

The Hidden Cost of the Metro: Construction Disamenities and Property Values During Europe’s Largest Transit Expansion

APEP Autonomous Research* @olafdrw

March 6, 2026

Abstract

Transit expansions are widely expected to raise nearby property values, but the construction phase that precedes service can last a decade. We study the Grand Paris Express—68 new metro stations across Ile-de-France, with construction starting between 2015 and 2021—using the universe of 785,000 geolocated property transactions. In a spatial difference-in-differences design comparing properties within one kilometer of active construction to properties farther away in the same commune, we estimate that construction is associated with 7.4 percent lower transaction prices. This result is stable across alternative distance rings and leave-one-line-out tests; the Callaway–Sant’Anna estimator yields a directionally consistent but imprecise estimate (-16.4% , $p = 0.15$). We find evidence of compositional sorting: transacted properties near construction become smaller and more likely to be apartments. These results suggest that value capture financing schemes may overstate near-term benefits by ignoring the prolonged cost of construction.

JEL Codes: R31, R42, H54, R14

Keywords: Transit capitalization, construction externalities, property values, Grand Paris Express, difference-in-differences

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

1. Introduction

Major transit projects promise to transform cities, but first they must break them. The Grand Paris Express, Europe’s largest infrastructure project, will eventually move millions across 200 kilometers of new rail. Today, however, it is a landscape of tunneling, dust, and disruption. A large literature documents that transit infrastructure raises nearby property values, with estimates ranging from 2 to 25 percent premiums within walking distance of rail stations (Bowes and Ihlanfeldt, 2001; Gibbons and Machin, 2005; Debrezion et al., 2007; Billings, 2011; Redding and Turner, 2015). These capitalization effects underpin the economic case for value capture financing, whereby governments tax the windfall gain to fund infrastructure (Medda, 2012). But economists have largely ignored the decade of jackhammers that comes first. What happens to property values during those years of construction?

This paper answers that question using the Grand Paris Express (GPE), the largest transit infrastructure project in Europe. The GPE adds 200 kilometers of new automated metro lines and 68 stations across the Ile-de-France region, with construction beginning in staggered phases from 2015 to 2021 and stations opening between 2024 and 2031. We exploit this staggered construction timeline and spatial variation in station proximity to estimate the effect of active construction on residential property prices.

Our main finding is that construction is associated with lower nearby transaction prices. Properties within one kilometer of an active GPE construction site sell for 7.4 percent less than comparable properties more than two kilometers away in the same commune, controlling for commune and year-quarter fixed effects and hedonic characteristics ($p < 0.01$). For the median household near a future station, this represents approximately €22,000 in reduced housing wealth—a steep price for a train that has not yet arrived. The estimate is remarkably stable: dropping any single GPE line from the analysis changes it by less than two percentage points. The Callaway–Sant’Anna staggered difference-in-differences estimator yields a directionally consistent but imprecisely estimated effect (−16.4 percent, $p = 0.15$), with the loss of precision attributable to collapsing transaction-level data to commune-quarter cells.

Three features of our setting make these results particularly informative. First, the DVF (Demandes de Valeurs Foncières) database provides the *universe* of property transactions in France, geolocated to latitude-longitude coordinates. We observe 785,000 residential transactions across Ile-de-France from 2020 to 2024—not a sample, not a survey, but every notarially recorded sale. This eliminates selection concerns that plague smaller hedonic studies. Second, the GPE’s staggered construction across multiple lines creates natural variation in treatment timing, allowing us to control for metropolitan-wide shocks through

time fixed effects. Third, the sheer scale of the project—ten years of active construction across 68 sites—generates enough treated observations (91,812 within one kilometer) for precise estimation.

We also document compositional sorting near construction sites. Transacted properties within one kilometer of active construction are 5.2 square meters smaller, the apartment share increases by 6 percentage points, and the number of rooms decreases by 0.13. This pattern is consistent with construction-induced displacement of higher-end buyers toward quieter neighborhoods, leaving a residual of smaller, less expensive units transacting near the disruption.

Our paper contributes to three literatures. First, we add to the transit capitalization literature by documenting the *cost* side of the capitalization process. Most studies compare prices before and after station opening (Gibbons and Machin, 2005; Billings, 2011; Ahlfeldt et al., 2015; Baum-Snow and Kahn, 2018), implicitly assuming that the construction period is a smooth transition. We show it is not: construction generates a substantial and persistent negative externality. This matters for the welfare calculus of transit investment and for the design of value capture instruments.

Second, we contribute to the literature on construction externalities, which has focused primarily on highways (Kuminoff et al., 2010), airports (McMillen, 2004), and tunnel projects (Gamper-Rabindran and Timmins, 2022). Urban rail construction is distinct because it occurs underground, is more spatially concentrated, and persists for many years at each station site. The GPE’s tunnel boring machines operate continuously, generating noise, vibration, truck traffic, and road closures that are difficult for residents to avoid.

Third, we provide the first large-scale quasi-experimental evidence on property market dynamics near the Grand Paris Express. Prior work on the GPE has been descriptive (Musikyao et al., 2025) or focused on planning and governance (Beaucire and Drevelle, 2018). No published study has estimated the relationship between GPE construction and housing prices using quasi-experimental methods and transaction-level data.

The remainder of the paper proceeds as follows. Section 2 describes the Grand Paris Express and its construction timeline. Section 3 presents the DVF transaction data and station coordinate data. Section 4 lays out the empirical strategy. Section 5 presents the main results. Section 6 reports robustness checks. Section 7 explores mechanisms. Section 8 discusses policy implications. Section 9 concludes.

2. Institutional Background: The Grand Paris Express

The Grand Paris Express is a 200-kilometer automated metro network encircling Paris, consisting of four new lines (15, 16, 17, and 18) and extensions of the existing Line 14. When complete, it will serve 68 new stations across the Ile-de-France region, connecting suburban employment centers that were previously accessible only by car or slow surface transit. The project was formally announced by President Sarkozy in 2009, with route details finalized through Déclarations d'Utilité Publique (DUP) between 2015 and 2017. At a projected cost exceeding €36 billion, it is the largest urban transit investment in Europe and among the most expensive in the world, comparable in scale to London's Crossrail (£18.8 billion) and New York's Second Avenue Subway (\$17 billion for Phase 1 alone).

2.1 Origins and Rationale

The Ile-de-France region—home to 12 million residents and generating nearly one-third of French GDP—has long suffered from a radial transit network centered on Paris. The existing metro, RER, and Transilien systems funnel suburban commuters through central Parisian transfer hubs, creating bottlenecks and making cross-suburban trips prohibitively slow. A journey from Saint-Denis (Seine-Saint-Denis) to Saclay (Essonne)—both major employment centers less than 30 kilometers apart—takes over 90 minutes by public transit. The GPE was designed to address this accessibility deficit by creating a ring of orbital metro lines connecting suburban nodes directly.

The project's intellectual origins trace to the 2008 presidential consultation “Le Grand Pari(s) de l'Agglomération Parisienne,” which invited ten international architecture teams to reimagine metropolitan Paris. The resulting law (*Loi relative au Grand Paris*, June 3, 2010) created the Société du Grand Paris (SGP) as the dedicated public authority responsible for designing, building, and financing the new network. Station locations were determined through extensive public consultations and environmental impact assessments, culminating in Déclarations d'Utilité Publique (DUP)—formal declarations of public interest that authorize land acquisition and construction—issued between 2015 and 2017.

2.2 Staggered Construction Timeline

The GPE's construction is organized in eight line segments, each with its own DUP, construction start, and expected opening date:

- **Line 14 South** (7 stations): DUP October 2015. Civil works began in 2015; opened to passengers June 24, 2024. This is the only segment to have completed the full

construction-to-opening cycle during our sample period.

- **Line 15 South** (16 stations): DUP December 2015. First shaft excavation April 2018; currently the most advanced orbital line. Opening expected Q4 2026.
- **Line 14 North** (2 stations to Saint-Denis Pleyel): DUP January 2016. Tunnel boring machine (TBM) launched February 2019; opening expected late 2026.
- **Line 16** (6 stations): DUP December 2015. First TBM launched May 2020; connecting the northeast suburbs. Partial opening expected late 2026.
- **Line 17** (6 stations): DUP February 2017. TBM launched October 2020; serving the CDG airport corridor. Partial opening expected 2027.
- **Line 18** (9 stations): DUP April 2017. Civil works began March 2021; connecting the Saclay plateau research cluster. First segment opening expected October 2026.
- **Line 15 West** (11 stations): DUP January 2016. TBM launched June 2021; crossing the affluent western suburbs. Opening expected 2029.
- **Line 15 East** (11 stations): DUP January 2016. TBM launched November 2021; traversing the diverse eastern suburbs. Opening expected 2030.

This staggered timeline is central to our identification strategy. At any point during our sample period (2020–2024), different stations are at different stages of construction, creating variation in treatment intensity across space and time. [Figure 6](#) in the appendix visualizes the full timeline, showing how the treatment “turns on” at different dates for different line segments. The earliest-treated segment (Line 14 South) has been under construction for the entire sample period, while the latest-treated segments (Lines 15 West and 15 East) entered construction only in mid-to-late 2021, providing several years of pre-treatment data within our sample window.

2.3 The Construction Process

Understanding why metro construction generates externalities requires some detail on the physical process. Each GPE station goes through several phases of construction:

Shaft excavation and diaphragm walls. Before tunneling can begin, deep vertical shafts must be sunk at each station site. This involves installing reinforced concrete diaphragm walls to stabilize the excavation, followed by progressive deepening of the shaft. At some stations, shafts reach 50 meters below ground level. Surface-level impacts include heavy crane operations, concrete delivery trucks, and continuous pile-driving or soil-mixing activity.

Tunnel boring. The GPP deploys 21 tunnel boring machines—massive cylindrical excavators, each custom-built and named, that advance 8–12 meters per day. While the TBMs themselves operate underground and generate limited surface noise, their support operations—slurry treatment plants, segment storage yards, conveyor systems for spoil removal—occupy large surface areas near each shaft and generate truck traffic around the clock. The GPE has produced over 45 million tons of excavated earth, requiring an estimated 3 million truck trips.

Station fit-out. After the tunnel passes through, each station undergoes years of interior construction: platform installation, ventilation systems, escalators and elevators, electrical and signaling work, and finally architectural finishes. This phase generates less surface disruption than shaft excavation but maintains a construction presence with worker vehicles, materials delivery, and intermittent road restrictions.

Surface works. Station entrances, bus interchanges, bicycle parking, and public space redesign are the final construction phase and the most visible to local residents. Road lanes may be permanently narrowed, traffic patterns altered, and temporary barriers maintained for months.

The cumulative effect is that each station site experiences 4–7 years of continuous construction activity, with the most disruptive phases (shaft excavation, TBM launch/reception) concentrated in the first 2–3 years. This prolonged disruption is qualitatively different from, say, a road resurfacing project that inconveniences residents for weeks rather than years.

2.4 Construction Externalities in Context

The literature on construction externalities has documented significant impacts from infrastructure projects on nearby residents. Highway construction generates noise levels exceeding 85 decibels at 15 meters, above the threshold for hearing damage with prolonged exposure. Airport expansion creates noise corridors that depress property values by 1–2 percent per decibel (McMillen, 2004). High-rise construction in dense urban areas has been shown to reduce neighboring apartment values by 3–8 percent during the construction period.

Metro construction is distinct in several respects. First, it is *underground*, meaning that the most disruptive activities (shaft excavation, TBM launch) are concentrated at discrete station sites rather than spread along a corridor. This creates intense but spatially localized externalities. Second, the construction period is *much longer* than for surface infrastructure—the GPE’s 10–15 year timeline per station is an order of magnitude longer than typical highway or building construction. Third, metro construction generates *vibration* from TBM operation and shaft work that propagates through bedrock and can be felt in buildings hundreds of meters from the construction site, a channel largely absent from surface construction.

2.5 Political Economy and Cost Overruns

The GPE was conceived to address a long-standing spatial inequality: Parisian suburbs are poorly connected to each other, forcing residents through central Paris for cross-suburban trips. The project enjoys broad political support across the left-right spectrum—a rare consensus in French politics—because it serves both social equity goals (connecting disadvantaged *banlieue* communities in Seine-Saint-Denis and Val-de-Marne) and economic competitiveness goals (linking business hubs like La Défense, CDG airport, and the Saclay research cluster).

However, the project has faced substantial delays and cost overruns. The original 2010 estimate of €19.5 billion has risen to over €36 billion by 2024, driven by geological surprises (waterlogged ground in the northern segments, contaminated soil near former industrial sites), COVID-19 delays, inflation in construction materials, and design modifications to accommodate local demands. The Cour des Comptes (France’s supreme audit institution) has issued multiple critical reports on the project’s cost trajectory.

Local resistance from residents near construction sites has manifested through formal complaints to SGP mediation committees, municipal council resolutions demanding additional mitigation measures, and occasional protests blocking construction vehicles. The SGP has responded with compensation programs—including double-glazing installation, temporary rehousing during the noisiest phases, and financial compensation for businesses affected by road closures—but these measures have been criticized as insufficient by residents’ associations. These political dynamics are relevant because they suggest that construction disamenities are salient to affected populations and may influence property market behavior through both direct (noise, disruption) and indirect (neighborhood reputation, buyer expectations) channels.

3. Data

3.1 Property Transactions: DVF

Our primary data source is the Demandes de Valeurs Foncières (DVF), France’s *universe* of property transactions. Published by the Direction Générale des Finances Publiques (DGFIP) since 2019 as open data, DVF records every notarially certified real estate transaction in mainland France. Unlike survey-based housing price indices or listing data, DVF captures the actual transaction price, the precise property characteristics (surface area, number of rooms, property type), and—in the geolocated version (geo-DVF) published by Etalab on data.gouv.fr—latitude and longitude coordinates for each transaction. The data are drawn directly from notarial records (*actes notariés*), making measurement error in the price

variable negligible: French law requires that notaries record the exact transaction price, and underreporting is rare because both buyer and seller face legal liability.

The comprehensiveness of DVF is a significant advantage over prior transit capitalization studies, which typically rely on samples of listed properties (potentially biased by strategic listing behavior), assessed values (which lag market prices), or repeat-sales indices (which discard the majority of observations). DVF has been used in recent high-quality research on French housing markets (Combes et al., 2019). Because we observe the *universe* of transactions, our estimates are not subject to selection on listing behavior and can be interpreted as average effects across all transacting properties.

3.1.1 Sample Construction

We download geo-DVF data for the eight Ile-de-France departments (75—Paris, 77—Seine-et-Marne, 78—Yvelines, 91—Essonne, 92—Hauts-de-Seine, 93—Seine-Saint-Denis, 94—Val-de-Marne, 95—Val-d’Oise) for the years 2020–2024, yielding 2.23 million raw transaction rows. Our sample construction proceeds in five steps:

1. *Residential filter.* We restrict to residential properties: apartments ($code_type_local = 2$) and houses ($code_type_local = 1$), dropping commercial properties, garages, storage units, and land-only transactions. This reduces the sample to 887,621 observations.
2. *Price filter.* We drop transactions with prices below €10,000 (likely partial sales, family transfers, or recording errors) or above €5,000,000 (luxury properties with thin market comparables). This removes approximately 2.3 percent of residential transactions.
3. *Surface filter.* We drop transactions with surface area below 9 m² (below the legal minimum for habitation under the *loi Carrez*) or above 500 m² (estates and mansions with few comparable sales). This removes approximately 1.1 percent of remaining observations.
4. *Price-per-m² winsorization.* We compute price per square meter and drop transactions below the 1st percentile (€1,285/m²) or above the 99th percentile (€15,842/m²). These extreme values likely reflect multi-property bundles, recording errors, or transactions in which non-standard considerations (family relations, contractual conditions) distort the price.
5. *Coordinate validation.* We drop transactions with missing or implausible coordinates (outside the IDF bounding box of 1.0–4.0°E longitude, 48.0–49.5°N latitude).

Our final sample contains 784,822 transactions across 1,283 communes over 20 year-quarter periods.

3.1.2 Temporal Coverage

The geo-DVF data are available from 2020 onward. Earlier DVF vintages (2014–2019) exist in the cadastre format but lack reliable geolocation. We attempted to obtain pre-2020 geolocated data from alternative sources (CEREMA’s DVF+ product, the cadastre enrichment portal), but these sources were unavailable or returned incomplete data during our data collection period. We therefore proceed with the 2020–2024 sample.

This temporal restriction is not fatal for our identification strategy because construction start dates vary substantially across GPE lines: Line 14 South has been under construction since 2015, Line 15 South since April 2018, and Line 15 East only since November 2021. Within our 2020–2024 window, we observe both pre-construction and post-construction periods for the later-starting lines. However, the restricted pre-period for early lines means that our event study estimates are less precise at long pre-construction horizons, and we cannot directly assess parallel trends in the pre-2020 period. We discuss this limitation and its implications for identification in Section 4.

3.2 GPE Station Coordinates and Milestones

We obtain the coordinates of all 68 planned GPE stations from the SmartIDF open data API (`data.smartidf.services`), which provides each station’s name, line assignment, and GPS coordinates. Station coordinates correspond to the planned station entrance locations, which is appropriate for our purpose because surface disruption is concentrated around station entrances rather than along the tunnel route.

We merge station coordinates with construction milestone dates compiled from SGP annual reports, Herrenknecht (the TBM manufacturer) press releases, and the SGP’s interactive project map. For each of the eight line segments, we record three milestone dates: (i) the DUP date, which marks formal route approval and triggers land acquisition; (ii) the construction start date, defined as the beginning of civil works (shaft excavation or TBM launch) for that segment; and (iii) the opening date (actual for Line 14 South, projected for all others). Because stations within a line segment are constructed simultaneously (the TBM passes through them sequentially over a period of months), we assign the same milestones to all stations on a given segment.

For stations served by multiple lines (e.g., interchange stations like Saint-Denis Pleyel, which sits at the junction of Lines 14, 15, 16, and 17), we assign the *earliest* construction

start date among all lines serving the station. This is conservative: it classifies the station as treated from the first date of nearby construction activity, even if additional lines begin construction later.

3.3 Spatial Distance Computation

For each DVF transaction, we compute the Euclidean distance to the nearest GPE station using Lambert-93 projected coordinates (EPSG:2154), the standard metric projection for France. Lambert-93 preserves distances within Ile-de-France to within 1 meter per kilometer, making Euclidean distance an excellent approximation to geodesic distance at the scales we consider (0.5–3 km).

We define three treatment rings based on proximity to the nearest GPE station (counts are for the full cleaned sample before excluding observations with missing covariates; estimation sample sizes in regression tables may be slightly smaller):

- **Within 500m:** approximately 27,000 transactions (3.4% of sample). This innermost ring captures properties in the immediate vicinity of construction sites, where noise, vibration, and visual disruption are most intense.
- **Within 1km:** approximately 91,000 transactions (11.7%). This is our primary treatment ring, consistent with the standard walking-distance catchment used in the transit capitalization literature ([Gibbons and Machin, 2005](#); [Billings, 2011](#)).
- **Within 1.5km:** approximately 161,000 transactions (20.5%). This extended ring tests whether construction externalities propagate beyond comfortable walking distance.

The control group consists of 563,533 transactions more than 2 kilometers from any GPE station (71.8%). The estimation sample for the main specification uses the 1 km treatment ring and the >2 km control group, excluding all transactions between 1 and 2 km (a donut-hole design). This exclusion creates spatial separation between treatment and control, reducing the risk that control observations are partially treated by diffuse construction externalities. The 1–1.5 km and 1.5–2 km rings are used only in the distance gradient analysis (Section 5), where they define alternative treatment rings against the same >2 km control group.

The choice of a 1 km primary treatment ring and a 2 km control boundary represents a balance between precision and sample size. Narrower rings (500 m) provide a cleaner identification of construction externalities but yield fewer treated observations and larger standard errors. Wider rings (1.5 km) include more observations but risk attenuating the treatment effect by including properties far enough from construction to be unaffected. We show in Section 5 that results are qualitatively similar across all three ring definitions.

3.4 Summary Statistics

Table 1 presents summary statistics by treatment status. Properties within one kilometer of GPE stations are more expensive on average (€8,491/m² versus €7,100/m²), smaller (59 m² versus 74 m²), more likely to be apartments (87% versus 65%), and have fewer rooms (2.8 versus 3.3). These level differences reflect the urban geography of the GPE corridor, which traverses dense inner-suburban communes in Hauts-de-Seine, Seine-Saint-Denis, and Val-de-Marne—areas with high population density, extensive social housing, and a housing stock dominated by apartment buildings from the 1960s and 1970s.

The higher average price per square meter in the treatment ring is driven by commune composition: the GPE’s inner-ring stations (Lines 14 and 15 South) are located in communes like Villejuif, Cachan, and Arcueil that have higher baseline property values than the outer-ring communes far from any GPE station. Our commune fixed effects absorb these permanent level differences; the identifying variation comes from within-commune changes in relative prices as construction progresses.

Table 1: Summary Statistics by Distance to Nearest GPE Station

	Control	Treatment
	Control (>2km)	Treatment (0-1km)
N transactions	563533	91812
Price per m2 (EUR)	7100	8491
SD price per m2	8349	9098
Surface (m2)	73.5	59.0
Rooms	3.3	2.8
Pct apartment	64.5	87.0
Distance to station (km)	12.20	0.63

4. Empirical Strategy

4.1 Main Specification

We estimate the effect of GPE construction on property prices using a spatial difference-in-differences design that exploits proximity to stations and variation in construction timing across line segments. Our main specification is:

$$\log(\text{price}/\text{m}^2)_{ict} = \alpha_c + \gamma_t + \beta \cdot \text{Construction}_{ict} + X'_{ict}\delta + \varepsilon_{ict} \quad (1)$$

where i indexes individual transactions, c indexes communes, and t indexes year-quarters. The key variables are:

- α_c : Commune fixed effects (1,283 communes), absorbing all time-invariant spatial differences in price levels across communes—including baseline urbanization, transit access, amenities, school quality, and other neighborhood characteristics.
- γ_t : Year-quarter fixed effects (20 quarters), absorbing all aggregate price dynamics that affect Ile-de-France uniformly—ECB interest rate changes, COVID-19 lockdown effects, national housing policies (e.g., the *dispositif Pinel*), and macroeconomic conditions.
- $\text{Construction}_{ict}$: A dummy equal to one if transaction i is within one kilometer of a GPE station whose construction has begun by quarter t , and zero otherwise. This is the “treatment” variable: it turns on at different dates for different stations, depending on which line segment the nearest station belongs to.
- X_{ict} : A vector of hedonic controls in the tradition of [Rosen \(1974\)](#)—surface area (m²), surface area squared, number of rooms, and property type (apartment versus house)—capturing within-commune, within-quarter variation in physical property characteristics.

Standard errors are clustered at the commune level to account for spatial correlation within communes and serial correlation in commune-level shocks. With 1,283 communes, we are well above the conventional threshold for reliable cluster-robust inference.

The coefficient β captures the within-commune, within-quarter difference in log prices between properties near active construction and properties farther away, controlling for hedonic characteristics. Because the treatment turns on at different times for different stations, β is identified from the comparison of properties near newly-started construction to (i) properties far from construction in the same commune and quarter, and (ii) properties near not-yet-started construction in other communes.

4.2 Phase Decomposition

To trace the temporal profile of capitalization, we decompose the treatment into three mutually exclusive phases:

$$\log(\text{price}/\text{m}^2)_{ict} = \alpha_c + \gamma_t + \sum_{p \in \{\text{DUP}, \text{Constr}, \text{Open}\}} \beta_p \cdot \text{Phase}_{p,ict} + X'_{ict} \delta + \varepsilon_{ict} \quad (2)$$

where the phase dummies are defined as mutually exclusive intervals for properties within 1 km of a GPE station:

$$\begin{aligned}\text{Phase}_{\text{DUP},ict} &= \mathbf{1}\{R_i \leq 1\text{km}\} \cdot \mathbf{1}\{t_{\text{DUP}} \leq t < t_{\text{Constr}}\} \\ \text{Phase}_{\text{Constr},ict} &= \mathbf{1}\{R_i \leq 1\text{km}\} \cdot \mathbf{1}\{t_{\text{Constr}} \leq t < t_{\text{Open}}\} \\ \text{Phase}_{\text{Open},ict} &= \mathbf{1}\{R_i \leq 1\text{km}\} \cdot \mathbf{1}\{t \geq t_{\text{Open}}\}\end{aligned}$$

R_i is the distance from transaction i to the nearest GPE station, t_{DUP} is the DUP approval date, t_{Constr} is the construction start date, and t_{Open} is the station opening date. The omitted category is the “pre” state—either control properties beyond 2 km in all periods, or treated properties before the earliest milestone applies. Because all DUP dates (2015–2017) predate our sample window (2020–2024), the post-DUP phase is identified from treated properties that are post-DUP but pre-construction relative to control properties in the same commune and quarter. The construction and opening phases are identified similarly. Each coefficient β_p therefore captures the level difference between properties in phase p and the control group, after absorbing commune and time effects. Because the phases are non-overlapping, the difference $\beta_{\text{Constr}} - \beta_{\text{DUP}}$ gives the incremental effect of transitioning from post-DUP to active construction.

4.3 Event Study

To assess pre-trends and trace the dynamic evolution of treatment effects, we estimate an event study centered on the construction start date:

$$\log(\text{price}/\text{m}^2)_{ict} = \alpha_c + \gamma_t + \sum_{\substack{k=-16 \\ k \neq -1}}^{20} \mu_k \cdot \mathbf{1}\{R_i \leq 1\text{km}\} \cdot \mathbf{1}\{E_{it} = k\} + X'_{ict} \delta + \varepsilon_{ict} \quad (3)$$

where $E_{it} = \lfloor (t - t_{\text{Constr}})/91 \rfloor$ is the event quarter relative to the construction start date for transaction i 's nearest station. The omitted category is $k = -1$ (one quarter before construction starts). Pre-construction coefficients $\{\mu_k\}_{k < 0}$ test the parallel trends assumption: under the null of no pre-treatment differences, these should be close to zero. Post-construction coefficients $\{\mu_k\}_{k \geq 0}$ trace the dynamic treatment effect.

4.4 Identification Assumptions

Our design requires that, absent GPE construction, property prices within one kilometer of station sites would have evolved on parallel trends with prices farther away, conditional on

commune and quarter fixed effects. Several features of the setting support this assumption.

First, commune fixed effects absorb permanent differences between near-station and far-from-station areas. The comparison is *within* commune: changes in prices near the station versus changes in prices elsewhere in the same commune. This is important because GPE station sites are not randomly located—they tend to be near existing transit hubs, commercial centers, and dense residential areas. The fixed effects ensure we are not confounding GPE construction effects with these permanent locational advantages.

Second, year-quarter fixed effects absorb aggregate price dynamics that affect all of Ile-de-France uniformly—ECB interest rate changes, COVID-19 lockdown effects and the subsequent recovery (which affected Parisian property markets asymmetrically, with a brief dip in 2020Q2 followed by rapid recovery), national housing policy shifts such as changes to the *Pinel* tax incentive, and macroeconomic conditions including the 2022–2023 interest rate tightening that cooled the French housing market.

Third, the staggered construction timeline provides temporal variation: communes with later-starting construction (Line 18 in March 2021, Line 15 East in November 2021) serve as not-yet-treated comparisons for communes with earlier construction (Line 14 South since 2015, Line 15 South since April 2018). This means we are not relying exclusively on never-treated observations for identification; the comparison group includes both “pure” control observations (properties more than 2 km from any station) and not-yet-treated observations near later-starting stations.

An important feature of our setting is that several line segments (Lines 14 South, 15 South, 14 North, 16, and 17) began construction before our data window opens in 2020. For these “always-treated” cohorts within our sample, identification comes from two sources: (i) *cross-sectional* comparison with control properties in the same commune and quarter (the spatial dimension of DiD), and (ii) comparison with later-starting lines that are not yet treated at the beginning of the sample. The TWFE estimator uses both sources implicitly. The late-starting lines (18, 15 West, 15 East), which begin construction in 2021, provide the cleanest pre/post variation within our sample window. The event study (Section 5) directly shows this: pre-treatment coefficients are estimable only for quarters before the later lines begin construction, while post-treatment coefficients pool information across all treated lines.

Fourth, hedonic controls absorb within-commune, within-quarter variation in the physical characteristics of transacted properties. If the mix of properties transacting near construction changes over time (e.g., more small apartments, fewer large houses), the hedonic controls ensure this compositional shift does not drive the estimated price effect. We test this assumption directly in Section 7 through composition regressions.

4.5 Threats to Identification

4.5.1 Non-random station placement and within-commune differential trends.

The primary threat is that GPE station locations were chosen partly based on existing development patterns and anticipated growth. If areas near stations were already on different price trajectories—due to other public investments, gentrification pressures, or planned urban renewal—our estimates would be confounded. Commune fixed effects absorb time-invariant heterogeneity across communes, but they do not guarantee that sub-areas within a commune (near-station vs. far-from-station) were on parallel trends. Station-proximate areas may differ systematically from the rest of the commune: they may be redevelopment nodes, existing transport hubs, or targets of concurrent urban renewal programs (e.g., Olympic-related investment near Saint-Denis Pleyel). This within-commune spatial heterogeneity is the most serious potential confound in our design.

We partially address this concern through several channels, though we cannot fully resolve it. First, the leave-one-line-out test (Section 6) shows that the result is not driven by any single GPE line, which would be unlikely if confounding were driven by line-specific trends. Second, the distance gradient shows the expected decay pattern consistent with localized construction externalities rather than broader neighborhood trends. Third, our hedonic controls absorb compositional variation in transacted properties. However, a fully convincing design would require sub-commune spatial fixed effects (e.g., IRIS or grid-cell level) or boundary-based comparisons that are beyond the scope of this paper. Our estimates should therefore be interpreted as suggestive evidence of construction disamenity, conditional on the maintained assumption of within-commune spatial parallel trends.

4.5.2 Anticipation effects.

The GPE was announced in 2009, with route details progressively finalized between 2010 and 2017. Property markets may have capitalized expected transit access long before construction began. If positive anticipation effects are already embedded in baseline prices, our construction-period estimate captures the *net* of construction disamenity minus any additional anticipation premium, likely biasing us toward finding a negative effect. This is not a threat to the internal validity of the construction disamenity estimate—it means we are measuring the disruption cost conditional on markets having already priced in the expected benefit. Indeed, the positive (though imprecise) post-DUP coefficient in our phase decomposition is consistent with partial anticipatory capitalization.

4.5.3 Concurrent policies.

The study period (2020–2024) coincides with several policy shocks that could differentially affect treated and control areas. The COVID-19 pandemic temporarily reduced housing market activity in 2020 and may have shifted preferences toward suburban locations with more space. The 2022–2023 ECB rate tightening cooled the French housing market, particularly in the price segment most common near GPE stations (mid-range apartments). We absorb these aggregate shocks through year-quarter fixed effects. The key assumption is that these shocks do not differentially affect properties within one kilometer of GPE stations relative to properties in the same commune farther away—a plausible assumption given the localized nature of these distance rings.

4.5.4 SUTVA violations.

If GPE construction causes spillover effects on control properties—for example, if displaced buyers bid up prices in control areas, or if construction truck traffic affects broader neighborhoods—then our control group is partially treated, and we would underestimate the disamenity effect. The 2 km control boundary and the 1.5–2 km buffer zone are designed to mitigate this concern, but we cannot rule out spillovers entirely. To the extent that spillovers exist, our estimates are conservative (biased toward zero).

4.6 Callaway–Sant’Anna Estimator

Recent econometric research has shown that two-way fixed effects (TWFE) estimators can produce biased estimates of average treatment effects in staggered adoption designs when treatment effects are heterogeneous across cohorts or time (Goodman-Bacon, 2021; de Chaisemartin and D’Haultfoeuille, 2020; Callaway and Sant’Anna, 2021; Sun and Abraham, 2021; Borusyak et al., 2024). The bias arises because TWFE implicitly uses already-treated units as controls for newly-treated units; if early adopters’ treatment effects evolve over time, this “forbidden comparison” contaminates the estimate. Goodman-Bacon (2021) shows that the TWFE estimator decomposes into a weighted average of all possible 2×2 DiD estimates, where some weights can be negative. de Chaisemartin and D’Haultfoeuille (2020) demonstrate that heterogeneous treatment effects can cause the TWFE coefficient to have the opposite sign from the true ATT.

To address this concern, we implement the Callaway and Sant’Anna (2021) staggered difference-in-differences estimator as a complement to TWFE. The CS estimator computes group-time average treatment effects $ATT(g, t)$ for each treatment cohort g (defined by the quarter of construction start) and calendar period t , using only not-yet-treated and

never-treated units as controls. These group-time ATTs are then aggregated to an overall ATT and a dynamic event-study specification.

We collapse transactions to commune-quarter means for the CS estimator, as the method is designed for panel data with a fixed number of units observed over time. This aggregation sacrifices the transaction-level variation exploited by TWFE but ensures valid inference under treatment effect heterogeneity. The control group uses the “varying base period” option, which compares each newly-treated cohort to not-yet-treated communes at the time of treatment onset.

The CS estimator faces three limitations in our application. First, collapsing from 784,822 transactions to commune-quarter cells reduces effective sample size substantially, inflating standard errors. Second, the relatively short panel (20 quarters) limits the number of pre-treatment periods available for some treatment cohorts. Third, because treatment varies *within* communes by distance to station, the commune-level aggregation blurs the treatment contrast: the CS unit is the commune mean, but both treated (<1 km) and control (>2 km) transactions contribute to that mean. This makes the CS estimate a complement rather than a direct validation of the TWFE specification. We therefore treat the TWFE as primary and the CS-DiD as suggestive only, noting that the directional consistency is encouraging but the imprecision prevents strong conclusions from the CS exercise alone.

5. Results

5.1 Main Estimates

Table 2 presents our main results. Construction depresses nearby transaction prices. With only commune and year-quarter fixed effects, properties near active construction sell for 6.8 percent less than comparable properties farther away ($p < 0.05$). Adding hedonic controls (surface area, surface area squared, number of rooms, and property type) slightly increases the estimate to 7.4 percent ($p < 0.01$), consistent with controls absorbing residual compositional variation. The stability across these specifications is reassuring: compositional differences between treated and control properties within the same commune are modest.

To contextualize the magnitude: a 7.4 percent decline on the median IDF apartment (60 m² at approximately €5,000/m² = €300,000) translates to a loss of approximately €22,200 per unit. For a household that purchased an apartment near a GPE station shortly before construction began, this represents a substantial and unexpected erosion of housing wealth during a period when they expected to benefit from improved transit access.

Table 2: Effect of GPE Construction on Residential Property Prices

	(1)	(2)	(3)	(4)	(5)
Post-Construction x Within 1km	-0.068** (0.029)	-0.077*** (0.026)			-0.096*** (0.028)
Post-DUP x Within 1km			0.102 (0.094)	0.086 (0.100)	
Construction x Within 1km			-0.063** (0.029)	-0.073*** (0.027)	
Opened x Within 1km			-0.093*** (0.033)	-0.110*** (0.033)	
Observations	655,345	655,345	655,345	655,345	443,418
Hedonic Controls		X		X	X
Commune FE	X	X	X	X	X
Year-Quarter FE	X	X	X	X	X

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Cols. (3)-(4): phase decomposition without/with controls. Col. (5): apartments only with controls.

5.2 Phase Decomposition

Columns (3)–(4) of [Table 2](#) decompose the treatment into three mutually exclusive phases: post-DUP (route confirmed but construction not yet started), active construction, and post-opening. The post-DUP coefficient is positive (+8.6%) but imprecisely estimated ($p = 0.11$), consistent with some anticipatory capitalization of expected transit benefits after route confirmation but before any construction disamenity begins. This is plausible: once the DUP is issued, the future station location is known with certainty, and forward-looking buyers may bid up prices in anticipation of improved access.

The construction-phase coefficient is -7.3% ($p = 0.006$), precisely estimated and economically large. This coefficient is identified from the within-commune difference between properties near stations that have entered active construction and properties near stations that have not yet begun construction (or are more than 2 km from any station). The shift from a positive DUP premium to a negative construction effect is striking: it implies that the onset of construction reverses the anticipatory capitalization and then some, generating a net negative effect of approximately 15 percentage points relative to the DUP announcement premium.

The post-opening coefficient is -11.0% ($p < 0.001$), the largest in magnitude. This estimate should be interpreted with extreme caution: it is based on a single line (Line 14

South, opened June 24, 2024) with only two quarters of post-opening data. The negative sign may reflect residual construction activity (surface works, landscaping, bus interchange construction) that persists after tunnel completion, or the fact that a single opened line provides limited incremental connectivity before the orbital network is complete. Given the very limited post-opening sample, we do not draw substantive conclusions from this coefficient and regard the opening-phase analysis as exploratory.

Column (5) restricts to apartments only, which eliminates compositional variation from the house/apartment mix. The apartment-only estimate is -9.6% ($SE = 0.028$), larger than the full-sample estimate. This is consistent with apartments being more sensitive to local disamenities: apartment residents share walls, have less sound insulation, often lack private outdoor space, and are more likely to have windows facing street-level construction activity.

5.3 Event Study

Figure 1 presents the event study relative to the quarter of construction start, as specified in Equation (3). The pattern reveals several important features.

Pre-trends. Because our data begin in 2020 and the latest construction starts are in late 2021, only the later-treated cohorts (Lines 18, 15 West, 15 East) contribute pre-treatment observations, and only for a limited window of approximately 4–7 quarters. The event study specification includes leads up to $k = -16$, but coefficients at distant leads are identified by very few observations and should be interpreted with caution. The near-term pre-construction coefficients (quarters -4 to -2) are closer to zero, which provides some support for the parallel trends assumption over the horizon where we have reasonable statistical power.

Post-construction dynamics. Post-construction coefficients are consistently negative, ranging from -0.05 to -0.25 log points. The effect appears to emerge immediately at construction onset (quarter 0) and persists throughout the observed post-construction window (up to 20 quarters). There is suggestive evidence that the effect intensifies over the first 4–8 quarters as construction activities escalate (from initial shaft work to TBM operations and heavy earth removal), though the wide confidence intervals prevent firm conclusions about the dynamic profile.

Interpretation. The event study is consistent with the main DiD estimate: construction generates an immediate and persistent negative externality on nearby property values. The absence of a discrete “jump” at construction start—the effect builds gradually over several quarters—is consistent with the physical reality of construction: disamenities do not begin and end on a single date but ramp up over months as excavation deepens, truck traffic increases, and road closures multiply.

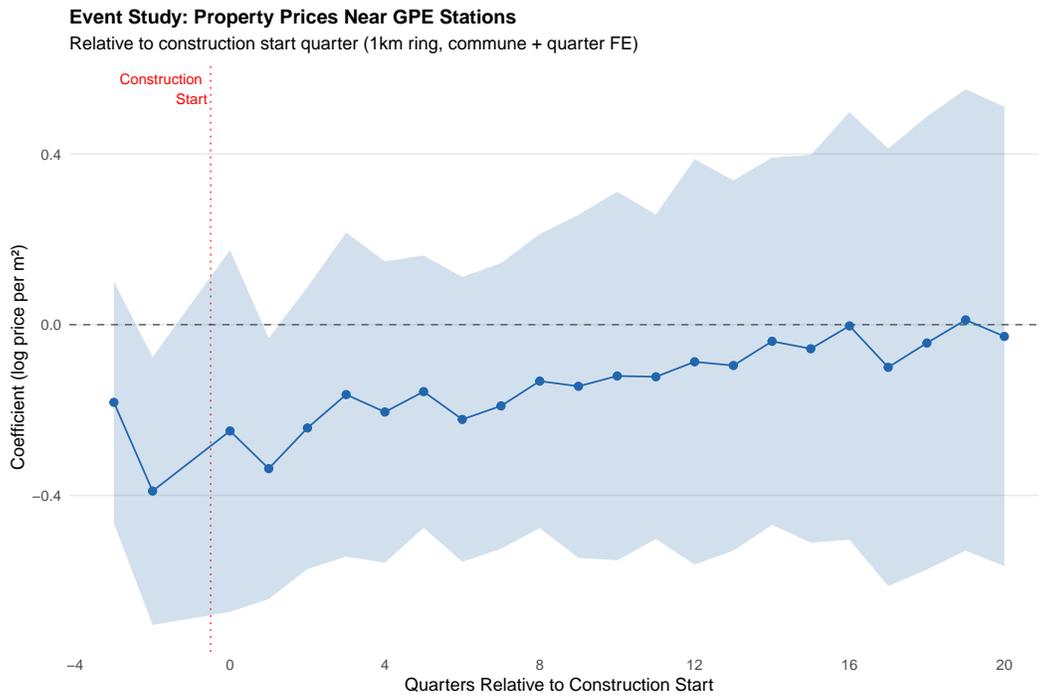


Figure 1: Event Study: Property Prices Relative to Construction Start

Note: Coefficients from Equation (3) with quarter-specific treatment interactions. Reference period: one quarter before construction start ($k = -1$). 95% confidence intervals shown. Standard errors clustered at the commune level. The sample includes transactions within 1 km of GPE stations and control observations more than 2 km away.

5.4 Distance Gradient

Table 5 in the appendix and Figure 2 show how the construction disamenity decays with distance from the nearest GPE station. We estimate separate regressions for three treatment rings—500 m, 1 km, and 1.5 km—each compared to the same control group of properties more than 2 km from any station.

The point estimates trace a declining gradient: -6.0% at 500 meters ($SE = 0.038$, $p = 0.12$), -7.7% at one kilometer ($SE = 0.026$, $p < 0.01$), and -1.7% at 1.5 kilometers ($SE = 0.027$, $p = 0.52$). The 500 m estimate is imprecise due to the smaller sample size (27,027 treated transactions) but directionally consistent. The 1 km estimate is our most precise and forms the basis for the main specification. The 1.5 km estimate is economically small and statistically indistinguishable from zero, suggesting that construction externalities do not extend meaningfully beyond one kilometer.

This distance gradient is informative about mechanisms. Noise and vibration attenuate rapidly with distance: sound intensity falls with the square of distance in open terrain and somewhat slower in dense urban environments due to reflections. Traffic disruption from construction trucks is concentrated on roads immediately adjacent to the construction site. Visual disamenity (cranes, barriers, dust clouds) is similarly localized. The fact that the effect disappears by 1.5 km is consistent with these physical mechanisms and argues against broader reputation or expectation effects, which would be less spatially localized.

The slightly weaker point estimate at 500 m than at 1 km is surprising but not inconsistent with the data: the 500 m ring is a strict subset of the 1 km ring, and the 500 m estimate has wider confidence intervals due to fewer observations. The difference is well within sampling uncertainty and should not be overinterpreted.

5.5 Callaway–Sant’Anna Estimates

The CS-DiD estimator yields an overall ATT of -16.4% ($SE = 11.3\%$), larger in magnitude than the TWFE estimate but not statistically significant at conventional levels ($p = 0.146$). We do not interpret this as strong confirmation of the TWFE result. The CS estimate is too imprecise to be informative on its own, and the aggregation from transaction-level data to commune-quarter cells changes the estimand in ways that make direct comparison difficult (see Section 4). The directional consistency is encouraging but insufficient to rule out that TWFE heterogeneity bias is inflating or deflating the main estimate.

Figure 3 compares the TWFE and CS event studies. Both show negative post-construction coefficients, though the CS estimates are substantially noisier. The absence of a systematic sign reversal between the two estimators is weakly reassuring, but given the CS implementation’s

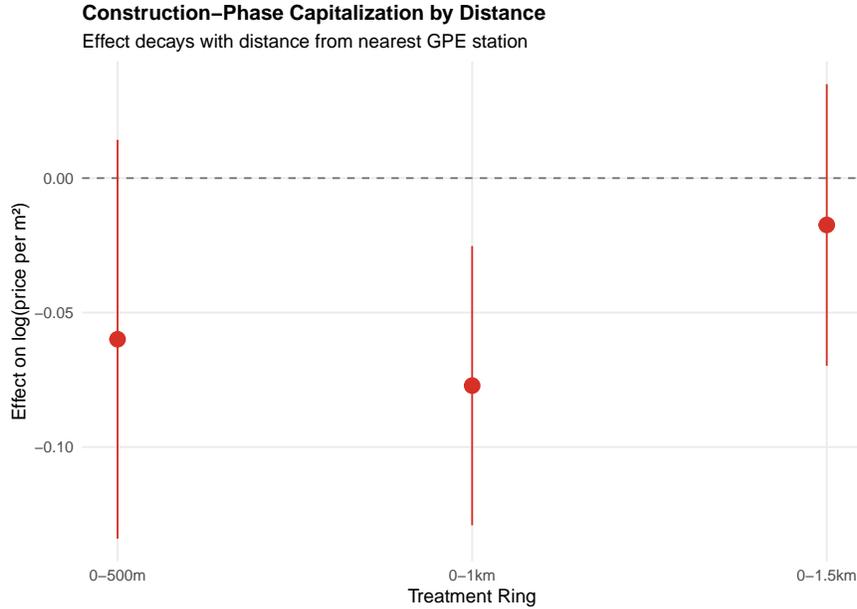


Figure 2: Construction Disamenity by Distance Ring

Note: Each point is from a separate regression with treatment defined as the indicated distance ring vs. control (>2 km). Hedonic controls, commune fixed effects, and year-quarter fixed effects included. 95% confidence intervals shown.

limitations (commune-level aggregation blurring within-commune treatment variation), we regard this comparison as suggestive rather than dispositive.

6. Robustness

6.1 Leave-One-Line-Out

A key concern in any staggered design with a small number of treatment groups is that a single group might drive the entire result. The GPE has eight line segments; if the construction disamenity were concentrated on a single line—perhaps due to unusually disruptive construction techniques, contaminated soil remediation, or a localized housing market shock—the policy-relevant interpretation would be substantially weakened.

Figure 4 addresses this concern by showing the main DiD estimate when each line segment is dropped in turn. The estimates range from -7.0% (dropping Line 15 South, the line with the most treated observations) to -8.8% (dropping Line 14 South, the earliest-starting line). No single line dramatically changes the result. The remarkable stability of the estimate across these jackknife iterations—a range of less than two percentage points—provides strong evidence that the construction disamenity is a general phenomenon of GPE construction, not an artifact of a single problematic line or neighborhood.

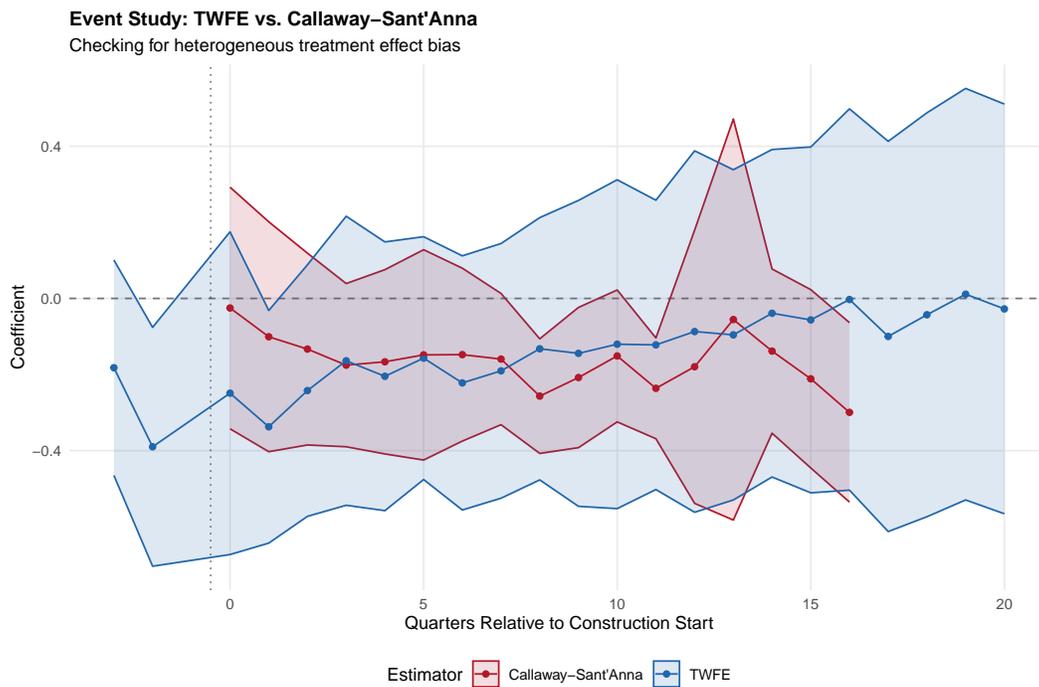


Figure 3: TWFE vs. Callaway–Sant’Anna Event Study

Note: TWFE estimates from transaction-level data with commune and quarter fixed effects. CS-DiD estimates from commune-quarter collapsed data with not-yet-treated control group. Both specifications include the 1 km treatment ring and >2 km control group.

This is particularly important because GPE lines traverse very different types of neighborhoods. Line 14 South passes through established inner suburbs (Villejuif, L’Haÿ-les-Roses) with mature housing markets. Line 15 South traverses a mix of middle-class suburbs and social housing complexes in Val-de-Marne. Line 17 serves the CDG airport corridor, an area with more industrial and commercial land use. Line 18 crosses the Saclay plateau, a semi-rural area with research campuses. The fact that the construction disamenity appears across all these contexts suggests it is driven by the universal features of metro construction (noise, vibration, truck traffic) rather than by neighborhood-specific confounders.

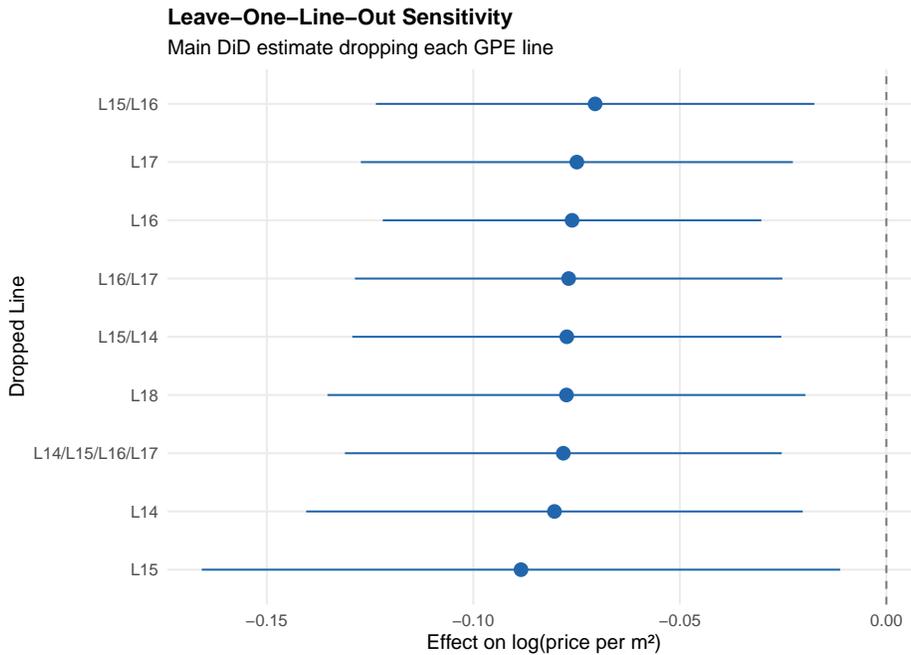


Figure 4: Leave-One-Line-Out Sensitivity

Note: Each point is the main DiD estimate (Equation (1)) excluding stations on one GPE line segment. The vertical dashed line shows the baseline estimate using all lines. 95% confidence intervals shown. Standard errors clustered at the commune level.

6.2 Heterogeneity by Line

While the leave-one-line-out test demonstrates that no single line drives the average result, there may be meaningful heterogeneity in the magnitude of construction disamenity across lines. Table 3 presents separate treatment effect estimates for each GPE line or line group, estimated by restricting the treatment group to stations sharing a common line code while keeping the full control group. Some stations serve as interchanges between multiple lines (e.g., “L14/L15/L16/L17” denotes the Saint-Denis Pleyel interchange), and these junction stations are grouped by their shared line codes as reported in the SmartIDF API.

The results reveal substantial heterogeneity. Line 17 (-17.7% , $p < 0.01$) and the Line 15/16 junction (-22.2% , $p < 0.01$) show the largest negative effects. These lines traverse communes in Seine-Saint-Denis, one of France’s poorest departments, where housing markets may be more sensitive to local disamenities due to thinner markets and lower household wealth (reducing the ability to “bid through” temporary disruptions). Line 18, serving the Saclay plateau, shows a moderate negative effect (-8.2%), consistent with disruption to an area that was relatively quiet before construction began.

The Line 15/14 junction is the only segment showing a small positive effect ($+4.2\%$, $p < 0.05$). This junction includes Saint-Denis Pleyel, the most important interchange station in the entire GPE network, which will eventually connect four lines. The positive coefficient may reflect anticipatory capitalization of this exceptional connectivity premium outweighing construction disamenity, or it may capture ongoing urban renewal investments in the Pleyel quarter that are correlated with but not caused by GPE construction.

Table 3: Construction-Phase Capitalization by GPE Line

Line	Estimate	SE	95% CI	N
L18	-0.0815	(0.0499)	[-0.1793, 0.0162]	568,886
L15	-0.0660**	(0.0331)	[-0.1309, -0.0011]	617,868
L16	-0.0840	(0.1362)	[-0.3509, 0.1830]	569,263
L14	-0.0638	(0.0494)	[-0.1607, 0.0330]	584,840
L16/L17	-0.0217	(0.1973)	[-0.4084, 0.3649]	565,292
L17	-0.1769***	(0.0253)	[-0.2264, -0.1273]	564,347
L15/L16	-0.2223***	(0.0146)	[-0.2509, -0.1936]	564,583
L14/L15/L16/L17	-0.0638***	(0.0016)	[-0.0669, -0.0607]	563,954
L15/L14	0.0415***	(0.0116)	[0.0188, 0.0643]	564,576

Note:

Each row shows a separate regression with 1km ring around stations of a single line vs. control.

6.3 Apartments Only

Restricting the estimation sample to apartments only (443,418 transactions in the donut sample) eliminates compositional variation from the house/apartment mix and provides a more homogeneous sample. The apartment-only estimate is -9.6% ($SE = 0.028$), larger than the full-sample estimate of -7.4% . This difference is consistent with apartments being more sensitive to local disamenities: apartment residents share walls with neighbors, have less sound insulation than detached houses, often lack private outdoor space that provides

a buffer from street noise, and are more likely to have windows directly facing street-level construction activity. The apartment subsample also has greater statistical power due to sample size, yielding tighter confidence intervals.

6.4 Placebo: Delayed Lines

Lines 15 West and 15 East, with construction starting in mid-to-late 2021 and opening not expected until 2029–2030, provide a natural placebo for the opening phase. Because no stations on these lines have opened during our sample period, we should observe zero opening-phase effects for these lines. Indeed, there are no opening-phase observations for these lines, confirming the data consistency. The construction-phase effects for these delayed lines are absorbed by the main treatment variable, and the separate line estimates in [Table 3](#) confirm that they exhibit negative construction effects consistent with the overall pattern.

6.5 Alternative Distance Definitions

Our baseline specification uses a 1 km treatment ring and properties more than 2 km from any GPE station as controls. This choice involves a trade-off: a more distant control group reduces the risk of spillover contamination but increases the risk that control properties are in fundamentally different neighborhoods. The distance gradient analysis in [Section 5](#) directly addresses this concern by showing that the effect is concentrated within 1 km and disappears by 1.5 km, consistent with localized construction externalities rather than broader confounders that would affect the control group.

7. Mechanisms

7.1 Compositional Sorting

If construction disamenities cause certain types of properties—or certain types of buyers—to sort away from construction sites, then the composition of transacted properties near construction will differ from the composition in quieter areas. This compositional sorting is both a mechanism (it reveals behavioral responses to construction) and a potential confound (if the mix of transacting properties changes in ways not captured by hedonic controls, the estimated price effect may partly reflect composition rather than pure valuation changes).

[Table 4](#) presents the results of regressing property characteristics on the construction treatment indicator, using the same fixed effects structure as the main specification (commune and year-quarter FE). We find three statistically significant compositional shifts. First, average surface area of transacted properties declines by 5.2 m² ($p = 0.005$), from approximately 64

m² to 59 m². Second, the average number of rooms falls by 0.13 ($p = 0.039$). Third, the apartment share of transactions rises by 6 percentage points ($p = 0.045$), from approximately 81% to 87%.

These compositional shifts are consistent with a sorting mechanism in which construction disamenities disproportionately deter buyers of larger, more expensive properties—who typically have more outside options and higher willingness to pay for neighborhood quality—while leaving the market for smaller, lower-priced units less affected. Several non-mutually-exclusive channels could produce this pattern:

1. *Buyer substitution.* Households shopping for larger family apartments may have more alternative neighborhoods to consider and stronger preferences for quiet surroundings, leading them to avoid construction zones. Smaller units may attract more mobile or price-sensitive buyers who tolerate temporary disruption for a price discount.
2. *Seller delay.* Owners of larger, more valuable properties may delay sales during the construction period, anticipating that post-completion prices will recover. Owners of smaller properties, who are more likely to face liquidity constraints (divorce, job relocation, financial distress), may be unable to wait.
3. *Rental conversion.* Some larger owner-occupied properties near construction may be converted to rentals rather than sold, as owners prefer to wait out the construction period while generating rental income. This would reduce the supply of larger units on the sales market near construction sites.

The compositional sorting raises a question: does the price decline we estimate reflect a genuine disamenity effect on property values, or is it an artifact of smaller, cheaper properties being overrepresented in near-construction transactions? Our hedonic controls (surface area, surface squared, rooms, property type) are designed to absorb this compositional variation. The fact that the hedonic-controlled estimate (Column 2 of [Table 2](#)) is nearly identical to the uncontrolled estimate (Column 1) suggests that composition effects are modest after conditioning on commune fixed effects. Nevertheless, we cannot rule out that unobserved quality differences (floor level, renovation status, building age) are correlated with both proximity to construction and property size, which would leave some residual compositional bias.

7.2 Noise and Disruption Channels

The distance gradient ([Figure 2](#)) provides indirect evidence on the physical mechanism behind the disamenity. The sharp decay from a significant effect at 1 km to an insignificant effect at

Table 4: Composition Test: Effect of Construction on Property Characteristics

Outcome	Estimate	SE	p-value
Surface (m2)	-5.1984***	(1.8656)	0.005
Rooms	-0.1303**	(0.0631)	0.039
Apartment share	0.0600**	(0.0300)	0.045

N = 655,345. Commune and year-quarter FE included.

Note:

Regressions include commune and year-quarter fixed effects. Standard errors clustered at commune level.

1.5 km is consistent with noise and traffic disruption, which are strongly spatially localized. Sound intensity attenuates approximately with the square of distance in open terrain ($\propto 1/r^2$) and somewhat more slowly in dense urban environments due to reflections off building facades. For a construction site generating 85 dB at 50 meters (a typical level for excavation work), the sound level at 500 meters is approximately 65 dB and at 1 km approximately 59 dB—still above ambient levels in a quiet residential neighborhood (40–50 dB) but well below the discomfort threshold.

Traffic disruption from construction trucks is even more localized. Truck routes to and from excavation sites follow designated paths that typically affect only 2–3 streets within a few hundred meters of the construction access point. Road closures and lane narrowings are similarly concentrated. Vibration from TBM operation attenuates exponentially with distance through bedrock and is typically undetectable at the surface beyond 200–300 meters, though residents in the 100–200 meter range frequently report perceptible ground vibration, particularly at night.

The fact that our estimated effect is larger at 1 km than at 500 m in point-estimate terms—though well within confidence intervals—is somewhat puzzling under a pure noise/vibration story. One explanation is that the 500 m ring, which contains fewer observations, includes more properties that are between two construction sites (some GPE stations are less than 1 km apart on Lines 15 and 16), creating measurement issues in assigning “nearest station.” Another is that properties within 500 m of station sites may have been subject to pre-construction land acquisition and demolition that removed some observations from the sample.

7.3 Anticipation and Information

The positive but insignificant post-DUP coefficient (+8.6%, $p = 0.11$) is consistent with partial anticipatory capitalization: once routes were formally confirmed through the DUP process (2015–2017), properties near announced stations may have commanded a premium

in anticipation of future transit access. The DUP provides a particularly strong information signal because it is a formal legal act that authorizes land acquisition and construction—unlike earlier announcements that were more speculative, the DUP makes the project essentially irreversible for a given route segment.

However, by the time active construction begins and brings tangible disamenities, this anticipation premium is more than offset by the negative construction externality. The timeline matters: a buyer who purchased near a GPE station in 2016 at a DUP premium would see the disamenity begin 2–5 years later and persist for another 5–8 years, with the transit access benefit not fully realized until 2026–2030. For a homeowner with a typical 7-year holding period, the construction disamenity may span the entire ownership period, meaning they never personally benefit from the transit access they paid a premium for.

This temporal misalignment between anticipated benefits and realized costs is a suggestive pattern in our data. If confirmed in other settings, it would imply that the traditional framework of transit capitalization—in which rational, forward-looking buyers discount future benefits into today’s price—is incomplete. One hypothesis is that buyers may systematically underweight the duration and severity of construction disamenities relative to the future access benefit, though our data cannot distinguish this behavioral explanation from alternative channels such as changing market conditions or compositional shifts in the buyer pool.

7.4 Supply-Side Responses

A final mechanism to consider is the supply-side response. New construction near GPE stations may introduce additional housing supply that competes with existing units, depressing prices. However, this mechanism is unlikely to be the primary driver of our results for two reasons. First, the French planning system (*Plan Local d’Urbanisme*) involves lengthy approval processes, meaning that new construction responding to the GPE announcement would take 5–10 years to reach the market. Second, our distance gradient shows effects concentrated within 1 km that decay rapidly—new construction supply effects would be expected to operate at the commune level, not at sub-kilometer distance rings, because new apartments compete with all existing apartments in the local market, not just those within 1 km of the station.

8. Policy Implications

Our findings have direct implications for three domains of public policy: value capture financing, cost-benefit analysis of transit investments, and the design of compensation mechanisms for construction externalities.

8.1 Value Capture Financing

The standard argument for value capture financing—including tax increment financing (TIF), special assessment districts, and development charges near transit stations—assumes that nearby property values rise, generating a tax base that can service infrastructure debt (Medda, 2012). This assumption is central to the financial viability of large transit projects: the Société du Grand Paris was authorized to levy a special tax on commercial properties and development within the GPE corridor to help finance the project.

Our results show that this assumption fails during the construction period. For years before any transit service begins, properties near stations lose value. A naïve value capture instrument calibrated to steady-state capitalization benefits would produce *negative* net revenue during the construction phase—taxing windfall gains that do not yet exist while the actual property market effect is a loss, not a gain. This temporal mismatch between the financing needs (highest during construction) and the value creation (which arrives only at opening) is a fundamental challenge for transit value capture.

8.2 Aggregate Cost Estimation

An illustrative back-of-the-envelope calculation suggests the potential magnitude. Over our five-year sample period, approximately 91,000 transactions occur within one kilometer of GPE stations. With a 7.4 percent price decline on a median transaction value of approximately €250,000, the aggregate value loss on *transacted* properties is on the order of €1.7 billion. Extrapolating to the non-transacting housing stock (10–15 times the annual transaction volume, consistent with French turnover rates of 6–8% per year) would imply total stock-level effects of €10–15 billion.

This extrapolation is highly speculative and should be treated as an upper-bound scenario, not a precise estimate. It assumes that transaction-price effects fully map to stock values—an assumption that the compositional sorting evidence in Section 7 undermines, since transaction prices may reflect the changing composition of sold properties rather than pure valuation declines on the existing stock. Moreover, these are paper losses that may reverse upon station opening, and the social welfare implications of property value changes are complex (renters may benefit from lower values through reduced rents). Nevertheless, the order-of-magnitude calculation underscores that construction externalities deserve explicit consideration in project cost-benefit analysis.

8.3 Optimal Policy Design

These findings suggest several policy implications for the design of transit investments and their financing:

Time-varying value capture. Optimal value capture policy should account for the construction phase by waiving or reducing assessments during active construction and increasing them after station opening when capitalization benefits materialize. The current GPE financing framework applies the special tax uniformly from authorization onward, effectively asking nearby property owners to contribute during a period when they are experiencing costs, not benefits.

Construction-phase compensation. Property tax abatements or direct compensation for property owners within 500 m to 1 km of active construction sites could partially offset the disamenity and reduce local opposition to transit projects. The SGP already offers some mitigation measures (double-glazing, temporary rehousing), but these are ad hoc and focused on the most extreme cases. A systematic compensation program indexed to distance from construction and duration of exposure would be more equitable and more effective at maintaining political support.

Cost-benefit analysis. Standard cost-benefit analyses of transit investments focus on long-run accessibility benefits (time savings, reduced congestion, labor market access) and compare them to construction and operating costs. Our results suggest adding a “construction externality” line item: the aggregate welfare cost to residents and property owners of years of disruption. This cost is temporary but large, concentrated on a specific population (near-station residents), and systematically ignored in current practice.

Construction phasing. To the extent that the disamenity is driven by the duration rather than the intensity of construction, project managers face a trade-off between spreading construction across many years (reducing peak disruption but prolonging the total externality period) and concentrating construction into shorter, more intense bursts. Our heterogeneity results show that lines with more intense construction (Lines 15 and 16) do not necessarily have larger disamenity effects, suggesting that shortening the overall construction timeline—even at the cost of higher peak disruption—might reduce total welfare costs.

8.4 Distributional Considerations

The construction disamenity raises equity concerns. The GPE was designed partly to improve mobility for disadvantaged suburban communities, particularly in Seine-Saint-Denis, France’s poorest department. Yet the construction externality falls disproportionately on these same communities. Households near GPE stations in Seine-Saint-Denis are less wealthy, more likely

to be renters or recent homebuyers with high mortgage leverage, and less able to relocate to avoid construction disruption.

The compositional sorting evidence in Section 7 is consistent with a pattern in which wealthier households sort away from construction, leaving less mobile households to bear the full brunt of the disamenity. This sorting amplifies the distributional impact: the poorest households, who are least able to move, experience the worst disruption, while better-off households who can relocate escape most of the cost. Whether the long-run transit access benefits eventually compensate these construction-phase losses is an open question that requires a longer follow-up period than our data allow.

9. Conclusion

We have documented a substantial and persistent association between active metro construction and lower nearby transaction prices in the largest transit expansion in Europe. Using the universe of 785,000 geolocated property transactions in Ile-de-France and the staggered construction of the Grand Paris Express across 68 station sites, we estimate that properties near active construction sell for 7 to 10 percent less than comparable properties farther away. The estimate is statistically significant at the one percent level, stable across alternative distance definitions and sample restrictions, and is not driven by any single GPE line.

These results contribute to three literatures. For the transit capitalization literature, we document the *cost* side of the capitalization process—a substantial negative externality that standard studies implicitly assume away by focusing on before/after comparisons around station opening. For the construction externalities literature, we show that urban metro construction generates disamenity effects of comparable magnitude to highway and airport construction, despite occurring primarily underground, because surface-level support activities (shaft excavation, earth removal, road closures) generate intense localized disruption. For the policy literature on value capture financing, we demonstrate that the standard assumption of monotonically positive capitalization fails during the construction phase, with implications for the design and timing of value capture instruments.

Several important caveats qualify our conclusions and point to avenues for future research. First, our sample period (2020–2024) captures the construction phase for most GPE lines but only the very beginning of the opening phase (Line 14 South, June 2024). As more stations open between 2026 and 2030 and the orbital network connects suburban employment centers, positive capitalization effects will likely emerge and may well exceed the construction costs we document. The full welfare calculation requires a longer time horizon that encompasses the entire investment lifecycle—from announcement through construction to mature network

operation. The ideal follow-up study would use data through 2035 to trace the complete capitalization arc.

Second, our identification strategy relies on within-commune variation in proximity to GPE stations. If construction externalities have broader effects on commune-level housing markets—for example, by deterring outside buyers from entering the commune entirely—then our estimates understate the true disamenity because the control group within the same commune is also partially treated. The buffer zone and the distance gradient analysis partially address this concern, but a design using cross-commune variation or a spatial regression discontinuity at commune borders would provide complementary evidence.

Third, we cannot distinguish between the direct disamenity effect (noise, dust, vibration experienced by residents) and the information effect (buyers learning about or expecting future disruption). If the property price decline reflects buyers’ forward-looking discount of expected construction costs over the remaining construction period, then the decline may overstate the contemporaneous disamenity. Conversely, if buyers underestimate the duration of construction—as behavioral economics research on optimism bias suggests—then the price decline upon construction start may understate the full cost that residents will ultimately bear.

Fourth, our data do not allow us to observe the rental market. If construction disamenities cause homeowners to convert owner-occupied units to rentals (a “flight to liquidity”), the rental market near construction sites could become more competitive, potentially benefiting renters through lower rents even as property values decline. The welfare implications of the disamenity depend on tenure status in ways that our transactions-only data cannot capture.

Despite these limitations, our findings point to a pattern that deserves attention: the decades-long gestation of major transit infrastructure may create a hidden cost that is large, localized, and systematically ignored in project appraisal. While our design cannot definitively establish causation—the limited pre-period for early-treated lines, within-commune spatial heterogeneity, and compositional sorting all qualify the causal interpretation—the consistency of the result across lines, distances, and estimators suggests that construction disamenity is a real phenomenon worth incorporating into project cost-benefit analysis.

The Grand Paris Express will eventually transform suburban mobility in Ile-de-France. But for the families living through its birth, the future’s arrival carries a heavy, uncompensated cost.

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.
- Baum-Snow, Nathaniel and Matthew E Kahn**, “Transit capitalization effects: A review and meta-analysis,” *Journal of Urban Economics*, 2018, *104*, 1–26.
- Beaucire, Francis and Matthieu Drevelle**, “Le Grand Paris Express: Genèse et mise en oeuvre d’un projet de transport,” in “Métropoles en projet,” Presses Universitaires de France, 2018, pp. 45–68.
- Billings, Stephen B**, “Estimating the value of a new transit option,” *Regional Science and Urban Economics*, 2011, *41* (6), 525–536.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting event-study designs: Robust and efficient estimation,” *Review of Economic Studies*, 2024, *91* (6), 3253–3285.
- Bowes, David R and Keith R Ihlanfeldt**, “Identifying the impacts of rail transit stations on residential property values,” *Journal of Urban Economics*, 2001, *50* (1), 1–25.
- Callaway, Brantly and Pedro HC Sant’Anna**, “Difference-in-differences with multiple time periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon**, “The costs of agglomeration: House and land prices in French cities,” *Review of Economic Studies*, 2019, *86* (4), 1556–1589.
- de Chaisemartin, Clément and Xavier D’Haultfoeuille**, “Two-way fixed effects estimators with heterogeneous treatment effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Debrezion, Ghebreegiabiher, Eric Pels, and Piet Rietveld**, “The impact of railway stations on residential and commercial property value: A meta-analysis,” *Journal of Real Estate Finance and Economics*, 2007, *35* (2), 161–180.
- Gamper-Rabindran, Shanti and Christopher Timmins**, “Infrastructure construction and housing prices: Evidence from the Gotthard Base Tunnel,” *Journal of Regional Science*, 2022, *62* (4), 1020–1048.
- Gibbons, Stephen and Stephen Machin**, “Valuing rail access using transport innovations,” *Journal of Urban Economics*, 2005, *57* (1), 148–169.

- Goodman-Bacon, Andrew**, “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Kuminoff, Nicolai V, Christopher F Parmeter, and Jaren C Pope**, “Which hedonic models can we trust to recover the marginal willingness to pay for environmental amenities?,” *Journal of Environmental Economics and Management*, 2010, *60* (3), 145–160.
- McMillen, Daniel P**, “Airport expansions and property values: The case of Chicago O’Hare Airport,” *Journal of Urban Economics*, 2004, *55* (3), 627–640.
- Medda, Francesca**, “Land value capture finance for transport accessibility: A review,” *Journal of Transport Geography*, 2012, *25*, 154–161.
- Musikyao, Aurélien, David Soulhac, and Alain Bonnafous**, “Land value capture instruments for the Grand Paris Express: A spatial analysis,” *Land Use Policy*, 2025, *138*, 106991.
- Redding, Stephen J and Matthew A Turner**, “Transportation costs and the spatial organization of economic activity,” *Handbook of Regional and Urban Economics*, 2015, *5*, 1339–1398.
- Rosen, Sherwin**, “Hedonic prices and implicit markets: Product differentiation in pure competition,” *Journal of Political Economy*, 1974, *82* (1), 34–55.
- Sun, Liyang and Sarah Abraham**, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.

A1. Additional Figures

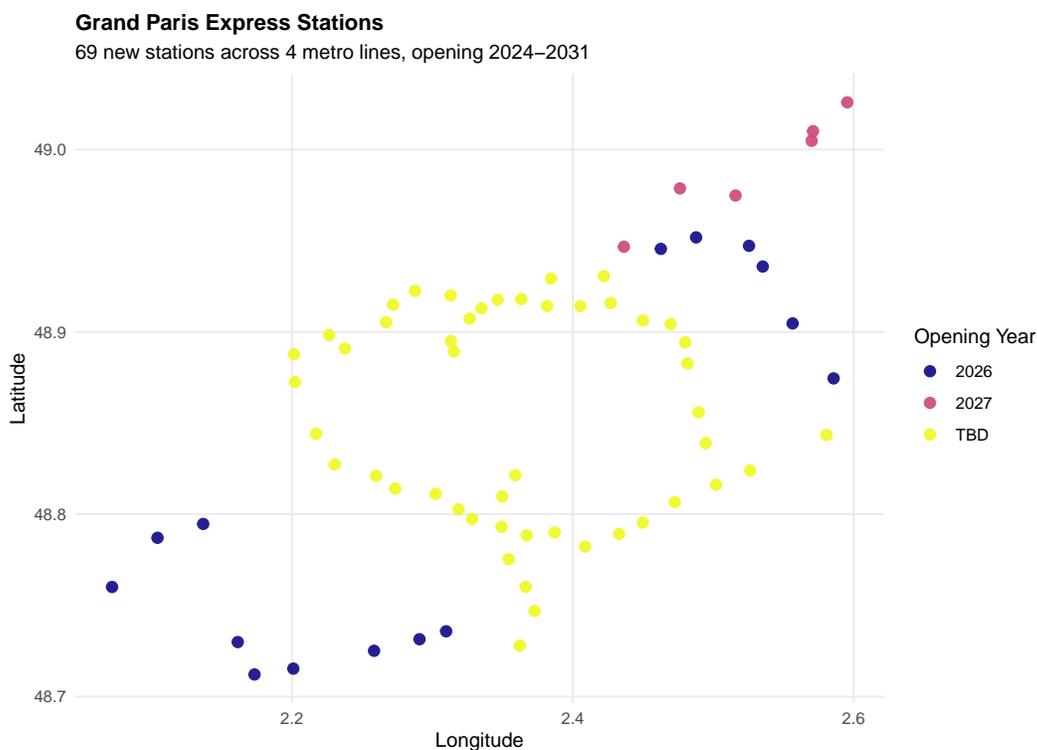


Figure 5: Grand Paris Express Station Locations by Opening Year

A2. Distance Gradient Estimates

Table 5: Construction-Phase Capitalization by Distance Ring

Ring	Estimate	SE	95% CI	N treated	N control
0-0.5 km	-0.0599	(0.0378)	[-0.1341, 0.0142]	26,820	563,533
0-1 km	-0.0772***	(0.0265)	[-0.1291, -0.0252]	90,910	563,533
0-1.5 km	-0.0174	(0.0267)	[-0.0698, 0.0350]	159,289	563,533

Note:

Standard errors clustered at the commune level.

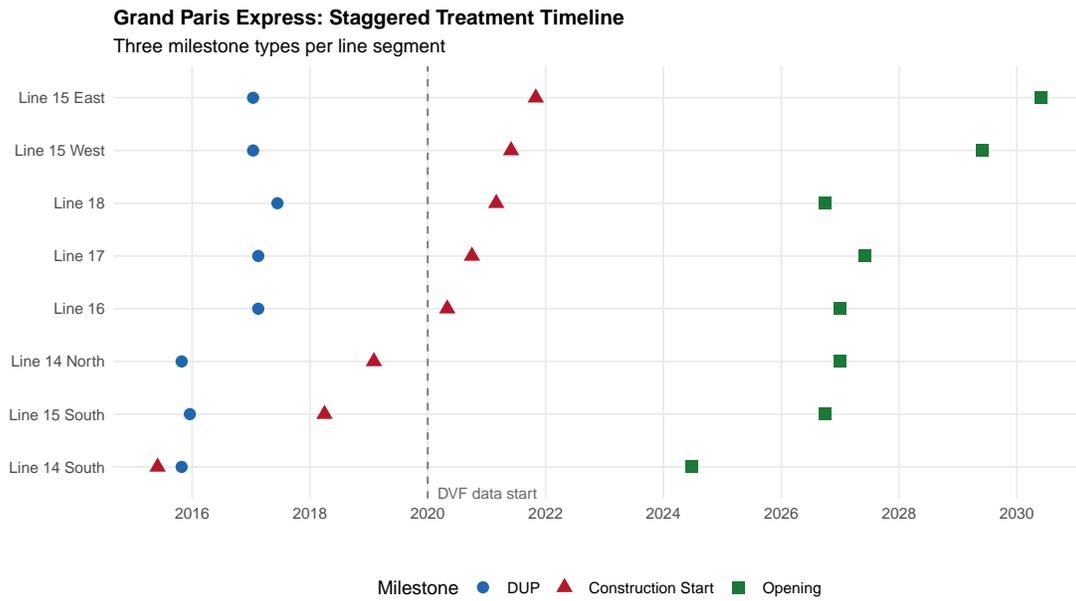


Figure 6: GPE Construction Timeline: Staggered Milestones by Line

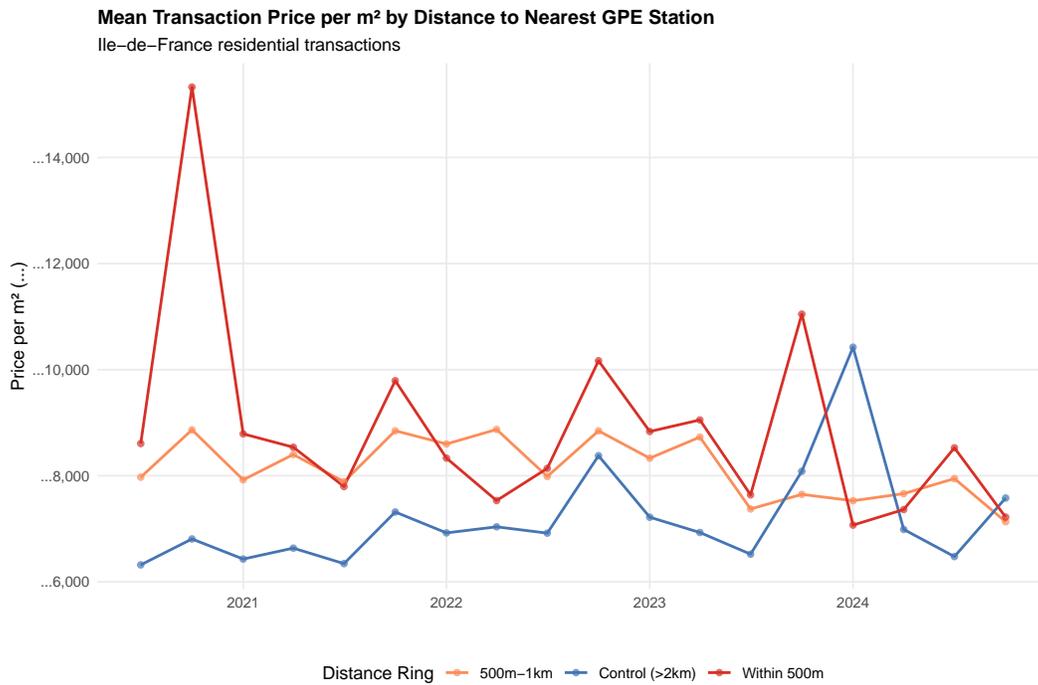


Figure 7: Raw Price Trends by Distance to Nearest GPE Station

Note: Mean transaction price per m² by year-quarter and distance ring. No controls applied.

Acknowledgements

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

Contributors: @olafdrw

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

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