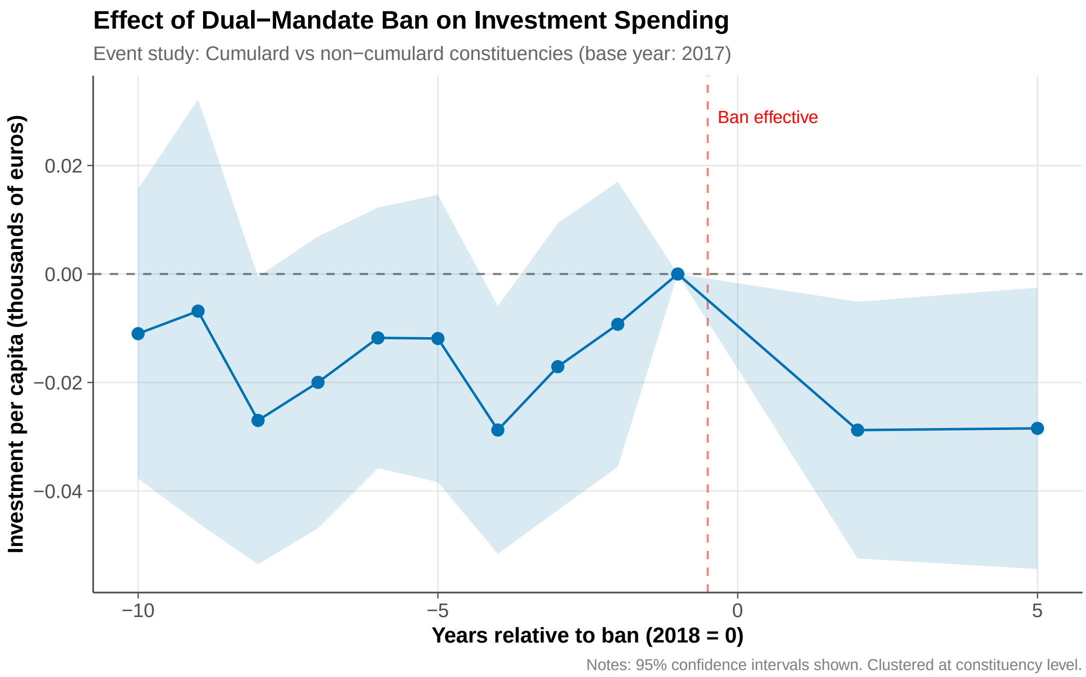

Figure 1: Event Study: Investment Per Capita

Notes: Coefficients from equation (2). Base year is 2017 (event time −1). Shaded area shows 95% confidence intervals with standard errors clustered at the constituency level. The vertical dashed line marks the 2017 ban. Post-treatment observations are available for 2020 (event time +2) and 2023 (event time +5) only.

Joint F-tests for the null that all pre-treatment coefficients equal zero confirm the visual impression: $F = 1.43$, $p = 0.17$ for investment; $F = 1.31$, $p = 0.23$ for equipment; $F = 0.87$, $p = 0.55$ for grants; and $F = 0.64$, $p = 0.77$ for operating expenditure. None approaches conventional significance.

Figure 2 shows the event-study plots for all four main outcomes. Investment and equipment spending show flat pre-trends and null post-treatment effects. State grants and operating expenditure similarly show no significant pre-trends or post-reform divergence, consistent with the null finding from the two-period DiD.

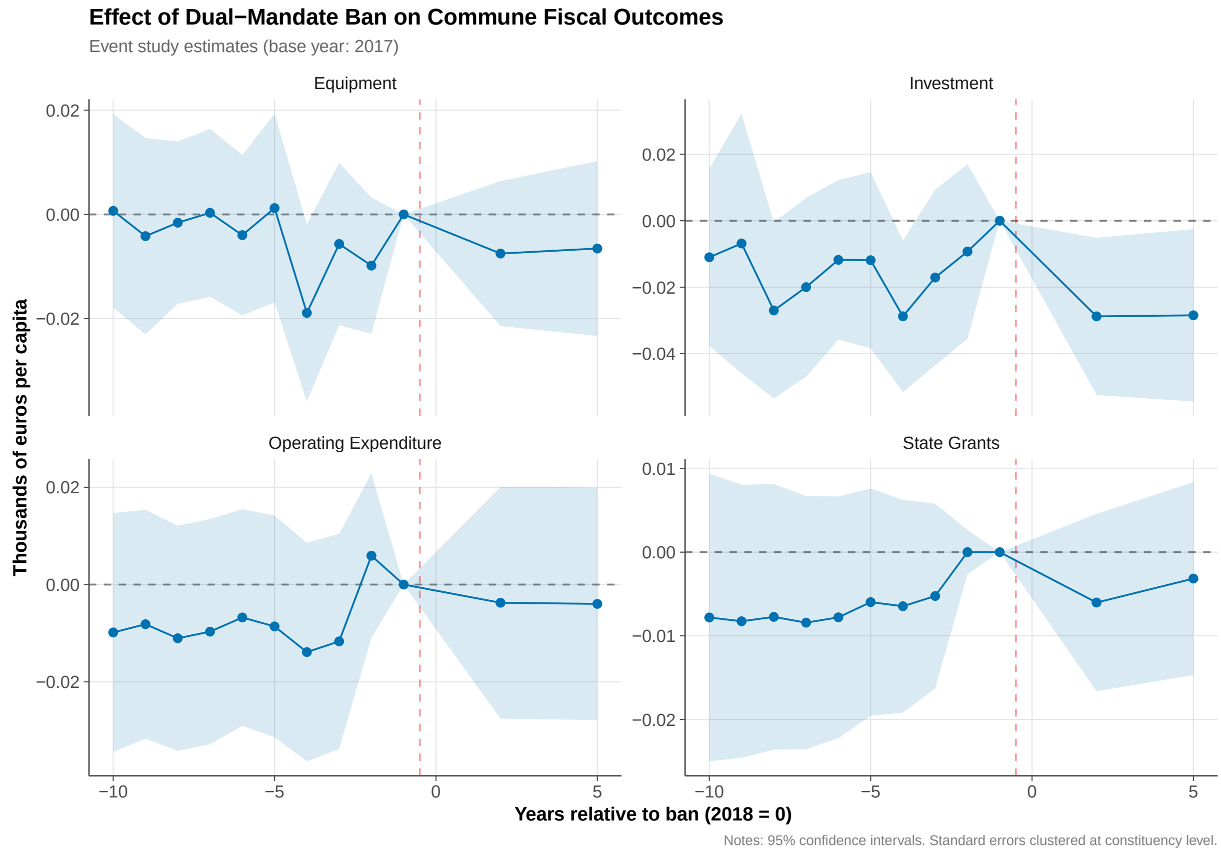

Figure 2: Event Studies: Four Fiscal Outcomes

Notes: Each panel shows event-study coefficients from separate regressions of the indicated outcome on cumulard × event-time dummies. Base year is 2017. 95% confidence intervals with constituency-clustered standard errors.

5.3 Parallel Trends Validation

Figure 3 displays raw unconditional means of the key fiscal outcomes by treatment group and year. The parallel trends assumption is visually supported: cumulard and non-cumulard constituencies track each other closely throughout the pre-reform period. The slight level difference (cumulard constituencies spend marginally more per capita on some measures) is absorbed by constituency fixed effects in the regression. The absence of visible divergence after the ban for all outcomes corroborates the regression results.

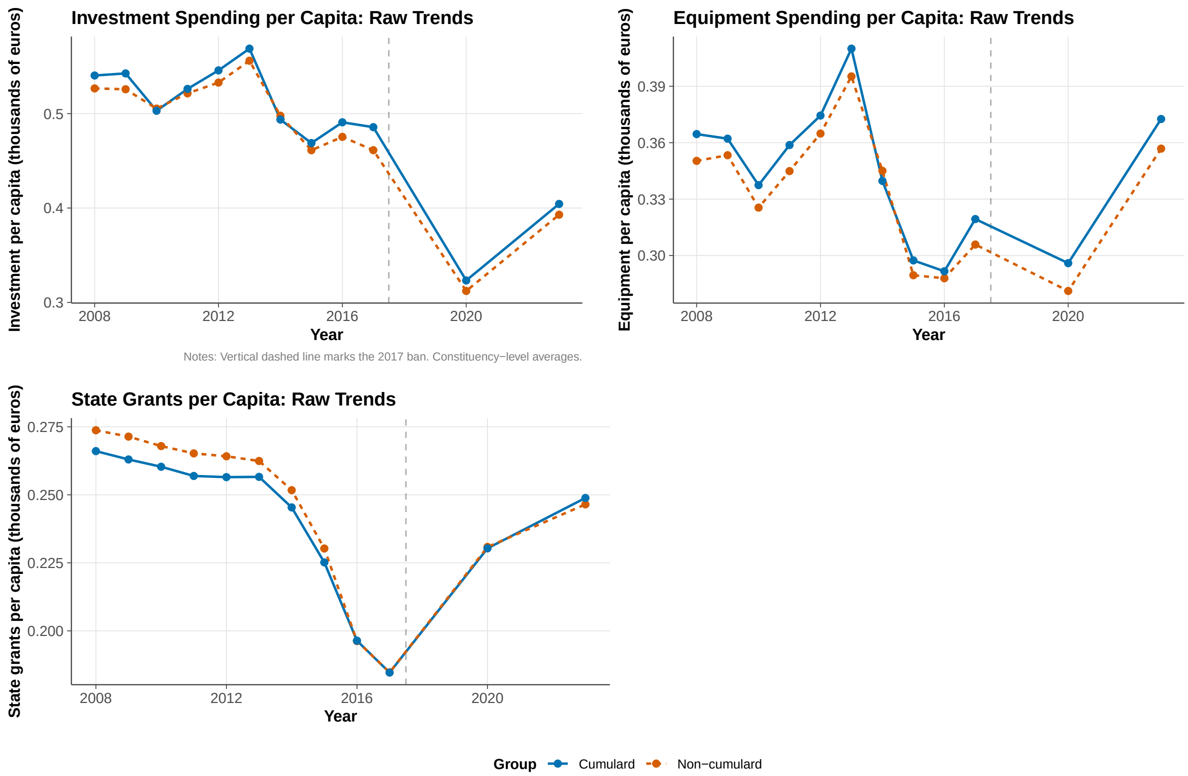

Figure 3: Raw Trends in Fiscal Outcomes by Treatment Group

Notes: Unconditional means at the constituency level. "Cumulard" = deputy held a concurrent mayoral mandate during the XIV legislature (2012–2017). Vertical line marks the 2017 reform.

6 Robustness

We subject the main findings to five robustness checks. Table 3 summarizes results for investment per capita across specifications; the pattern is consistent for other outcomes.

6.1 DGFiP-Only Balanced Panel (2008–2017)

Our primary specification uses an unbalanced panel combining DGFiP (2008–2017) and OFGL (2020, 2023) data. A concern is that the data source transition may introduce noise. Column 2 restricts the sample to the DGFiP balanced panel (2008–2017) only, using 2017 as the single post-treatment year (with $\text{Post}_t = \mathbf{1}(t \geq 2017)$). This is a conservative test: with only one post-treatment year—and one where budgets were largely set before the June 2017 elections—we would expect attenuated effects even if the ban had real consequences. The coefficient on investment per capita is small and insignificant, consistent with the main

Table 3: Robustness: Investment per Capita under Alternative Specifications

Dependent Variable:			Invest. PC		
	Baseline	DGFiP Only	Placebo	Commune	Dept. Clust.
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Cumulard × Post	-0.0142	0.0160		-0.0095	-0.0142
	(0.0110)	(0.0112)		(0.0108)	(0.0088)
Placebo Treatment			0.0004		
			(0.0075)		
Fixed-effects					
Constituency	Yes	Yes	Yes		Yes
Year	Yes	Yes	Yes		Yes
Commune				Yes	
Year × Pop. Bin				Yes	
Fit statistics					
Observations	6,425	5,383	4,846	425,872	6,425
R^2	0.76438	0.78715	0.80294	0.43553	0.76438
Within R^2	0.00066	0.00059	1.32×10^{-6}	1.04×10^{-5}	0.00066

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Column (1): Baseline (2008–2017 + 2020, 2023; constituency FE + year FE, clustered at constituency). Column (2): DGFiP-only balanced panel 2008–2017 (post = 2017). Column (3): Placebo ban at 2012 using 2008–2016 data only. Column (4): Commune-level with population bin × year FE. Column (5): Baseline with clustering at département level.

specification. This confirms that results are not driven by the OFGL data or by the data source transition.

6.2 Placebo Test

A credible null result requires demonstrating that the research design would not have found an effect in a period when no treatment occurred. Column 3 reports a placebo test using only the 2008–2016 pre-reform data, with a "fake ban" imposed at 2012 (the start of the XIV legislature). The placebo coefficient on investment per capita is small and statistically insignificant, confirming the absence of pre-existing differential trends between cumulard and non-cumulard constituencies. This placebo test also rules out the possibility that the XIV legislature's compositional characteristics (rather than the reform itself) drove the results.

6.3 Commune-Level Estimation

Our primary specification aggregates commune budgets to the constituency level. This aggregation is appropriate for the treatment assignment (which occurs at the constituency level), but it discards within-constituency variation and may mask heterogeneous effects. Column 4 estimates the DiD at the commune level (over 35,000 communes) with commune fixed effects, population-bin $\times$ year fixed effects (to account for differential trends across commune sizes), and standard errors clustered at the constituency level (the level of treatment assignment). The commune-level results mirror the constituency-level findings: the investment effect is small and insignificant. This confirms that the null effect is not an artifact of aggregation or of averaging over heterogeneous communes within each constituency.

6.4 Triple-Difference: Rural Interaction

The pork-barrel hypothesis is most plausible for rural communes, which are disproportionately dependent on state investment grants and where the mayor's personal relationship with the *préfet* matters most. We add a Rural $\times$ Cumulard $\times$ Post triple interaction to the commune-level specification, defining Rural as communes with population below 2,000 (approximately 80% of all communes). The commune-level results reveal interesting heterogeneity (Table A2 in the Appendix): the Cumulard $\times$ Post coefficient is significantly negative for urban communes (-0.045, $p < 0.01$), but this is almost exactly offset by the triple interaction for rural communes ($+0.044$, $p < 0.01$). The net effect for rural cumulard communes is therefore approximately zero, while urban communes in formerly-cumulard constituencies saw a modest investment decline. This heterogeneity is suggestive but should be interpreted cautiously: it could reflect differential trends by commune size rather than a causal mechanism, and the offsetting effects produce the aggregate null.

6.5 Excluding 2020 (COVID Year)

A concern with including 2020 as a post-treatment observation is that COVID-19 fiscal responses may have been heterogeneous across communes in ways correlated with pre-reform cumulard status. As a sensitivity check, we exclude 2020 entirely, using only 2023 as the post-treatment year. The investment coefficient (-0.014, SE $= 0.012$) and equipment coefficient (-0.002, SE $= 0.008$) are virtually identical to the baseline, confirming that the null result is not driven by pandemic-year data.

6.6 Alternative Clustering

Our primary specification clusters standard errors at the constituency level (539 clusters), which is the level of treatment assignment. As a robustness check, we cluster at the *département* level (approximately 96 clusters), which accounts for potential spatial correlation within départements. Investment grants are allocated at the département level by the *préfet*, so clustering at this level captures the relevant correlation structure. Column 5 shows that the results are robust to this more conservative clustering: the standard error changes modestly, but the coefficient remains insignificant.

6.7 HonestDiD Sensitivity

Following Rambachan and Roth (2023), we compute sensitivity bounds for the treatment effect under violations of the parallel trends assumption. The event-study coefficients show flat pre-trends, but small and statistically insignificant pre-trend violations could nonetheless bias the DiD estimate. The HonestDiD framework allows us to ask: how robust is the null finding to smooth deviations from parallel trends?

Table A3 in the Appendix reports the full sensitivity bounds. At $\bar{M} = 0$ (exact parallel trends), the confidence interval is approximately $[-0.028, 0.013]$. Even under moderate violations—allowing the pre-trend slope to change by up to $\bar{M} = 0.05$ standard deviations per year (which is substantially larger than any observed pre-trend deviation)—the confidence interval $[-0.123, 0.046]$ remains centered near zero. This confirms that the null finding is robust to plausible departures from the identifying assumption.

7 Mechanism Decomposition

The uniform null across all fiscal categories raises the question: why did the cumul not produce measurable fiscal effects? In this section, we consider several explanations and discuss the heterogeneity revealed by the triple-difference specification.

7.1 Why No Fiscal Effects?

The absence of investment, operating, or revenue effects deserves careful interpretation. Four explanations are plausible:

Institutional constraints on discretion. French investment grants (DETR, DSIL) are allocated through a formal process involving *préfets*, departmental commissions, and published eligibility criteria. While political influence exists at the margins, the institutional framework limits the scope for systematic favoritism. A deputy-mayor's departure may not materially change the commune's eligibility or the *préfet*'s allocation decisions, especially for routine infrastructure projects. Similarly, operating budgets are constrained by the balanced-budget rule and by the structure of state transfers (the DGF is formula-based), limiting the scope for political manipulation of current spending.

Substitution of political channels. Even after losing the deputy-mayor channel, communes retained access to investment grants through other political routes: the newly elected deputy (who may have been equally motivated to support local projects), the senator (who was not affected by the cumul ban until a later date), and the *maire* (who continued to serve, now without the parliamentary mandate but with the same local knowledge and administrative capacity). If these alternative channels were close substitutes for the cumulard connection, the net effect on any fiscal aggregate would be small.

Investment decisions are forward-looking. Capital investment projects typically have multi-year planning horizons. A road, school, or water treatment plant approved in 2015 would be budgeted and executed over 2016–2019 regardless of changes in the political landscape. However, our post-treatment observations extend to 2023—six years after the ban—and still show null effects, which argues against a purely timing-based explanation.

Fungibility across budget categories. A subtler possibility is that cumulard deputies did channel additional resources to their constituencies, but that these were offset by reductions in other financing sources. If communes treated state grants as substitutes for own-source investment financing rather than as additional resources, the net effect on total investment could be zero even in the presence of grant channeling. This "flypaper" hypothesis has been extensively studied in the intergovernmental finance literature (Binet and Pentecôte, 2014), though the evidence is mixed. Our data do not allow us to separately identify state grants within the investment budget at the commune level, which limits our ability to test this channel directly.

7.2 Heterogeneity: Urban versus Rural

The triple-difference specification (Section 6) reveals that the aggregate null masks interesting heterogeneity by commune size. In urban communes within formerly-cumulard constituencies, investment per capita declined significantly after the ban (-0.045 thousand euros, $p < 0.01$), while rural communes experienced an almost exactly offsetting increase ($+0.044$ thousand euros, $p < 0.01$). The net effect for rural communes is approximately zero.

This pattern has several possible interpretations. First, it could reflect reallocation of investment resources within constituencies: the departure of the cumulard deputy may have shifted the political equilibrium away from urban centers (where the deputy-mayor's commune was typically located) toward smaller rural communes. Second, it could reflect differential exposure to other post-2017 shocks: the DSIL grant program, created in 2016 and expanded thereafter, disproportionately targeted medium-sized investment projects in rural areas. Third, the offsetting effects could be coincidental—reflecting independent trends by commune size that happen to cancel at the constituency level. We regard this heterogeneity as suggestive rather than conclusive, noting that the constituency-level analysis (which is the appropriate level for the treatment assignment) shows a clear null.

7.3 The Fiscal Non-Cost of the Cumul

The uniform null across all fiscal categories implies that the *cumul des mandats*—whatever its effects on parliamentary productivity, legislative oversight, or political accountability—did not generate measurable fiscal rents or efficiencies in local public finance. Neither the pork-barrel hypothesis (cumulards steered investment to their communes) nor the rent-extraction hypothesis (cumulards inflated operating budgets) finds support in the data.

This finding has methodological implications for the broader literature on political connections and local spending. A common approach is to examine total spending or total transfers, but our results show that disaggregation by budget category is important for understanding the null: the absence of effects is uniform rather than reflecting offsetting positive and negative effects across categories. The null is a genuine null, not a masking artifact.

We emphasize that the absence of fiscal effects should not be interpreted as evidence that the cumul had no consequences. The dual mandate may have affected outcomes that do not appear in commune budgets: parliamentary voting behavior, legislative productivity, policy coordination between national and local government, or the quality of public services. Our analysis captures only the fiscal dimension of a multifaceted institutional arrangement. The finding that severing the local–national connection carried no measurable fiscal cost is nonetheless informative for policymakers considering similar reforms.

8 Discussion and Conclusion

France's 2017 ban on dual mandates was the most significant reform of political institutions in the Fifth Republic since the introduction of the five-year presidential term in 2000. It severed the institutional connection between 248 national legislators and the local communes they governed as mayors. Proponents argued it would professionalize Parliament and reduce conflicts of interest. Opponents warned it would cut off communes from the national resources that deputy-mayors could mobilize. Our analysis adjudicates between these claims empirically.

The most striking finding is a non-finding. The ban had no detectable effect on local public investment, equipment spending, state grant receipts, operating expenditure, revenue, or debt. The feared collapse of locally-directed spending did not materialize—nor did the hoped-for rationalization of inflated budgets.

This comprehensive null is itself informative. It implies that the "pork-barrel" channel of the cumul was either nonexistent, too small to detect, or quickly substitutable through alternative political connections. Several interpretations are consistent with the evidence.

First, French investment grant allocation may be more bureaucratic than political. The DETR and DSIL are formally allocated by *préfets* based on eligibility criteria (population, fiscal capacity, project quality), with advisory input from a local commission. While political influence on this process has been documented (Enikolopov, 2014), the deputy-mayor channel may have been one among many, and its removal was compensated by other routes of influence.

Second, the extensive turnover in the 2017 elections—driven primarily by the *En Marche!* wave rather than the cumul ban itself—may have disrupted political connections for both treated and control constituencies. If the new deputies in non-cumulard constituencies were equally inexperienced as those replacing cumulards, the control group's fiscal trajectory would also have been affected, biasing the DiD toward zero. However, the event study shows no common break for either group, making this explanation less compelling.

Third, the institutional constraints on commune budgets—the balanced-budget rule for operating sections, the formula-based DGF, and the multi-year nature of investment planning—may have limited the scope for political manipulation in the first place. The cumul may have affected softer outcomes (political attention, coordination quality, responsiveness) without leaving a trace in fiscal aggregates.

This paper contributes to a broader literature on the design of political institutions. Besley and Coate (2004) argue that accountability mechanisms and selection effects interact to determine policy outcomes. The cumul ban was designed to improve accountability by forcing politicians to focus on a single mandate, but our results suggest that its fiscal consequences were negligible across all budget categories.

Several caveats apply. First, and most importantly, our outcome variables are aggregated to the constituency level, which averages across all communes in each constituency. If the cumulard deputy-mayor channeled resources specifically to their own "home" commune (the one where they served as mayor), the constituency-level estimates would dilute this commune-specific effect across many unaffected communes. A direct test of the pork-barrel hypothesis would require identifying each deputy-mayor's specific commune and comparing it to appropriate counterfactual communes—an analysis we leave to future work with more granular treatment data.

Second, our treatment variable captures only the deputy-mayor form of the cumul. Deputies who held other local executive mandates (departmental or regional council presidents, EPCI presidents) are classified as non-cumulards in our primary specification, which may attenuate the treatment effect. Third, the 2017 elections coincided with a massive realignment of French politics around the *En Marche!* movement, which may have disrupted political networks in ways that interact with the cumul ban. Fourth, the post-treatment period is observed at only two points (2020 and 2023) due to OFGL data quality limitations, which reduces statistical power relative to a continuously observed panel.

Several avenues for future research emerge from our findings. First, examining the effect of the cumul ban on legislative productivity—absenteeism, bill sponsorship, committee participation—would complement our fiscal analysis and speak directly to the "divided attention" argument that motivated the reform. Second, the 2020 municipal elections provide an opportunity to study whether the departure of cumulard deputies affected electoral competition and incumbent advantage at the local level. Third, comparing the French experience with countries that maintain or have recently restricted dual mandates (Belgium, Italy) could yield cross-national evidence on the fiscal consequences of political institution design. Fourth, the heterogeneity between urban and rural communes revealed by our triple-difference specification deserves further investigation with commune-level data that can identify the deputy-mayor's own commune specifically.

The broader lesson is that political institutions may shape governance in ways that are less fiscally consequential than the pork-barrel narrative suggests. The French cumul did not produce the dramatic fiscal distortions that its critics feared, nor did its abolition produce the disruption that its defenders predicted. The institutional separation of national and local mandates, in the French case, appears to have been fiscally neutral.

For policymakers considering similar reforms—such as restrictions on dual mandates in Belgium, Italy, or Germany—our findings offer a reassuring message: the fiscal disruption predicted by opponents did not materialize. But they also caution that the fiscal benefits promised by proponents are equally elusive. If the goal is to improve local public investment, the evidence points toward reforming the grant allocation system itself (reducing *préfet* discretion, increasing formula-based allocation) rather than targeting the political connections of individual legislators.

References

Bach, Laurent, "What Happens When Politicians Cannot Hold Multiple Offices? Evidence from French Municipalities," *European Journal of Political Economy*, 2019, *56*, 55–73.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-In-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Besley, Timothy and Stephen Coate, "Can Political Competition Reduce the Quality of Regulation?," *Journal of Public Economics*, 2004, *88* (12), 2677–2697.

Binet, Marie-Estelle and Jean-Sébastien Pentecôte, "Fiscal Transfers and Local Spending: Evidence from France," *Applied Economics*, 2014, *46* (10), 1108–1122.

Breuillé, Marie-Laure, Pascale Duran-Vigneron, and Anne-Laure Samson, "The French Inter-Municipal Revolution," *Regional Science and Urban Economics*, 2018, *72*, 107–121.

Cruz, Cesi, Julien Labonne, and Pablo Querubin, "Politician Family Networks and Electoral Outcomes: Evidence from the Philippines," *American Economic Review*, 2017, *107* (10), 3006–3037.

Díaz, Mónica Solé, "Cumul des mandats: le French Exception Revisited," *Parliamentary Affairs*, 2013, *66* (3), 561–578.

Enikolopov, Ruben, "Politicians, Bureaucrats and Targeted Redistribution," *Journal of Public Economics*, 2014, *120*, 74–83.

Ferraz, Claudio and Frederico Finan, "Motivating Politicians: The Impacts of Monetary Incentives on Quality and Performance," *American Economic Review*, 2011, *101* (1), 431–461.

Fiva, Jon H. and Askill H. Halse, "Power Sharing and Technology Improvement: Evidence from Renewable Energy," *Journal of Public Economics*, 2018, *159*, 89–105.

Fouirnaies, Alexander and Andrew B. Hall, "How Do Electoral Incentives Affect Legislator Behavior?," *American Political Science Review*, 2022, *116* (2), 626–642.

François, Abel and Raul Magni-Berton, "Le cumul des mandats en France: causes et conséquences," *Revue française de science politique*, 2006, *56* (4), 597–612.

Gagliarducci, Stefano and Tommaso Nannicini, "The Cost of Running for Office," *Journal of the European Economic Association*, 2013, *11* (2), 322–356.

Golden, Miriam A. and Lucio Picci, "Pork-Barrel Politics in Postwar Italy, 1953–94," *American Journal of Political Science*, 2008, *52* (2), 268–289.

Hodler, Roland and Paul A. Raschky, "Regional Favoritism," *Quarterly Journal of Economics*, 2014, *129* (2), 995–1033.

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

Appendix

A. Treatment Classification Details

Table A1: Cumulard Classification of XIV Legislature Deputies

	N	Share (%)
Total constituencies	577	100.0
Cumulard (deputy-mayor)	259	44.9
Non-cumulard	318	55.1

Notes: Classification based on Wikidata records of XIV legislature deputies (2012–2017) who simultaneously held the office of *maire* in a French commune. The 2017 organic law (Loi organique n°2014-125) prohibited the cumul of parliamentary and local executive mandates effective from the June 2017 legislative elections.

The cumulard classification is based on Wikidata records of XIV-legislature deputies who simultaneously held a mayoral office (*maire de [commune]*). Of the 577 metropolitan constituencies, Wikidata identifies 259 cumulard and 318 non-cumulard deputies. After merging with commune budget data, 38 constituencies (primarily overseas territories or those with incomplete fiscal records) drop from the analysis, yielding a final sample of 539 constituencies: 248 cumulard and 291 non-cumulard.

Our classification captures the primary form of the cumul (deputy + mayor), which was both the most common and the most controversial. We do not classify deputies who held lesser local mandates (e.g., simple municipal councillor, departmental councillor without executive function) as cumulards, since these mandates were not affected by the 2014 law's prohibition on *executive* functions.

B. Event Study: Equipment and Grants

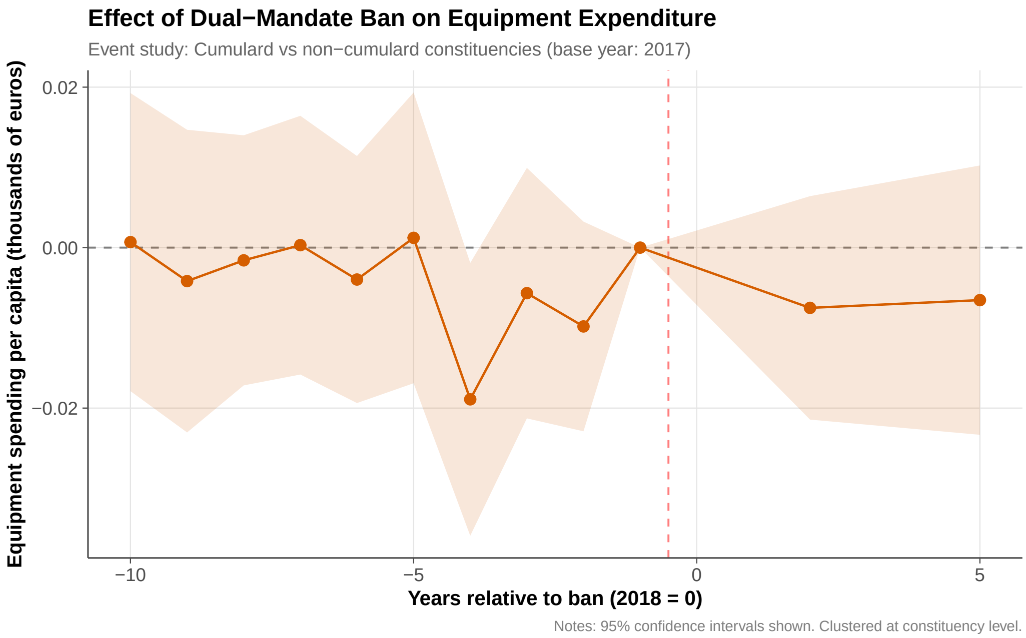

Figure A1: Event Study: Equipment Expenditure Per Capita

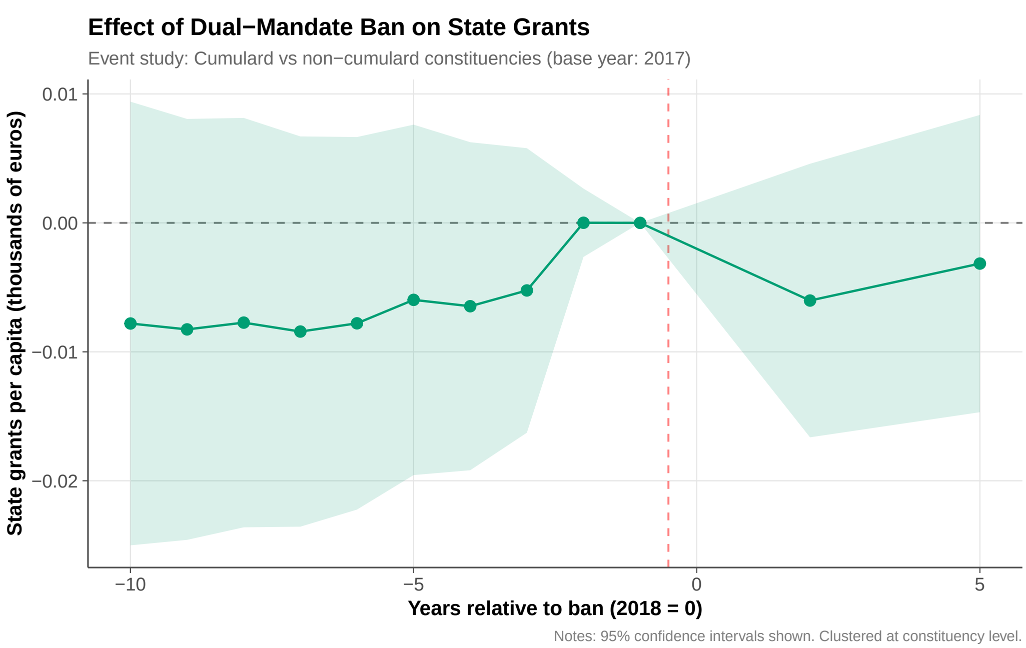

Figure A2: Event Study: State Grants Per Capita

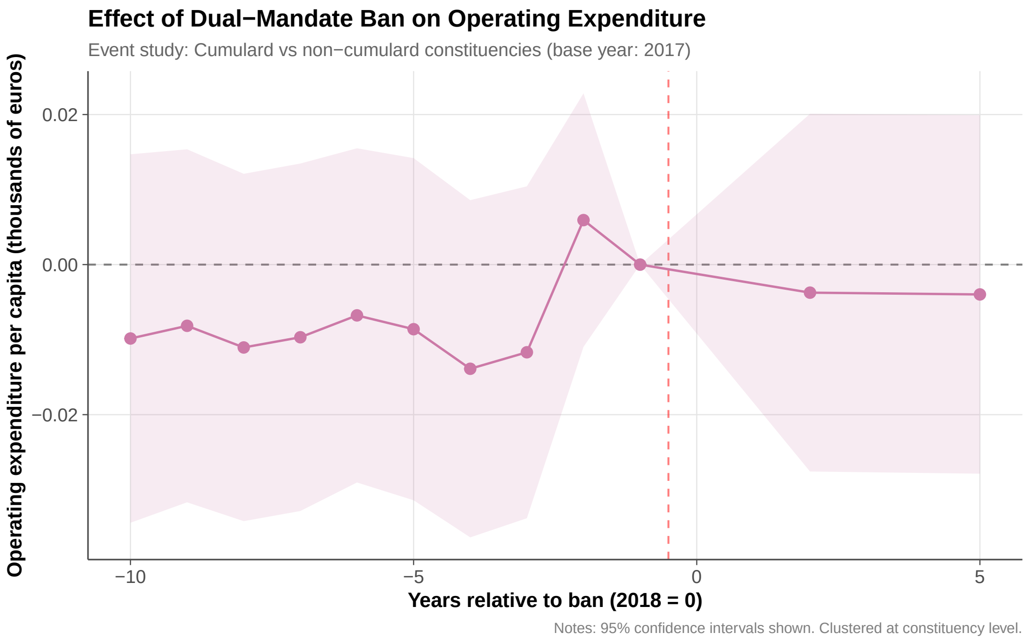

Figure A3: Event Study: Operating Expenditure Per Capita

Table A2: Triple-Difference: Rural Interaction (Commune Level)

Dependent Variables:	Invest. PC	Equip. PC
Model:	(1)	(2)
Variables		
Cumulard × Post	-0.0446***	-0.0137*
	(0.0082)	(0.0078)
Cumulard × Post × Rural	0.0438***	0.0096
	(0.0112)	(0.0091)
Fixed-effects		
Commune	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	425,872	425,872
R^2	0.43520	0.38395
Within R^2	6.69×10^{-5}	7.79×10^{-6}

Clustered (Constituency) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Commune-level regressions with commune and year fixed effects. Rural = commune population below 2,000. "Cumulard × Post" captures the urban effect; "Cumulard × Post × Rural" captures the additional rural differential. Standard errors clustered at the constituency level. "PC" denotes per capita (thousands of euros per inhabitant).

Table A3: HonestDiD Sensitivity Analysis: Investment Per Capita

$\bar{M}$	Lower Bound	Upper Bound
0.00	-0.028	0.013
0.01	-0.078	0.003
0.02	-0.091	0.016
0.03	-0.103	0.026
0.04	-0.113	0.036
0.05	-0.123	0.046

Notes: Sensitivity bounds from Rambachan and Roth (2023). $\bar{M}$ controls the maximum change in slope of the pre-trend violation between consecutive periods. At $\bar{M} = 0$, exact parallel trends are assumed. Positive $\bar{M}$ allows smooth deviations. The dependent variable is investment per capita (thousands of euros per inhabitant). All intervals include zero, confirming robustness of the null finding.

C. Triple-Difference Results

D. HonestDiD Sensitivity Bounds

E. Treatment Distribution by Département

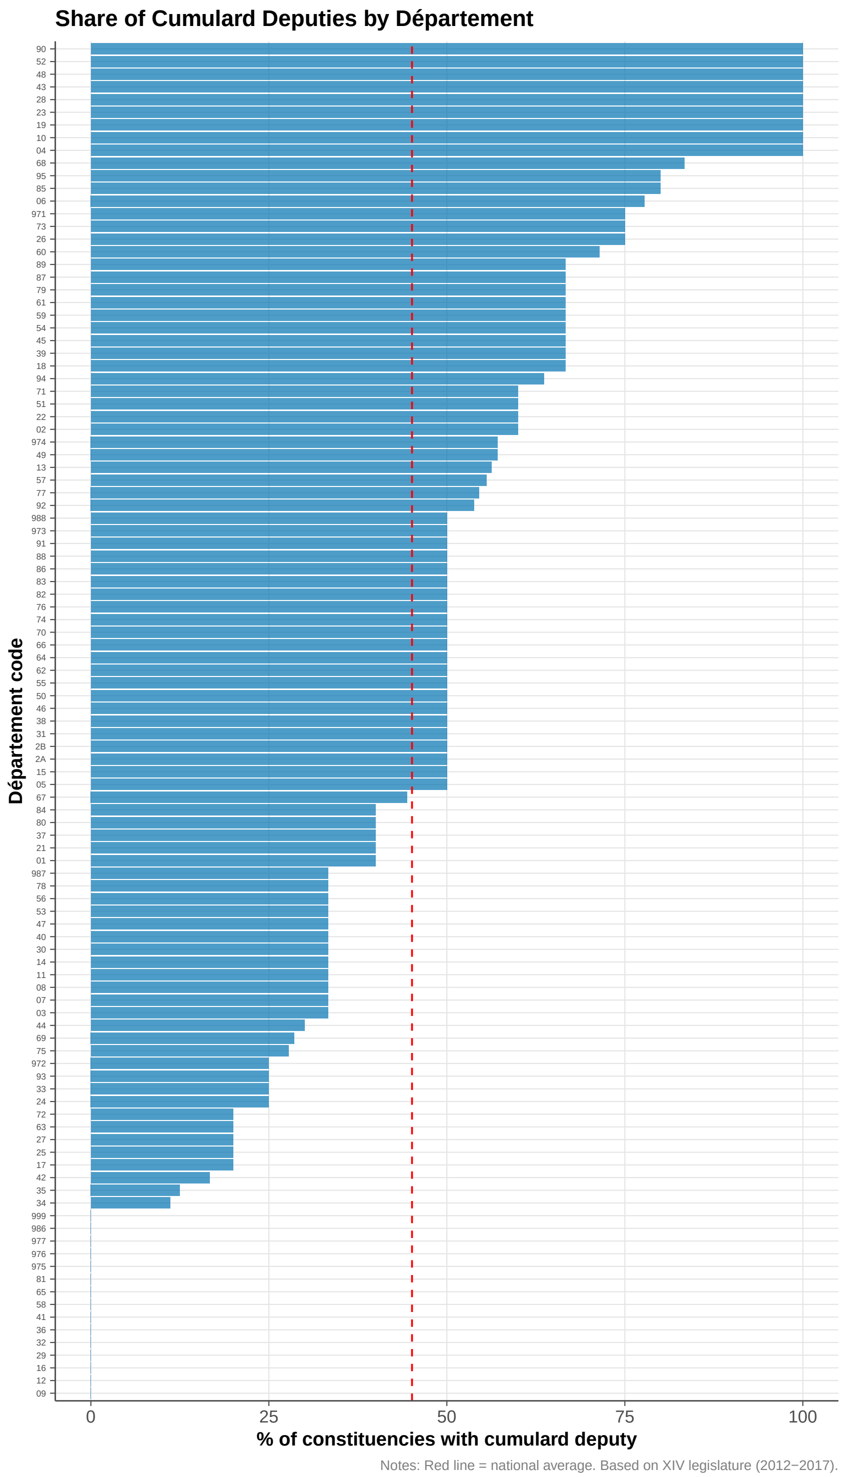

Figure A4: Share of Cumulard Constituencies by Département

Acknowledgements

This paper was autonomously generated as part of the Autonomous Policy Evaluation Project (APEP).

Contributors: @olafdrw

First Contributor: https://github.com/olafdrw

Project Repository: https://github.com/SocialCatalystLab/ape-papers