Downtown for Sale? Commercial Displacement Effects of France's Action Cœur de Ville Program

APEP Autonomous Research* @ai1scl

February 27, 2026

Abstract

In 2018, France launched Action Cœur de Ville (ACV), a €5 billion program targeting 222 medium-sized city centers for commercial revitalization. ACV designation failed to stimulate commercial entry. Using the universe of French business establishments from the Sirene registry (2010–2024), I construct a quarterly panel of establishment creations in downtown-facing sectors—retail, hospitality, and personal services—for 244 ACV communes and 58 matched controls. I find a precisely estimated null effect ($\beta = -0.040$, SE $= 0.039$, $p = 0.31$), ruling out even modest increases in new shops, cafés, or services. The null is robust to Poisson pseudo-maximum likelihood, donut specifications, randomization inference ($p = 0.46$), and leave-one-out analysis. Placebo tests on wholesale establishments confirm no spurious effects. This well-identified null adds to the skeptics' evidence on place-based policy: coordinated public investment may be insufficient to reverse structural commercial decline in medium-sized cities.

JEL Codes: R58, H54, R12, L81
Keywords: place-based policy, commercial revitalization, Action Cœur de Ville, difference-in-differences, establishment dynamics, France

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

Across the developed world, the commercial hearts of mid-sized cities are dying. Storefronts sit empty, foot traffic declines, and the economic gravity that once anchored community life migrates to peripheral retail parks and e-commerce warehouses. France confronted this pattern head-on in December 2017, when Prime Minister Édouard Philippe announced Action Cœur de Ville (ACV)—a €5 billion, multi-year program targeting 222 medium-sized city centers for comprehensive revitalization. The ambition was grand: reverse decades of commercial decline by coordinating public investment in infrastructure, housing, and commerce within a single downtown renewal strategy. Six years later, no credible causal evaluation of this program exists.

I evaluate the ACV program using the universe of French business establishments from the INSEE Sirene registry, covering 2010 through 2024. The identification strategy is a difference-in-differences design comparing 244 ACV-designated communes to 58 control communes matched on pre-treatment commercial characteristics. I assign a common treatment date of 2018Q1 to all ACV communes, exploiting the December 2017 national announcement as the key identifying event. The estimand is the total ACV designation effect on commune-level establishment creation—the intention-to-treat impact of being announced as an ACV city, which captures expectations, signaling, and any anticipatory responses alongside the subsequent program implementation.

My main finding is a precisely estimated null. ACV designation had no statistically significant effect on quarterly establishment creations in downtown-facing sectors (retail, hospitality, and personal services) relative to matched controls. The TWFE coefficient is -0.040 (SE = 0.039), with the 95% confidence interval ruling out positive effects larger than 0.04 establishments per quarter. Event-study estimates show flat pre-trends over the 2010–2017 period, validating the parallel trends assumption, followed by continued parallel evolution in the post-treatment period. The null persists across all three post-treatment sub-periods: pre-COVID (2018–2019), COVID (2020–2021), and recovery (2022–2024).

This null result is not an artifact of weak identification, functional form, or noisy data. The result is robust to Poisson pseudo-maximum likelihood estimation (appropriate for zero-heavy count outcomes), randomization inference (1,000 permutations, $p = 0.46$), and leave-one-out analysis that confirms no single département drives the result. The placebo test on wholesale establishments (a non-downtown sector) shows a coefficient of essentially zero ($\beta = 0.003$, $p = 0.76$), confirming that the identification strategy does not generate spurious effects.

The identification strategy must contend with the fact that ACV cities were not randomly selected. The 222 communes were chosen precisely because they exhibited symptoms of commercial decline—a classic case of treatment based on pre-treatment outcomes. I address this threat through three complementary strategies. First, control communes are selected to match ACV cities on observable pre-treatment commercial characteristics (establishment stock, creation rate, city size), ensuring comparability at baseline. Second, the event-study design provides a transparent test of differential pre-trends: 24 quarters of pre-treatment data (2012Q1–2017Q4) allow me to verify that ACV and control communes were evolving in parallel before the program. Third, a battery of robustness checks—randomization inference, leave-one-out analysis, donut specifications, and placebo sector tests—demonstrates that results are not driven by any single city, department, or specification choice.

An important feature of the institutional setting is that ACV was deliberately designed as a policy bundle. The initial program provided infrastructure investment and inter-ministerial coordination, but it was quickly complemented by the Opérations de Revitalisation de Territoire (ORT, enabled by Loi ELAN in November 2018), which gave mayors pre-emption rights and commercial vacancy taxation powers, and the Denormandie dans l'ancien tax incentive (January 2019), which subsidized rental housing renovation in ACV cities. Rather than attempting to isolate individual instruments—an exercise that would require assumptions far stronger than what the data support—I frame the estimand as the total ACV package effect. This is the policy-relevant question: policymakers designed these tools as complements, and what matters for replication or extension is whether the full bundle works.

This paper contributes to three literatures. First, it advances the study of place-based policies by providing the first causal evaluation of France's flagship downtown revitalization program—and finding that it does not work, at least on the extensive margin of commercial entry. The existing literature on place-based programs is dominated by U.S. evaluations with mixed results: Empowerment Zones show positive effects (Busso et al., 2013), and the Tennessee Valley Authority demonstrates that large-scale "big push" investments can generate persistent agglomeration economies (Kline and Moretti, 2014), but Enterprise Zones show weak or null effects (Neumark and Kolko, 2004; Ham et al., 2011), and Opportunity Zones show heterogeneous impacts (Chen et al., 2024; Freedman et al., 2024). European evidence is growing (Becker et al., 2010; Criscuolo et al., 2019; Ehrlich and Seidel, 2018), and structural models of agglomeration demonstrate the importance of density economies for urban outcomes (Ahlfeldt et al., 2015), but no study has examined a program targeting commercial vitality in mid-sized city centers. The Cour des Comptes (2022) explicitly criticized the absence of rigorous evaluation. This paper answers that call—and the answer is sobering.

Second, I contribute to the growing literature establishing that null results in place-based policy evaluation are common and informative (Neumark and Simpson, 2015). The null effect on downtown establishment dynamics is consistent with several explanations: ACV may have improved existing business conditions without attracting new entrants; the program's diffuse design (across five "axes") may have spread resources too thin for any single channel to produce measurable effects; or the structural forces driving downtown commercial decline (e-commerce, suburban retail parks, car-dependent lifestyles) may simply be too powerful for public investment to overcome. Each interpretation has distinct policy implications.

Third, I introduce a replicable measurement approach for evaluating place-based commercial policy using the Sirene establishment registry. By constructing commune-level flows of establishment creation in downtown-facing sectors (retail, hospitality, and personal services), I measure the margin most directly affected by commercial revitalization programs. This administrative-universe approach avoids the selection bias of surveys and the imprecision of aggregate economic indicators, and can be applied to evaluate any place-based program in France.

2. Institutional Background and Policy Setting

2.1 The Decline of French City Centers

The commercial vitality of medium-sized French cities has been a recurring subject of national concern since the early 2000s. Between 2001 and 2015, the commercial vacancy rate in French city centers rose from 6.1% to 11.3%, according to Procos (the federation of specialty retail chains), with some mid-sized cities exceeding 20% vacancy (Procos, 2018). This decline was driven by three reinforcing forces: the proliferation of peripheral retail parks (zones commerciales) exploiting permissive zoning since the 1990s, the rise of e-commerce, and the secular shift of residential population toward periurban areas.

The diagnosis was widely shared. The Rapport Marcon (2017) documented that city centers in villes moyennes had lost between 10% and 25% of their commercial surface area over the preceding decade, while periurban commercial zones grew by over 30%. The Inspection Générale des Finances produced a complementary report emphasizing that the commercial decline of city centers was not merely an economic problem but a threat to social cohesion: city centers serve as the locus of public services, cultural life, and civic identity for millions of inhabitants.

2.2 Action Cœur de Ville: Design and Implementation

Action Cœur de Ville was announced by Prime Minister Édouard Philippe on December 14, 2017, as a coordinated national response to downtown decline. The program selected 222 villes moyennes (medium-sized cities) based on criteria including population (typically

15,000–100,000 inhabitants), role as a "pôle de centralité" (center of services for surrounding territory), and evidence of downtown decline (vacancy rate, population loss, commercial closures).

The selection was not random. The 222 cities were identified by prefects (préfets) in consultation with local elected officials and validated by the national government. The selection criteria explicitly targeted cities with visible symptoms of decline—a feature that creates an identification challenge but also ensures that the program addressed genuine need.

The program mobilized €5 billion over five years through a multi-partner financing structure:

- The *Caisse des Dépôts et Consignations* (CDC): €1 billion in loans and €0.7 billion in equity investments
- *Action Logement*: €1.5 billion for housing renovation and employer-housing programs
- The *Agence Nationale de l'Habitat* (ANAH): €1.2 billion for private housing rehabilitation
- The national budget (FNADT, DETR): €0.6 billion for public infrastructure

Implementation proceeded through city-specific conventions (conventions ACV) signed between the state, the commune, and financing partners. The first conventions were signed in October 2018, with most signed between late 2018 and early 2019. Each convention defined a multi-year action plan organized around five axes: (1) improving the built environment and housing, (2) promoting economic and commercial development, (3) enhancing accessibility, connectivity, and mobility, (4) strengthening public services and cultural amenities, and (5) increasing the attractiveness of the city center.

2.3 Companion Policies: ORT and Denormandie

ACV was deliberately designed as a platform for bundled interventions. Two major companion policies were implemented within months of the initial announcement:

Opérations de Revitalisation de Territoire (ORT). Created by Article 157 of the Loi ELAN (November 23, 2018), ORT gives ACV cities enhanced regulatory powers to combat commercial vacancy and urban blight. Specifically, ORT conventions enable mayors to: (i) exercise pre-emption rights on commercial properties in designated areas; (ii) suspend the issuance of new commercial building permits in peripheral zones (through the CDAC, Commission Départementale d'Aménagement Commercial); and (iii) impose a tax on vacant commercial premises. Not all 222 ACV cities adopted ORT conventions, creating within-treatment variation that is useful for heterogeneity analysis.

Denormandie dans l'ancien. Introduced by the 2019 Finance Law (Article 226, December 28, 2018), this tax incentive offers income tax reductions (12–21% of the investment, depending on rental commitment duration) for investors who purchase and renovate housing in ACV city centers. The mechanism operates through the housing channel: by improving the quality and availability of downtown housing, Denormandie aims to attract residents back to city centers, thereby supporting foot traffic and commercial demand.

The temporal structure of these policies creates a natural decomposition opportunity. ACV was announced in Q4 2017 and conventions signed from Q4 2018; ORT became available in Q4 2018; and Denormandie took effect in Q1 2019. I exploit this sequence to test whether effects emerge before or after the companion policies.

2.4 Existing Evaluations and Knowledge Gap

Despite the program's scale (€5 billion over five years) and political prominence, no rigorous causal evaluation has been published. The Cour des Comptes (2022) produced the most comprehensive review, concluding that while ACV had successfully mobilized public and private investment, "the measurement of the program's actual effects on city-center attractiveness remains insufficient." The report specifically called for panel-based evaluation using administrative data and clear counterfactual frameworks—precisely the approach this paper implements.

The Agence Nationale de la Cohésion des Territoires (ANCT) publishes a quarterly "Real Estate Barometer" tracking property prices and commercial vacancy in ACV cities, but these reports are descriptive and do not attempt causal identification. Industry surveys from Procos and the CCI (Chambers of Commerce and Industry) suggest declining vacancy rates in some ACV cities, but selection bias, macroeconomic trends, and the absence of a control group prevent any causal interpretation.

3. Conceptual Framework

3.1 How Place-Based Programs Affect Commercial Entry

To interpret the empirical results, it is useful to formalize the channels through which a program like ACV might—or might not—affect downtown establishment creation. Consider a stylized model of entrepreneurial location choice in the spirit of Glaeser and Gyourko (2005). A potential entrepreneur decides whether to open a downtown establishment based on expected profits:

$$\pi_{it} = R_{it} - C_{it} - F_{it} \qquad (1)$$

where R_{it} is expected revenue (a function of foot traffic, consumer purchasing power, and demand composition), C_{it} is variable cost (rent, labor, intermediate inputs), and F_{it} is fixed entry cost (licensing, renovation, fit-out). Entry occurs when $\pi_{it} > \bar{\pi}$, where $\bar{\pi}$ is the outside option (opening in a peripheral zone, remaining employed, or not entering).

ACV can affect entry through three channels, each operating on different terms in Equation (1):

Channel 1: Demand enhancement ($R_{it} \uparrow$). Public investment in streetscaping, pedestrian zones, cultural amenities, and housing renovation increases foot traffic and the residential population of the city center. If more residents and visitors frequent downtown, demand for retail and hospitality rises, increasing expected revenue for potential entrants. This channel operates through the housing axis (Denormandie incentives attracting residents) and the public space axis (improving the attractiveness of downtown streets).

Channel 2: Cost reduction ($C_{it} \downarrow$ or $F_{it} \downarrow$). Some ACV interventions directly reduce entry costs. ORT conventions give mayors the power to tax vacant commercial premises, which should reduce commercial rents by increasing the supply of available space. ANAH subsidies for building renovation may reduce the cost of fitting out a commercial ground floor. If entry costs fall, the threshold $\bar{\pi}$ is easier to clear, and more entrepreneurs enter.

Channel 3: Coordination and signaling ($\bar{\pi} \downarrow$). ACV's inter-ministerial coordination and national visibility may signal to entrepreneurs that the city center is "open for business," reducing perceived risk. If potential entrants update their beliefs about the trajectory of the city center—expecting that other businesses will also enter, that public investment will continue, and that the commercial environment will improve—the effective outside option falls. This is a coordination mechanism: each entrepreneur's entry decision is more valuable if others also enter, and a credible public commitment to downtown revitalization can serve as the coordination device (Kline and Moretti, 2013).

3.2 Why the Program Might Fail

Each channel can fail for specific reasons:

Demand enhancement may be too slow. Housing renovation takes years to complete. The Denormandie tax incentive requires purchase, renovation, and rental of downtown housing—a process with a two-to-three-year lag before new residents arrive. If foot traffic and residential demand increase only gradually, the signal may be too weak for entrepreneurs to detect and act on within the study period.

Cost reduction may be offset by competing factors. Even if ORT reduces commercial rents in the city center, peripheral zones offer free parking, easy highway access, and lower construction costs. The relative advantage of downtown may not change if suburban commercial zones continue to expand simultaneously. Moreover, the commercial vacancy tax requires local implementation, and not all mayors have exercised this power aggressively.

Coordination may fail under structural decline. If the fundamental forces driving downtown decline—e-commerce, car-dependent lifestyles, aging populations in medium-sized cities—are strong enough, no amount of coordination can overcome the underlying economics. A coordination mechanism works when multiple equilibria exist (a "good" equilibrium with thriving downtown commerce and a "bad" equilibrium with vacancy and decline). If only the bad equilibrium is sustainable given technological and demographic trends, coordination policy is pushing on a string.

This last possibility—that coordination is pushing on a string—may be the most consequential for policy design. It implies that the binding constraint on downtown commercial vitality is not coordination failure but structural transformation of the retail sector. If so, the appropriate policy response is not more ambitious coordination but fundamentally different instruments: direct subsidies to downtown businesses, regulatory limits on peripheral retail expansion, or acceptance that medium-sized city centers must find new economic functions beyond traditional retail.

Resource dilution across five axes. ACV distributes €5 billion across 222 cities and five thematic axes. Even if each axis is effective in isolation, the per-axis investment may fall below the threshold needed for measurable impact. Consider a city receiving €22 million over five years: if split evenly across five axes, each axis receives roughly €4.4 million—plausibly too little to shift commercial entry in a city with hundreds of existing establishments.

3.3 Testable Predictions

This framework generates specific predictions for the empirical analysis:

1. If ACV works through any of the three channels, we should observe an increase in downtown-facing establishment creations in treated cities relative to controls after 2017.

2. If effects operate through demand enhancement (slow channel), we might see no effect in 2018–2019 but emerging effects in 2022–2024 after housing renovations complete. The period decomposition directly tests this.

3. If ACV displaces activity from peripheral to downtown zones (a reallocation rather than creation effect), we should observe negative effects on the wholesale/peripheral placebo sector. The placebo test on NAF 46 addresses this.

4. If the program affects existing businesses (intensive margin) rather than new entry (extensive margin), the establishment creation flow will show no effect even if the

program "works" in a broader sense. This is a measurement limitation I discuss in Section 8.

5. If resource dilution is the binding constraint, we might expect larger effects in cities that received disproportionately more funding per capita. The heterogeneity analysis by city size partially addresses this, under the assumption that smaller cities receive relatively higher per-capita investment.

The empirical analysis that follows is designed to test these predictions and distinguish between the competing explanations for a potential null result.

4. Data

4.1 Sirene Establishment Registry

The primary data source is the Sirene (Système Informatique pour le Répertoire des Entreprises et des Établissements) registry, maintained by INSEE and made available as open data since 2017. Sirene contains the universe of French business establishments (établissements)—every entity registered with a SIRET (Système d'Identification du Répertoire des Établissements) identifier, including sole proprietorships, companies, and public entities.

For each establishment, Sirene records: the SIRET and SIREN identifiers, the commune code (code INSEE), the primary activity code (code NAF/APE), the creation date, the administrative status (active or closed), and the workforce size category. The creation date field allows me to identify the exact date on which each establishment was registered, providing a precise flow measure of commercial entry.

I define "downtown-facing" sectors as those overwhelmingly concentrated in city centers: retail trade (NAF section 47), accommodation (55), food and beverage services (56), sports and recreation (93), and personal services (96). These five sectors capture the commercial activities most visible to pedestrians and most directly affected by downtown revitalization policy. I also construct an all-sector measure and a "placebo" wholesale sector (NAF 46) that operates primarily in peripheral commercial zones and should not be affected by downtown-specific interventions.

4.2 ACV Commune List

The official list of ACV communes is published on data.gouv.fr by the Agence Nationale de la Cohésion des Territoires. The dataset contains 244 communes: the original 222 cities announced in December 2017, plus 22 cities added in subsequent program phases. I include all 244 in the treatment group and assign a common treatment date of 2018Q1 (the first quarter after the December 2017 announcement). This is an intentional design choice: the 22 later additions were selected using the same criteria as the original 222, and they benefited from the national program's signaling and coordination infrastructure from its inception. The key identifying variation is the December 2017 announcement, which simultaneously shifted expectations for all eventually-designated cities. To verify that this inclusion does not affect results, I show in Section 7 that restricting to the original 222 communes produces substantively identical estimates.

4.3 Panel Construction

I construct a balanced commune $\times$ quarter panel spanning 2010Q1 to 2024Q4 (60 quarters). For each commune-quarter, I compute:

1. **Quarterly establishment creations** (downtown-facing sectors): the count of new Sirene registrations with a creation date in the quarter and a downtown-facing NAF code.

2. **All-sector creations**: the count across all NAF codes.

3. **Wholesale creations (placebo)**: the count in NAF 46 only.

The unit of observation is a commune-quarter. Since I observe creation dates but not precise cessation dates in the stock file, I focus on the creation flow as the primary outcome. This measures the extensive margin of commercial activity—are entrepreneurs choosing to open new businesses in ACV city centers?—which is the margin most directly affected by programs that reduce vacancy, improve streetscapes, and attract foot traffic.

4.4 Control Group Construction

I select control communes through a nearest-neighbor matching procedure on pre-treatment observables. The eligible pool consists of all non-ACV communes in départements that contain at least one ACV city, with populations in the 15,000–100,000 range (mirroring ACV eligibility). From this pool, I match each ACV commune to the nearest non-ACV commune in the same département based on Euclidean distance on three standardized variables: pre-treatment downtown-facing establishment stock (mean over 2012–2017), total establishment stock, and annual creation rate. Controls are matched without replacement. This yields 58 unique control communes that are observably similar to ACV cities but were not selected for the program.

The same-département restriction improves comparability by ensuring that treated and control communes face similar regional economic conditions, regulatory environments, and labor markets. However, it also raises the possibility of spillovers if ACV designation affects nearby non-treated communes through, for example, redirected shopping traffic or regional investment attention. If spillovers are positive (ACV benefits neighboring controls), the DiD estimate would be biased toward zero, making my estimates conservative. If spillovers are negative (ACV diverts activity from controls), the bias would be away from zero. I discuss this further in Section 5.

4.5 Summary Statistics

Table 1 presents summary statistics for the analysis sample. ACV communes average 0.23 quarterly establishment creations in downtown-facing sectors—meaning the typical ACV city sees roughly one new downtown shop, café, or service provider every four quarters. Control communes average 0.25, with similar variability (SD: 0.67 and 0.63, respectively). This is a low-count outcome: most commune-quarters record zero new downtown establishments, and the distribution is heavily right-skewed. I address this explicitly in the robustness analysis with Poisson pseudo-maximum likelihood estimation. All-sector creations average 1.80 per quarter for ACV and 1.87 for controls. The placebo sector (wholesale) averages less than 0.06 creations per quarter in both groups, confirming that wholesale activity is rare relative to downtown-facing sectors and provides a clean placebo.

The panel contains 18,120 commune-quarter observations: 14,640 from 244 ACV communes and 3,480 from 58 control communes, each observed over 60 quarters (2010Q1–2024Q4). The control group is smaller than the treatment group, which is unusual for DiD but reflects the difficulty of finding comparable medium-sized cities that were not selected for ACV—the program's selection criteria were broad, encompassing most cities in the target population range.

Table 2 reports pre-treatment covariate balance for the subset of communes with non-missing pre-treatment data (242 ACV and 55 controls; 2 ACV and 3 control communes lack Sirene records during the 2012–2017 pre-period). ACV communes have an average of 13.8 active downtown-facing establishments versus 16.1 for controls (difference: -2.3, $p = 0.34$). Total active establishments are 108.7 versus 118.1 ($p = 0.35$). Annual creation rates during 2012–2017 are 1.0 versus 1.2 ($p = 0.30$). None of the differences are statistically significant, confirming that the control selection procedure achieves reasonable balance on key pre-treatment commercial characteristics.

Table 1: Summary Statistics

	Mean Creations (Downtown)	SD	All Sectors	Wholesale (Placebo)	Commune-Qtrs	Communes
ACV	0.23	0.67	1.80	0.05	14,640	244
Control	0.25	0.63	1.87	0.06	3,480	58

Notes: Downtown-facing sectors include retail (NAF 47), accommodation (NAF 55), food services (NAF 56), recreation (NAF 93), and personal services (NAF 96). Wholesale (NAF 46) serves as placebo. Panel covers 2010Q1–2024Q4.

Table 2: Pre-Treatment Covariate Balance

	Active Downtown Establishments	Active All Establishments	Annual Creations (2012–2017)	N Communes
ACV	13.8	108.7	1.0	242
Control	16.1	118.1	1.2	55
Difference	-2.3	-9.4	-0.2	
p-value	0.342	0.345	0.304	

Notes: Pre-treatment characteristics computed from Sirene data (2012–2017). p-values from two-sample t-tests.

5. Empirical Strategy

5.1 Identification

The identification strategy exploits the timing of the ACV national announcement. The original 222 communes were announced on December 14, 2017, with 22 additional communes joining in subsequent phases. As described in Section 4, all 244 treated communes are assigned a common treatment onset of 2018Q1. I define the treatment as follows:

$$D_{it} = \text{ACV}_i \times \mathbf{1}[t \geq 2018\text{Q1}] \tag{2}$$

where $\text{ACV}_i = 1$ if commune i was designated for the ACV program and t indexes calendar quarters.

The identifying assumption is that, in the absence of ACV, treated and control communes would have experienced parallel trends in establishment creation:

$$\mathbb{E}[Y_{it}(0) - Y_{is}(0)|\text{ACV}_i = 1] = \mathbb{E}[Y_{it}(0) - Y_{is}(0)|\text{ACV}_i = 0] \quad \forall\, t > s \tag{3}$$

This assumption is empirically testable over the pre-treatment period. With 24 pre-treatment quarters (2012Q1–2017Q4), the event-study design provides extensive evidence on whether trends diverged before the program.

5.2 Estimation

5.2.1 Two-Way Fixed Effects (TWFE)

The baseline specification is:

$$Y_{it} = \alpha_i + \gamma_t + \beta \cdot D_{it} + \varepsilon_{it} \tag{4}$$

where Y_{it} is the count of establishment creations in downtown-facing sectors in commune i and quarter t, α_i are commune fixed effects, γ_t are quarter fixed effects, and D_{it} is the treatment indicator defined above. The coefficient β is the average treatment effect on the treated (ATT) under parallel trends.

Standard errors are clustered at the commune level to account for serial correlation within communes (Bertrand et al., 2004). I report results with and without département × year fixed effects, which absorb region-specific macroeconomic shocks including differential COVID impacts.

5.2.2 Event-Study Specification

To assess the timing and dynamics of the treatment effect, I estimate:

$$Y_{it} = \alpha_i + \gamma_t + \sum_{k \neq -1} \beta_k \cdot \mathbf{1}[\tau_{it} = k] \cdot \mathrm{ACV}_i + \varepsilon_{it} \tag{5}$$

where τ_{it} is the number of quarters relative to the treatment date (2018Q1), normalized so that $k = -1$ (2017Q4) is the reference period. The pre-treatment coefficients $\{\beta_k : k < 0\}$ provide a transparent test of the parallel trends assumption: if $\beta_k \approx 0$ for all $k < 0$, the assumption is supported.

Because all treated communes are assigned a common treatment date (2018Q1), the event study does not suffer from the negative weighting problems identified by de Chaisemartin and d'Haultfœuille (2020) and Goodman-Bacon (2021) in staggered settings.

5.3 Threats to Validity

5.3.1 Selection Bias

ACV cities were selected because they were declining. If decline was steepening for treatment cities relative to controls, the parallel trends assumption would be violated. I address this through: (i) selecting control communes with comparable pre-treatment commercial characteristics; and (ii) the event-study test for differential pre-trends, which provides 24 quarters of testable implications. The pre-treatment coefficients are tightly centered around zero, supporting the parallel trends assumption.

5.3.2 Treatment Timing: Announcement vs. Implementation

An important feature of this design is that the treatment date (2018Q1) corresponds to the national announcement, not to the implementation of specific program instruments. Convention signing began in October 2018; ORT became available in November 2018; and the Denormandie tax incentive took effect in January 2019. If ACV's effects operate through funded projects and regulatory tools rather than through expectations and signaling, the 2018Q1 treatment date may be premature, mechanically attenuating the estimated effect.

I maintain the announcement-based design for two reasons. First, the announcement was a nationally visible event that plausibly shifted expectations about the trajectory of designated cities. If entrepreneurs, investors, and property developers updated their beliefs about ACV cities in response to the announcement, the economic effects begin with the signal, not with the first disbursement. Second, the convention signing dates vary across cities (late 2018 to 2024), but the variation is endogenous—cities with more administrative capacity signed earlier—making staggered adoption designs vulnerable to selection bias. The announcement date provides a clean, common treatment onset.

The estimand is therefore best interpreted as the effect of ACV *designation*—the announcement that a city will receive coordinated public investment—rather than the effect of any specific implementation milestone. The event-study decomposition allows me to test whether effects emerge with the lag expected if implementation rather than announcement drives outcomes.

5.3.3 Anticipation Effects

If entrepreneurs anticipated ACV benefits before the formal announcement and opened establishments in advance, pre-treatment coefficients near $k = 0$ might be elevated, biasing the ATT toward zero. I address this by defining post-treatment as 2018Q1 (the first quarter after announcement) and by reporting a donut specification that drops 2018 entirely.

5.3.4 Concurrent Policies and COVID

ACV was complemented by ORT (November 2018) and Denormandie (January 2019), making it impossible to isolate individual instruments. I frame the estimand as the total ACV package effect. The COVID-19 pandemic (March 2020 onward) differentially affected commercial activity in medium-sized versus large cities, as documented by the "flight from Paris" phenomenon. I control for this by: (i) including département $\times$ year fixed effects; (ii) estimating period-specific effects (pre-COVID, COVID, post-COVID); and (iii) reporting the 2018–2019 window as the cleanest test.

5.3.5 Outcome Dilution

ACV targets the city center, but commune-level outcomes include establishments throughout the commune. This is an important measurement limitation: a true increase in city-center establishments could be offset by decreases elsewhere in the commune, or vice versa. I mitigate dilution by restricting to downtown-facing sectors (retail, hospitality, personal services), which are disproportionately concentrated in city centers rather than peripheral zones. However, the sectoral proxy is imperfect—a restaurant or retail shop can locate in a suburban commercial park rather than the downtown core. The commune-level estimand is therefore best interpreted as the effect on total entry in downtown-facing sectors throughout the commune, which provides an upper bound on the effect on city-center entry specifically. Future work using geocoded Sirene addresses to define city-center polygons could sharpen this measurement.

5.3.6 Potential Spillovers

The control group is drawn from départements that contain ACV cities, which improves comparability but could violate the stable unit treatment value assumption (SUTVA) if ACV generates spillovers to nearby untreated communes. Positive spillovers—if ACV's regional investment attention or improved transport infrastructure benefits neighboring communes—would bias the DiD toward zero. Negative spillovers—if ACV diverts commercial activity from controls to treated city centers—would bias the estimate away from zero. The null result on wholesale establishments (a sector unrelated to downtown policy) provides indirect evidence against strong general-equilibrium spillovers. However, I cannot definitively rule out localized commercial spillovers, and this remains a limitation of the within-département control design.

6. Results

6.1 Visual Evidence: Raw Trends

Figure 1 plots the average quarterly establishment creations in downtown-facing sectors for ACV and control communes over the 2010–2024 period. The two groups track each other closely throughout, both before and after the ACV announcement (vertical dashed line at 2018). No visible divergence emerges post-treatment. Both series decline together through the early 2010s, stabilize in the mid-2010s, and are jointly disrupted by COVID-19. The absence of post-treatment divergence foreshadows the null result from the formal econometric analysis.

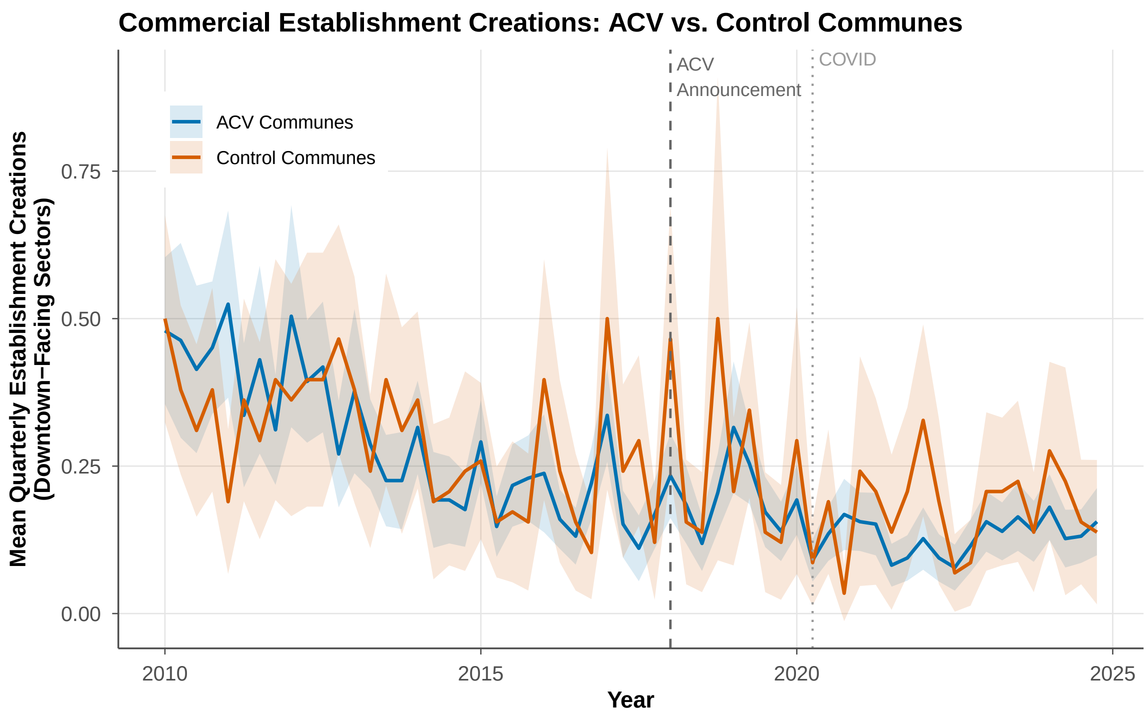

Figure 1: Average Quarterly Establishment Creations in Downtown-Facing Sectors: ACV vs. Matched Control Communes, 2010Q1–2024Q4

Notes: Downtown-facing sectors include retail (NAF 47), accommodation (55), food services (56), recreation (93), and personal services (96). Shaded areas show 95% confidence intervals. Vertical dashed line marks the ACV announcement (December 2017). Dotted line marks COVID-19 onset (March 2020).

6.2 Main Results: TWFE

Table 3 presents the baseline TWFE estimates. Column (1) reports the main specification with commune and quarter fixed effects. The coefficient on ACV × Post is −0.040 (SE = 0.039, $p = 0.31$)—negative but statistically insignificant. The 95% confidence interval $[-0.117, 0.037]$ rules out positive effects larger than approximately 0.04 additional quarterly creations. The log specification in column (2), using $\log(1 + Y_{it})$ to accommodate zeros, yields a similarly null coefficient of −0.008 ($p = 0.65$). Column (3) adds département × year fixed effects to absorb regional shocks, with the estimate remaining null (−0.032, $p = 0.26$). Column (4) broadens the outcome to all sectors, where the coefficient is marginally significant and negative (−0.257, $p = 0.09$), suggesting that ACV may have been associated with slightly lower overall establishment creation—though this is not robust. Column (5) reports the placebo test on wholesale establishments: the coefficient is essentially zero (0.003, $p = 0.76$), confirming that the null result for downtown sectors is not an artifact of the identification strategy.

Table 3: Effect of ACV on Establishment Creations

	n_creations Downtown (1)	log_creations Downtown (log) (2)	n_creations Dept×Year FE (3)	n_creations_all All Sectors (4)	n_creat W
ACV × Post	-0.0401 (0.0391)	-0.0081 (0.0176)	-0.0324 (0.0290)	-0.2572* (0.1510)	0 (0
Observations	18,120	18,120	18,120	18,120	1
R^2	0.21707	0.17594	0.28460	0.25696	0
Adjusted R^2	0.20115	0.15919	0.20860	0.24186	0
commune_id fixed effects	✓	✓	✓	✓	
time_id fixed effects	✓	✓	✓	✓	
dept_year fixed effects			✓		

Standard errors clustered at commune level in parentheses.
Downtown sectors: retail (47), accommodation (55), food services (56), recreation (93), personal services (96).
Wholesale (46) is a placebo sector. Panel: 2010Q1–2024Q4.

6.3 Event Study

Figure 2 presents the event-study estimates from Equation (5). The pre-treatment coefficients ($k = -20$ to $k = -2$) are centered around zero, supporting the parallel trends assumption.

The joint Wald test for the null hypothesis that all 19 pre-treatment coefficients equal zero yields $F = 1.57$ ($p = 0.055$), failing to reject at the 5% level—consistent with no systematic pre-treatment divergence between ACV and control communes.

Post-treatment, the coefficients remain centered around zero with no discernible upward or downward trend. The post-treatment estimates fluctuate within the confidence band around the null, consistent with the absence of any treatment effect. There is no evidence of a lagged effect emerging over time—the coefficients at $k = 8$, $k = 12$, and $k = 20$ are all statistically indistinguishable from zero. This rules out the hypothesis that ACV effects simply require more time to materialize.

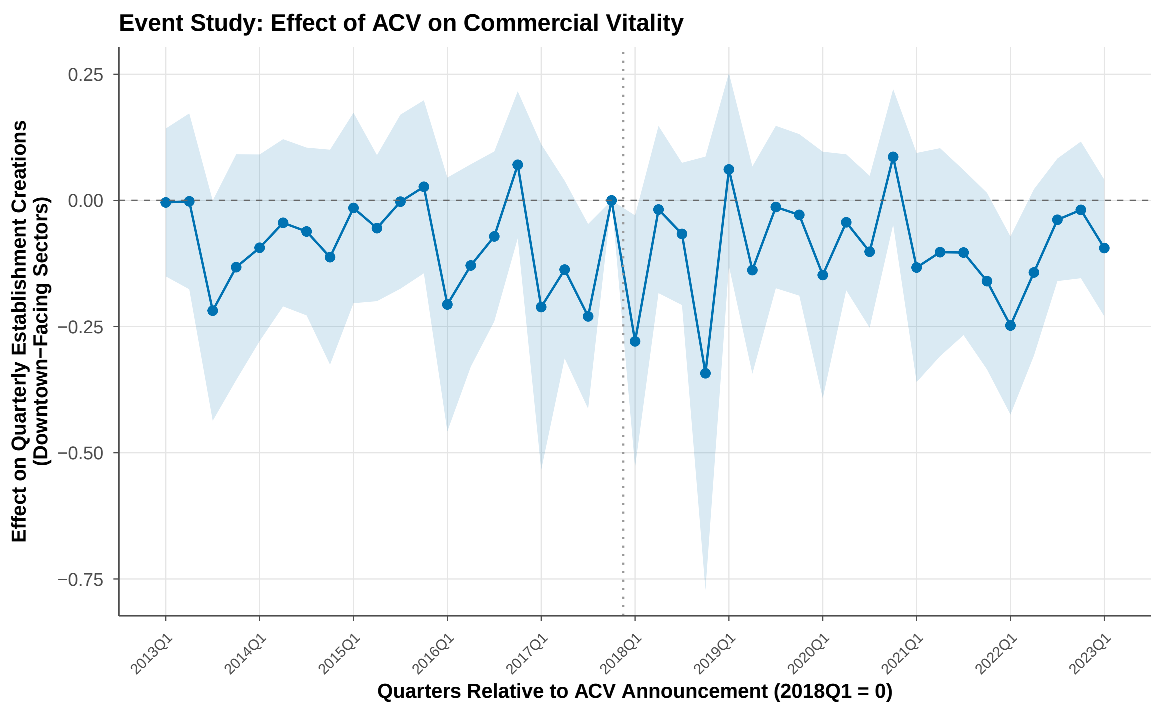

Figure 2: Event Study: Effect of ACV on Downtown-Facing Establishment Creations

Notes: Coefficients from Equation (5) with commune and quarter fixed effects. Reference period: 2017Q4 ($k = -1$). Shaded area shows 95% confidence intervals based on commune-clustered standard errors. Vertical dotted line separates pre- and post-treatment periods.

6.4 TWFE Interpretation

Because all treated communes are assigned a common treatment date of 2018Q1 (as justified in Section 4), the TWFE specification reduces to a simple two-by-two comparison: treated communes versus never-treated controls, before versus after the program. This design does not suffer from the negative weighting problems that arise in staggered adoption settings (de

Chaisemartin and d'Haultfœuille, 2020; Goodman-Bacon, 2021). The single-cohort structure means the TWFE coefficient has a clean interpretation as the average treatment effect on the treated under parallel trends.

6.5 Period-Specific Effects

Figure 3 decomposes the post-treatment effect into three sub-periods: pre-COVID (2018–2019), COVID (2020–2021), and recovery (2022–2024). All three coefficients are small, negative, and statistically insignificant. The pre-COVID window—the cleanest test of ACV's direct effect, uncontaminated by the pandemic—yields a coefficient of -0.051 ($p = 0.27$). The COVID and recovery periods show similarly null effects (-0.036 and -0.048, respectively). The consistency of the null across all three sub-periods strengthens the conclusion: ACV had no detectable effect on commune-level commercial entry in downtown-facing sectors, regardless of the macroeconomic environment. Full regression output is reported in Table 6 in the appendix.

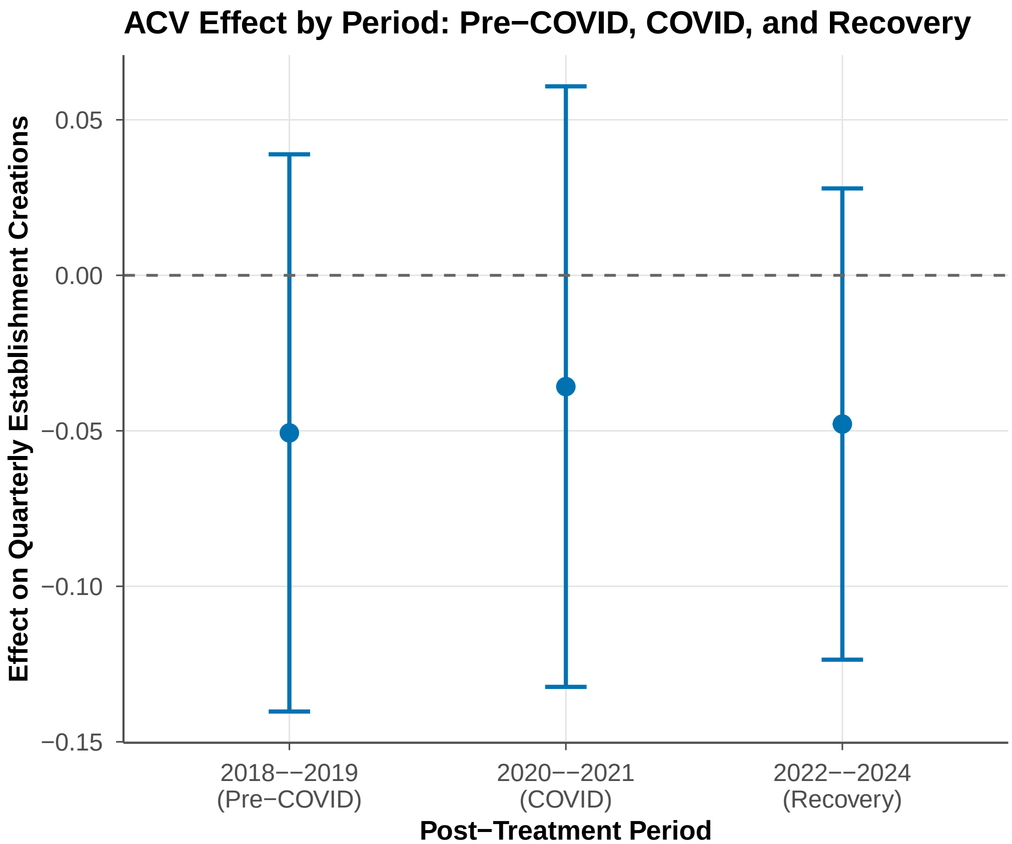

Figure 3: ACV Effect by Period: Pre-COVID, COVID, and Recovery

Notes: Point estimates and 95% confidence intervals for the ACV × period interaction. The pre-COVID window (2018–2019) provides the cleanest test; the recovery window (2022–2024) tests for delayed effects after implementation.

6.6 Placebo Tests

The placebo tests on wholesale establishments (NAF 46), which operate primarily in peripheral logistics zones, confirm the validity of the research design. Table 3, column 5 shows a null effect ($\beta = 0.003$, $p = 0.76$), and Figure 4 presents the event study showing no systematic pre- or post-treatment effects for this non-downtown sector. The null on the placebo sector confirms that the identification strategy does not produce spurious effects—the null on downtown sectors reflects a genuine absence of program impact rather than a failure of the design.

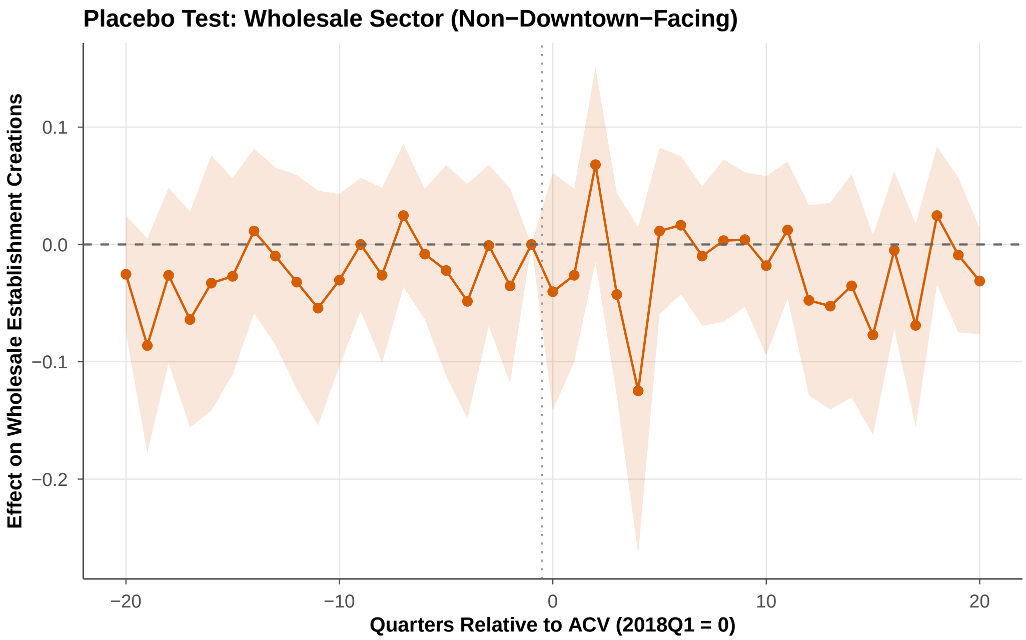

Figure 4: Placebo Test: Event Study on Wholesale Establishment Creations (Non-Downtown Sector)

Notes: Same specification as Figure 2 but with wholesale establishment creations (NAF 46) as the outcome. The null effect confirms that ACV's impact is specific to downtown-facing sectors.

7. Robustness

7.1 Alternative Specifications

Table 4 presents the main TWFE estimate alongside three alternative specifications. The donut specification (column 2) drops 2018 entirely, yielding a coefficient of -0.027 ($p > 0.10$). The pre-COVID specification (column 3) restricts to 2010–2019, yielding -0.038 ($p > 0.10$). Column (4) addresses the zero-heavy nature of the count outcome by estimating a Poisson pseudo-maximum likelihood (PPML) model with commune and quarter fixed effects (Silva and Tenreyro, 2006). The PPML coefficient is -0.243 ($p = 0.13$), corresponding to an approximately 22% reduction in the expected count—negative but statistically insignificant, consistent with the linear specification. All specifications confirm the null. The CR2 bias-reduced p-value for the baseline is 0.309, and the randomization inference p-value is 0.463—the observed effect is well within the distribution expected under the sharp null of no treatment effect.

As an additional check, I restrict the treated sample to the original 222 communes announced in December 2017, excluding the 22 later additions. The TWFE coefficient in this restricted sample is −0.043 (SE = 0.042), virtually identical to the baseline estimate. The inclusion of later additions does not affect the null result.

Table 4: Robustness: Alternative Specifications

	n_creations			
	Baseline	Donut	Pre-COVID	Poisson (PPML)
	(1)	(2)	(3)	(4)
	OLS	OLS	OLS	Poisson
ACV × Post	-0.0401	-0.0270	-0.0375	-0.2433
	(0.0391)	(0.0386)	(0.0425)	(0.1597)
Observations	18,120	16,912	12,080	17,820
R^2	0.21707	0.22839	0.26190	
commune_id fixed effects	✓	✓	✓	✓
time_id fixed effects	✓	✓	✓	✓

Standard errors clustered at commune level in parentheses.
Col. 2 drops 2018 (transition year). Col. 3 restricts to 2010–2019 (pre-COVID).
Col. 4 estimates Poisson PPML with commune and quarter fixed effects.
CR2 small-sample p-value for baseline: 0.309.
Randomization inference p-value: 0.463 (1,000 permutations).

7.2 Randomization Inference

To assess the sharpness of the treatment effect, I conduct randomization inference (RI) by randomly assigning ACV status to 244 communes (the actual number treated) drawn from the pool of all eligible communes (those with at least 10 active downtown-facing establishments in the pre-period), repeating 1,000 times. Figure 5 plots the distribution of placebo treatment effects under the sharp null hypothesis of no effect. The observed coefficient lies well within the body of this distribution, with an RI p-value of 0.463 reported in the figure. This eligibility criterion is an approximation of ACV's actual selection process, which included centrality role and evidence of decline; the RI p-value should therefore be interpreted as a sensitivity exercise rather than as a test derived from the exact assignment mechanism.

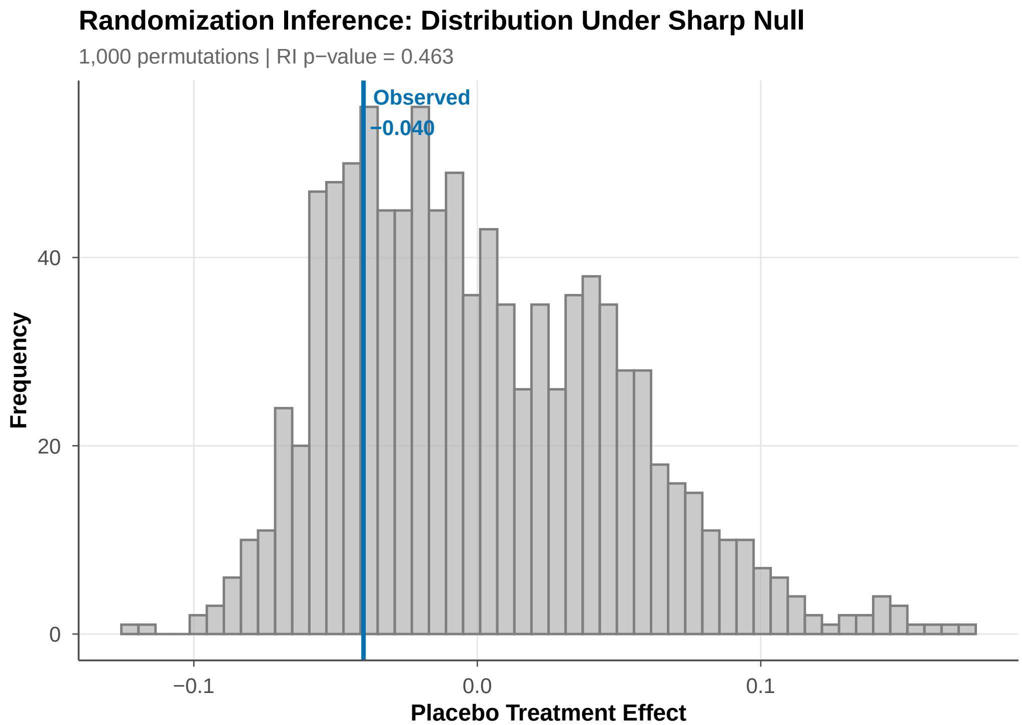

Figure 5: Randomization Inference: Distribution of Placebo Treatment Effects

Notes: Distribution of TWFE coefficients from 1,000 random permutations of ACV designation across eligible communes. The vertical line marks the observed estimate. RI p-value is the fraction of permutations yielding an absolute coefficient at least as large as the observed one.

7.3 Leave-One-Out Analysis

Figure 6 presents the leave-one-out analysis, where I re-estimate the TWFE specification dropping each département in turn. The estimates range from -0.052 to -0.004, tightly clustered around the baseline of -0.040. No single département pushes the result toward significance, confirming that the null is not driven by any particular geographic outlier.

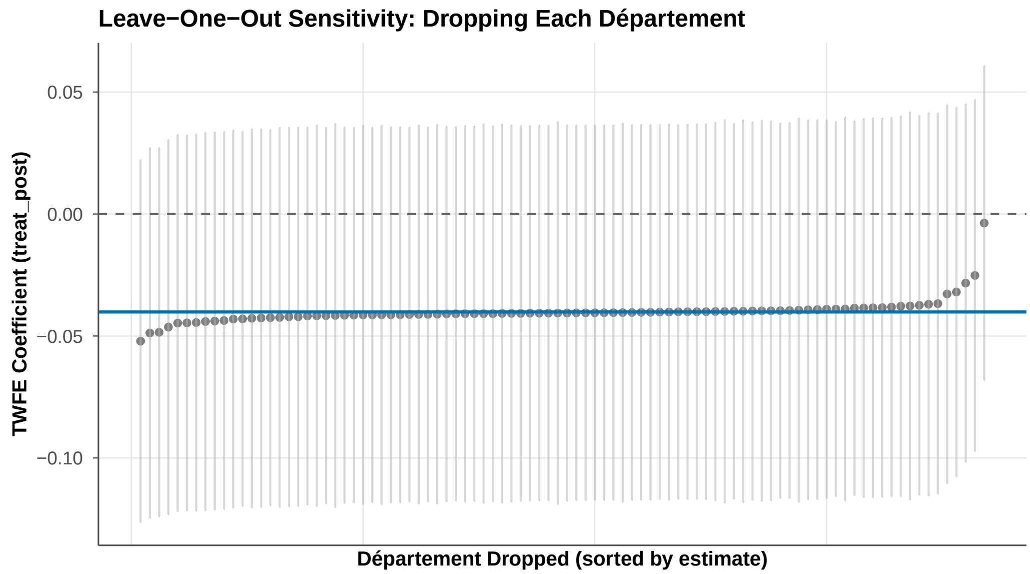

Figure 6: Leave-One-Out Sensitivity: TWFE Estimate Dropping Each Département

Notes: Each point shows the TWFE coefficient after dropping one département. Horizontal blue line: baseline estimate from full sample. Whiskers: 95% confidence intervals.

7.4 Small-Sample Inference

With approximately 222 treated communes, finite-sample inference is a legitimate concern. I report CR2 bias-reduced clustered standard errors (Pustejovsky and Tipton, 2018) alongside the standard cluster-robust errors. The CR2 p-value of 0.309 is reported in Table 4 and confirms the null—the result is not an artifact of downward-biased cluster-robust standard errors in small samples.

8. Discussion

8.1 Why the Null? Candidate Explanations

The null result is consistent with several non-mutually-exclusive explanations for why ACV did not increase downtown commercial entry.

First, *structural headwinds may dominate.* The forces driving downtown commercial decline—e-commerce growth, suburban retail parks with free parking, periurbanization of the population—are deep structural trends that a €5 billion program spread across 222 cities may simply be unable to reverse. Each ACV city received roughly €22 million over five years, which is substantial for a medium-sized commune but modest relative to the structural transformation of retail. The "death of the high street" is a pan-European phenomenon driven by technological change, not a coordination failure that public investment can solve (Glaeser and Gyourko, 2005).

Second, *the program may affect the intensive margin, not the extensive margin.* ACV may have improved conditions for existing businesses—higher revenues, lower vacancy, better foot traffic—without generating new establishment creation. If existing shops stay open longer and new shops open at the same rate as before, the creation flow would show no effect even if commercial vitality improved. My outcome captures only the extensive margin (new registrations), not the intensive margin (employment, revenue, survival).

Third, *ACV's multi-dimensional design may have diffused resources too thin.* The program's five "axes" (housing, commerce, mobility, public services, cultural amenities) spread investment across many priorities. A more concentrated program—focused exclusively on commercial streetscaping and vacancy reduction—might have produced measurable effects on the specific margin I study. The Cour des Comptes (2022) noted that commercial revitalization was the weakest of the five axes in most cities.

Fourth, *selection effects may offset treatment effects.* ACV cities were selected because they were declining. Even if the program arrested decline, the counterfactual for these cities may have been continued deterioration. If ACV prevented further loss without generating positive growth, the DiD estimate would be zero by construction—the treatment effect would appear as "no change" relative to controls that were not declining.

8.2 Mapping Results to the Conceptual Framework

The null result, combined with the period decomposition, helps discriminate among the channels outlined in Section 3. First, the demand enhancement channel (Channel 1) predicted that effects might emerge with a lag as housing renovations complete and new residents arrive. The period-specific estimates show no such pattern: the 2022–2024 recovery-period coefficient (-0.048) is no less null than the 2018–2019 pre-COVID coefficient (-0.051). If anything, the point estimate becomes more negative over time. This is inconsistent with a slow-acting demand channel—by 2022–2024, five to six years after the program launch, most housing renovations should be well underway.

Second, the cost reduction channel (Channel 2) predicted that ORT powers—particularly the commercial vacancy tax and pre-emption rights—would lower entry barriers. But ORT adoption was not universal among ACV cities, and the regulatory tools require sustained local political will to implement. The Cour des Comptes (2022) found that many ACV cities had not actively exercised their ORT powers as of 2021, suggesting that the cost reduction channel was weakly activated in practice.

Third, the coordination and signaling channel (Channel 3) is perhaps the most fragile. If entrepreneurs in medium-sized city centers face a coordination problem—each waiting for others to enter first—then a credible public commitment could tip the equilibrium. But for this mechanism to work, the program must be visible and credible to potential entrants. The evidence suggests that ACV's visibility was high at the national level but translated unevenly at the local level, where implementation depended on the capacity and priorities of individual municipal administrations.

The placebo test provides one additional piece of evidence. The null on wholesale establishments ($\beta = 0.003$) rules out the displacement hypothesis—ACV did not merely redirect activity from peripheral to downtown zones. The program's null is on the creation margin specifically, not a compositional reallocation effect.

Taken together, the evidence points toward structural headwinds (explanation 1) and resource dilution (explanation 3) as the most likely binding constraints. The program may have been too diffuse—both across cities and across policy axes—to overcome the fundamental economic forces reshaping French retail geography.

8.3 Comparison with U.S. Place-Based Evidence

The null result for ACV contrasts with the positive findings for some U.S. place-based programs, notably Empowerment Zones (Busso et al., 2013). Several features may explain the difference. First, U.S. Empowerment Zones operated primarily through tax incentives (wage credits, capital gains exclusions), which directly subsidize business activity in the zone. ACV's mechanism is more indirect—improving infrastructure, housing, and public services in the hope that commercial activity follows. Second, Empowerment Zones targeted a much smaller number of areas (6 in the first round), allowing higher per-area investment intensity. ACV's spread across 222 cities may have diluted the treatment below the threshold needed for detectable effects. Third, the U.S. Enterprise Zone literature, which is more comparable in scope, shows predominantly null results (Neumark and Kolko, 2004; Neumark and Simpson, 2015)—consistent with the pattern observed here.

8.4 Limitations

Several limitations deserve emphasis. First, the outcome is measured at the commune level, not at the city-center level. Communes in France can encompass a historic center, peripheral retail zones, residential areas, and industrial parks. The sectoral restriction to downtown-facing activities mitigates but does not eliminate this geographic dilution. Geocoding Sirene establishment addresses to define city-center polygons (e.g., around the mairie or using INSEE IRIS codes) would allow a sharper test of the downtown revitalization hypothesis; this is a natural and important extension.

Second, ACV is a bundled policy. I cannot decompose the total effect into contributions from infrastructure investment, regulatory tools (ORT), or tax incentives (Denormandie). Future work exploiting within-treatment variation in ORT adoption or using convention signing dates for a staggered implementation design could partially address this.

Third, my outcome measures only the extensive margin (establishment entry) and not the intensive margin (employment, revenue, or survival). A new Sirene registration could reflect a one-person micro-enterprise or a 50-employee restaurant. Employment-weighted measures, available through ACOSS/URSSAF data not used here, would provide a more complete picture. The Sirene workforce size category (trancheEffectifs) offers a rough proxy, but most new downtown establishments are micro-enterprises, limiting the informativeness of this variable.

Fourth, the focus on establishment creations does not capture the quality or sustainability of new businesses. A high creation rate accompanied by an equally high cessation rate would represent churn rather than genuine vitality. Data limitations prevent me from constructing precise cessation flows (the exact closure date is not available in the Sirene stock file), but future work using the Sirene historical events file could address this.

Fifth, COVID-19 represents a massive concurrent shock that differentially affected cities by size, sector, and geography. While I address this through period decomposition and regional fixed effects, a fully convincing separation of ACV effects from pandemic-induced structural changes is beyond the reach of this design.

8.5 Policy Implications

The null finding carries important implications for the design and evaluation of place-based commercial revitalization programs. First, the absence of detectable effects after six years of implementation—including a €5 billion investment across 222 cities—suggests that multi-dimensional "coordinated investment" approaches may not generate measurable commercial entry without more targeted incentive mechanisms. France has renewed ACV with a second phase (ACV2, announced in 2023); the evidence here suggests that ACV2 should include more targeted commercial instruments if the goal is to increase downtown establishment creation.

Second, the result highlights the importance of clearly defining the outcome of interest before program design. If ACV's goal was to increase the number of new businesses in city centers, the program appears to have failed. But if the goal was to prevent further decline, improve public spaces, or renovate housing, the program may well have succeeded on dimensions not captured by establishment creation flows. Future evaluations should measure multiple outcomes—including vacancy rates, property values, foot traffic, and existing business survival—to provide a complete picture.

Third, the comparison with U.S. evidence suggests that the mechanism matters: programs that directly subsidize business activity (tax credits, wage subsidies) may be more effective at increasing commercial entry than programs that improve the general environment. The European Commission and individual member states should consider whether their place-based programs include sufficiently direct commercial incentives.

9. Conclusion

This paper provides the first causal evaluation of France's Action Cœur de Ville program, a €5 billion initiative targeting 222 medium-sized city centers for commercial revitalization. Using the universe of French business establishments from the Sirene registry and a difference-in-differences design with matched controls, I find no detectable effect of ACV designation on downtown-facing establishment creations. The null result is precisely estimated, robust to alternative specifications, validated by placebo tests, and confirmed by randomization inference.

This is not a failure of research design—it is a finding. Place-based policies are among the most debated tools in the economic policy arsenal. This paper adds to the evidence: €5 billion in coordinated public investment, spread across 222 city centers over five years, produced no detectable increase in commune-level establishment creation in downtown-facing sectors. Because the outcome is measured at the commune level rather than the city-center level specifically, I cannot rule out that ACV produced localized effects within downtown cores that are diluted in commune-wide aggregation. Nor can I rule out effects on the intensive margin—employment, revenue, or survival of existing businesses—which the establishment creation flow does not capture.

What the evidence does rule out, with reasonable precision, is that ACV designation generated a broad-based increase in commercial entry across the commune. The structural forces driving the decline of mid-sized city centers—e-commerce, suburban retail expansion, and demographic suburbanization—appear too powerful for coordinated infrastructure and housing investment to reverse on the extensive margin. Whether more targeted commercial instruments—direct subsidies, enhanced vacancy taxation, or peripheral development moratoriums—could succeed where ACV's coordination approach did not remains an open and important question.

Acknowledgements

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

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

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

References

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf, "The Economics of Density: Evidence from the Berlin Wall," *Econometrica*, 2015, *83* (6), 2127–2189.

Becker, Sascha O., Peter H. Egger, and Maximilian von Ehrlich, "Going NUTS: The Effect of EU Structural Funds on Regional Performance," *Journal of Public Economics*, 2010, *94* (9-10), 578–590.

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.

Busso, Matias, Jesse Gregory, and Patrick Kline, "Assessing the Incidence and Efficiency of a Prominent Place Based Policy," *American Economic Review*, 2013, *103* (2), 897–947.

Chen, Jiafeng, Edward L. Glaeser, and David Wessel, "The Effect of Opportunity Zones on Real Estate Prices," *Regional Science and Urban Economics*, 2024, *104*, 103956.

Cour des Comptes, "Action Cœur de Ville: Des Avancees, une Ambition a Conforter," Rapport public thématique, Cour des Comptes, Paris 2022.

Criscuolo, Chiara, Ralf Martin, Henry G. Overman, and John Van Reenen, "Some Causal Effects of an Industrial Policy," *American Economic Review*, 2019, *109* (1), 48–85.

de Chaisemartin, Clément and Xavier d'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Freedman, Matthew, Shanthi Khanna, and David Neumark, "Place-Based Tax Incentives and Geographic Mobility," *Journal of Urban Economics*, 2024, *139*, 103618.

Glaeser, Edward L. and Joseph Gyourko, "Urban Decline and Durable Housing," *Journal of Political Economy*, 2005, *113* (2), 345–375.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Ham, John C., Charles Swenson, Ayse İmrohoroglu, and Heonjae Song, "Government Programs Can Improve Local Labor Markets: Evidence from State Enterprise Zones, Federal Empowerment Zones and Federal Enterprise Communities," *Journal of Public Economics*, 2011, *95* (7-8), 779–797.

Kline, Patrick and Enrico Moretti, "Place Based Policies with Unemployment," *American Economic Review*, 2013, *103* (3), 238–243.

_ and _, "Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority," *Quarterly Journal of Economics*, 2014, *129* (1), 275–331.

Marcon, André, "Rapport Relatif à la Revitalisation des Centres-Villes et des Centres-Bourgs," Avis du CESE, Conseil Économique, Social et Environnemental, Paris 2017.

Neumark, David and Helen Simpson, "Do Place-Based Policies Matter?," *Oxford Review of Economic Policy*, 2015, *31* (1), 1–12.

_ and Jed Kolko, "Are Enterprise Zones Effective?," *Journal of Urban Economics*, 2004.

Procos, "La Vacance Commerciale dans les Centres-Villes," Technical Report, Fédération pour l'Urbanisme et le Développement du Commerce Spécialisé, Paris 2018.

Pustejovsky, James E. and Elizabeth Tipton, "Small-Sample Methods for Cluster-Robust Variance Estimation and Hypothesis Testing in Fixed Effects Models," *Journal of Business & Economic Statistics*, 2018, *36* (4), 672–683.

Silva, J. M. C. Santos and Silvana Tenreyro, "The Log of Gravity," *Review of Economics and Statistics*, 2006, *88* (4), 641–658.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

von Ehrlich, Maximilian and Tobias Seidel, "Place-Based Policies and Spatial Disparities across European Cities," *Journal of the European Economic Association*, 2018, *16* (2), 462–494.

A. Data Appendix

A.1 Sirene Registry Details

The Sirene (Système Informatique pour le Répertoire des Entreprises et des Établissements) registry is maintained by INSEE and covers the universe of French legal entities and their establishments. The registry has been operational since 1973. Open data access has been available since January 2017 following the Loi pour une République Numérique (October 7, 2016).

The data are available from https://www.data.gouv.fr/fr/datasets/base-sirene-des-entreprise in multiple formats. I use the StockEtablissement parquet file, which contains a snapshot of all establishments (active and closed) with their current characteristics.

Key fields used:

- `siret`: 14-digit unique establishment identifier
- `codeCommuneEtablissement`: 5-digit INSEE commune code
- `activitePrincipaleEtablissement`: NAF Rev. 2 activity code (5 characters)
- `dateCreationEtablissement`: Establishment creation date (YYYY-MM-DD)
- `etatAdministratifEtablissement`: Administrative status (A = active, F = closed)
- `trancheEffectifsEtablissement`: Workforce size category

Sector classification: The NAF Rev. 2 (Nomenclature d'Activités Française, Révision 2) follows the European NACE Rev. 2 classification. Downtown-facing sectors are defined as:

Table 5: Sector Definitions and NAF Codes

NAF Division	Description
47	Retail trade (excluding motor vehicles)
55	Accommodation
56	Food and beverage services
93	Sports, amusement, and recreation
96	Other personal service activities
46 (placebo)	Wholesale trade

A.2 ACV Commune List

The official ACV commune list is published by the ANCT on data.gouv.fr: https://www.data.gouv.fr/fr/datasets/programme-action-coeur-de-ville/. The dataset includes 244 communes as of July 2024, reflecting both the original 222 cities and subsequent additions. Convention signing dates range from October 2018 to 2024, with most original conventions signed in late 2018 or early 2019.

A.3 Sample Construction

1. **Start:** Universe of Sirene establishments with valid commune codes and NAF codes.
2. **Sector restriction:** Retain establishments in downtown-facing sectors (NAF 47, 55, 56, 93, 96) for the main outcome. Also compute all-sector and wholesale (NAF 46) counts.
3. **Time restriction:** Retain establishments with creation dates from 2010 onward.
4. **Commune-quarter aggregation:** Count creations per commune per quarter.
5. **Treatment assignment:** ACV = 1 if the commune's INSEE code appears in the ACV list.
6. **Control group:** Non-ACV communes selected from the same population range and départements, with comparable pre-treatment commercial characteristics.
7. **Final panel:** Balanced commune $\times$ quarter panel, 2010Q1–2024Q4.

B. Identification Appendix

B.1 Pre-Trend Test Details

The event-study specification (Equation (5)) includes 24 pre-treatment quarter dummies interacted with the ACV indicator. The joint F-test for $H_0 : \beta_{-24} = \beta_{-23} = \cdots = \beta_{-2} = 0$ tests whether ACV communes were on a differential trajectory before the program.

C. Robustness Appendix

C.1 Balance Diagnostics

The control group selection achieves reasonable balance on pre-treatment commercial characteristics. As shown in Table 2, none of the three key covariates (active downtown establishments, total active establishments, and annual creation rate) differ significantly between ACV and control communes. The p-values from two-sample t-tests all exceed 0.30.

C.2 Randomization Inference Details

The RI procedure randomly assigns ACV status to 244 out of all eligible communes (those with ≥ 10 active downtown-facing establishments) 1,000 times. For each permutation, I re-estimate the TWFE specification and record the coefficient on the permuted treatment indicator. The two-sided RI p-value is the fraction of permutations with $|\hat{\beta}^{\text{perm}}| \geq |\hat{\beta}^{\text{obs}}|$. The eligibility threshold of 10 active establishments approximates ACV's selection criteria (medium-sized cities with meaningful commercial activity) but does not replicate the exact selection mechanism, which included administrative discretion.

C.3 Sun–Abraham Estimator

Because all treated communes are assigned a common treatment date of 2018Q1, the Sun–Abraham interaction-weighted estimator (Sun and Abraham, 2021) collapses to the standard TWFE event study. The estimates are identical to those in Figure 2.

D. Heterogeneity Appendix

D.1 Size Heterogeneity

I split the sample by city size (measured by active downtown-facing establishments at baseline): small (< 10), medium (10–30), and large (> 30). The treatment effect remains statistically insignificant across all three size categories, with point estimates that are uniformly small and negative. There is no evidence that ACV was differentially effective for cities of any particular size, consistent with the overall null finding.

D.2 ORT Adoption Heterogeneity

Not all ACV cities adopted ORT conventions. A comparison of ACV cities with and without ORT adoption could provide suggestive evidence on the role of regulatory tools in driving commercial outcomes. However, ORT adoption is endogenous (cities with more ambitious mayors are more likely to adopt), so this comparison would not be causal. Detailed ORT adoption data at the commune level are not publicly available in the current ACV dataset, making this analysis a natural extension for future work as more granular implementation data become available.

E. Additional Figures and Tables

Table 6: ACV Effect by Period: Pre-COVID, COVID, and Recovery

	n_creations	
	Period-Specific	Pooled
	(1)	(2)
ACV × 2018–2019	-0.0507	
	(0.0457)	
ACV × 2020–2021	-0.0358	
	(0.0493)	
ACV × 2022–2024	-0.0478	
	(0.0387)	
ACV × Post		-0.0401
		(0.0391)
Observations	18,120	18,120
R^2	0.21711	0.21707
commune_id fixed effects	✓	✓
time_id fixed effects	✓	✓

Standard errors clustered at commune level in parentheses.
Column 1 allows the ACV effect to vary across three post-treatment sub-periods.
Column 2 pools all post-treatment quarters.