Public procurement and educational outcomes: Evidence from primary education in France (2015-2024)

Vincent Jardine & Camille Naudy

2025

Abstract

Public investment in education is typically studied through aggregate spending or school finance reforms, with limited attention to the administrative channels through which these investments are implemented. In many education systems, however, local governments play a central role in translating public resources into material learning conditions through public procurement.

This paper examines whether education-related municipal procurement affects student outcomes in France. While curricula and teachers are centrally managed by the national government, municipalities are responsible for primary school infrastructure, equipment, and services, most of which are delivered through procurement contracts. We combine the BeauAMP database of public contracts (2015–2023) with administrative data on standardized test scores at entry into lower secondary school (2017–2024) to construct a municipality-level panel.

Using staggered difference-in-differences estimators robust to heterogeneous treatment timing, we analyze the dynamic effects of procurement activity on student performance. The results provide limited evidence that education-related procurement leads to improvements in average test scores. While some specifications suggest negative effects, these estimates are sensitive to the choice of estimator and are not consistently robust across specifications.

By contrast, we document a clear and robust increase in the number of students taking standardized tests following procurement activity. This suggests that procurement primarily affects school capacity or the attractiveness of local schooling rather than measured learning outcomes. Additional heterogeneity analyses reveal that the negative score effects are concentrated among urban municipalities, while fiscal capacity strongly predicts procurement participation but does not systematically moderate its effects on learning outcomes.

Taken together, these findings highlight the importance of distinguishing between different channels through which public investment affects education systems. They suggest that municipal procurement plays a role in shaping access to schooling and enrollment dynamics, while its short-run impact on measured academic performance remains uncertain.

1 Introduction

Understanding how public investment translates into improvements in educational outcomes is a central question in the economics of education. A large body of research has examined the effects of school spending, finance reforms, and resource allocation on student achievement. However, much less attention has been paid to the administrative mechanisms through which these investments are implemented in practice. In many education systems, a substantial share of education spending is operationalized through local procurement decisions that determine how infrastructure, equipment, and services are delivered to schools.

France provides a particularly relevant setting to study these mechanisms. While curricula and teachers are centrally managed by the Ministry of Education, municipalities are responsible for the material conditions of primary schooling, including school buildings, maintenance, equipment, and a range of complementary services. These responsibilities are largely implemented through public procurement, which constitutes the main operational channel through which local governments translate public resources into concrete educational investments.

Public procurement represents a significant share of public spending and plays a key role in shaping local public service provision. In the education sector, procurement decisions determine the quality of school facilities, the availability of digital and pedagogical equipment, and the provision of ancillary services such as catering or transport. Despite its importance, the role of procurement as an implementation channel for education investment has received limited empirical attention. Most of the education economics literature focuses on funding levels, teacher quality, or class size, with comparatively little emphasis on how public investments are operationalized and how procurement activity may shape the learning environment.

This issue is particularly salient in the French institutional context. Although educational standards are nationally defined, municipalities differ widely in fiscal capacity, administrative resources, and investment strategies. As a result, similar formal responsibilities may translate into heterogeneous material conditions across territories. At the same time, spatial disparities in student performance persist despite a highly centralized education system. Understanding whether variation in municipal procurement activity contributes to these differences is therefore of both academic and policy relevance.

This paper investigates whether education-related municipal procurement affects student outcomes at the end of primary school. We focus on standardized test scores measured at entry into lower secondary education, which capture cumulative learning during primary schooling. We assemble a new municipality-level panel combining the BeauAMP database of public con-

tracts (2015–2023) with administrative test score data from the French Ministry of Education (2017–2024).

Our empirical strategy exploits the staggered timing of procurement activity across municipalities. Because procurement occurs at different points in time and its effects are likely to materialize gradually, we implement staggered difference-in-differences estimators robust to heterogeneous treatment timing following Callaway and Sant'Anna (2021). We complement this approach with alternative estimators to assess the robustness of the results.

The findings provide limited evidence that education-related procurement improves average test scores. While some specifications suggest negative effects, these estimates are sensitive to the choice of estimator and are not consistently robust. In contrast, we document a clear and robust increase in the number of students taking standardized tests following procurement activity. This pattern suggests that procurement may expand school capacity or increase the attractiveness of local schooling options.

We further investigate whether these enrollment dynamics can explain the relationship between procurement and student performance. The results provide little support for strong compositional effects, indicating that changes in student composition are unlikely to fully account for the observed patterns in test scores. We document heterogeneous effects across municipality types: the negative score effects are concentrated in urban areas, while rural and semi-urban municipalities display estimates close to zero. Analyses by fiscal capacity reveal that procurement participation is strongly increasing in municipal wealth, while the effects of procurement conditional on participation do not vary systematically with fiscal resources. Overall, the evidence suggests that education-related procurement primarily affects access to schooling and enrollment dynamics rather than generating immediate improvements in measured academic performance.

By providing nationwide evidence on education-related procurement and student outcomes in France, this paper contributes to several strands of literature. First, it complements research on education finance by focusing on the administrative channels through which investments are implemented. Second, it contributes to the literature on public procurement by examining how procurement activity affects social outcomes in the education sector. Third, the heterogeneity analyses contribute to the literature on fiscal decentralization by documenting that fiscal capacity is a primary determinant of procurement participation in education, rather than a moderator of its effectiveness. More broadly, it highlights the importance of distinguishing between access, capacity, and learning outcomes when evaluating the impact of public investment in education.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature. Section 3 presents the data. Section 4 describes the empirical strategy. Section 5 reports the

results and robustness checks. Section 8 concludes.

2 Literature review

This paper relates to two complementary strands of literature. The first examines whether targeted investments in school environments improve educational outcomes. The second analyzes public procurement as a policy instrument shaping welfare outcomes through the quality, efficiency, and equity of public service provision. Together, these literatures provide the conceptual foundations for studying education-oriented public procurement as a mechanism linking local public investment to student achievement.

2.1 Targeted School-Level Investments and Educational Outcomes

A large body of research in the economics of education has investigated the relationship between public investment and student performance. Early indications, such as (Card and Krueger, 1992), documented positive associations between school quality—measured through class size, instructional resources, or facilities—and long-term educational and labor market outcomes. These findings helped establish the view that material learning conditions matter and that targeted investments can contribute to reducing educational inequalities.

This interpretation was challenged by influential reviews, most notably (Hanushek, 1997), who argued that the empirical relationship between education spending and student achievement was often weak or inconsistent. According to this perspective, increasing resources alone may generate limited returns when institutional incentives are misaligned or when funds are poorly allocated. This debate shifted attention away from aggregate spending levels toward the institutional and managerial conditions under which resources are effectively transformed into learning gains.

More recent quasi-experimental evidence has substantially refined this discussion. Exploiting variation from school finance reforms in the United States, (Jackson et al., 2016) show that sustained and targeted increases in school spending lead to significant improvements in test scores, educational attainment, and adult earnings, particularly for students from disadvantaged backgrounds. Similarly, (Lafortune et al., 2018) document that contemporary finance reforms reallocating resources toward low-income districts reduced achievement gaps. A comprehensive synthesis by (Jackson and Mackevicius, 2021) confirms that increases in public school spending generally improve academic performance, while highlighting substantial heterogeneity linked to spending composition and local administrative capacity.

A growing literature focuses specifically on investments in school infrastructure and learning en-

vironments. Using quasi-experimental variation from large-scale facility upgrades, (Lafortune and Schönholzer, 2022) show that improvements in school buildings enhance student achievement, reduce absenteeism, and generate positive spillovers for surrounding neighborhoods. Complementary evidence reviewed by (Escueta et al., 2020) suggests that investments in digital tools and pedagogical equipment can yield meaningful learning gains when embedded in coherent organizational strategies, whereas isolated or poorly coordinated investments often produce limited effects.

Recent contributions further emphasize the importance of implementation and local capacity. For instance, (Jackson, 2024) highlights that the effectiveness of education spending critically depends on administrative quality and institutional alignment, reinforcing the idea that how resources are deployed matters as much as how much is spent.

2.2 Public Procurement as a Welfare-Enhancing Policy Instrument

A second strand of literature examines public procurement as a central component of public sector performance and welfare outcomes. Traditionally viewed as a technical and regulatory function aimed at ensuring compliance, transparency, and cost-efficiency, procurement is increasingly recognized as a strategic policy instrument influencing the quality and distribution of public services.

Recent empirical work emphasizes that procurement outcomes depend not only on spending levels but also on institutional design and administrative capacity. In a landmark cross-country study, (Bosio et al., 2022) show that procurement systems characterized by greater transparency, professionalization, and managerial discretion achieve better outcomes in terms of cost, quality, and timeliness.

International organizations have reinforced these insights. The (World Bank, 2020, 2023) identify procurement as a critical lever for government effectiveness and inclusive growth. Similarly, recent OECD reports (OECD, 2023, 2024) emphasize the role of procurement capacity and data systems in improving public service delivery and reducing territorial inequalities.

A growing applied literature also examines procurement as a channel shaping welfare outcomes beyond efficiency. (Fazekas and Blum, 2021) review the evidence on procurement reforms and find that only a limited set of tools consistently improve outcomes. More recent work by (Fazekas et al., 2023) further stresses that procurement performance reflects broader state capacity and administrative learning, rather than formal rules alone.

2.3 Contribution of the present study

This paper contributes to the literature by examining public procurement as a key implementation channel through which education investments are translated into local learning environments. While existing research has extensively documented the effects of school spending and resource allocation on student outcomes, much less attention has been paid to the administrative processes that govern how these resources are deployed in practice.

First, the paper provides new empirical evidence on the relationship between education-related municipal procurement and student performance using comprehensive nationwide data from France. By combining detailed procurement information with standardized test scores, it offers a novel perspective on how local investment decisions are operationalized and how they relate to educational outcomes.

Second, the paper contributes to the literature by documenting that education-related procurement primarily affects school capacity rather than measured academic performance. We find limited evidence that procurement leads to improvements in average test scores, but strong and consistent evidence that it increases the number of students enrolled. This result highlights an important distinction between the effects of public investment on access to education and its effects on learning outcomes.

Third, the paper investigates the mechanisms underlying this relationship. While procurement activity is associated with higher enrollment, we show that changes in student composition do not appear to explain the observed patterns in test scores. This suggests that the link between procurement and educational outcomes operates through channels that are not directly captured by standardized performance measures, or that potential benefits may materialize only in the longer run.

More broadly, these findings contribute to the literature on education finance, decentralization, and public procurement by emphasizing the importance of implementation processes. They suggest that evaluating education policies requires not only examining how much is spent, but also understanding how investments are delivered and through which channels they affect different dimensions of educational systems.

3 Data

To analyze these relationships between public investments and territorial results regarding indicators of sustainable development, two data sources have emerged as particularly valuable.

First, the BeauAMP database (*Base Étendue, Améliorée et Unifiée des Annonces des Marchés*

Publics) offers an especially valuable resource. Covering the period from 2015 to 2023, this dataset was specifically designed to capture the diversity and complexity of public procurement practices in France. Drawing on official sources such as the *Bulletin officiel des annonces des marchés publics* (BOAMP) and the *Journal officiel de l'Union européenne* (JOUE), BeauAMP provides detailed information on awarded contracts, including contract value, procurement procedure, duration, and the nature of both buyers and suppliers.

A key strength of BeauAMP lies in its capacity to illuminate the spatial and institutional dynamics of public procurement. By geolocating public buyers—and, when available, contract awardees—the database allows researchers to track how public spending is distributed across territories and evolves over time. Contracts are categorized using CPV codes, which facilitates sectoral analyses in areas such as education, health, environmental services, and social care. BeauAMP also differentiates between levels of government (central, regional, municipal), making it possible to examine how various layers of the public sector use procurement as an instrument of policy implementation. Through its longitudinal coverage, institutional granularity, and geographic precision, BeauAMP allows for a more grounded and contextualized understanding of how public procurement shapes local development trajectories. Although still underutilized in academic research, this dataset opens new avenues for evaluating the strategic role of procurement in advancing sustainability goals across territories.

The second source of data relates to educational performance, drawn from the national standardized assessments administered by the Direction de l'évaluation, de la prospective et de la performance (DEPP) within the French Ministry of Education. Since 2017, these assessments have measured students' foundational competencies upon entry into sixième—the first year of lower secondary education—providing a consistent and comparable indicator of learning outcomes across municipalities. They constitute one of the few longitudinal datasets enabling the study of educational performance at a fine territorial scale.

The assessments cover two main domains, French language and mathematics, and report average test scores at the school level, disaggregated by gender, sector (public or private), and within-school performance groups. Based on these results, a municipality-level composite index of educational performance is constructed by aggregating average school scores within each commune and weighting them by the number of students enrolled. This aggregation yields a coherent measure of educational attainment at the end of primary education, capturing learning outcomes at a critical transition point toward secondary schooling.

A key strength of the DEPP assessments lies in their standardized and exhaustive coverage. The use of identical testing frameworks and scoring methods ensures comparability over time and

space, while their nationwide scope allows for a detailed exploration of spatial disparities and temporal trends in student performance. As such, these data provide a reliable proxy for assessing the territorial distribution of educational quality and equity.

When combined with the BeauAMP public procurement database (2015–2023), the DEPP data make it possible to construct an annual municipality-level panel (2017–2024) linking investment behavior to educational outcomes. Each observation includes:

- the number and total value of education-related procurement contracts (infrastructure, digital equipment, pedagogical services);
- the average test scores
- a set of socio-economic and fiscal controls capturing municipal characteristics.

This integrated framework enables the analysis of whether greater education-oriented procurement intensity is associated with improvements in student achievement at the entrance to lower secondary school, and whether these effects vary according to local institutional quality, governance capacity, and socio-economic conditions.

3.1 Defining education-related procurement

A central methodological step in this study consists in identifying which municipal procurement contracts can be considered education-related. Since primary education infrastructure and equipment fall under municipal responsibility in France, we restrict our measure of educational investment to procurement awarded directly by municipalities. This ensures both conceptual consistency and complete coverage in the BeauAMP database, where municipal buyers are systematically identified.

To isolate education-oriented procurement within the broader set of municipal contracts, we rely on the *object* variable, which contains a short textual description of the purchased good or service. We implement a keyword-based screening procedure using the term "school" (*école*) and its variants, which allows us to identify contracts explicitly intended for primary school buildings, equipment, and services. This procedure yields a total of 13,453 school-related lots awarded by 1,110 municipalities between 2015 and 2023.

To assess the reliability of this classification, we manually reviewed a subsample of contract descriptions. This verification indicates that the keyword-based approach produces no apparent false positives: all inspected contracts clearly concerned goods, works, or services for primary schools. However, the method may generate false negatives, since relevant contracts might not

systematically mention the term "school" in their textual description. Such omissions would lead to under-identification of education-related procurement and thus introduce a potential selection bias, whereby our estimates capture the impact of a subset—rather than the full universe—of municipal education investments.

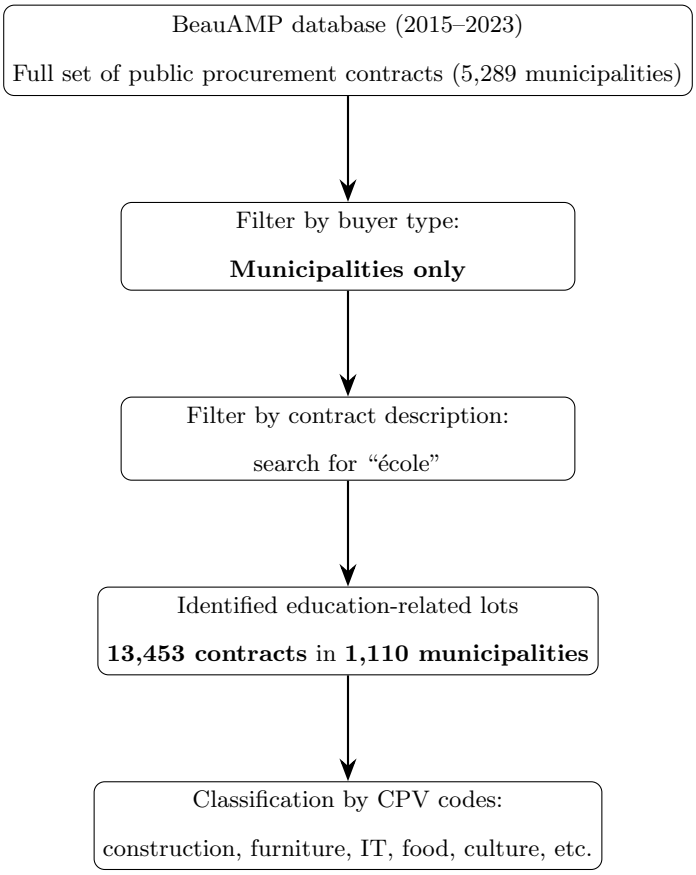

Figure 1: Identification process of education-related municipal procurement

The resulting corpus of education-related procurement is overwhelmingly dominated by construction and architectural services, which account for around 40% of all lots. These contracts primarily concern building maintenance, renovation, extension, or compliance upgrades for primary schools. Other categories—such as furniture, information technology, and cultural or food-related services—represent a much smaller share of municipal investment in schools (see Table 1). This pattern indicates that education-related procurement in municipalities is driven first and foremost by physical infrastructure needs, with complementary investments in equipment and services playing a secondary role.

Table 1: Top ten CPV sectors in education-related procurement (2015–2023)

CPV Category	Number of lots	Share (%)
Construction (45)	5,227	38.8
Furniture (39)	879	6.5
IT equipment (30)	506	3.8
Architecture (71)	490	3.6
Leisure and culture (92)	430	3.2
Food (15)	403	3.0
Books and publications (22)	272	2.0
Sewage and waste services (90)	255	1.9
Hotels and restaurants (55)	248	1.8
Transport services (60)	227	1.7

3.2 Aggregating the data at the city x year level

To perform the empirical analysis, we aggregate the data into a panel at the municipality $\times$ year level. We create two main types of variables to account for the amount of procurement performed:

- number of lots awarded
- total procurement spendings (based on the prices)

We create a general variable and a set of industry specific variables for both of these metrics. The main reasons for including a "number" of lots variable is to complement prices which are imperfectly reported. The primary issue we have is that a large share (30 percent) of lots have no associated price. A second issue is potential misreporting of prices. We attempt to mitigate this problem in two ways:

- Removing contracts with extremely high prices by trimming the top 1% off the price distribution
- Taking into account potential confusion between total contract price and lot price. In some multi-lot contracts, all lots have exactly the same price. We suspect that in this case, the total price was assigned to every lot. Thus, when computing total procurement spendings, we consider that the amount was spent only once for these types of contracts (while in other contracts we use the sum of prices accross lots).

Figure 2: Secondary schools according to treatment status

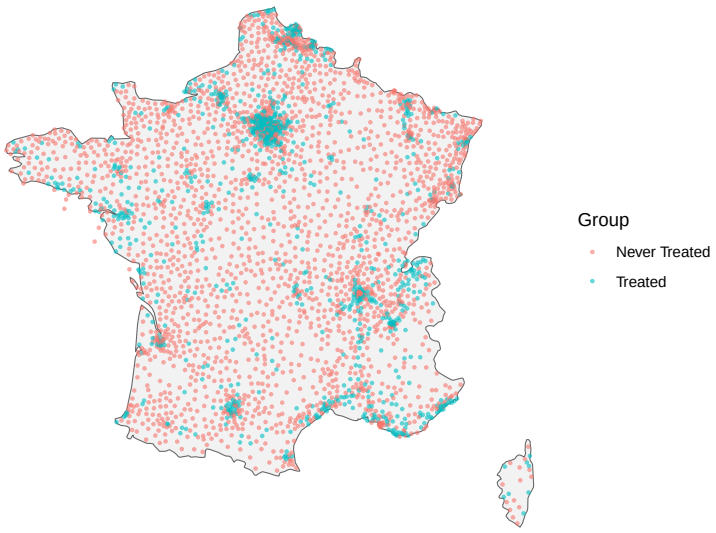

After these data precautions, the data on procurement spendings appears reliable[1], with average yearly spendings of 973,000 euros, and median yearly spendings of 380,000 euros per municipality. Moreover, when ranking cities by total spending, we find that the top spots are occupied by large French cities such as Toulouse, Marseille and Montpellier.

The unit of observation in the empirical analysis is the municipality-year. Procurement contracts are aggregated at this level and combined with annual data on student outcomes. Consequently, the number of observations in the regression sample reflects the size of the municipality-year panel rather than the number of contracts observed in the raw procurement data.

To investigate heterogeneous treatment effects, we supplement the main panel with two additional data sources. First, municipal population density is measured using the seven-level communal density grid (grille communale de densité) produced by INSEE, which classifies each municipality along a continuum from densely populated urban centers to isolated rural communes. We aggregate this classification into three broad categories: urban, semi-urban, and rural. Second, municipal fiscal capacity is measured using the fiscal potential per capita (potentiel fiscal par habitant) drawn from the 2017 criteria for the allocation of the Dotation Globale de Fonctionnement (DGF), published annually by the Direction Générale des Collectivités Locales (DGCL). This indicator captures the structural tax base of each municipality independently of local tax rate choices, and constitutes the standard measure of fiscal capacity in the French decentralization literature.

[1]The estimation sample excludes 63 municipalities (1.7%) for which the 2017 baseline enrollment is missing. These municipalities tend to be slightly smaller (6.1 primary schools on average versus 7.5 in the main sample) and are less likely to engage in procurement (24% versus 34%). Results are robust to alternative sample definitions.

4 Empirical Strategy

4.1 Identification strategy

We estimate the effect of education-related municipal procurement on student performance using a staggered difference-in-differences framework. This approach exploits variation in the timing of procurement activity across municipalities.

Several features of our empirical context require particular attention.

Lagged treatment effects. Procurement investments occur at the primary school level, while the outcome variable—standardized test scores—is measured when students enter lower secondary school. As a result, procurement implemented in year t may only affect cohorts entering secondary school several years later. Infrastructure improvements or equipment investments may also require time to translate into observable changes in learning outcomes. Consequently, we expect procurement effects to emerge gradually over time rather than immediately.

Unobserved past procurement. Procurement activity is observed starting in 2015, while the first test scores in our dataset correspond to students entering secondary school in 2017. Investments performed before 2015 may therefore affect early cohorts in our sample but remain unobserved. This issue is particularly relevant when interpreting the effects for municipalities treated early in the observation window. A related concern is that municipalities engaged in procurement before the start of the education score panel (2017) cannot be assigned a clean pre-treatment period. In practice, no municipality in the final sample has a recorded treatment date prior to 2017, as the procurement and education panels overlap from that year onward. The 470 units flagged as "already treated in the first period" by the estimation procedure correspond to municipalities dropped during panel balancing due to missing observations across periods, rather than to units with unobserved prior treatment. This does not affect the interpretation of the estimates.

Staggered adoption of treatment. Municipalities engage in procurement at different points in time. To account for this staggered rollout of treatment, we rely on the estimator proposed by Callaway and Sant'Anna (2021), which allows for heterogeneous treatment timing and dynamic treatment effects. In this framework, municipalities that have not yet received procurement constitute the control group for municipalities treated earlier.

Treatment intensity. Procurement investments vary substantially in magnitude across municipalities. To capture this variation, we classify treated municipalities into three groups according to terciles of total procurement spending per student: small, medium, and large spenders. This allows us to explore heterogeneity in treatment effects across different levels of procurement inten-

sity.

In addition to the baseline estimates, we examine heterogeneous treatment effects along two dimensions. First, we split the sample by municipality type (urban, semi-urban, rural) based on the INSEE density grid, to assess whether procurement effects differ across local contexts. Second, we stratify municipalities into terciles of fiscal potential per capita to explore whether fiscal capacity moderates the effects of procurement on student outcomes. For these subgroup analyses, we replace the academic region fixed effect with a département fixed effect to preserve sufficient within-group variation.

4.2 Baseline specification

Our baseline specification can be written as follows:

$$Y_{i,t} = \alpha + \tau D_{i,t} + \boldsymbol{\beta X_{i,t}} + \delta_i + \gamma_t + \varepsilon_{i,t}$$

where $Y_{i,t}$ denotes the average test score in municipality i at time t, and $D_{i,t}$ is a treatment indicator equal to one once the municipality is exposed to education-related procurement. The vector $X_{i,t}$ includes control variables capturing municipal characteristics, including the number of primary schools, the average distance between primary and secondary schools, the academic region (*académie*), the number of students taking the test in 2017, and the average test score in 2017.

The model is estimated using the Callaway and Sant'Anna (2021) difference-in-differences estimator with not-yet-treated municipalities as the control group. Standard errors are obtained using bootstrap procedures clustered at the municipality level.

In addition to average treatment effects, we estimate dynamic treatment effects through event-study specifications that trace the evolution of outcomes relative to the year of first procurement activity.

5 Results

5.1 Baseline effects of education-related procurement on test scores

We begin by examining the baseline relationship between education-related procurement and student performance. Table 2 reports the average treatment effects estimated using the staggered difference-in-differences framework described in Section 4.

Table 2: Effect of education-related procurement on average test scores

	All spenders	Small Spenders	Mid Spenders	Large Spenders
ATT	−0.773*	−0.792*	−1.265*	-0.889
	(0.32)	(0.384)	(0.417)	(0.49)
N	2883	2531	2506	2450
P (pre-trends)	0.55	0.52	0.84	0.13

*Note: * indicates that 0 is excluded from the 95% confidence interval.*

Across most specifications, procurement activity is associated with modest negative effects on average test scores. Estimated effects range between approximately −0.8 and −1.3 points depending on procurement intensity. These effects are statistically significant for the overall sample and for municipalities in the lower and middle terciles of procurement spending.

Figure 3 presents the corresponding dynamic treatment effects. The event-study plots show no evidence of differential pre-treatment trends, supporting the validity of the identifying assumptions. After procurement activity, coefficients gradually become negative, although the magnitude remains modest.

Figure 3: Event study plots: Effect of school related procurement spendings on test scores, by length of exposure

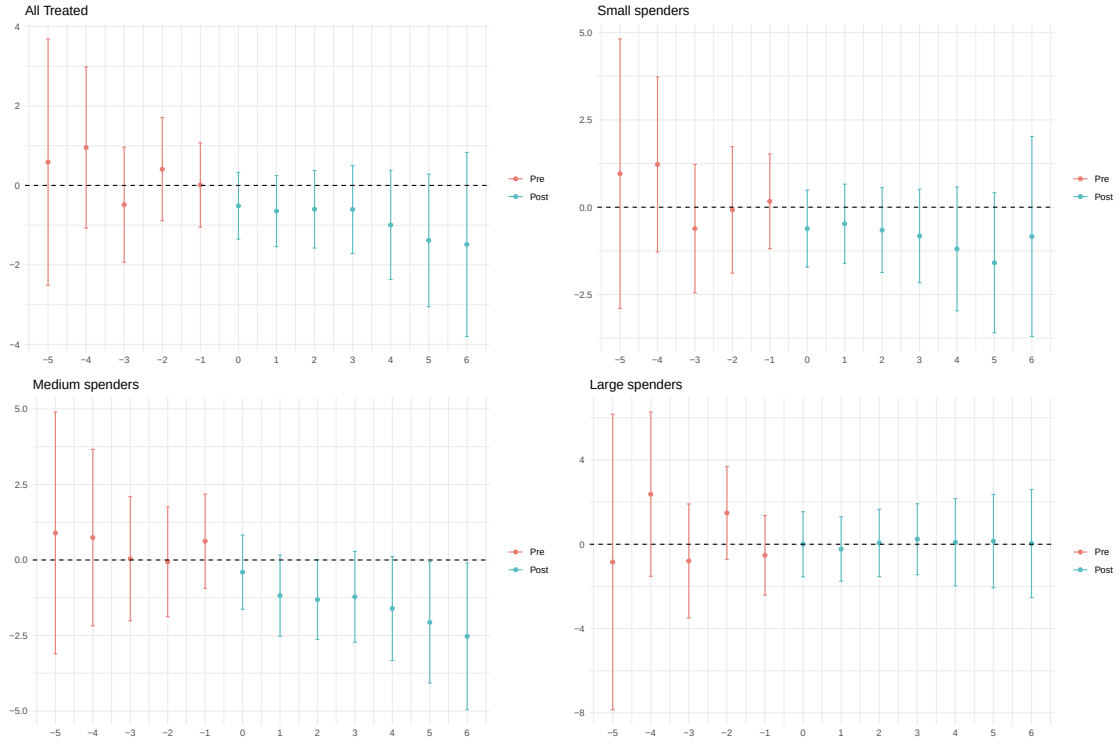

5.2 Procurement and student enrollment

To better understand the mechanisms underlying these results, we examine whether procurement activity affects the number of students entering secondary school. Investments in school infrastructure may increase school capacity or make local schools more attractive, potentially leading to higher enrollment.

Table 3 reports the estimated effects of procurement activity on the logarithm of the number of students taking the standardized test.

Table 3: Effect of education-related procurement on (log) number of students

	All spenders	Small Spenders	Mid Spenders	Large Spenders
ATT	0.024*	0.028*	0.019*	0.025*
	(0.007)	(0.008)	(0.008)	(0.01)
N	2887	2535	2510	2454
P (pre-trends)	0.01	0.44	0.29	0.06

*Note: * indicates that 0 is excluded from the 95% confidence interval.*

Table 4: Effect of education-related procurement on number of students

	All spenders	Small Spenders	Mid Spenders	Large Spenders
ATT	16.492*	12.362*	10.165	14.495*
	(7.854)	(4.941)	(6.551)	(6.325)
N	2887	2535	2510	2454
P (pre-trends)	0.23	0.59	0.64	0.36

*Note: * indicates that 0 is excluded from the 95% confidence interval.*

Across most specifications, procurement activity is associated with a statistically significant increase in the number of students taking the test. The estimates suggest an increase of approximately two to three percent on average.

Expressed in levels, the results indicate an increase of roughly 10 to 16 additional students per municipality following procurement activity.

5.3 Interpreting the mechanism: capacity and composition effects

Section 5.2 shows that education-related procurement is associated with a statistically significant increase in the number of students taking standardized tests. This finding suggests that procurement activity may expand school capacity or increase the attractiveness of local schools. In this section, we investigate whether this increase in enrollment can account for the effects observed on average test scores.

5.3.1 Changes in the distribution of student performance

If procurement activity leads to an inflow of additional students, the composition of the student population may change. Such compositional effects could affect average test scores without reflecting changes in underlying learning conditions. To explore this channel, we examine whether procurement affects the dispersion of test scores.

Figure 4 reports the dynamic effects of procurement on the standard deviation of test scores. The estimates suggest a modest decline in score dispersion in the years following treatment, particularly in the medium run. However, confidence intervals remain relatively wide, and the estimates are not statistically precise across all periods. While this pattern is consistent with a change in the composition of the student population, the evidence remains suggestive and does not provide strong support for large compositional shifts.

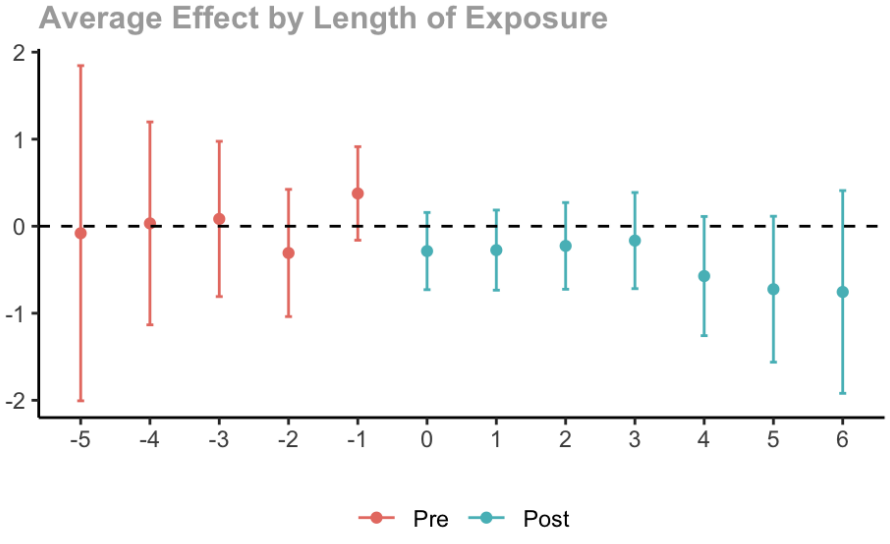

Figure 4: Dynamic effects of procurement on score dispersion

5.3.2 Does enrollment mediate the effect on test scores?

We next assess whether the relationship between procurement and average test scores is driven by changes in student enrollment. To do so, we re-estimate the baseline specification including contemporaneous enrollment as an additional control variable.

Table 5 compares the baseline estimates with those obtained when controlling for the (log) number of students. The estimated coefficients remain very similar across specifications, both in magnitude and statistical significance. In particular, the inclusion of enrollment does not substantially attenuate the estimated effect of procurement on test scores.

Table 5: Effect of procurement on test scores: role of enrollment

	Baseline	+ Enrollment
	(1)	(2)
Average procurement effect (ATT)	-0.036	-0.057
	(0.265)	(0.266)
Log number of students		1.84**
		(0.771)
Municipality fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
Observations	28,291	28,291

Notes: Column (1) reports the baseline specification. Column (2) adds contemporaneous enrollment as a control. Standard errors clustered at the municipality level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

These results indicate that variations in student enrollment do not fully account for the relationship between procurement activity and average test scores. While procurement clearly affects the number of students, this channel does not appear to be the primary driver of the observed effects on performance.

6 Heterogeneity

6.1 Heterogeneity by municipality type

The baseline estimates reported in Section 5.1 pool all treated municipalities together, potentially masking important heterogeneity in the effects of procurement across different types of local con-

texts. In particular, the mechanisms through which procurement may affect student outcomes — school capacity expansion, disruption from construction works, or changes in the attractiveness of local schooling — are likely to operate differently in urban, semi-urban, and rural settings. Urban municipalities tend to have larger and more complex school systems, face greater competition from private schools, and may be more exposed to short-run disruption effects from infrastructure works. Rural municipalities, by contrast, operate smaller schools with less capacity to absorb additional enrollment, and procurement investments may primarily reflect maintenance needs rather than capacity expansions.

To explore this heterogeneity, we classify municipalities according to the seven-level population density grid (grille communale de densité) produced by INSEE, aggregated into three broad categories: urban (density levels 1–2), semi-urban (levels 3–4), and rural (levels 5–7). We then re-estimate the baseline staggered difference-in-differences specification separately for each group, replacing the academic region fixed effect with a département fixed effect to preserve sufficient within-group variation.

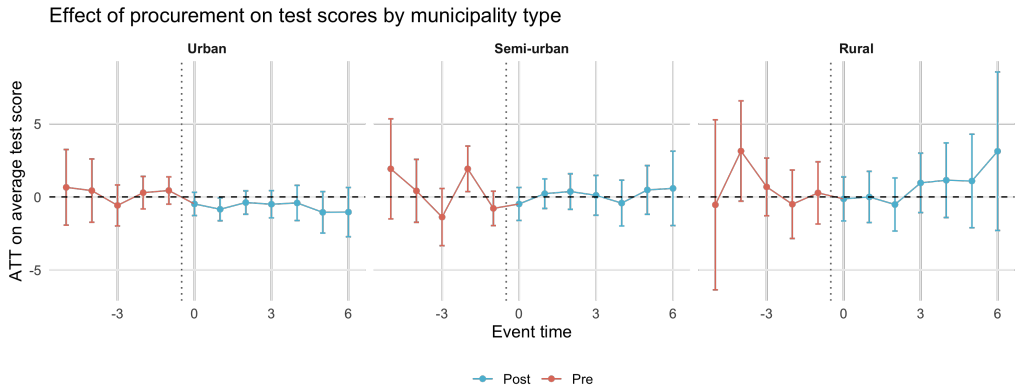

Figure 5: Dynamic effects of procurement on average test scores by municipality type

Notes: The figure reports event-study estimates from the Callaway and Sant'Anna (2021) estimator estimated separately for urban, semi-urban, and rural municipalities. The dotted vertical line indicates the year of first procurement. Pre-treatment periods are shown in red, post-treatment periods in blue. Confidence intervals are at the 95% level.

Figure 5 reports the dynamic treatment effects by municipality type. The results reveal a pattern consistent with the overall baseline estimates being driven primarily by urban municipalities. In urban areas, post-treatment coefficients are systematically negative and relatively stable over the event window, with point estimates ranging from approximately -0.4 to -1.5 points. By contrast, estimated effects in semi-urban and rural municipalities are centered around zero throughout the post-treatment period, with wide confidence intervals reflecting the smaller number of treated

units in these subsamples.

Table 6 summarizes the average treatment effects across municipality types. The estimated ATT for urban municipalities is -0.62 (SE = 0.43), compared to 0.06 (SE = 0.53) for semi-urban and 0.36 (SE = 0.82) for rural municipalities. None of these estimates is statistically significant at conventional levels, which partly reflects the reduction in statistical power resulting from splitting the sample. Nevertheless, the sign and magnitude of the urban estimate is consistent with the baseline result, while the rural and semi-urban estimates are close to zero and imprecisely estimated.

Table 6: Effect of procurement on test scores by municipality type

	Urban	Semi.urban	Rural
ATT	-0.625	0.062	0.365
Std. Error	(0.432)	(0.527)	(0.818)
95% CI	[-1.472, 0.222]	[-0.971, 1.094]	[-1.239, 1.969]
N communes	1117	818	1664

These patterns are consistent with two non-exclusive interpretations. First, procurement-related disruption — particularly from construction works — may be more consequential in urban schools, where larger student populations and denser school networks amplify short-run negative effects on the learning environment. Second, the null results in rural areas may reflect the fact that procurement in these municipalities is less likely to involve large-scale infrastructure projects and more likely to concern routine maintenance, with limited impact on measured academic performance. Disentangling these channels would require more granular data on the nature and scale of procurement works, which we leave for future research.

6.2 Fiscal capacity

Municipal fiscal capacity may shape the effects of procurement on student outcomes through several channels. Wealthier municipalities may procure higher-quality goods and services, implement investments more efficiently, or concentrate procurement in types of projects — such as large-scale infrastructure works — that are less directly linked to short-run learning improvements. At the same time, fiscally constrained municipalities may undertake procurement precisely in response to acute infrastructure deficits, potentially generating larger marginal returns.

To explore this heterogeneity, we classify municipalities into three groups based on terciles of their fiscal potential per capita (potentiel fiscal par habitant) drawn from the 2017 DGCL criteria for the allocation of the Dotation Globale de Fonctionnement. This measure captures the structural

tax base of each municipality independently of local tax rate choices, and constitutes the standard indicator of fiscal capacity in the French decentralization literature.

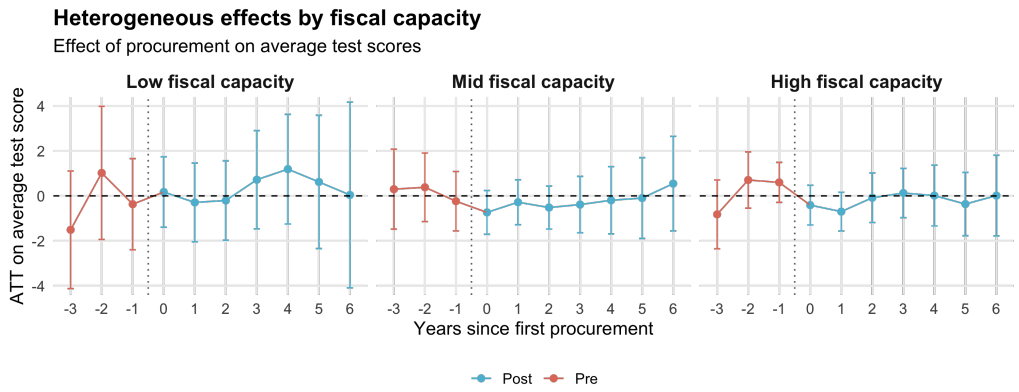

Figure 6: Dynamic effects of procurement on average test scores by fiscal capacity

Notes: The figure reports event-study estimates from the Callaway and Sant'Anna (2021) estimator estimated separately for low, medium and high fiscal capacity municipalities. The dotted vertical line indicates the year of first procurement. Pre-treatment periods are shown in red, post-treatment periods in blue. Confidence intervals are at the 95% level.

Table 7 and Figure 6 report the estimated average treatment effects and dynamic event-study plots by fiscal capacity group. None of the three estimates is statistically significant at conventional levels, reflecting the reduction in statistical power associated with sample splitting. The point estimates suggest a modest positive effect for low-capacity municipalities (ATT = 0.26, SE = 0.78) and small negative effects for mid- and high-capacity municipalities (ATT = -0.34 and -0.24 respectively), but these differences are not precisely estimated and should be interpreted with caution.

Table 7: Effect of procurement on test scores: heterogeneity by fiscal capacity

	Low	Mid	High
ATT	0.262	-0.345	-0.239
Std. Error	(0.782)	(0.463)	(0.451)
95% CI	[-1.271, 1.796]	[-1.252, 0.563]	[-1.124, 0.646]
N communes	1117	1133	1097
N traitées	155	392	626
% early adopters	29.7%	32.1%	44.4%

Two features of the data warrant attention. First, fiscal capacity is strongly associated with

procurement activity: 57% of high-capacity municipalities engage in procurement compared to only 14% of low-capacity municipalities. This pattern suggests that fiscal resources are a primary determinant of whether municipalities undertake education-related procurement at all, rather than a moderator of its effects conditional on procurement. Second, high-capacity municipalities exhibit a higher share of early adopters (44% treated by 2017 versus 30% for low-capacity municipalities), which reduces the pool of valid comparison units and contributes to the imprecision of the high-capacity estimates.

Taken together, these results do not provide evidence that fiscal capacity systematically moderates the effect of procurement on test scores. The more striking finding is that fiscal capacity strongly predicts procurement participation, suggesting that observed heterogeneity in educational outcomes across municipalities may partly reflect differences in procurement intensity driven by fiscal constraints rather than differences in the effectiveness of procurement investments.

6.3 Construction versus non-construction procurement

Education-related municipal procurement covers heterogeneous types of investments. A substantial share of identified contracts corresponds to construction and renovation works, while the remainder concerns equipment, furniture, digital tools, or school-related services.

To examine this heterogeneity, we split procurement activity into construction contracts (CPV code 45) and non-construction contracts.

Figure 13 reports the dynamic treatment effects for the two categories. Overall, neither construction nor non-construction procurement displays strong or persistent effects on average test scores. Estimated coefficients fluctuate around zero and confidence intervals remain wide.

Table 8 summarizes the average treatment effects.

Table 8: Average treatment effects: construction versus non-construction procurement

Procurement type	ATT	Std. Error	CI low	CI high
Construction	0.00	0.49	-0.97	0.97
Non-construction	-0.49	0.36	-1.18	0.21

Notes: The table reports average treatment effects aggregated across treatment cohorts and post-treatment periods using the Callaway and Sant'Anna (2021) estimator. Confidence intervals are computed at the 95% level.

These results suggest that disaggregating procurement by broad category does not reveal strong differences in educational impacts. The absence of statistically significant effects is likely partly driven by reduced statistical power when the sample is split into smaller treatment groups.

Education-related municipal procurement covers heterogeneous types of investment. A substantial share of identified contracts corresponds to construction and renovation works, while the remainder concerns non-construction items such as equipment, furniture, digital tools, or school-related services. Distinguishing between these categories is useful because the mechanisms through which they may affect learning outcomes are likely to differ. Construction projects may generate delayed benefits and short-run disruption effects, whereas non-construction procurement may affect the school environment more rapidly but on a smaller scale.

To examine this heterogeneity, we split education-related procurement into two categories based on CPV codes. Construction procurement corresponds to CPV code 45, while non-construction procurement includes all other education-related contracts. For each category, we construct a separate treatment timing variable equal to the first year in which the municipality is exposed to that type of procurement. We then estimate staggered difference-in-differences models using the same specification as in the baseline analysis.

Figure 13 reports the dynamic treatment effects for the two categories. Overall, neither construction nor non-construction procurement displays a strong or persistent effect on average test scores. The estimated coefficients fluctuate around zero throughout the event window, and confidence intervals remain wide. For construction procurement, the pattern is broadly flat, with a small negative estimate around the year of treatment followed by coefficients close to zero. Non-construction procurement exhibits somewhat more negative point estimates on average, but again without statistically robust post-treatment effects. These patterns suggest that once education-related procurement is decomposed by category, no clear dynamic effect emerges for either type of investment. Figure 13 is consistent with this interpretation, as the two series remain close to zero across the event window.

Table 8 summarizes the average post-treatment effects aggregated across cohorts and post-treatment periods. The estimated average treatment effect for construction procurement is equal to 0.00, with a standard error of 0.49 and a 95% confidence interval ranging from -0.97 to 0.97. For non-construction procurement, the estimated ATT is -0.49, with a standard error of 0.36 and a confidence interval ranging from -1.18 to 0.21. In both cases, the confidence intervals include zero, indicating that the average effects are not statistically distinguishable from no effect.

Taken together, these results suggest that the overall relationship between education-related procurement and student achievement is not driven by a specific category of contracts. If anything, non-construction procurement is associated with slightly more negative point estimates, but the magnitude remains modest and imprecisely estimated. More generally, the absence of strong differences between construction and non-construction procurement indicates that disaggregating

procurement by broad category does not reveal sharply contrasting educational effects in the data.

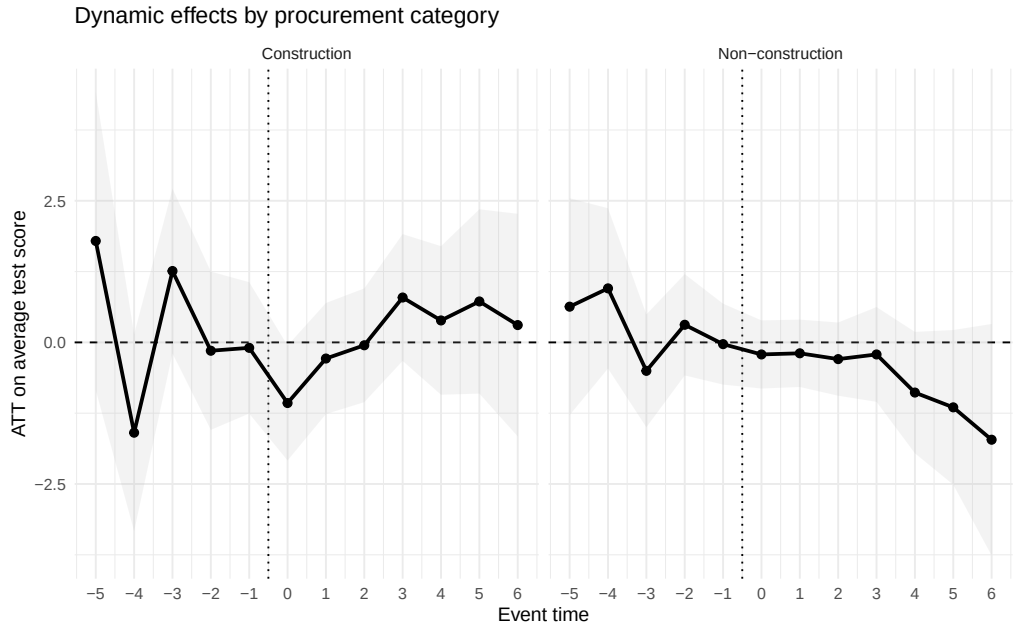

Figure 7: Dynamic effects by procurement category

The absence of statistically significant effects once procurement is disaggregated by category is likely driven by the reduced statistical power of the estimates, as the number of treated observations becomes substantially smaller in each category. Importantly, the point estimates remain consistent with the baseline results: non-construction procurement is associated with slightly negative effects, while construction procurement displays coefficients close to zero.

7 Robustness checks

7.1 Alternative staggered DiD estimators

As a robustness check, we re-estimate the baseline specification using alternative staggered difference-in-differences estimators. In addition to the main Callaway and Sant'Anna (2021) specification with not-yet-treated municipalities as the control group, we consider two alternatives: the Sun and Abraham (2021) event-study estimator and a Callaway and Sant'Anna specification using never-treated municipalities as the control group.

Table 9 reports the corresponding average treatment effects. The two Callaway and Sant'Anna specifications yield similar estimates, ranging from -0.75 to -0.84, and both remain statistically significant. This suggests that the negative baseline effect is robust to alternative control-group definitions within the same estimation framework.

By contrast, the Sun and Abraham estimator produces an estimate close to zero and statistically insignificant. This indicates that the estimated effect of procurement on test scores is sensitive to the choice of staggered DiD estimator. Taken together, these results suggest that the data do not provide robust evidence that education-related procurement improves student test scores, and that the sign of the average score effect should be interpreted with caution.

Importantly, this sensitivity concerns the estimated effect on test scores rather than the broader conclusion of the paper. Across specifications, we consistently find that procurement activity is associated with an increase in student enrollment, while evidence of short-run improvements in measured academic performance remains limited.

Table 9: Robustness checks using alternative staggered DiD estimators

Estimator	ATT	Std. Error	CI low	CI high
Callaway–Sant'Anna (not-yet-treated)	-0.745	0.328	-1.387	-0.103
Sun and Abraham	0.022	0.283	-0.533	0.576
Callaway–Sant'Anna (never-treated)	-0.835	0.365	-1.551	-0.119

Notes: The table reports average treatment effects obtained using alternative staggered difference-in-differences estimators and control groups. Confidence intervals are computed at the 95% level.

8 Discussion and conclusion

This paper examines whether education-related municipal public procurement affects student outcomes in France. Using a newly constructed municipality-level panel combining procurement data and standardized test scores, and relying on staggered difference-in-differences estimators, we analyze the dynamic effects of procurement activity on educational performance.

The results provide limited evidence that education-related procurement leads to improvements in average test scores. While some specifications suggest negative effects, these estimates are sensitive to the choice of estimator and are not consistently robust across alternative identification strategies. This suggests that the data do not provide strong evidence of short-run improvements in measured academic performance following procurement activity.

By contrast, we document a clear and robust increase in the number of students taking standardized tests following procurement activity. This finding indicates that procurement investments are associated with changes in school capacity or the attractiveness of local schooling options. Additional analyses provide limited support for strong compositional effects, suggesting that changes

in student composition are unlikely to fully explain the observed patterns in test scores.

Taken together, these results highlight the importance of distinguishing between different channels through which public investment affects education systems. In this context, municipal procurement appears to play a role in shaping access to schooling and enrollment dynamics, while its short-run impact on measured learning outcomes remains uncertain.

These findings have several implications for policy. First, they suggest that evaluating education policies solely on the basis of test scores may overlook important effects related to access and capacity. Second, they highlight the role of local governments and administrative processes in shaping educational environments within centralized systems. Finally, they point to the need for further research on the long-run effects of procurement investments, as well as on the specific types of procurement that may be more directly linked to learning outcomes.

More broadly, this paper contributes to a growing literature emphasizing that the effectiveness of public investment depends not only on the amount of resources allocated, but also on how these resources are implemented and translated into local service provision.

References

Bosio, E., Djankov, S., Glaeser, E., and Shleifer, A. (2022). Public procurement in law and practice. *American Economic Review*, 112(4):1091–1117.

Callaway, B. and Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Card, D. and Krueger, A. B. (1992). Does school quality matter? returns to education and the characteristics of public schools in the united states. *Journal of Political Economy*, 100(1):1–40.

Escueta, M., Quan, V., Nickow, A. J., and Oreopoulos, P. (2020). Education technology: An evidence-based review. *Journal of Economic Literature*, 58(4):897–996.

Fazekas, M. and Blum, J. R. (2021). Improving public procurement outcomes: Review of tools and the state of the evidence base. *World Bank Policy Research Working Paper*, (9690).

Fazekas, M., Wachs, J., et al. (2023). State capacity and public procurement performance. *Public Administration Review*.

Hanushek, E. A. (1997). Assessing the effects of school resources on student performance: An update. *Educational Evaluation and Policy Analysis*, 19(2):141–164.

Jackson, C. K. (2024). Education finance and the distribution of school quality. *Annual Review of Economics*, 16:453–481.

Jackson, C. K., Johnson, R. C., and Persico, C. (2016). The effects of school spending on educational and economic outcomes: Evidence from school finance reforms. *Quarterly Journal of Economics*, 131(1):157–218.

Jackson, C. K. and Mackevicius, C. (2021). The distribution of school spending impacts. *NBER Working Paper*, (28517).

Lafortune, J., Rothstein, J., and Schanzenbach, D. W. (2018). School finance reform and the distribution of student achievement. *American Economic Journal: Applied Economics*, 10(2):1–26.

Lafortune, J. and Schönholzer, D. (2022). The effects of school infrastructure investments on student outcomes. *American Economic Journal: Economic Policy*.

OECD (2023). Government at a glance 2023.

OECD (2024). Public procurement in oecd countries.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

World Bank (2020). Enhancing government effectiveness and transparency: The fight against corruption.

World Bank (2023). Public procurement and governance review.

A Additional Results

A.1 Event-study plots for student enrollment

Figure 8: Event study plots: Effect of procurement spending on number of students

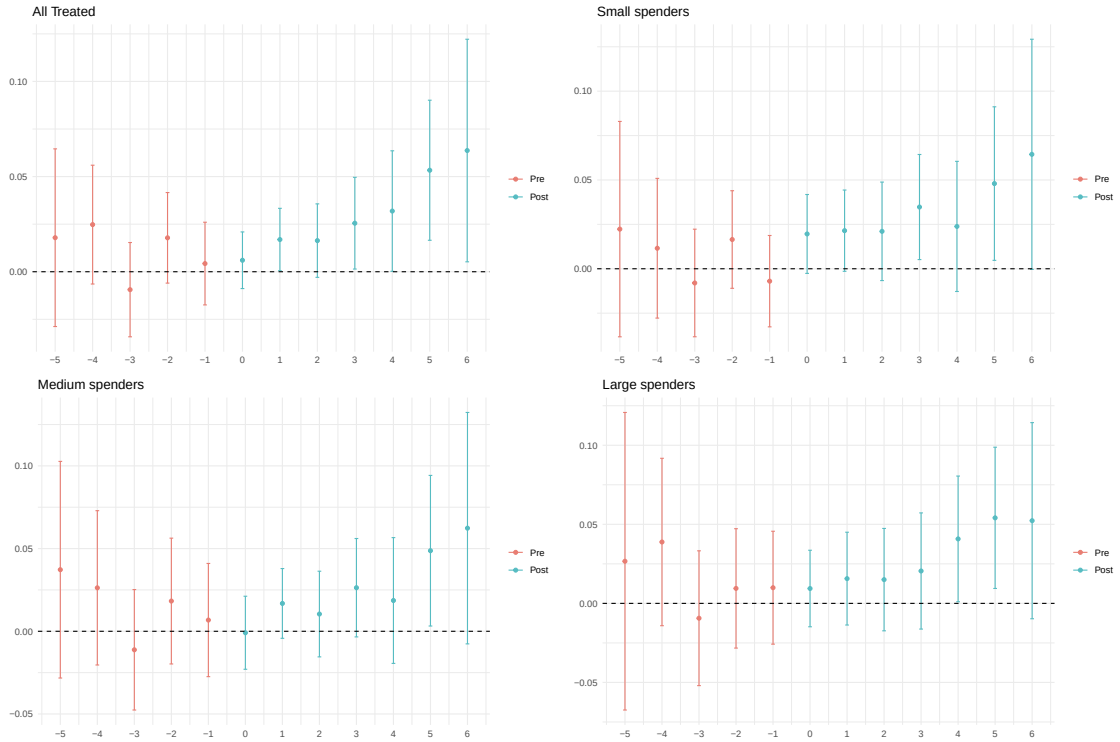

A.2 Enrollment

Table 10: Effect of procurement on test scores: role of enrollment

Dependent Variable:	score_avg	
	Baseline	+ Enrollment
Model:	(1)	(2)
Variables		
year = -6	-2.10**	-2.05**
	(0.898)	(0.897)
year = -5	-0.845	-0.810
	(0.726)	(0.726)
year = -4	-0.608	-0.616
	(0.508)	(0.507)
year = -3	-0.988**	-0.974**
	(0.469)	(0.470)
year = -2	-0.165	-0.159
	(0.349)	(0.349)
year = 0	-0.366	-0.373
	(0.279)	(0.280)
year = 1	-0.282	-0.305
	(0.284)	(0.285)
year = 2	-0.085	-0.097
	(0.307)	(0.308)
year = 3	0.168	0.150
	(0.342)	(0.344)
year = 4	0.384	0.363
	(0.401)	(0.402)
year = 5	0.125	0.087
	(0.448)	(0.448)
year = 6	0.287	0.214
	(0.555)	(0.556)
Log number of students		1.84**
		(0.771)
Fixed-effects	29	
codecommune	Yes	Yes
year	Yes	Yes

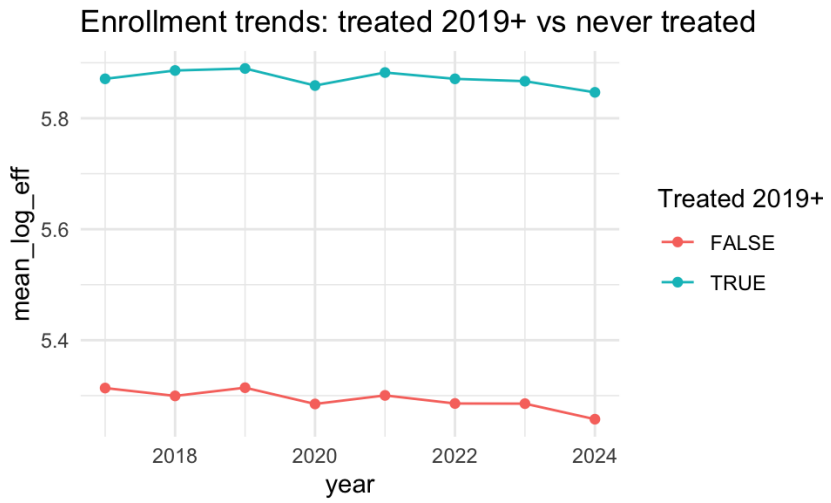

Figure 9: Enrollment trends by treatment status (2017–2024)

Notes: The figure reports the average log number of students for municipalities treated from 2019 onward and never-treated municipalities. Trends are parallel throughout the pre-treatment period, supporting the parallel trends assumption underlying the enrollment estimates. The level difference between the two groups reflects the larger average size of treated municipalities, which is absorbed by municipality fixed effects in the estimation.

B Additional Descriptive Statistics

B.1 Test score evolution

A first descriptive plot examines the average evolution of test scores between control and treated municipalities over the period 2017–2024, indexing test scores to 100 in 2017. A municipality is considered treated if it experiences education-related procurement at any point during the observation period, including not-yet-treated units.

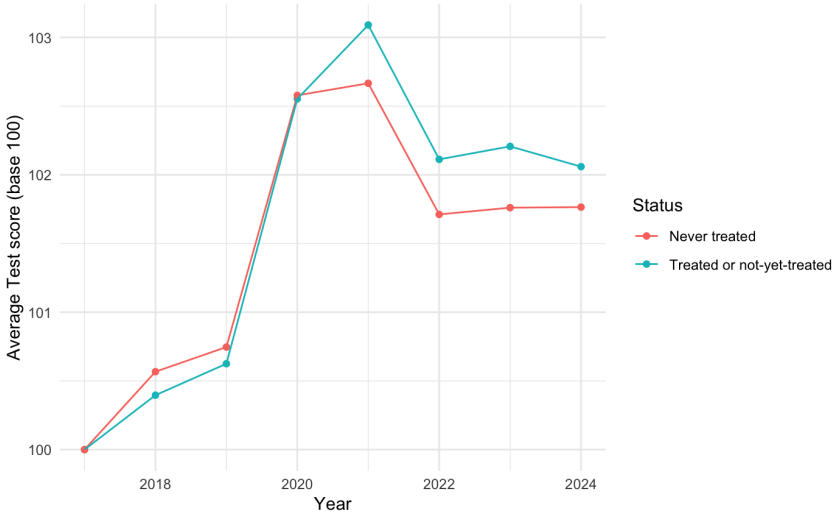

Figure 10: Evolution of average test scores by treatment status (base 100 in 2017)

Notes: The figure reports the evolution of average standardized test scores for treated (including not-yet-treated) and never-treated municipalities. Test scores are indexed to 100 in 2017.

B.2 Evolution of the amount of education-related procurement per year

We next examine the evolution of education-related procurement activity over time, measured as the average number of procurement lots per municipality-year.

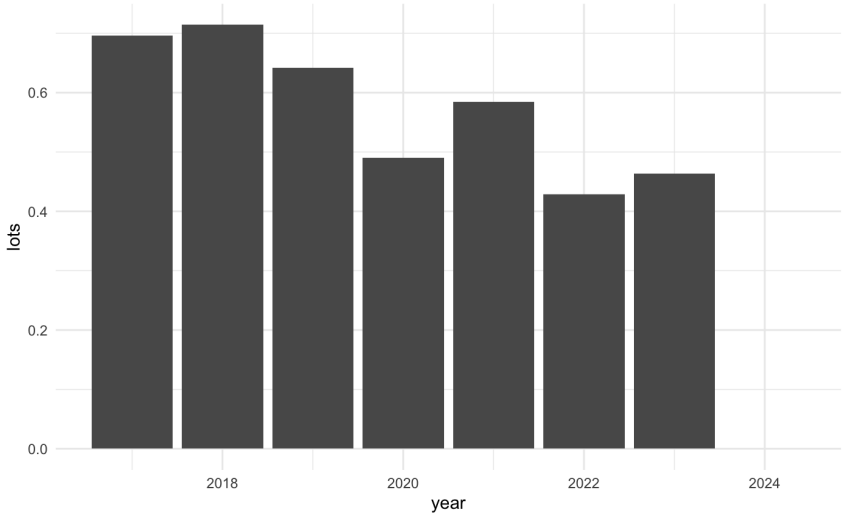

Figure 11: Average number of education-related procurement lots per municipality-year

Notes: The figure reports the average number of education-related procurement contracts (lots) per municipality and year over the period 2017–2024.

C Robustness checks

C.1 Sun & Abraham Event-Study

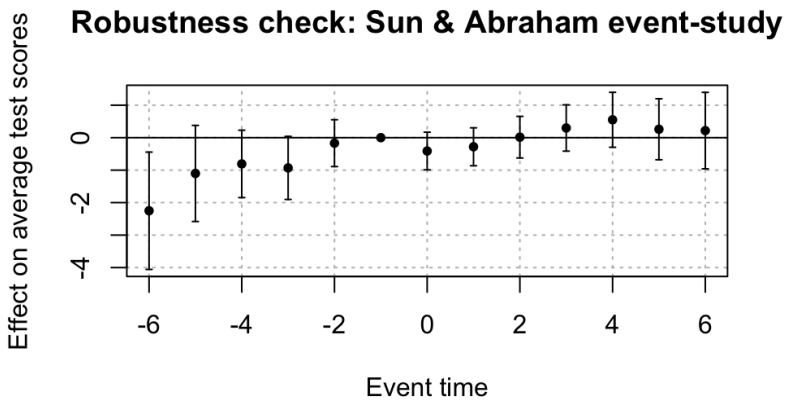

Figure 12: Event-study plot

C.2 Callaway & Sant'Anna using never-treated municipalities as the control group

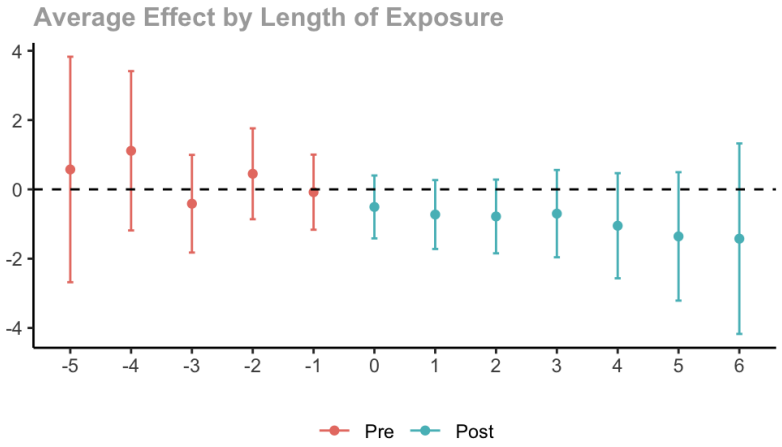

Figure 13: Dynamic effects by procurement category