The Hidden Costs of Urban Renewal[*]

Matéo Moglia[†]

September 10, 2026

Preliminary and incomplete, results may still change.

Job Market Paper

Abstract

Why does housing supply stay low where rents are highest? Linking permit-level survey data to building permits and precise parcel-level land occupation in the Paris area, I directly measure construction costs rather than inferring them from equilibrium outcomes. Thanks to an option-value model of redevelopment in which construction takes time and land is bargained, I recover parcel-level fixed costs, net of both frictions, and decompose them: demolition accounts for 15% of their cross-sectional variance, observed land-use regulation for 21%, and 37% remains unexplained by either. Two counterfactual results follow. First, setting fixed costs to zero raises supply by 104.1% and reduces rent by 22.7%. Entirely removing the density cap attains only 41% of that gain. Second, the friction is large: the median parcel requires a transfer of 1,839 €/m^2 to be profitable for redevelopment. In all, the binding constraint is not land-use regulation but the construction fixed costs of housing supply.

Keywords: Urban renewal, land use, teardowns, housing supply elasticity, capital investment with uncertainty

JEL: R14, R31, R33, Q53

[*]I am extremely grateful to Pierre Boyer and Benoit Schmutz-Bloch for their invaluable guidance. This paper greatly benefited from discussions with Suzanne Bellue, Gilles Duranton, Bertrand Garbinti, Jessie Handbury, Yurou He, Jonathan Goupille-Lebret, Alain Trannoy, Etienne Wasmer, and seminar and workshop participants at AFSE (Nantes), Aix-Marseille School of Economics, Humboldt University, CREST, HEC Paris, UEA (Barcelona), the Wharton School, Università della Svizzera italiana. This research was supported by subsidies from the Agence Nationale de la Recherche (Labex ECODEC No. ANR-11-LABEX-047), the Chaire Professorale Jean Marjoulet, and the Chaire DDX. I also thank the Berlin Quantitative Spatial Economics Group and the Wharton School for their hospitality during the writing of the paper. Data are accessed through the Centre d'accès sécurisé aux données.

[†]Centre de Recherche en Economie et Statistiques (CREST). Ecole polytechnique, Institut Polytechnique de Paris, Palaiseau, France. *Contact*: mateo.moglia@gmail.com.

1. Introduction

The world's most productive cities now grow inward. Vacant land in their cores is scarce, so additional floorspace increasingly comes from replacing existing structures rather than from expansion at the fringe (Frolking et al., 2024). Redevelopment has become the dominant margin of urban growth and the one we understand least. Across advanced economies, environmental constraints and land-preservation policies push cities inward.[1] The relevant question is therefore no longer how much housing can be built on vacant land, it is why existing urban land is not redeveloped when expected returns are high.

Redevelopment is rare even where returns appear large. In the Paris metropolitan area, valuable land still sits under low-density housing, ageing commercial buildings, warehouses, and car parks. Standard urban models predict that rising demand pushes land towards its highest-valued use. The same puzzle appears wherever inward growth dominates. Metropolitan areas like Tokyo, New York, San Francisco, Seoul, or Toronto still exhibit areas of low density.

The dominant explanation is regulation. A large literature argues that zoning, density limits, permitting requirements, political opposition, prevent developers from responding to market signals (Gyourko and Molloy, 2015; Glaeser and Gyourko, 2018; Chapelle et al., 2025). In this view, regulation is a tax on development and accounts for the low supply elasticities of high-price cities. But redevelopment requires more than permission.

In this paper, I study an alternative explanation to low housing supply elasticity: fixed costs. Developers must acquire land, bargain with incumbent owners, evict occupants, demolish structures, remediate sites, and wait for construction to finish. These costs are sunk and irreversible, so redevelopment is not only production decision but an option-exercise decision: the developer can build today or wait. A low redevelopment rate therefore does not necessarily itself imply stringent regulation (Baum-Snow, 2023). Even absent regulatory barriers, redevelopment may remain unattractive when clearance is costly or the option value of waiting is high.

Assessing the relative importance of regulation and fixed costs requires measuring redevelopment costs. A growing structural literature recognises that fixed redevelopment costs shape housing supply (Murphy, 2018; Baum-Snow and Han, 2024; Rollet, 2025), yet these costs enter as latent residuals, recovered from equilibrium outcomes and from parametric assumptions on cost distribution. We consequently know little about their level, their composition, or their size relative to the regulatory constraints em-

[1] Since 2021, France has made zero net land take a national objective, explicitly steering development away from greenfield expansion (Salin and de L'Estoile, 2024).

phasised elsewhere. Regulation is observable and sometimes treated as the whole problem. Redevelopment costs are acknowledged but remain a black box. This paper asks how large redevelopment costs are, what they are made of, and how much of the low supply response in dense cities they explain.

To answer those questions, I assemble a parcel-level panel covering more than 550,000 developed parcels, linking cadastral records, the universe of land and property transactions, building permits, demolition declarations and local land-use plans. I focus on the Paris metropolitan area, one of Europe's largest and most supply-constrained housing markets. I develop a parcel-level real-options model in which developers choose whether to redevelop immediately or wait, and in which uncertainty comes from future housing prices rather than transitory profitability shocks. Investment is irreversible, construction takes time, landowners bargain with developers. Redevelopment costs are then recovered directly, parcel by parcel, as the wedge required to rationalise observed land prices and observed redevelopment outcomes jointly. This delivers a distribution of costs instead of a scalar.

Recovered redevelopment fixed costs are large, highly heterogeneous, and decisive for redevelopment decisions. To understand the composition of this cost, I exploit a unique assembly of parcel-level data. In particular, I recover developer-reported demolition and construction, and match them to exhaustive parcel and project characteristics. These costs are substantial and depend on the removed structure. Redevelopment also carries a time cost: conditional on location and prices, non-residential parcels take significantly longer to redevelop than residential ones. To my knowledge, no previous structural study pairs recovered redevelopment costs with direct measurement of their components.

The composition of these costs departs sharply from the barriers emphasised in policy debate. Land-use regulation explains only 21% of their cross-sectional variation; measured demolition costs, 15%; surrounding built density, 12%. The dominant obstacles to redevelopment are not primarily regulatory. A large share arises from the physical and economic cost of replacing what already stands. These magnitudes reverse the ranking of policies implied by the regulatory view.

Eliminating redevelopment costs raises housing supply by 104.1% in general equilibrium. Removing density restrictions reaches only 41% of that gain, equivalent to a uniform redevelopment subsidy of 1,343 €/m^2. The two instruments act on different margins. Density relaxations affect projects already close to profitable, so their effect is concentrated where parcels would have redeveloped anyway. Redevelopment costs determine whether the option is exercised at all. In high-rent cities, supply is governed less by the intensity of permitted development than by the extensive margin of redevelopment itself.

The paper is set in the Paris metropolitan area, which offers an unusually clean setting to conduct such analyses. Between 1999 and 2021 its population grew by 11% while land take-up rose by less than a percent, implying redevelopment instead of sprawling. I assemble a unique parcel-level panel covering more than 550,000 developed parcels from 2014 to 2025. It links the fiscal universe of built units, the three-dimensional building envelope from satellite data, the universe of building permits from administrative sources, the universe of land and property transactions, municipal land-use plans, various soil quality and connectivity measures from geological sources, and a developer cost survey. I am able to track redevelopment projects, developers characteristics, and the price at which land and units are traded.

Four facts about that panel motivate the analysis. First, redevelopment is the supply margin. Of the housing units authorised over the period, 64% come from demolishing a standing structure and rebuilding, 20% from transformation within an existing envelope, and only 10% from empty parcels. Second, redevelopment reallocates land across uses, as a large share of units are supplied on former non-housing parcels. Third, the parcels that should redevelop first often do not. Around 37,000 units could be built on vacant parcels, brownfields, on urban infills. Fourth, the two frictions I later add to the model are visible in the raw data. Construction time differs sharply by use of origin, with tertiary and secondary redevelopments significantly slower than housing ones, and land acquisition is locally segmented, with a median municipal market served by an effective seventeen developers.

I build on the canonical frameworks of irreversible investments under uncertainty (Capozza and Li, 1994; McDonald and Siegel, 1986). A demand shifter follows a geometric Brownian motion and rents clear the market against the standing stock, so the rent observed by developers inherits that diffusion. Data confirm that landowners are atomistic and therefore take prices as given while aggregate supply still feeds back on the level of rents. Redevelopment is irreversible, so the developer waits until the local rent crosses a threshold equal to the naive break-even rent scaled by the option multiplier. Conditional on exercising, the developer chooses density up to the regulatory envelope. Two frictions separate this from the canonical stopping problem (Dixit and Pindyck, 1994). Construction takes time, indexed by the use of origin, which delays the revenue stream but also spreads the construction cost, so the net wedge is smaller than the raw discount on rents. Land changes hands bilaterally, so I replace the zero-profit condition with Nash bargaining over the redevelopment surplus and calibrate the landowner's weight from local concentration. I also depart from the quantitative literature in the source of uncertainty. In the latter, redevelopment timing is driven by idiosyncratic shock on costs (and, isomorphically, on profits). Here it is driven by the price process itself, since projects of this scale do not receive random jumps in returns between periods.

Inverting the observed transaction price at the observed rent then yields the fixed cost as a fully determined parcel-level "residual", not a free parameter chosen to reproduce an aggregate development rate. The inversion also delivers an exact accounting of what that residual is not. The gap between the frictionless redevelopment value and the price at which land trades splits into a delay wedge and a developer bargaining rent, and the fixed cost as the remainder. Because nothing in the procedure targets the model's supply elasticity, the sign of the fixed cost, or the wedge between fundamental and perceived rent growth, all three serve as over-identification checks. The recovered elasticity of 0.6 sits against an external estimate identified from aggregate price and quantity variation the model never saw, and the fixed cost is negative for well under one percent of parcels.

The fixed cost can then be confronted with parcel-level characteristics, which is what no previous structural study has been able to do. The *Prix de revient des logements neufs* survey asks developers what they spent on a stratified sample of permits, separating construction from demolition, and reports the permit identifier, which lets me match it back to parcels. Demolition costs are fixed in exactly the sense the model requires. They are uncorrelated with the floorspace authorised and scale instead with the structure removed, at 35 to 56 euros per square metre depending on destination, on top of a mobilisation cost of around 25,000 euros incurred whatever the quantity. I use the hedonic coefficients from this survey to predict demolition cost for every parcel in the panel, and then regress the recovered fixed cost per square metre on seven blocks of observables. The blocks are demolition, ground condition, grid connectivity, built surroundings, parcel scale, out-of-model land-use regulation, and municipal controls. Specifically, regulation enters through permitting refusal rates, architectural protection and zoning type rather than the density cap.

A Shapley decomposition apportions the explained variance across these blocks by averaging each one's marginal contribution over all orderings of entry. Land-use regulation is the largest single block at 21%, ahead of demolition at 15%, built surroundings at 12%, and ground condition at 14%. Grid connectivity is indistinguishable from zero and 37% is explained by none of the blocks. Two features of this result matter for what follows. Regulation is the one observable that survives differencing out the municipality, since zoning and architectural protection vary parcel by parcel where a municipal refusal rate cannot. And the ranking is stable across the full parameter band and across every alternative first-stage price specification, with the unexplained share far more sensitive to the calibration of the option problem than to the hedonic prediction of the land price. The decomposition concerns dispersion. Regulation is what makes fixed costs differ across parcels. It is not what makes them large.

Counterfactual analysis acts on two sets of instruments. First, I subsidize fixed costs with unconditional transfers to developers. Second, I relax the density caps. I also separate the schedule shift from the market response, reporting the partial-equilibrium supply deviation alongside the general-equilibrium one, since new supply depresses the clearing rent and pushes marginal parcels back below their thresholds. Setting the fixed cost to zero is the benchmark rather than a policy, since it bounds what any instrument acting on costs can achieve. It raises supply by 104.1% and lowers rent by 22.7%, with 30% of the direct effect clawed back once the market re-clears. Removing the density cap entirely attains only 41% of that increase. A uniform subsidy of 1,343 €/m^2 would deliver the same supply. Increasing the density cap leaves most thresholds where they were, whereas lowering the fixed cost lowers every threshold at once.

Inverting the threshold condition parcel by parcel then narrows the counterfactual analysis. For each parcel I solve for the smallest transfer that would bring its threshold below the prevailing rent, which partitions the stock into parcels that redevelop anyway, parcels a finite subsidy would reach, parcels blocked by the value of the standing structure, and parcels where even an unconstrained envelope would deliver a smaller building than the one already there. Among reachable parcels the median transfer is 1,839 €/m^2 of land, which is roughly the variable construction cost. Programmes at the scale of the *Fonds Friches*, a nationwide programme for brownfield renewal, reach only the very left tail. Ranking reachable parcels by that transfer traces a frontier with no inframarginal spending by construction.

Contribution to the literature This paper contributes to the micro-foundations of housing supply literature (Baum-Snow and Han, 2024) by opening the "black box" of redevelopment fixed costs. Glaeser and Gyourko (2003) and the subsequent body of works (Glaeser and Gyourko, 2018; Gyourko and Molloy, 2015; Diamond et al., 2019) think of regulations as a tax on supply. I complement this approach by focusing on the role of fixed costs in shaping construction. Murphy (2018), Peng (2023), and Rollet (2025) estimate high fixed costs structurally from redevelopment patterns but remain agnostic about their composition. Grovenstein et al. (2011) allows for heterogeneous land costs in development. My direct measurement using developer survey data reveals that transformation costs vary systematically with past land use, a distinction their unified fixed-cost parameters do not capture. Methodologically, I depart from their choice (an idiosyncratic profitability shock to motivate redevelopment timing). Large-scale projects do not experience random jumps in returns between periods. Instead, I build on the real options tradition (McDonald and Siegel, 1986; Capozza and Helsley, 1990; Dixit and Pindyck, 1994), surveyed in Womack (2015), where developers face irreversible investment under price uncertainty. Lange and Teulings (2024) and Lebret et al. (2026) recently showed that predictable demand growth creates option value to delaying construction. I extend this literature in three ways. First, I incorporate heterogeneous fixed costs tied to parcel history, generating a distribution of redevelopment thresholds that explains within-city variation in supply responses. Second, I let construction take time, which is standard in the theoretical real-options literature since Majd and Pindyck (1987) and Bar-Ilan and Strange (1996) but rarely taken to parcel-level data, and which turns aggregate supply into a predetermined object that responds to rents with a lag. Third, I document that the developer market is concentrated, as Favilukis et al. (2026) and Quintero (2023) find for the United States: I replace the zero-profit condition with Nash bargaining over the redevelopment surplus, with the landowner's weight calibrated from local developer concentration. This delivers the incidence of supply-side policy directly, since a grant on fixed costs is capitalised into land at a rate equal to that weight.

The paper also speaks to quantitative urban models that typically impose city-level supply elasticities (Redding, 2024). I provide micro-foundations for within-city heterogeneity: even uniform price increases produce differential responses depending on where parcels sit relative to their transformation cost thresholds. By doing so, I am in line with Baum-Snow and Han (2024), who explicitly model area-specific housing supply elasticity. However, they parametrize the fixed costs distributions where I provide a direct measure. Similarly, Combes et al. (2021) relies on the zero-profit condition to derive supply elasticity for individual homes in France, using similar data. I set the analysis in a dense, already-built context, where the relevant margin is replacement rather than first occupation.

Roadmap The remainder of the paper proceeds as follows. Section 2 describes the institutional context, the data sources, and presents motivating facts. Section 3 develops the theoretical model and derives optimal thresholds. Section 4 details the model calibration and Section 5 presents the baseline results. Section 6 presents counterfactual policy simulations. Section 7 discusses extensions. Section 8 concludes.

2. Data and context

2.1. Institutional context

The analysis is set in the Paris area, consisting of the city of Paris (*département* 75), and three conterminous *départements*: Hauts-de-Seine (92), in the West, richer and oriented towards the tertiary sector, Seine-Saint-Denis (93), in the North and the East, one of the poorest area of continental France with an important industrial past, and Val-de-Marne (94), in the South, with stark spatial inequalities. Paris area, also dubbed as *Petite couronne*, is dense, with only two municipalities classified as "intermediate dense", the 141 others being "dense", the highest density classification by INSEE. Historically, the area has been a powerhouse of the French industrial power. After the successive waves of deindustrialization, the Paris area is now geared towards the tertiary sector. Around 10% of the French population lives there and the area generates 12% of the French GDP. The area is attractive, for firms and for households, creating an important tension on the commercial and housing real estate markets.

The area is split into 143 municipalities (including 20 Paris districts, or *arrondissements*). Municipal councils levy local taxes (especially on establishments and on property) and have the power over land use regulations. They decide on zoning plans (*plan local d'urbanisme*) and deliver construction permits (see Chapelle et al. (2023, 2025), and Lafont, 2026 for more details on French context).

2.2. Data

I rely on a rich parcel-level dataset covering the Paris area from 2014 to 2024. I restrict on parcels where zoning allows housing. Appendix A.1 presents extensively the data used. Usually, a parcel hosts one building (but large parcels can host several). A building hosts one or more units, where a unit can be a housing unit or an activity unit. Activity units are distinguished depending on their destination. They can be a tertiary unit, a secondary unit, or an annex. Each unit appears in the dataset as long as taxes are collected. A parcel can also host a vacant building (a physical structure without a taxed unit) or no building at all (empty parcel). The *Fichiers Fonciers* provide the universe of units in France, from 2014 to 2025. For each parcel, I observe the characteristics at the unit level: construction date, floorspace (only for non-industrial units), and destination. I aggregate the units at the parcel level and define the parcel's destination as the destination of at least 90% of units, otherwise the parcel is said to mixed use. Figure A1 summarizes the parcels' destination as of 2025.

New developments are identified thanks to the *Sitadel* database which contains the universe of authorized building permits in France since 2011. I restrict the analysis to projects where (i) at least one housing unit is added on the parcel and (ii) the contractor is a private firm or a physical person (to remove pure social housing projects, which are outside of the scope of my analysis). *Sitadel* provides information on the additional floorspace supplied and the issuance date. Project data may also include the demolished floorspace, contractor characteristics, as well as the dates at which the construction starts and ends.

Housing cost is measured by both rents and sale prices. Rents are measured at the municipal level by the *Observatoire des Loyers* from 2015 onwards. It uses online ads to build a hedonic monthly rent per square metre. Sales data are from the *Demande de Valeurs Foncières* (DVF). I apply a similar hedonic measurement to obtain the sale price per square metre at the municipal level for different housing unit types.[2]

Other sources include the *BD Topo*, a yearly 3D representation of all buildings in France, which measures changes in the building envelope (footprint × height) at the parcel level. Population and income are measured with the *Fidéli* database. Land use regulations are drawn from *Apur*, which provides shapefiles on zoning. In absence of precise floor-area-ratio data at the parcel level, I infer them based on pre-existing building structure.

2.3. Motivating evidence

To provide context on the Paris area and justify the subsequent model's assumptions, I present some empirical facts.

Artificialization reached a steady-state As opposed to other contexts where cities can grow at the extensive margin (i.e, by taking up more land at the outskirts), the Paris area is land constrained. Figure 1 plots the change (in base 1999) in population and land take-up in the Paris area from 1999 to 2021. Land take-up increased by less than a percent over the period, driven by the most rural parts of the Seine-Saint-Denis *département*, in the north of Paris. Over the same period, population grew by 11%: it confirms the densification of the area over the period.

[2]Following Chapelle et al. (2023), I regress the logarithm of the price per square meter on observed housing characteristics (number of rooms, type of dwellings, construction date), municipal fixed effects μ_m, and time fixed effects μ_t. We obtain the hedonic price index by computing $P_{mt} = \exp\left(\mu_m + \mu_t + \sum_i \frac{1}{n_{m,t}} \beta X_i + \frac{1}{2}\sigma^2\right)$ where $n_{m,t}$ is the number of transactions in municipality m at time t and σ is the root mean square error.

Housing is mainly supplied through demolition-reconstruction I match the universe of authorized building permits (2014-2024) with the envelope of building before (2013) and after (2026) the permit was issued. Only permits creating at least one unit and strictly more than 9 square meters (the legal minimum floorspace for a rented unit) are considered. I distinguish three types of constructions: the new units, on empty parcels; the within-transformation, i.e. without demolition of the existing structure; and the demolition-reconstruction, i.e. with a change in the building envelope. 18,873 permits creating housing units were authorized, creating 298,329 new units and 17,981,231 square meters of residential floorspace. 64% of the new units is provided through demolition-construction, 20% is provided through within-building transformation (retrofit, unit splits, office-to-housing transitions), and only 10% is supplied on empty parcels (Figure 2).

It is worth noticing that most building permits delivered over the period create housing (69%). On Figure C4, I compare the sample of permits creating housing units with the ones creating only non-housing units. 79% of permits involving a demolition are supplying new housing units, whereas only 56% of permits without demolition are creating new housing units. This is driven by the within-building transformation of mixed-use and tertiary building, a margin I am not investigating in this paper.

Non-housing parcels produce more housing unit when redeveloped A substantial share of the housing supply is provided on non-housing parcels (10%) and on mixed-use parcels (20%). Demolition-reconstruction and non-housing-to-housing transitions appear to be important margins for housing

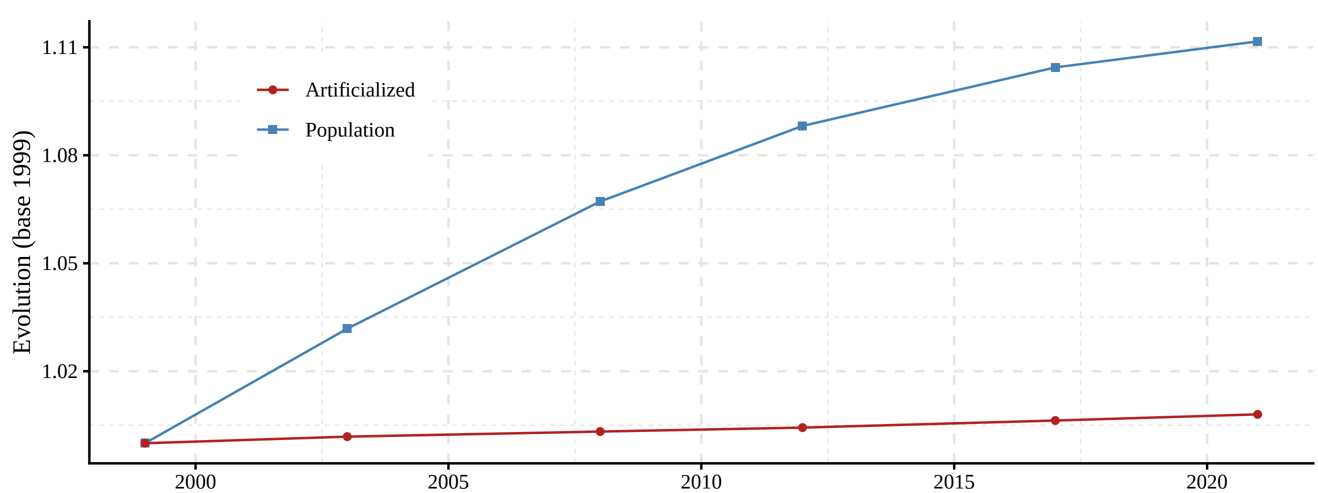

Data: *Mode d'Occupation des Sols*, *Recensement Permanent*. Evolution in base 1999 (=1). Artificialized land is the share of land affected by human actions, except parks, forests, and artificial waterbodies over total area. Population is the total Paris area population.

Figure 1: Evolution of population and land take-up (1999-2021)

densification. Using the sample of permits and the transition typology presented *supra*, I compute the floorspace-to-area ratio (FAR), as total housing floorspace over parcel area, and the units-to-area ratio (UAR), as the total units over parcel area. Higher FAR/UAR indicates more intense housing supply for a given parcel size. Table C3 reports the mean and median ratios, units, and floorspace supplied by project type (new units, demolition, no demolition), and previous land use (house, mixed-use, multi-unit housing, new units, and non-housing). Projects involving demolition and projects set on former non-housing parcels supply more housing in absolute terms (median of 225 and 800 units), but also per square meter (1.12 and 1.60 FAR and 0.17 and 0.30 UAR).

A large number of parcels could be redeveloped I distinguish three types of parcels, all located in areas zoned for housing. First, the *empty* parcels. Those parcels do not host a physical building. They can be a parking lot, an unused plot, a wasteland. Second, the *vacant* parcels. Those parcels still host a physical building but no tax is collected. I focus on parcels that used to host economic activities (i.e, all establishments located on the parcel are closed as of 2024) and not all vacant parcels to limit the count of false positive (a physical building can host economic activities or housing units and be considered as fiscally empty. Appendix A provides more details on the building definition and composition determination.). Third, the *infill* parcels. Those parcels host a building that is significantly lower than the contiguous buildings. For each parcel, I measure the development density around it (average

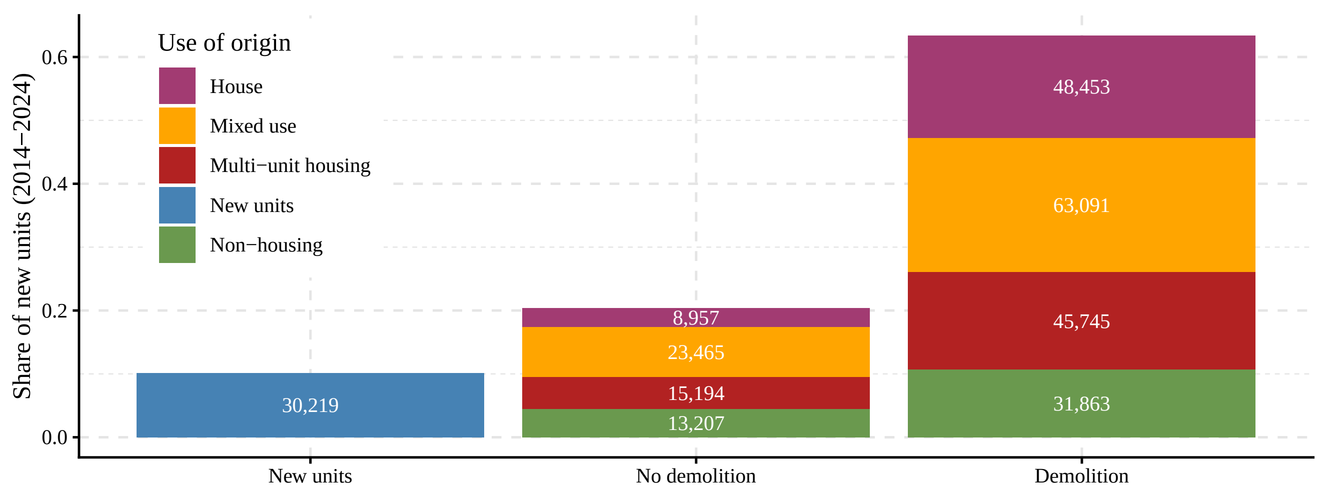

Data: *Sitadel*, *Fichiers Fonciers*, *BD Topo* (2014-2025). New units are housing units built on empty lots; No demolition are housing units provided in building without change in envelope (footprint × height); Demolition are housing units provided in demolished and rebuilt buildings. Use of origin is defined as the main occupation of the parcel before the permit was issued. The numbers correspond to the count of new units by transition type × use of origin.

Figure 2: Housing supply by type of transformation and use of origin

building-coverage ratio and average height in a 100-meter radius). A back-of-the-envelope computation indicates a substantial margin for redevelopment if those parcels were similar to the ones in the radius. Their demolition-reconstructions could increase housing supply, especially through higher height.

Empty and vacant parcels are mainly located in the closest Paris suburbs (Fig. C3). Note that all projects are located in the 98th percentile of rents at the France level, indicating high potential profitability. Table 1 summarizes the potential supply. The absence of redevelopment of those parcels, even with high demand and high price, suggests important frictions in the extensive margin of redevelopment.

Table 1: Potential housing supply on empty and vacant parcels

Type	Parcels	Potential floorspace	Potential units
Empty	1,460	720,504	11,289
Vacant	265	330,279	5,367
Vertical potential	1,490	1,233,345	20,556
Total	3,215	2,284,128	37,212

Data: Fichiers Fonciers, BD TOPO. Potential parcels are located on residentially zoned areas. Empty parcels are parcels without physical building. Vacant parcels are parcel with no taxed units and past economic activity. Data as of 2024. Potential floorspace is measured as the median building coverage ratio × parcel area × median height, divided by a 2.5 meter ceiling height. Potential units assume a size of 60 sqm per unit.

Project length is heterogeneous across projects Thanks to the building permit data, I can track projects timeline. For each permit, I observe the date at which it was validated. For around 60% of permits, I observe the construction start and end date. The missing values are either due because construction has not started yet or due to the absence of reporting from the developers. For projects where the data are available, I measure the construction length, as the difference between the start and end date, and the project length, from the permit validation to the construction end (by law, permits must be rejected or accepted 3 months after the application). The difference between the project and the construction lengths may capture construction delays (credit constraint, capital acquisition, labour acquisition) or other frictions (extreme climate events). These measures highlight an important mechanism: projects which take more time to be completed may bear higher fixed costs and uncertainty.

Table 2 reports both an omnibus test and the full set of pairwise comparisons. The Kruskal–Wallis rank sum test rejects the null of equal duration distributions across uses of origin ($\chi^2 = 118.7$, df $= 7$, $p < 2.2 \times 10^{-16}$). Because the omnibus test only establishes that some category differs, I rely on the pairwise Wilcoxon comparisons (Benjamini–Hochberg adjusted) to locate the differences. They reveal two poles. Projects originating from tertiary and secondary activity uses form the slow extreme: tertiary-activity conversions differ from nearly every other category ($p < 0.01$), while the two activity categories are statistically indistinguishable from one another. At the fast end, new units differ sharply from mixed-use, no-unit, secondary- and tertiary-activity origins. The remaining categories (annex, mixed

Table 2: Construction time across uses of origin: Kruskal–Wallis test and pairwise Wilcoxon rank sum comparisons

	Annex	Mixed	Multi	New unit	No unit	Secondary	Single
Mixed	0.383						
Multi	0.620	0.058*					
New unit	0.258	<0.001***	0.602				
No unit	0.804	0.089*	0.316	0.003***			
Secondary	0.038**	0.051*	0.010**	<0.001***	0.020**		
Single	0.634	<0.001***	0.804	0.089*	0.049**	0.003***	
Tertiary	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	0.708	<0.001***
Kruskal–Wallis: $\chi^2 = 118.7$, df $= 7$, $p < 2.2 \times 10^{-16}$							

Data: Sitadel. Cells report Benjamini–Hochberg adjusted p-values from pairwise Wilcoxon rank sum tests of the null that construction time is equally distributed between the two uses of origin (column and line). The Kruskal–Wallis statistic tests the joint null of equal distributions across all uses of origin. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Abbreviations: Ann = Annex, Mix = Mixed use, MFU = Multi-family unit, New = New unit, NoU = No unit, Sec = Secondary activity, SFU = Single-family unit, Ter = Tertiary activity.

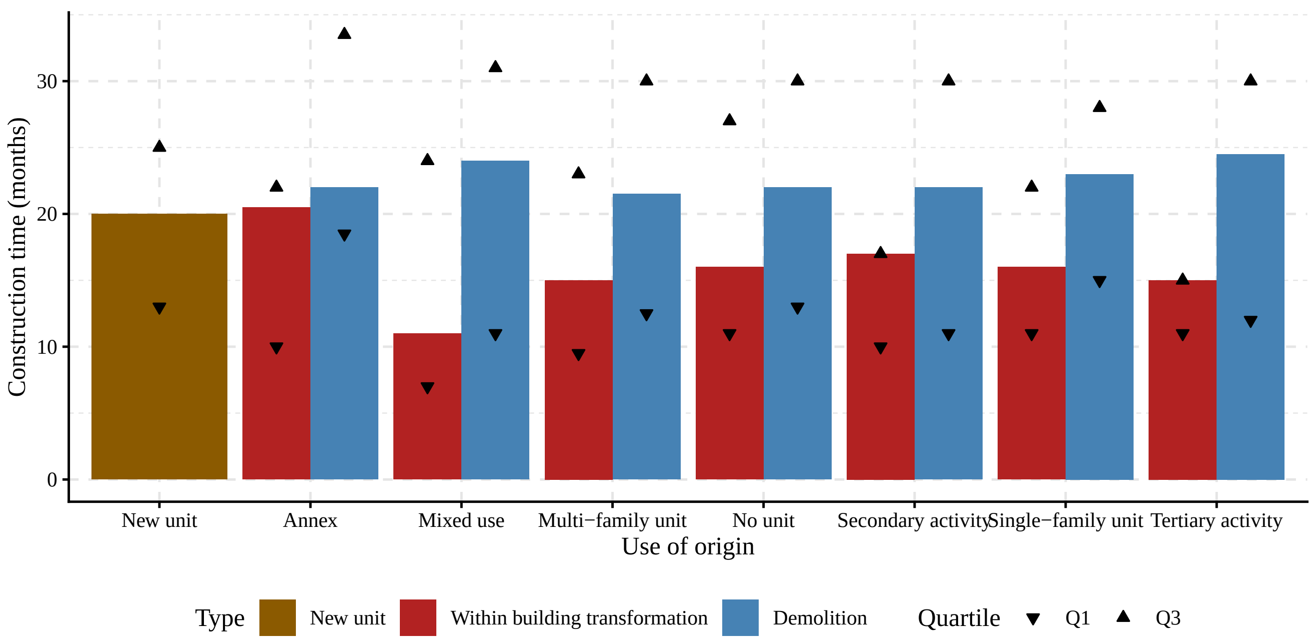

Data: Sitadel. First quartile (down triangle), median (bar), and third quartile (up triangle) of construction time (in months) by project type (new unit, within-building transformation, demolition) and parcel's use of origin. Construction time is the time difference between the construction start date and the building completion date, for permits creating at least one new housing unit.

Figure 3: Construction time across uses of origin

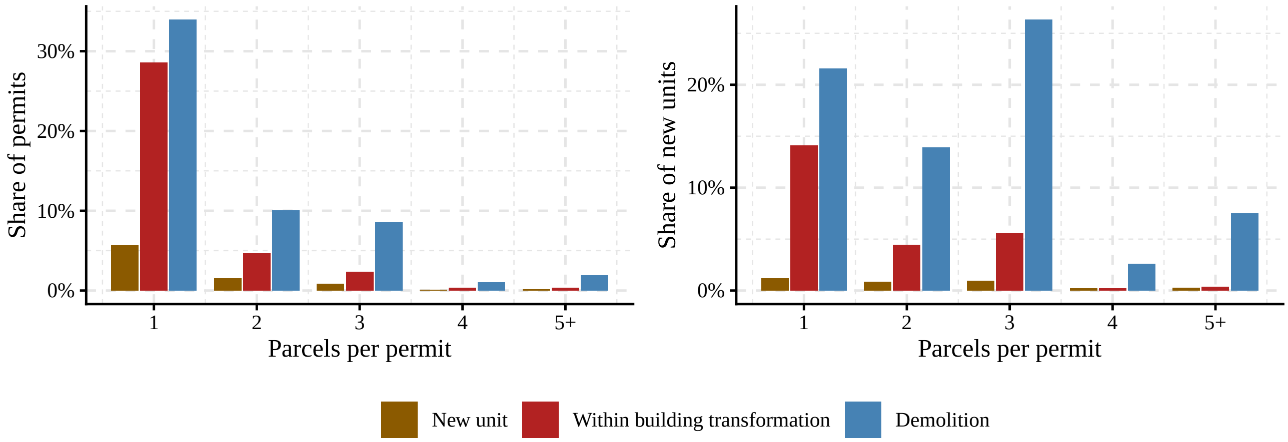

Data: Sitadel. Share of permits (left panel) and new units (right panel) by number of parcels involved in a permit, ventilated by project type (new unit, within-building transformation, and demolition).

Figure 4: Share of projects and new units, by count of parcels

use, multi-family, no unit, and single-family) exhibit indistinguishable construction time.

Figure 3 corroborates this ordering. For a given use of origin, demolition-based redevelopment (blue) is systematically slower than within-building transformation (red), consistent with building demolition. The activity-origin categories (secondary and tertiary) both exhibit the highest median construction times. New-unit construction (orange) sits at the low end of the distribution, consistent with the easiest construction processes. Similar results are found by Glancy et al. (2026) using US data, who also account for the planning time.

Land assembly is an important redevelopment margin but without observable premium Large projects require several parcels (Figure C5). The median project creating more than 150 units requires 5 parcels. Above 11 new units, most projects require strictly more than 1 parcel. 63% of new units are supplied in multi-parcel projects and 31% of permits span several parcels (Figure 4).

Is there a land price premium for assembled parcels? I retain only projects for which all parcels are acquired before the permit is authorised. If a parcel is sold several times before the permit, I keep only the transaction closest to the permit date. For each project, I classify the acquisition pattern into three mutually exclusive types: single (one-off transaction), naive (several mutations, all within a window of 3 months of one another), and Sequential (mutations spread beyond that window). The dependent variable is the transaction price per square metre of land, trimmed at the 1st and 99th percentiles, sold and bought by private actors. The estimating equation is:

$$\text{Price per square meter}_i = \beta_1 \text{Naive}_{p(i)} + \beta_2 \text{Sequential}_{p(i)} + X_i + \gamma_t + \gamma_c + \varepsilon_i, \tag{1}$$

where Single is the omitted category, X_i collects parcel controls (building floorspace, project parcel count, use of origin) and transaction controls (buyer and seller type fixed effects), as well as municipal and time fixed effects. Standard errors are clustered at the municipal level.

The identification logic is that, conditional on the set of controls, β_1 and β_2 isolate the price differential associated with parcels bought as part of a multi-parcel assembly rather than as stand-alone acquisitions. The interpretation follows the hold-up reasoning of Brooks and Lutz (2016) and the land-assembly literature (Miceli and Sirmans, 2007): once a developer has begun assembling land for a project, the marginal landowner holds a complementary input and can extract a premium. Assembled parcels should be sold at a higher price than stand-alone parcels. The results give qualified support to the hold-up hypothesis, specifically in naive (near-simultaneous) assembly rather than sequential purchasing. In the preferred single-gap table (Table 3), naive-assembly parcels carry a premium of roughly 205–236€ per m^2 relative to stand-alone parcels, economically large against a mean price of 1,984€/m^2 (about a 10–12% premium). The estimate is stable and significant at the 5% level across all five columns as controls and fixed effects are added. The sequential coefficient, by contrast, is small and never significant, suggesting no measurable premium once acquisitions are spread out in time. Table C5 reports robustness tests for different time windows (0, 1, 3, and 6 months). Effects are sensitive to the choice of the windows to define a sequential assembly.

To investigate the mechanism further, I exploit within-assembly variation: restricting to sequential projects, I rank each mutation by its sale date and estimate a project-fixed-effect specification. The position coefficients are identified off the ordering of purchases inside a given assembly, differencing out any project-level selection into assembling. The within-assembly position figure (Figure C6) shows a flat, statistically indistinguishable from zero marginal effect of sale position on land price. I find no evidence of the rising, pivotal-seller premium that a sequential hold-up story would predict. Taken together, the evidence points to a premium on bundling parcels rather than to escalating extraction due to hold–out. The data cannot rule out *ex-ante* negotiation procedures leading to buying parcels in block. This effect does not come as a surprise as sequential assembly is marginal among multi-parcel projects; 90% of parcels used in a project are bought within less than a quarter, 95% within less than a year.

Developer market is relatively fragmented and individual projects are small The housing supply literature has long relied on the perfect competition assumption on the developer market. However, empirical evidence (Favilukis et al., 2026; Lefebvre and Daly, 2026; Charreire et al., 2025; Quintero, 2023) points to a relatively concentrated market, where only a few actors are supplying most of new units. Developers are identified with their SIREN identifier (a firm identifier that nests all establishments' identifiers).[3]

Table 4 and Figure 5 characterise the degree of concentration in land acquisition across the Paris area region. Two patterns stand out. First, concentration is a local phenomenon that dissolves sharply with spatial aggregation. At the municipal level the median market is served by an effective number of $1/\text{HHI} \approx 17$ developers once self-builds are excluded (*With dev.* sample), falling to roughly 13 among projects of five dwellings or more. The same statistic rises by an order of magnitude at the departmental level ($\approx$ 170) and again at the regional level ($\approx$ 330), and the four-firm concentration ratio contracts from 35.9% to 11.4% to 7.2% along the same progression. This gradient suggests spatially segmented land markets. Second, the Lorenz curves' curvatures are pronounced: a small share of developers accounts for a large share of developed units, with the top decile of developers developing the majority

Table 3: Land-price premium in land assembly

	(1)	(2)	(3)	(4)	(5)
Naive assembly	235.787**	234.827**	219.760**	229.435**	204.484**
	(98.821)	(98.988)	(90.048)	(88.396)	(87.811)
Sequential assembly	-3.618	2.601	83.677	102.256	73.244
	(114.670)	(114.537)	(98.702)	(96.897)	(101.757)
Year FE	✓	✓	✓	✓	✓
Municipality FE	✓	✓	✓	✓	✓
Built area		✓	✓	✓	✓
Seller			✓	✓	✓
Buyer				✓	✓
Parcel count					✓
Mean Dep. Var	1,984€	1,984€	1,984€	1,984€	1,984€
Observations	9,482	9,482	9,482	9,482	9,482
R^2	0.533	0.534	0.547	0.554	0.554

Data: Sitadel, DVF. OLS regression (1). All models include year and municipal fixed effects. Built area and parcel count are continuous variables. Seller and Buyer are categorical variables stating if the seller (resp. buyer) is a physical person, a private firm, a public actor, or if its type is unknown. Sample is the set of construction projects for which all permitted parcels are traded prior the permit authorisation date. Naive assembly means that several parcels were bought within a 3-month window. Sequential assembly means that several parcels were bought outside a 3-month window. Single, the case where only one parcel is traded, is the omitted category. Municipal-level clustered standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

[3] Some firms can have several SIREN identifiers even though they are the same. The results presented are an *lower bound* of the real market concentration.

of parcels. Taken together, the evidence indicates that market power in local land acquisition does not arise from region-wide oligopoly (no single developer is large relative to aggregate supply) but from the conjunction of sharp local segmentation.

Table 4: Developer concentration in land acquisition

	Sample		
Statistic	All	With dev.	Dev. & large (≥5)
Municipal			
Active developers	101	35	25
HHI	2.00	5.79	7.44
CR4	18.6%	35.9%	42.3%
Effective devs.	50.0	17.3	13.4
Department			
Active developers	4,784	1,360	916
HHI	0.15	0.60	0.85
CR4	5.5%	11.4%	13.8%
Effective devs.	661.2	171.5	120.5
Regional			
Active developers	16,810	4,249	2,789
HHI	0.07	0.30	0.43
CR4	3.4%	7.2%	8.7%
Effective devs.	1452.2	332.8	233.4

Data: Sitadel. Median concentration indices across spatial levels. Active developers in the count of unique developer identifiers and developer-less projects. HHI is a Herfindahl-Hirschmann Index (0-1 scale). CR4 is the share of parcels developed by the 4 largest developers. Effective developers is $1/\text{HHI}_i$, i.e. the number of equal-sized developers that would produce the same HHI. The first column is the universe of permits creating at least one housing unit. The second column restricts the sample for permits with a referenced developer. The third column restricts to permits creating more than 5 units and with a developer.

Individual projects are atomistic with respect to the housing stock, and in turn to the price gradient (5). Even the single largest project represents a median of 1.35% of the existing dwelling stock in its municipality, rising to 6.2% at the 99th percentile and a maximum of 11.8%. These shares fall at the departmental level (at most 0.12%) and the regional level (0.02%). No developer, therefore, is large enough to shift the stock of housing, either locally or in aggregate, which justifies treating each developer as a price-taker with respect to the gradient. The picture differs for the annual flow of new units: because construction is lumpy and municipalities may contain few projects in a given period, the largest project accounts for a median of 14.6% of municipal new supply and up to 77% in the most concentrated case. This lumpiness is not relevant for the price gradient determination. Developers take as given the price gradient, which is pinned down by the stock of dwellings rather than by a single period's construction. Atomicity and local market power are properties of two different markets and are not in tension. A developer is atomistic with respect to the housing market, in the sense that no project moves the stock

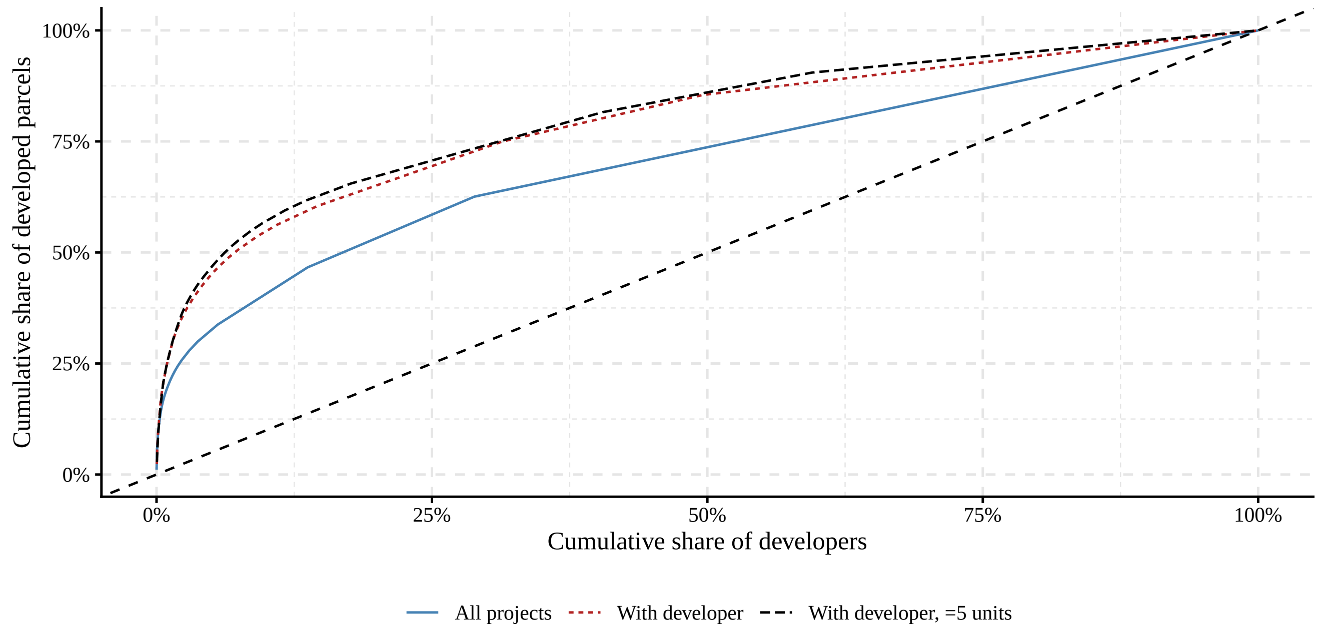

Figure 5: Cumulative distribution of developers

Data: Sitadel. Lorenz curves of developed parcels against the cumulative share of developers, ranked largest first. Developers are identified by SIREN. The plain blue curve is the universe of permits creating at least one housing unit. The small-dashed red line restricts the sample for permits with a referenced developer. The long-dashed black line restricts to permits creating more than 5 units and with a developer. The dashed 45° line denotes perfect equality, greater curvature indicates greater concentration.

that pins down the rent gradient, which is what Table 5 establishes and what the model uses when it lets agents take the price path as given. She is not atomistic with respect to a given landowner, because acquiring a specific parcel is a bilateral negotiation in a locally segmented market, which is what Table 4 establishes and what the model uses when it splits the surplus. The first fact governs the price process, the second governs who captures the gains from a transaction at that price.

Table 5: Largest project as a share of aggregate stock and flow

	Share of aggregate stock			Share of aggregate flow		
Level	Median	P99	Max	Median	P99	Max
Municipal	1.35%	6.19%	11.83%	14.60%	74.91%	76.91%
Department	0.10%	0.12%	0.12%	0.98%	2.09%	2.12%
Regional	0.02%	0.02%	0.02%	0.29%	0.29%	0.29%

Data: Sitadel, Fichiers Fonciers. Each cell reports the largest project's share, computed within spatial units. The stock is defined as the number of existing dwellings. The aggregate flow is the pooled (2014-2025) authorized housing units. The median, the 99th percentile, and the max are computed across considered spatial units. Interpretation: the largest project represents a median share of 1.35% of the existing units at the municipal level.

3. Dynamic model of housing supply

I depart from the quantitative spatial model literature and assume that profits and costs are not idiosyncratically shocked. Developers instead observe and condition their actions on rental prices R_t, which adjust endogenously to clear the housing market. The model belongs to the tradition of option value in investment decisions with an optimal stopping time (Dixit and Pindyck, 1994). Recent developments include Lebret et al. (2026), from which I also gain inspiration. Once a parcel is bought for redevelopment, the sunk cost is paid and the decision is irreversible. For simplicity and to remain within a relevant time horizon, I assume that new structures do not depreciate and cannot be further redeveloped. The model characterises a transition path of housing supply over a finite horizon, usually twenty years.

The framework adds two frictions to this benchmark. First, construction is not instantaneous. A project started today delivers floorspace one to two years later, depending on the use of origin, so revenue is discounted relative to the date at which the option is exercised. Second, the parcel changes ownership through a bilateral transaction rather than through frictionless bidding. The landowner and the developer split the surplus by generalised Nash bargaining. The competitive benchmark in which the landowner captures everything is the limiting case in which her bargaining weight equals one. Both frictions drive a wedge between the value a redevelopment creates and the price at which land is traded.

Developers make two decisions. They first decide whether to acquire a parcel from an incumbent landowner. Conditional on acquisition they choose the optimal housing supply. One key characteristic of the model is that developers do not inherit a parcel but must buy it before any development takes place. In stage one, taking acquisition as given, the developer chooses density h^{new} to maximise the value created net of construction costs. In stage two the parties bargain over the parcel, and the transaction price splits the resulting surplus. The density choice remains independent of the price at which the parcel trades, because it maximises the total surplus and the split is a constant share of that surplus.

Finally, I account for heterogeneity in current land use. Parcels are used for housing or for non-housing purposes, or are vacant. The continuation value from waiting differs across uses, which creates variation in redevelopment timing that is independent of fixed costs. The time-to-build also differs across uses. The fixed cost is identified as the difference that rationalises the gap between the delay-adjusted redevelopment value and the observed land price, net of the developer's bargaining rent. I further decompose that fixed cost into a construction component and a regulatory component.

3.1. Demand and price determination

The housing market clears through rental price adjustment. The demand shifter A_t captures exogenous changes in housing demand, such as population growth, income and amenities, and is the sole source of aggregate uncertainty. It evolves according to a geometric Brownian motion:

$$dA_t = \mu A_t \, dt + \sigma A_t \, dW_t \tag{2}$$

where μ is the expected growth rate of demand, σ its volatility and dW_t a standard Wiener process. Denoting by Q_t the aggregate housing supply at time t, the inverse demand function is:

$$R_t = A_t \cdot Q_t^{-\alpha} \tag{3}$$

where $\alpha > 0$ is the inverse demand elasticity. Higher housing supply reduces rental prices, which creates a general equilibrium feedback on redevelopment incentives. Section 3.6 shows that the construction lag turns Q_t into a predetermined object, so that this feedback arrives with a delay.

Rental price dynamics and the atomistic developer In practice developers and landowners observe the rental price R_t directly rather than the underlying demand shifter A_t. From equation (3) the equilibrium rent is the product of a diffusion, A_t, and the aggregate supply path $Q_t^{-\alpha}$, which moves discretely each time a project is delivered. The realised process R_t is therefore not a geometric Brownian motion. It is a diffusion punctuated by downward jumps at delivery dates.

The central assumption is that each parcel i is atomistic. It has measure zero in the continuum of parcels, so any individual redevelopment decision has a negligible effect on aggregate supply, as established in the motivating facts. Each landowner and developer therefore takes the prevailing supply level as given and does not forecast its future path, internalising neither their own effect on aggregate supply nor its stochastic dynamics. Conditional on Q_t being treated as fixed, the only perceived source of variation is A_t, so the perceived rental price process is a geometric Brownian motion:

$$dR_t = \delta R_t \, dt + \sigma R_t \, dW_t \tag{4}$$

where σ is the volatility of A_t and δ is the expected growth rate of rental prices from the individual agent's perspective. I calibrate δ externally from the net residential yield in Section 4 and treat it as a primitive of the agent's price expectations. Assumption (4) also delivers the expectation that prices the construction lag, since $\mathbb{E}_t[R_{i,t+s}] = R_{it}e^{\delta s}$ for any horizon s.

This atomistic assumption keeps the individual stopping problem tractable while preserving general equilibrium feedback at the aggregate level. The wedge between the perceived and the realised process is exactly the supply channel. Each agent perceives a constant drift δ and solves a standard stopping problem, but the equilibrium price level responds endogenously to supply. As projects are delivered, Q_t rises and the market-clearing rent falls for all parcels, so agents re-optimise from a lower price level than they had anticipated. This dampening operates entirely through the level of rents in equation (3) and not through the perceived drift.

Spatial heterogeneity Equation (3) implies a single price R_t across all locations. To allow a spatially heterogeneous price gradient I define ϕ_i, a time-invariant location quality shifter capturing amenities, access to Paris and neighbourhood quality, such that $R_{it} = \phi_i R_t$ is the rental price in location i at time t. Aggregate shocks then affect all neighbourhoods proportionally, while ϕ_i determines persistent spatial differences in rent levels. I calibrate ϕ_i from observed baseline rents.

3.2. Intensive housing supply with time-to-build

The construction lag A project started at date t on parcel i is completed at $t + \ell_i$. The lag $\ell_i = \ell_{u(i)}$ is deterministic and indexed by the use of origin $u(i)$. I make four assumptions. Density is chosen at the start date and cannot be revised. The project cannot be abandoned once started[4]. The standing structure is emptied at the start date, so its rental flow is forgone from t rather than from $t + \ell_i$. Fixed costs are sunk at the start date.

Two discount factors summarise the lag. Rental income begins at completion, so under the perceived process (4) the present value of a perpetual flow starting at $t + \ell_i$ is the undelayed value scaled by λ_i^R. Construction outlays are spread at a constant rate over the build, so their present value is the undelayed cost scaled by λ_i^C:

$$\lambda_i^R = e^{-(\rho-\delta)\ell_i} \qquad \text{and} \qquad \lambda_i^C = \frac{1 - e^{-\rho \ell_i}}{\rho \ell_i}$$

where ρ is the discount rate.

[4]It rules out the compound option studied by Majd and Pindyck (1987) and the reversal of the uncertainty comparative static obtained by Bar-Ilan and Strange (1996). No precise data exist on stopped construction projects in France.

The two factors enter through their ratio,

$$\kappa_i \equiv \frac{\lambda_i^R}{\lambda_i^C}$$

which is the effective discount applied to the net cash flow. Delaying the cost alongside the revenue offsets most of the delay. κ_i is therefore much closer to one than λ_i^R is. Under the calibration $\kappa_i < 1$, so the lag can be seen as a (moderate) tax on the project.

Project value and cost A developer acquiring parcel i at time t and building h^{new} obtains the present value of perpetual rental income from completion onwards:

$$\mathcal{V}_i(R_{it}, h^{\text{new}}) = \lambda_i^R \cdot \frac{R_{it} \cdot h^{\text{new}}}{\rho - \delta}$$

Note that $\rho > \delta$ is required for the present value to be finite, which the calibration ensures. Development costs $C_i(h^{\text{new}})$ are the sum of convex construction costs and a fixed cost F_i. Assuming a Cobb-Douglas relation between land and capital in housing production (Combes et al., 2021), the unit cost is $c_i = c_0 T_i^{-1/\zeta}$, where T_i is parcel area in square metres and c_0 is the capital cost parameter:

$$C_i(h^{\text{new}}) = F_i + \lambda_i^C \cdot c_0 T_i^{-1/\zeta} (h^{\text{new}})^{1+1/\zeta} \tag{5}$$

where F_i is the fixed cost. It collects all costs that do not vary with project scale. The developer chooses density to maximise the value created:

$$\pi_i(R_{it}) = \max_{h^{\text{new}} \leq \bar{h}_i} \lambda_i^R \frac{R_{it} h^{\text{new}}}{\rho - \delta} - \lambda_i^C c_0 T_i^{-1/\zeta} (h^{\text{new}})^{1+1/\zeta} \tag{6}$$

where $\bar{h}_i$ is the regulatory density limit. I write π_i^0 for the profit evaluated at $\ell_i = 0$, which is the frictionless profit of the model.

Lemma 1 *The construction lag rescales the redevelopment payoff. For every parcel and every rent,* $\pi_i(R) = \lambda_i^C \, \pi_i^0(\kappa_i R)$*, in the interior regime as well as in the constrained regime.*

The lemma follows from factoring λ_i^C out of (6), which leaves the frictionless problem evaluated at the delay-adjusted rent $\kappa_i R$.

Optimal density The first-order condition of (6) gives the unconstrained optimal density:

$$(h^{\text{new}})^{\star} = \left[\frac{\zeta\, \kappa_i R_{it}}{c_0(\zeta+1)(\rho-\delta)}\right]^{\zeta} T_i \tag{7}$$

Optimal density is linear in parcel area, so developers build at a constant floorspace-to-land ratio when unconstrained. That ratio is increasing in the rental price and in the expected rental price growth rate. It is decreasing in the construction lag, and falls by the factor κ_i^{ζ}. The intuition is direct. Delay "taxes" the revenue stream more heavily than the construction spend. The marginal square metre is worth less at the moment the density is chosen. Time-to-build makes projects smaller and delays them.

Regulatory constraint Equation (7) is an interior solution provided $h^{\star}$ lies below the maximum allowed supply $\bar{h}_i$. Define the saturation rent as the rental price at which the unconstrained optimum equals the regulatory limit:

$$h^{\star} = \bar{h}_i \implies \bar{R}_i = \frac{1}{\kappa_i} \cdot \frac{c_0(\zeta+1)(\rho-\delta)}{\zeta} \left(\frac{\bar{h}_i}{T_i}\right)^{1/\zeta} \tag{8}$$

Parcels with tighter regulatory constraints relative to their area reach saturation at lower rental prices. The lag raises the saturation rent by the factor κ_i^{-1}, so it widens the interior regime. It shifts parcels away from the corner at which the density cap binds and towards the region in which density responds to prices. The optimal density is then:

$$h_i^{\text{new}}(R_t) = \begin{cases} \left[\dfrac{\zeta\, \kappa_i R_{it}}{c_0(\zeta+1)(\rho-\delta)}\right]^{\zeta} T_i & \text{if } R_t < \bar{R}_i \qquad \text{(Interior)} \\ \bar{h}_i & \text{if } R_t \geq \bar{R}_i \qquad \text{(Constrained)} \end{cases} \tag{9}$$

Redevelopment value Substituting (9) into (6) gives the value created by redevelopment, gross of the fixed cost. In the interior regime:

$$\pi_i(R_{it}) = \frac{\lambda_i^C}{\zeta} \left[\frac{\zeta\, \kappa_i R_{it}}{(\zeta+1)(\rho-\delta)}\right]^{\zeta+1} c_0^{-\zeta} T_i$$

and in the constrained regime:

$$\pi_i(R_{it}) = \lambda_i^R \frac{R_{it}\bar{h}_i}{\rho-\delta} - \lambda_i^C c_0 T_i^{-1/\zeta} \bar{h}_i^{(\zeta+1)/\zeta}$$

The redevelopment value is convex in the rental price and linear in parcel area in the interior regime, and linear in the rental price in the constrained regime. Larger parcels are more valuable because they benefit from lower unit construction costs and accommodate more housing at the optimal density. Relative to the frictionless benchmark, the interior value is scaled by $\lambda_i^C \kappa_i^{\zeta+1}$, which compounds a smaller project with a discounted cash flow.

3.3. The land transaction

Surplus and disagreement points The parcel changes hands at the moment of redevelopment. The landowner's disagreement point is the value of keeping the parcel in its current use, V_i^w, defined in Section 3.4. The developer's disagreement point is the value of her next best project, which I normalise to zero, as she holds nothing else. The total surplus created by the transaction is therefore

$$S_i(R_t) = \pi_i(R_{it}) - F_i - V_i^w \tag{10}$$

which is the value the redevelopment adds over leaving the standing structure in place. Trade occurs whenever $S_i \geq 0$.

Nash bargaining The parties split the surplus according to the generalised Nash solution with landowner weight $\theta \in (0, 1]$. The transaction price and the developer's rent are

$$p_i^L(R_t) = V_i^w + \theta\, S_i(R_t), \qquad \Pi_i^D(R_t) = (1-\theta)\, S_i(R_t) \tag{11}$$

Setting $\theta = 1$ recovers the competitive model in which developers bid the price up to the full fixed cost value and earn zero expected profit. Values of θ below one describe a thin local market in which a landowner faces few credible counterparties. Two properties of (11) matter for what follows. The price is increasing in the redevelopment value and decreasing in the fixed cost at the rate $\partial p_i^L / \partial F_i = -\theta$, so landowners bear only a fraction of transition costs rather than their entirety. The density choice (9) is unaffected: it maximises S_i and both parties receive a constant share of the maximised surplus.

3.4. Extensive housing supply with heterogeneous current uses

Landowners observe the current rental price and, from the atomistic perspective, treat R_t as a geometric Brownian motion with drift δ and volatility σ. They also observe the state of their parcel and the price function $p_i^L(R_t)$ implied by (11). All agents discount with the same rate ρ. At each instant the landowner chooses between waiting, which means continuing the current use, or selling at the bargained price.

3.4.1 Continuation values by current use

Housing parcels The continuation value from renting the existing structure is

$$V_i^w(R_{it}) = \frac{R_{it} \cdot h_i^{\text{old}}}{\rho - \delta} \tag{12}$$

where the denominator reflects discounting at ρ and expected rental growth at δ.

Commercial and industrial parcels Denoting non-housing rental value by v_i^C and its expected growth rate by δ^C, the continuation value is

$$V_i^w = \frac{v_i^C}{\rho - \delta^C} \tag{13}$$

If non-residential demand grows more slowly than residential demand, these parcels have lower continuation values and redevelop earlier, all else equal. I remain agnostic about the actual destination of non-housing parcels.

Vacant parcels

$$V_i^w = 0 \tag{14}$$

No current revenue implies no opportunity cost of redevelopment. Vacant parcels redevelop first.

3.4.2 Optimal stopping problem

Under uncertainty and irreversibility the landowner faces a real options problem. The optimal stopping time is

$$\tau_i^\star = \inf\{t \geq 0 : R_t \geq R_i^\star\}$$

and the optimal threshold satisfies

$$R_i^\star = \frac{\beta}{\beta - 1} \cdot R_i^{\text{naive}} \tag{15}$$

where $\beta > 1$ is the positive root of the fundamental quadratic

$$\frac{1}{2}\sigma^2\beta(\beta - 1) + \delta\beta - \rho = 0 \tag{16}$$

that is,

$$\beta = \frac{1}{2} - \frac{\delta}{\sigma^2} + \sqrt{\left(\frac{\delta}{\sigma^2} - \frac{1}{2}\right)^2 + \frac{2\rho}{\sigma^2}}$$

The multiplier $\beta/(\beta - 1) > 1$ captures the option value of waiting. Higher volatility raises β and hence the threshold, which delays redevelopment. A higher discount rate does the same, while faster expected rental growth lowers the threshold. The naive threshold depends on the current use through V_i^w and on the construction lag through κ_i and λ_i^C. In the interior regime it has the closed form

$$R_i^{\text{naive}} = \frac{1}{\kappa_i} \cdot \frac{(\zeta + 1)(\rho - \delta)}{\zeta} \left[\frac{\zeta \left(F_i + V_i^w\right)}{\lambda_i^C \, c_0^{-\zeta} T_i} \right]^{\frac{1}{\zeta+1}} \tag{17}$$

Two properties of (17) organise the rest of the paper.

Proposition 1 *The construction lag does not affect the option value multiplier. It enters the exercise payoff multiplicatively by Lemma 1, so the fundamental quadratic (16) and its root β are unchanged. Holding F_i and V_i^w fixed, the lag rescales the threshold by*

$$\frac{R_i^\star(\ell_i)}{R_i^\star(0)} = \frac{1}{\kappa_i \, (\lambda_i^C)^{1/(\zeta+1)}} \; > \; 1$$

The proposition isolates what time-to-build does and does not do in this framework. It raises the trigger, and by a magnitude an order smaller than the raw discount λ_i^R would suggest, because the cost is delayed alongside the revenue. It does not change the value of waiting, since the lag is a deterministic multiplicative wedge and carries no additional optionality. Recovering the reversal of Bar-Ilan and Strange (1996), under which longer lags can make the trigger fall in σ, would require an abandonment option during construction, which the irreversibility assumption excludes.

Proposition 2 *Conditional on the fixed cost, bargaining is neutral. The developer solves $\max_\tau \mathbb{E}\left[e^{-\rho\tau}(1-\theta)S_i(R_\tau)\right]$, and $(1-\theta)$ is a positive constant, so the stopping rule coincides with the one that maximises the joint surplus. The threshold (15) is therefore invariant to θ.*

3.4.3 Comparative statics across parcel types

The model predicts heterogeneous redevelopment patterns by current use. Vacant parcels have the lowest thresholds, since only the fixed cost and the lag delay them, and they redevelop first. Non-housing parcels with high current revenue have high continuation values and redevelop only when housing values substantially exceed the value in current use. Dense housing parcels rarely redevelop unless the densification potential $\bar{h}_i - h_i^{\text{old}}$ is large. Across all types, redevelopment timing depends heavily on the distribution of fixed costs.

3.5. Identification of fixed costs

The residual All structural parameters other than F_i are calibrated externally, and the construction lag and the bargaining weight are calibrated externally as well. Inverting the price relation (11) at the observed rent then recovers the fixed cost as a fully determined residual:

$$F_i = \pi_i(R_{i0}) - V_i^w - \frac{p_i^L - V_i^w}{\theta} \tag{18}$$

which collapses to $F_i = \pi_i^0 - p_i^L$ when $\ell_i = 0$ and $\theta = 1$ — the case in which the landowner appropriates the whole surplus and construction is instantaneous. Equation (18) makes the identification problem explicit. A longer lag lowers π_i, and a lower landowner weight raises the surplus imputed from a given observed price. Both reduce the recovered fixed cost, and neither is separately identified from within the residual.

Proposition 3 *Let π_i^0 denote the frictionless redevelopment value. For any parcel with positive surplus, the gap between the frictionless value and the observed land price decomposes exactly into three terms:*

$$\underbrace{\pi_i^0 - p_i^L}_{\text{observed gap}} = \underbrace{\pi_i^0 - \pi_i}_{\text{delay wedge}} + \underbrace{(1-\theta)S_i}_{\text{developer rent}} + \underbrace{F_i}_{\text{fixed cost}}$$

What the data discipline Proposition 2 implies that timing and prices identify different parameters. Only the fixed cost enters the threshold, so observed redevelopment behaviour disciplines F_i. The bargaining weight enters only the price, so the cross-section of land prices disciplines θ given F_i. The cross-section also bounds θ directly, because a fixed cost cannot be negative:

$$F_i \geq 0 \iff \theta \geq \theta_i \equiv \frac{p_i^L - V_i^w}{\pi_i(R_{i0}) - V_i^w} \tag{19}$$

The right-hand side is computable without knowledge of F_i and delivers a data-driven floor for the calibration.

Incidence Equation (11) also gives the incidence of any policy that lowers the fixed cost. A grant of one euro raises the surplus by one euro, of which θ is capitalised into the land price and $1-\theta$ accrues to the developer. Under efficient bargaining the effect on the exercise decision is independent of θ, so the split governs who receives the transfer without governing whether the project happens.

3.6. General equilibrium with staggered delivery

Irreversibility and the running maximum Redevelopment is irreversible, so the state variable governing the extensive margin is not the current price but the running maximum

$$\mathcal{M}_t \equiv \max_{s\leq t} R_s$$

Parcel i has started by time t if and only if $R_i^\star \leq \mathcal{M}_t$. With time-to-build, starting and delivering are distinct events. Parcel i has delivered by time t if and only if it started by $t-\ell_i$, that is if and only if $R_i^\star \leq \mathcal{M}_{t-\ell_i}$. When prices fall below a previous peak no new parcels are triggered and the stock of started projects is frozen, while deliveries continue to arrive from the pipeline.

Density at the moment of starting Density is chosen once, at the start date, and fixed thereafter:

$$h_i^{\text{new}} \equiv h_i^{\text{new}}(R_{\tau_i^\star})$$

This is a direct consequence of irreversibility. Post-development density adjustment is not permitted, so h_i^{new} depends on the price prevailing when construction begins and not on the price at delivery. In the cross-section, parcels started at high prices are denser than those started at low prices, and the lag drives a wedge between the price that determined a building's size and the price it faces on completion.

Aggregate supply and market clearing Let $\bar{Q} \equiv \int_{\{i:u_i=\text{housing}\}} h_i^{\text{old}}\,di$ denote the standing housing stock and let

$$\Delta_u(m) \equiv \int_{\{i:u(i)=u,\ R_i^\star\leq m\}} \left[h_i^{\text{new}}(R_{\tau_i^\star}) - h_i^{\text{old}}\right] di \tag{20}$$

denote the net floorspace added by parcels of use u whose threshold lies below m. Aggregate supply is then additively separable across lag groups:

$$Q_t = \bar{Q} + \sum_u \Delta_u(\mathcal{M}_{t-\ell_u}) \tag{21}$$

The first term is the standing stock, invariant to the price path. Each term of the sum is the pipeline of one use of origin, evaluated at the running maximum that prevailed one construction lag ago. Non-housing and vacant parcels contribute $h_i^{\text{old}} = 0$ before redevelopment, and a housing-to-housing redevelopment contributes only the densification gain, which is positive only if the new structure exceeds the demolished one. The market-clearing condition is

$$R_t = A_t \cdot Q_t^{-\alpha} \quad \text{with } Q_t \text{ given by (21)}$$

Aggregate supply at date t is predetermined, since it depends on running maxima strictly earlier than t. The equilibrium rent is therefore obtained explicitly rather than as a fixed point, which is the sense in which time-to-build simplifies the aggregate problem while complicating the dynamics.

Overshooting The delay changes the character of the transition. A positive demand shock raises rents, which triggers starts, but the supply response that would dampen the shock arrives only ℓ_u years later. Rents therefore overshoot the path they would follow under contemporaneous clearing, and the running maximum ratchets to a higher level. The counterfactual exercises in Section 6 compare supply at different permanent price levels, at which $R_t = \mathcal{M}_t = \mathcal{M}_{t-\ell_u}$ for all u, so the overshoot affects the transition path without affecting the permanent-level comparisons.

3.7. Implications for housing supply

3.7.1 Local housing supply elasticities

The aggregate supply elasticity with respect to the rental price, evaluated at a moment when the price is at its running maximum, is

$$\varepsilon(R_t) = \frac{dQ_t}{dR_t} \cdot \frac{R_t}{Q_t} \tag{22}$$

The central object governing it is the cross-sectional distribution of redevelopment thresholds. Define

$$F(r) \equiv \mu\big(\{i : R_i^\star \leq r\}\big)$$

as the measure of parcels whose threshold lies below r, with density $f(r) = F'(r)$. The shape of f is endogenously determined by the joint distribution of $(F_i, \bar{h}_i, \phi_i, v_i^C, h_i^{\text{old}}, \ell_i)$ through (17). A region with many thresholds clustered near the current price has a highly elastic extensive margin, and a region with a sparse threshold distribution near the current price has a rigid one.

At a moment when the price is at its running maximum, a marginal increase dR triggers all parcels with $R_i^\star \in [R_t, R_t + dR)$. Differentiating the pipeline terms of (21) under the integral sign gives

$$\frac{d}{dR}\int h_i^{\text{new}} \cdot \mathbf{1}(R_i^\star \leq R)\, di = \mathbb{E}[h^{\text{new}} \mid R^\star = R] \cdot f(R)$$

so that

$$\frac{dQ_t}{dR_t} = \underbrace{\mathbb{E}\left[\left.\frac{\partial h_i^{\text{new}}}{\partial R}\right|_{R_t} \middle| R_i^\star = R_t \right] \cdot f(R_t)}_{\text{Intensive margin of newly triggered parcels}} + \underbrace{\mathbb{E}[h_i^{\text{new}}(R_t) \mid R_i^\star = R_t] \cdot f(R_t)}_{\text{Extensive margin}} \tag{23}$$

Both margins are governed by the threshold density. There is no intensive margin from parcels already redeveloped, since density is chosen irreversibly at the start date. The expression holds for any price process with a well-defined running maximum, and the geometric Brownian motion assumption is needed only for the threshold formula (15). Two remarks follow from the extensions. The derivative $\partial h_i^{\text{new}}/\partial R$ is scaled by κ_i^ζ, so the intensive margin of a newly triggered parcel is weaker for slower projects. Equation (23) is the elasticity of starts, whereas the elasticity of deliveries at date t is the same object evaluated at $\mathcal{M}_{t-\ell_u}$.

3.7.2 Role of fixed costs and frictions in supply elasticity

Channel 1: rightward shift of the threshold distribution Higher fixed costs shift thresholds rightward, and the effect is amplified by the option value multiplier:

$$\frac{\partial R_i^\star}{\partial F_i} = \frac{\beta}{\beta - 1} \cdot \frac{\partial R_i^{\text{naive}}}{\partial F_i} > 0$$

Fewer parcels then have thresholds near the current price, so the extensive margin is thinner. Since the multiplier is constant across parcels, the effect is uniform in proportional terms.

Channel 2: current use heterogeneity For non-housing parcels the continuation value depends on current revenue, and $\partial R_i^{\text{naive}}/\partial v_i^C > 0$. High current revenues push thresholds rightward and reduce the threshold density near current prices. The interaction of current revenue, the fixed cost and the construction lag determines both the location and the spread of the threshold distribution for each use of origin, so f may be multimodal, with clusters corresponding to land use categories.

Channel 3: density at the threshold Even conditional on crossing, the density built depends on the fixed cost. A higher fixed cost delays crossing to a higher price, which implies denser construction at crossing. In general equilibrium, however, parcels with low fixed costs cross earlier, expand aggregate supply and reduce the market-clearing rent, so parcels crossing later face lower prices than a partial equilibrium calculation would predict. The delivery lag weakens this feedback further, since the supply expansion that depresses rents arrives after the later parcels have already committed to a density.

Channel 4: regulatory constraint interaction Where the density cap binds, density is capped regardless of the fixed cost, and the extensive margin contribution is $\bar{h}_i f(R_t)$ irrespective of the price level. The construction lag raises the saturation rent by κ_i^{-1} and therefore moves parcels out of this regime, which raises the weight of the intensive margin in (23). Even under loose regulation, high fixed costs impede redevelopment, so the interaction of the two determines the intensive margin's contribution to aggregate elasticity.

Channel 5: bargaining By Proposition 2 the bargaining weight leaves the elasticity unchanged, conditional on the fixed cost. It nonetheless affects the estimated elasticity, because the fixed cost is recovered from the price by (18) and is decreasing in the surplus imputed to the developer. A lower landowner weight therefore implies lower fixed costs, lower thresholds and a more elastic supply.

Proposition 4 *The aggregate supply elasticity $\varepsilon(R_t)$ is governed by the cross-sectional density of redevelopment thresholds $f(R_t)$ and by the density built at threshold crossing $\mathbb{E}[h^{new} \mid R^\star = R_t]$. In areas with higher fixed transition costs F_i or longer construction lags ℓ_i:*

1. *The threshold distribution shifts rightward, which reduces $f(R_t)$ at any given price, with the effect amplified by the option value multiplier $\beta/(\beta-1)$ in the case of fixed costs and by the factor $\kappa_i^{-1}(\lambda_i^C)^{-1/(\zeta+1)}$ in the case of lags.*

2. *The density built at threshold crossing is lower, mechanically so under a lag through the factor κ_i^{ζ}, and in general equilibrium because supply expansion depresses the market-clearing rent.*

3. *The supply response arrives later, since deliveries at date t reflect the running maximum at $t-\ell_i$.*

4. *Aggregate supply elasticity $\varepsilon(R_t)$ is lower.*

Bargaining power affects the elasticity only through the identification of F_i.

4. Model calibration and estimation

4.1. Calibration summary

The calibration involves six blocks: demand dynamics (μ, σ, α), financial parameters (δ, ρ), construction technology (ζ, c_0), the two frictions of Section 3 (the construction lag ℓ_u and the bargaining weight θ), fixed costs and land use regulation (F_i, $\bar{h}_i$), and parcel-level variables (R_t, ϕ_i, T_i, v_i^C, h_i^{old}). Every structural parameter is fixed externally, either from the literature or from direct estimation on auxiliary data. Table 6 reports values and sources. All variables are measured in 2023 unless stated otherwise. Appendix B develops each block, reports the auxiliary estimates and documents the sensitivity of the calibrated objects.

Two results dictate the strategy. Proposition 2 states that the bargaining weight is absent from the stopping rule, so no timing moment identifies it. Proposition 3 states that the lag and the fixed cost enter the land price through the same wedge, so no price moment separates them. Both frictions must therefore be pinned down outside the land-price equation. I take the lag from permit-to-completion durations in Sitadel, by use of origin, and the weight from the concentration of land acquisition, subject to a floor implied by $F_i \geq 0$. The option-value block is unaffected by either choice: by Proposition 1 the multiplier $\beta/(\beta - 1) = 1.50$ is invariant to the lag, and neither μ nor α enters β.

The fixed cost is then not a free parameter but a fully determined parcel-level residual, obtained by inverting (18) at the calibrated parameters. Three objects are consequently over-identification checks rather than calibration targets: the implied housing supply elasticity, the fit of predicted to observed land prices, and the sign restriction $F_i \geq 0$. Section 5.1 reports them.

Demolition component can be measured directly rather than inferred. I use the *Prix de revient des logements neufs* survey, matched to Sitadel permits and to the parcels they cover, and estimate a hedonic decomposition of project-level demolition cost on demolished floorspace by destination. Marginal costs cluster in the 35-56€ per square metre range across destinations and the mobilisation intercept is 25,269€. Parcel area is insignificant throughout, so demolition cost scales with the structures removed and not with the land. I predict out of sample onto the full parcel sample using destination-specific floorspace from the Fichiers Fonciers. This yields a measured floor for F_i, against which Section 5.6 reports the share of the fixed cost the model does not name. Appendix B.6 gives the estimates and the robustness exercises.

Table 6: Parameters used in simulation

Par.	Name	Source	Value
μ	Expected growth rate of demand	Calibration (Observatoires des Loyers, FF)	0.01
σ	Volatility of demand	Calibration (Observatoires des Loyers, FF)	0.04
α	Inverse demand elasticity	Chapelle et al. (2023)	0.35
δ	Rental price growth rate	$\rho -$ net_cap (net yield, IDF)	0.02
ρ	Discount rate	Banque de France (+ risk premium)	0.065
β	Option value	Computation	2.98
ζ	Construction cost convexity	Combes et al. (2021)	1.17
c_0	Average construction cost	PRLN, EPTB	1,186
ℓ_{hous}	Construction lag, housing	Sitadel (permit to completion)	2.0
ℓ_{mixed}	Construction lag, mixed	Sitadel (permit to completion)	2.2
ℓ_{nonhous}	Construction lag, non-housing	Sitadel (permit to completion)	2.2
ℓ_{vacant}	Construction lag, vacant	Sitadel (permit to completion)	1.8
θ	Landowner bargaining weight	$1 -$ HHI (Sitadel), floor from $F_i \geq 0$	1.00
F_i	Fixed costs	Structural residual	–
$\bar{h}_i$	Maximum floorspace (LUR)	90th percentile envelope	–
ϕ_i	Location quality shifter	Observatoires des Loyers	–
T_i	Land area	Fichiers Fonciers	–
v_i^C	Non-housing revenue	DGFiP (valeur locative cadastrale)	–
h_i^{old}	Existing housing floorspace	Fichiers Fonciers	–

4.2. Estimation procedure

With all structural parameters and both frictions fixed externally, the estimation reduces to inverting the land transaction parcel by parcel to recover the fixed cost, then decomposing that fixed cost onto its economic determinants. The procedure has five steps.

Step 1: model primitives Given the calibrated parameters of the perceived price process (ρ, δ, σ), I compute the option value root β from the fundamental quadratic (16) and the associated multiplier $\beta/(\beta-1)$. For each parcel I form the regulatory cap $\bar{h}_i$ from (28) and the local rent $R_{it} = \phi_i R_t$. Because β depends only on perceived-process parameters, the multiplier is common to all parcels and invariant to the aggregate supply path, to the construction lag and to the bargaining weight.

Step 2: construction wedges For each use of origin I form $\lambda_u^R = e^{-(\rho-\delta)\ell_u}$, $\lambda_u^C = (1 - e^{-\rho\ell_u})/(\rho\ell_u)$ and $\kappa_u = \lambda_u^R/\lambda_u^C$, and attach them to every parcel through its use of origin. These two scalars are the entire footprint of the construction lag in the estimation, by Lemma 1.

Step 3: redevelopment value and the regime split At the calibrated parameters I evaluate the redevelopment value under the counterfactual of zero fixed cost, distinguishing the interior and corner regimes by comparing the delay-adjusted rent $\kappa_i R_{it}$ to the saturation rent

$$\bar{R}_i = \frac{1}{\kappa_i} \cdot \frac{c_0(\zeta+1)(\rho-\delta)}{\zeta} \left(\frac{\bar{h}_i}{T_i}\right)^{1/\zeta}$$

When the constraint is slack I use (3.2), otherwise the corner expression (3.2). I record the regime flag for each parcel, and I also record the frictionless value π_i^0, obtained by setting $\lambda_i^C = \kappa_i = 1$, which is needed for the decomposition of Proposition 3.

Step 4: fixed cost as a determined residual The continuation value V_i^w is formed use by use as in (12)–(14). Because no parameter is free, the fixed cost is exactly identified from (18) as

$$F_i = \pi_i(R_{i0}) - V_i^w - \frac{\hat{p}_i^L - V_i^w}{\theta}$$

The parcel-level threshold is then $R_i^\star = \frac{\beta}{\beta-1} R_i^{\text{naive}}$ with R_i^{naive} given by (17). I trim the top and bottom percentiles of the $R_i^\star$ distribution to remove parcels where the hedonic prediction is unreliable, and I report separately the parcels priced below their standing value, for which $\hat{p}_i^L < V_i^w$ and the fixed cost is not interpretable.

Step 5: decomposition and over-identification I confront the model-implied supply elasticity with the external estimate of α, I report the three-way decomposition of Proposition 3, and I decompose F_i onto predicted demolition cost and out-of-model land-use regulation with commune fixed effects (Section 5.6).

5. Baseline results

5.1. Model validation

The model is not fitted to any moment in the estimation sample, so its ability to reproduce independent quantities is a genuine test. It now faces four such tests rather than three, since the sign restriction on F_i has become informative about θ.

First, the model predicts an aggregate housing supply elasticity, evaluated at the 2023 rent level, of 0.6, against the external estimate of 0.35 from Chapelle et al. (2023) for the Paris area over 1999–2019. This is a cross-object test rather than an untargeted moment: the external estimate is identified from aggregate price–quantity variation, whereas the model elasticity aggregates parcel-level thresholds built from parameters that never saw that variation. The construction lag moves this object through two offsetting channels, since it raises triggers, which thins the extensive margin, and lowers desired density, which moves parcels out of the corner regime and thickens the intensive margin. The net movement is small, and I emphasise that ℓ_u is not calibrated to this moment: doing so would destroy the over-identification it provides.

Second, the wedge between fundamental and perceived growth has the sign the supply channel requires. The demand process grows at $\hat{\mu} = 0.01$ while perceived rents grow at $\delta = 0.02$. The two would coincide if aggregate supply were fixed, and the gap is the rate at which supply expansion erodes the market-clearing rent through (3). It is not imposed anywhere in the calibration.

Third, the recovered fixed costs are economically sensible. Only a minimal share of parcels (0.1% of housing and 4.7% of non-housing parcels) exhibit negative fixed costs, or 1125 parcels out of 543,974. For virtually all parcels there is therefore a substantial and positive wedge between the redevelopment value and the acquisition cost, as the theory requires. Under a correctly specified model this share would be exactly zero, so it bounds the combined contribution of specification error and hedonic prediction error in $\hat{p}_i^L$.

Fourth, the share of negative fixed costs is monotone in θ and therefore disciplines it. A lower landowner weight attributes more of the gap between the redevelopment value and the land price to the developer's rent and less to cost, so it lowers F_i and raises the share of parcels the model cannot rationalise. The adopted θ sits above the floor implied by (19), and Table C11 reports how the share of negative fixed costs, the median threshold and the elasticity move across the (ℓ, θ) grid. The reader can read off directly how far the concentration mapping could be pushed before the model breaks.

5.2. Decomposing the wedge

Before decomposing F_i itself, it is worth reporting what F_i does not contain. Proposition 3 splits the gap between the frictionless redevelopment value and the observed land price into a delay wedge, a developer rent, and a fixed cost. Table 7 reports the three components by use of origin, in euros per square metre of parcel and as shares of the aggregate gap. Across the sample the construction lag accounts for 16.8% of the gap, the developer's bargaining rent for 0.0%, and the fixed cost for the remaining 83.2%. The delay wedge is largest for the non-housing origins, which combine the longest lags with the largest redevelopment values, and the developer rent scales with the surplus and is therefore concentrated in the same parcels.

The exercise is not a decomposition in the statistical sense, since the first two components are calibrated rather than estimated. The fixed cost that Section 5.6 decomposes is the third term alone, and the magnitudes reported there should be read against a smaller total than the competitive, instantaneous-construction benchmark would deliver.

Table 7: Decomposition of the gap between the frictionless payoff and the land price

		Median, EUR per sqm of parcel			Share of aggregate gap		
Use	N	Delay	Developer rent	Fixed cost	Delay	Developer rent	Fixed cost
Housing	104,760	608	0	3,282	16.0%	0.0%	84.0%
Mixed	16,339	1,020	0	4,457	17.5%	0.0%	82.5%
Non-housing	4,950	891	0	4,279	17.4%	0.0%	82.6%
Vacant	1,577	529	0	3,372	13.4%	0.0%	86.6%

Data: Model estimation. The gap is $\pi_i^0 - p_i^L$, the difference between the redevelopment value under instantaneous construction and the observed land price, decomposed as in Proposition 3 into a delay wedge $\pi_i^0 - \pi_i$, a developer rent $(1-\theta)S_i$, and the fixed cost F_i. The first block reports the median of each component in euros per square metre of parcel area, the second reports each component summed across parcels as a share of the aggregate gap, so the three shares add to one hundred. N is the number of parcels in each use of origin with a positive surplus S_i and a positive gap. The delay wedge is evaluated at the lags of Table B1 and the developer rent at $\theta = 1.00$.

5.3. Incidence of the surplus

The bargaining extension delivers one object the competitive benchmark cannot: the incidence of a policy acting on fixed costs. A grant that lowers F_i by one euro raises the surplus by one euro, of which θ is capitalised into the land price and $1-\theta$ accrues to the developer. At the adopted weight, 100% of any transfer is therefore capitalised into land regardless of whether it induces redevelopment. By Proposition 2 the exercise decision is unaffected by the split, so the incidence result is separable from the quantity results of Section 6 and can be applied to any of the subsidy scenarios there. At the 2023 rent, developers retain 0.0% of the aggregate redevelopment surplus, or 0million euros.

5.4. Threshold distribution

Figure 6 plots the redevelopment threshold by parcel use of origin, as well as the 2023 rent level. Unsurprisingly, vacant parcels exhibit a low threshold, where some parcels should have already redeveloped. The density of housing and non-housing parcels are qualitatively similar. The mixed-use parcel, mostly located in city centres and high-attractivity areas exhibit the highest thresholds. Even if rent were to double, a substantial share of the latter would not redevelop. Consistent with the threshold crossings, supply is expected to gradually increase (Figure 7) over the next 50 years.

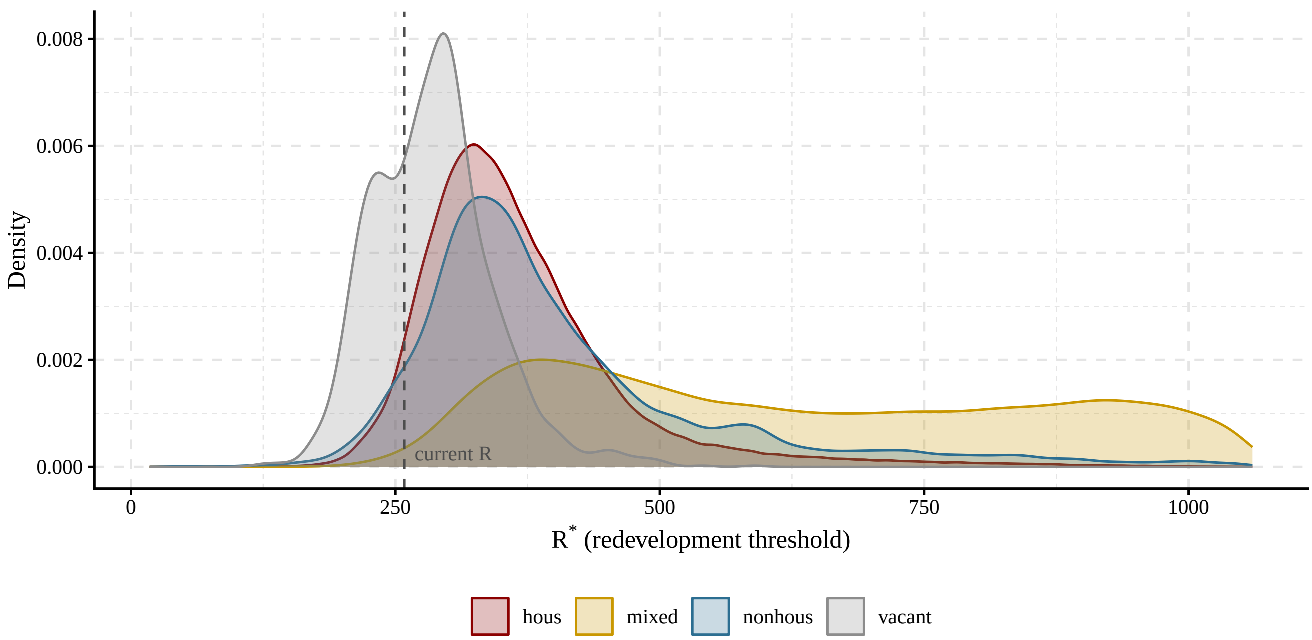

Data: Model estimation. Distribution, by type (housing, mix-use, non-housing, and vacant parcels), of redevelopment thresold $R_i^\star$. Vertical dashed line is 2023 rent level. 1st and 99th percentile are dropped.

Figure 6: Redevelopment threshold $R_i^\star$ distribution

5.5. Housing supply elasticity decomposition

With variable and fixed costs in hand, the model predicts housing supply for a given rent level. I compute housing supply elasticity as the supply response for a 1% increase in aggregate rent level R_t. As mentioned before, the model predicts a 0.6 elasticity.

I decompose elasticity along two dimensions. First, I decompose the elasticity into an intensive and an extensive margin. The total supply response to a change in the aggregate rent decomposes as in (23).

Both margins operate on redeveloping parcels, but through distinct channels. The extensive margin counts how many parcels cross their threshold $R_i^\star$ as rent rises: it is governed by the density of thresholds $f(R^\star)$ around the current rent and, through the dispersion of $R_i^\star$, by the heterogeneity of fixed costs. The intensive margin measures how much taller the parcels that were already redeveloping choose to build: it is governed by the cost convexity ζ, and it is mechanically shut down for parcels sitting at the regulatory cap, where $\partial h_i^{\text{new}} / \partial R_t = 0$.

I evaluate (23) numerically by bumping the aggregate rent and recomputing supply from crossed parcels and their optimal densities. At the 2023 rent, the aggregate elasticity of 0.63 is almost exclusively extensive in Paris: the intensive margin is near zero because the density cap binds for the large majority of parcels, leaving the redevelopment decision, not the density choice, as the operative channel (Table C7). At prevailing rents, additional housing supply comes almost entirely from new redevelopments crossing their threshold rather than from taller buildings on parcels already redeveloping. Outside of Paris, the results are more nuanced: a substantial share of parcels, especially mixed and non-housing parcels, hit the regulatory cap.

Second, I allow the elasticity to vary across municipalities in the sample. I compute the elasticity com-

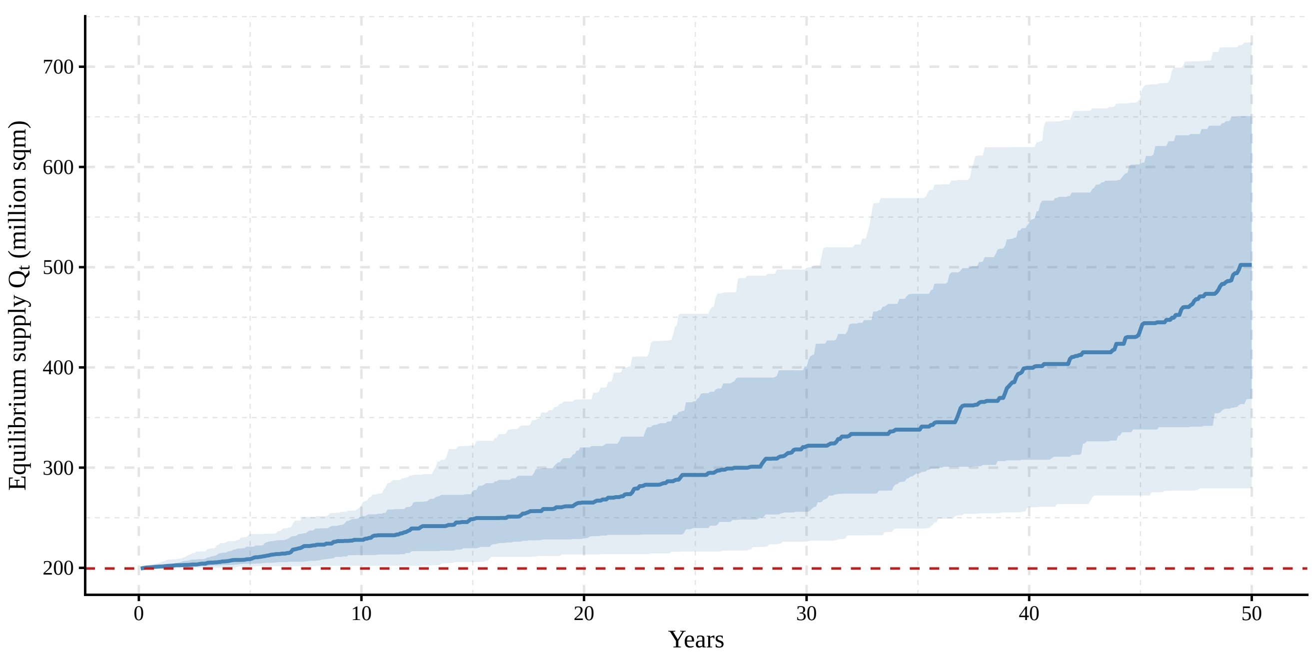

Data: Model estimation. Blue solid line is the housing supply evaluated at median demand level. Dark blue ribbon is the housing supply at the top and bottom 25th percentiles of the demand level. Pale blue ribbon is the housing supply at the top and bottom 10th percentiles of the demand level. Dashed horizontal line is supply evaluated in 2023. Each demand draw (500 draws) is following a monthly geometric Brownian motion from 2023 to 2073.

Figure 7: Housing supply evaluated along the demand gradient

mune by commune and map it across the Paris area. High-ϕ inner-Paris municipalities, where rents already sit well above most thresholds and density caps bind tightly, display the most inelastic supply, the intensive margin is exhausted and almost no parcels remain below threshold. Peripheral, lower-ϕ municipalities retain a larger mass of marginal parcels and a more elastic supply. Figure C8 maps the various municipal-level elasticities evaluated in 2023.

5.6. Fixed costs decomposition

I predict the parcel price $\hat{p}_i^L$ and recover F_i from (18). I decompose the fixed cost, which by Proposition 3 is already net of the construction delay and of the developer's bargaining rent, so the magnitudes below are shares of a smaller total than a competitive, instantaneous-construction benchmark would deliver. Figure C9 maps the recovered cost per square metre at the cadastral section level. To explain the fixed cost, I regress its logarithm per square metre on seven blocks of observables: demolition, ground condition, grid connectivity, built surroundings, parcel size, out-of-model land-use regulation[5], and municipal controls. The baseline decomposition is

$$\log \frac{F_i}{T_i} = a + \mathbf{Dem}_i + \mathbf{Soil}_i + \mathbf{Conn}_i + \mathbf{Surr}_i + b \log T_i + \mathbf{Reg}_i + \mathbf{Mun}_i + e_i \tag{24}$$

The dependent variable is expressed per square metre because F_i scales mechanically with parcel size, its decomposition would recover size rather than economic content. Predicted demolition cost and log parcel area still correlate at 0.451, so the two are not orthogonal, but the normalisation impedes parcel size from absorbing the decomposition (parcel size ends up accounting for 0.6% of the variance). Standard errors are clustered at the municipal level.

Sample I restrict the estimation sample. Parcels with a non-positive fixed cost are dropped, since the dependent variable is in logs and a negative fixed cost is not interpretable in any case. Vacant parcels are dropped, 1,584 in total, because they carry no structure to remove and the demolition block is undefined for them. The fixed cost per square metre is trimmed at the first and ninety-ninth percentiles, which removes the parcels where the hedonic price prediction is least reliable. The decomposition is estimated on 530,436 parcels.

[5] The distinction between in-model and out-of-model regulation is essential for a clean decomposition. The density cap $\bar{h}_i$ is in-model: it already enters the redevelopment value through the corner solution (3.2), so regressing F_i on it would recover a mechanical relationship. I therefore exclude $\bar{h}_i$ and use alternative measures of land-use constraints: permitting restrictions, architectural protection, and zoning.

Pre-conditions Two conditions must hold so the decomposition can be read as economic content rather than as a restatement of the model. The first is orthogonality. The fixed cost must not be a mechanical function of the two objects that already enter the redevelopment value, parcel area and the density cap. Regressing F_i/T_i on both yields coefficients that are statistically significant at this sample size but economically negligible, and an adjusted R^2 of 2%. The correlation with parcel area is essentially zero and the correlation with the density cap is 0.098. The cap is therefore not perfectly excluded from the fixed cost.

The second is separability, since (24) assumes the fixed cost is additively separable across blocks. If regulation raises demolition cost, through asbestos-removal rules on older buildings for example, the two blocks are not cleanly separable and the linear decomposition misattributes the shared component. I test this by summarising each block by its fitted linear index from the additive model, so that the test carries one term per block pair rather than hundreds of dummy-by-dummy products, and adding the interactions to the additive baseline. The targeted demolition-by-regulation interaction is rejected at the one percent level ($p = 0.003$), as is the omnibus test of all pairwise block interactions ($p = 0.001$). Strict separability therefore fails. It fails by a negligible margin: the interactions raise the R^2 by 0.0046, on more than half a million observations. I read this as separability being a relatively good approximation.

Sequential estimates Table C8 reports the estimates. The individual coefficients are not interpretable in isolation. The regressors are correlated, so no single coefficient captures the marginal contribution of its variable to the cross-sectional variance of fixed costs. The columns differ in how the demolition block enters and in whether regulation is included. The first omits demolition entirely and serves as a reference, the next two enter the demolished floorspace disaggregated by destination, the following two replace it with the single predicted cost obtained from the hedonic coefficients of Table C6, and the last adds municipality fixed effects. The disaggregated form is the more demanding of the two, since it lets the data choose the weight on each destination rather than imposing the weights estimated on the PRLN sample.

Shapley decomposition To measure each block's contribution I use a Shapley decomposition. It apportions the model R^2 across groups of regressors by averaging each group's marginal contribution over all orderings of entry. Unlike the sequential R^2 gains from adding blocks one at a time, the Shapley value splits variance shared between correlated groups symmetrically. The resulting shares are order-invariant and sum to the explained variance, with the complement reported as the unexplained share.[6]

The seven blocks are as follows. *Demolition* collects the demolished floorspace by destination, in inverse hyperbolic sine, together with the height and the construction date of the structure to be removed, both of which proxy the difficulty of the operation rather than its scale. *Ground condition* collects an indicator for a former industrial or potentially polluted site and the soil geological type of the parcel. *Grid connectivity* is the distance to the gas network and to the underground electricity grid. *Built surroundings* is the number of contiguous parcels and their mean building height, which proxy the difficulty of operating in a dense environment. *Parcel scale* is log parcel area. *Land-use regulation* is the municipal permit refusal rate, an indicator for architectural protection covering *Architectes des Bâtiments de France* perimeters and historical monument buffers, and zoning type. *Municipal controls* are population density, median income, and the distance to the Paris border. The four middle blocks are additive under the Shapley decomposition, so I also report their sum as a single parcel-characteristics share.

Table 8 reports the decomposition in three columns, which differ only in the treatment of location. The first includes the municipal controls, the second omits them, and the third removes the municipal mean from the dependent variable and from every regressor before computing the shares. The third column is the identification-cleanest of the three, since no municipal variable can be mechanically present in F_i and the shares are identified off parcel-level variation alone. I read the middle column as the main specification, for two reasons. The municipal controls partly re-project the spatial variation already embedded in the hedonic land price and are not interpretable as a cost component *per se*. And the within-commune column discards the between-municipality variation that a planner setting policy at the metropolitan level would care about. Shares are expressed as fractions of the total variance of F_i, so the unexplained row is one minus the R^2 and each column sums to one hundred.

Results Four findings follow. First, out-of-model land-use regulation is the largest identified component, at 20.6% of the cross-sectional variance, ahead of demolition at 14.8%. The ordering is not close and it does not depend on the specification: regulation leads demolition by 15 to 11 with municipal controls and by 10 to 2 within commune, and the two ranges do not overlap in any cell of the parameter band. This concerns the dispersion of the fixed cost, not its level, and it is consistent with the counterfactuals of Section 6, which act on the level and find the density cap a second-order instrument. Regulation is

[6] This property is essential here, since demolition and regulation are not orthogonal and a naive sequential attribution would assign the shared component to whichever block enters first. It also carries a cost. Where two blocks share a latent factor, each receives half of it by construction rather than by evidence, which is why the built-surroundings share below should be read alongside the demolition share and not independently of it.

what makes fixed costs differ across parcels. It is not what makes them large.

Second, the parcel's own physical situation accounts for 27.6% of the variance, the second largest identified block, and it is concentrated in two components. Ground condition contributes 14.1% and built surroundings 12.5%, while grid connectivity contributes 0.4% with a bootstrap standard error of 0.8 and is indistinguishable from zero. What raises the cost of putting a parcel back into housing use is what is under it and what is around it, not how far it sits from a utility network. The built-surroundings share should be read jointly with demolition rather than on its own, since the mean height of neighbouring buildings and the height of the structure being removed correlate at 0.89 and the Shapley value splits that common factor evenly between the two blocks by construction.

Third, 37.0% of the variance is explained by none of the blocks. It is the least precisely determined number in the exercise, ranging from 36 to 51% across the parameter band.

Fourth, the explanatory power of the whole exercise is largely between-municipality. Within commune the unexplained share rises to 84%, so roughly three quarters of what the other columns explain is variation across municipalities rather than across parcels within them. Regulation is the exception. It retains 10% of total variance within commune, which is close to three fifths of everything the within specification explains, and it does so on parcel-level measures alone, since the municipal refusal rate is collinear with the fixed effect and drops out. Zoning and architectural protection, assigned parcel by parcel inside a single municipality, are the one observable that continues to discriminate once location is differenced away.

Standard errors come from a cluster bootstrap that resamples municipalities rather than parcels, so as to respect the clustered structure of the data, with 500 replications.

Robustness Two exercises check that the shares do not hinge on the first stage. Table C9 re-recovers F_i across the parameter band (ρ, ζ, c_0) and re-estimates the shares on each cell. Only these three parameters are varied, since δ is tied to ρ by the net-yield identity of Section 4 and moves with it, as do the construction wedges λ_u^C and κ_u and the continuation value V_i^w, all of which are recomputed inside each cell. Table C10 instead holds the parameters at their calibrated values and swaps the hedonic prediction of the parcel price, comparing the baseline specification with a parsimonious one, a log specification, a saturated one, and the bootstrap mean of the hedonic prediction.

The ranking is stable throughout. Regulation ranges from 18 to 24 % across the band and demolition from 11to 14 %, so the two intervals never meet and regulation leads in all cells.[7] The two exercises differ sharply in how much they move the unexplained share. Varying the structural parameters moves it from 36 to 51 %, whereas varying the first-stage price specification moves it only from 37 to 37 %. What the fixed cost contains is therefore far more sensitive to the calibration of the option-value problem than to the hedonic prediction of the land price, which is the reverse of the usual concern with structural residuals and is worth stating plainly.

Neither exercise varies the two frictions of Section 3. The construction lag and the bargaining weight shift the level of F_i rather than its cross-sectional composition, since both enter through use-of-origin averages and a single scalar respectively rather than through parcel-level variation. Table C11 reports their effect on the level directly.

Table 8: Fixed cost: Shapley decomposition

	Share of total variance (%)		
Component	With controls	Without controls	Within commune
Demolition	11.4	14.8	2.0
	(1.0)	(1.3)	(0.2)
Ground condition	9.2	14.1	3.3
	(1.5)	(2.1)	(1.0)
Grid connectivity	0.3	0.4	0.2
	(0.6)	(0.8)	(0.3)
Built surroundings	9.7	12.5	1.0
	(1.1)	(1.3)	(0.1)
Parcel size	0.4	0.6	0.0
	(0.1)	(0.1)	(0.0)
Land-use regulation	14.9	20.6	9.8
	(1.1)	(1.6)	(0.9)
Municipal (location)	24.4	–	–
	(1.4)		
Unexplained	29.7	37.0	83.6
	(3.5)	(3.8)	(1.2)

Data: Model estimation. Grouped Shapley (LMG) decomposition of the cross-sectional variance of $\log(F_i/T_i)$, estimated on 530,436 non-vacant parcels with a positive fixed cost, trimmed at the first and ninety-ninth percentiles. Each cell is the share of the *total* variance attributable to that block, so columns sum to one hundred and the unexplained row is one minus the R^2. Blocks are defined in Section 5.6. Column (1) includes the municipal controls, column (2) omits them, and column (3) removes the municipal mean from the dependent variable and from every regressor before computing the shares, which mechanically eliminates the municipal block. Column (2) is the specification quoted in the text. The implied R^2 is 70%, 63% and 16% respectively. Cluster bootstrap standard errors in parentheses, resampling municipalities, 500 replications.

[7] These figures are not directly comparable to the main table. Both robustness exercises use a coarser four-block partition and enter demolition as the single predicted cost of Table C6 rather than as the disaggregated surfaces, which mechanically lowers the demolition share by roughly one percentage point relative to Table 8.

6. Counterfactual housing supply

6.1. Counterfactual scenarios

To highlight the mechanisms at play, I run two series of counterfactual exercises. The first acts on the fixed cost F_i, the second on the density cap $\overline{h}_i$. The two series operate on different margins of the redevelopment decision: comparing them tells which friction binds.[8] Although these scenarios do not all have real-life counterparts, several policy recommendations follow.

A policy shifts the parcel-level threshold $R_i^\star$ and hence the aggregate supply schedule $Q(\cdot)$. The equilibrium is re-solved under the policy schedule using the same draws of the primitive process A_t, so that every reported deviation nets out the common demand shock. Most results are estimated over a 20-year horizon.

6.1.1 Fixed cost subsidies

Removing the fixed cost The first counterfactual measures how many parcels would redevelop absent their fixed cost. I re-solve the model imposing $F_i = 0$ for every parcel. This is a benchmark rather than a policy: it bounds what any instrument acting on F_i can achieve. The non-redevelopment it leaves is attributable to: the continuation value V_i^w, the density cap $\overline{h}_i$, or to the option value.

Targeting under-built parcels The second counterfactual subsidises parcels that are at or below 50% of their density cap $\overline{h}_i$ and lie within one kilometre of the Paris border. These represent 7% of the parcel stock (around 38,000 parcels). The subsidy is a transfer linear in parcel size, s euros per square metre, capped at the parcel's fixed cost, so that the outlay on parcel i is $\min(s \cdot \text{area}_i, \max(F_i, 0))$ and the policy threshold solves

$$R_i^{\star,\text{pol}} = \frac{\beta}{\beta - 1} \cdot R_i^{\text{naive}} \left(\max(F_i - s \times \text{area}_i,\, 0)\right)$$

I set $s = 50\,€/\text{m}^2$, as the per-square-metre cost of a nationwide brownfield renewal programme in France (*Fonds Friches*), and then increase it to 100, 500 and 1,000 $€/\text{m}^2$ in order to establish how large a transfer the friction requires.

[8] An additional series I aim to add in a future iteration of the papers relates to the option-value part. By reducing the interest rate ρ, through preferred-rate loan for instance, or reducing uncertainty σ, a policy leader may also trigger more redevelopment.

Targeting non-residential parcels The third counterfactual applies the same schedule to vacant and non-housing parcels (10,000 parcels, 1.8% of the stock). These differ from the previous group as their continuation value V_i^w is lower, so a given transfer moves $R_i^\star$ further, but they would also redevelop more readily absent the policy.

The minimum inducing subsidy The fourth exercise inverts the problem. Let $R_i^{\text{tgt}}(T) = \mathcal{M}_T \phi_i$ denote the local rent prevailing at horizon T. For each parcel I solve the threshold condition for the fixed cost that would just tip parcel i into redevelopment at that rent,

$$\tilde{F}_i(T) = \lambda_i^C \left(\frac{\kappa_i R_i^{\text{tgt}}(T)}{\frac{\beta}{\beta-1} \cdot \frac{(\zeta+1)(\rho-\delta)}{\zeta}} \right)^{\zeta+1} \frac{c_0^{-\zeta}\, T_i}{\zeta} \;-\; V_i^w,$$

and define the minimum inducing subsidy as $\tilde{s}_i(T) = \big(F_i - \tilde{F}_i(T)\big)/T_i$. This is a partial-equilibrium object: it holds the rent path at its baseline value and therefore bounds what any targeting rule can achieve, since a policy that treats many parcels lowers the equilibrium rent and raises the subsidy each one requires. By Proposition 2 it is invariant to the bargaining weight, since θ does not enter the threshold. The weight governs instead who receives the money, an incidence question taken up below.

Four exclusive groups partition the parcel stock. *Already* parcels satisfy $F_i \leq \tilde{F}_i$ and redevelop under the baseline rent, so subsidising them is a pure transfer, of which a share θ is capitalised into the land price and the remainder goes to the developer. *Reachable* parcels satisfy $0 < \tilde{F}_i < F_i$: a finite subsidy induces them, and $\tilde{s}_i$ measures how large it must be. Two groups are blocked for distinct reasons. Where $\tilde{F}_i < 0$, the existing structure is valuable enough that no reduction in the fixed cost brings $R_i^\star$ below the prevailing rent. These parcels are blocked by V_i^w and are beyond the reach of any cost subsidy. Where $h_i^{\text{new}} < h_i^{\text{old}}$, redevelopment would replace the existing structure with a smaller one and is never undertaken. Because $\overline{h}_i$ is defined as the 90th percentile of the building density distribution, some parcels lay above this limit.

Ranking reachable parcels by $\tilde{s}_i$ and treating the cheapest first traces a budget frontier with no inframarginal transfer by construction. Comparing the cost per induced parcel along that frontier with the cost per induced parcel under the observable targeting rules above measures how much of the potential gain is lost by conditioning on characteristics a planner can actually observe.

6.1.2 Less stringent land-use regulations

Parcels capped by land use regulations point to a second instrument. Raising the density cap $\overline{h}_i$ acts on two margins simultaneously. Parcels that already redevelop build larger, and the higher payoff lowers $R_i^\star$ so that additional parcels cross their thresholds. This differs from a fixed-cost subsidy, which operates only on the extensive margin, and it carries no public outlay.

I consider three exercises. First, I remove the density cap entirely, which bounds what any relaxation can deliver (no corner solution). Second, I solve for the uniform multiplier m such that $Q(m\,\overline{h}_i)$ matches the supply obtained under $F_i = 0$ at the evaluation horizon. The multiplier m expresses how much additional permitted density substitutes for eliminating the fixed cost altogether. Because the two schedules have different shapes, the match is exact only at a reference date. Third, I raise every municipality whose median density cap lies below the Paris area median up to that median, holding the others fixed. The Paris area median FAR is 1.9 and exactly half of municipalities are below this median, accounting for 58.2% of parcels. Only the last of the three is implementable, the first two are benchmarks against which to read it.

6.2. Counterfactual results

Results are expressed as deviations from the baseline path, solved under the same draws of A_t. For each policy I report the partial-equilibrium deviation ΔQ^{PE}, which evaluates the policy's supply at the baseline equilibrium state and so isolates the policy-implied supply shift, and the general-equilibrium deviation ΔQ^{GE}, which evaluates it at the state that clears the market under the policy. Their difference $\Delta Q^{\mathrm{rent}} \leq 0$ measures the share of the direct effect absorbed by the market re-clearing at a lower rent.

6.2.1 Removing the fixed cost

Table 9 reports the headline effects at a twenty-year horizon. Eliminating the fixed cost raises aggregate supply by 104.1% and lowers the equilibrium rent by 22.7%. The partial-equilibrium shift is larger, at 148.1%: roughly 30% of the direct effect is clawed back once rents adjust. The number of redevelopment projects rises by 193,331, and net new floorspace by 282.59 million square metres.

Two features of this benchmark deserve emphasis. First, the effect is overwhelmingly on the extensive margin: the fixed cost is what keeps parcels from redeveloping at all, not what limits how much is built on them once they do. Second, and despite its magnitude, the benchmark leaves a substantial share of the stock untouched. Only 78.2% of parcels redevelop after 50 years, even with F_i set to zero throughout, so the remaining friction lies elsewhere (in the continuation value V_i^w, in the density cap, or in the option value itself).

Figure 8 shows the decomposition over time (50 years). The schedule shift peaks early and declines thereafter: as the baseline rent rises on its own, more parcels would have crossed in any case, and the policy's marginal contribution shrinks. The rent adjustment tracks it with a lag, since the additional supply must accumulate before it depresses the clearing rent. The intervention therefore matters most in the first decade, which argues for a front-loaded rather than a permanent instrument.

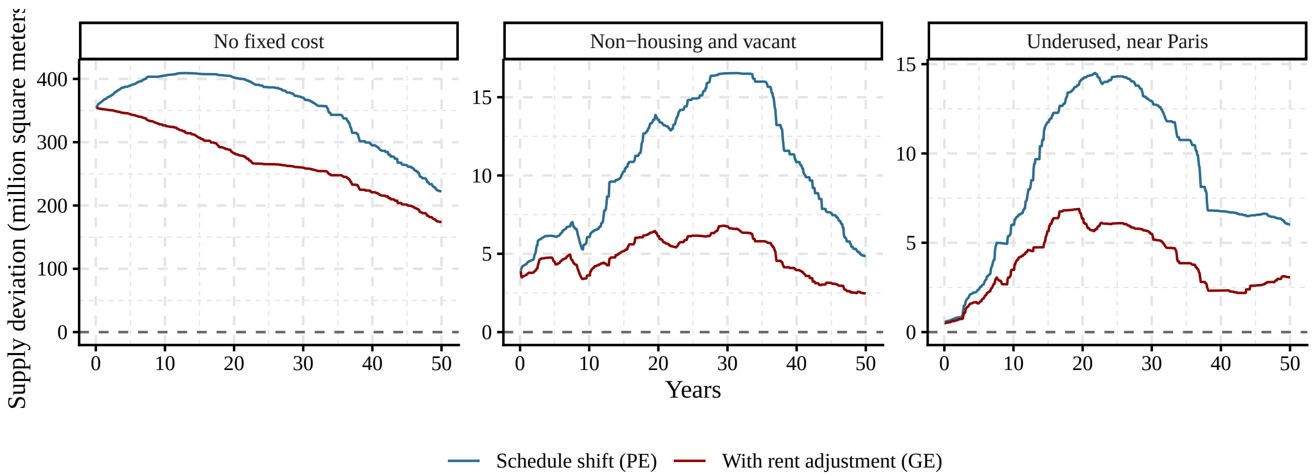

Counterfactual estimation under 3 counterfactual scenarios: Integral fixed cost removal; subsidize non-housing and vacant parcels; subsidize parcels built 50% or below the density cap located up to 1km around Paris. Blue line is the partial equilibrium supply path (without rent adjustment) and the red line is the general equilibrium path (with rent adjustment).

Figure 8: Decomposition of the counterfactual supply response

6.2.2 Targeted subsidies

The two targeted schemes deliver very little. At the programme-realistic rate of 50 €/m^2, the underused-parcel scheme induces 145 redevelopments out of 37,897 parcels targeted, a conversion rate of 0.38%, the non-residential scheme induces 67 out of 10,002, or 0.67%. Escalating the subsidy twentyfold, to 1,000 €/m^2, raises these to 8.6% and 19.6% respectively. The negative change in project count is a general equilibrium effect: parcels that cross their thresholds increase the net floorspace, decreasing the rent, and pushing some parcels further away from their threshold.

Figure 9 makes the mechanism transparent. Take-up (the share of targeted parcels that redevelop) is substantially higher than conversion at every subsidy level, and the gap is the inframarginal transfer (i.e. the parcels who would have redeveloped absent the policy due to price increase). The number of inframarginal parcels is invariant to the subsidy by construction, since lowering $R_i^\star$ cannot prevent a parcel from crossing. The share of the outlay that is wasted therefore falls with the subsidy only because the induced count grows, not because targeting improves. At 50 €/m^2, 95.7% of spending under the underused scheme accrues to parcels that would have redeveloped anyway.

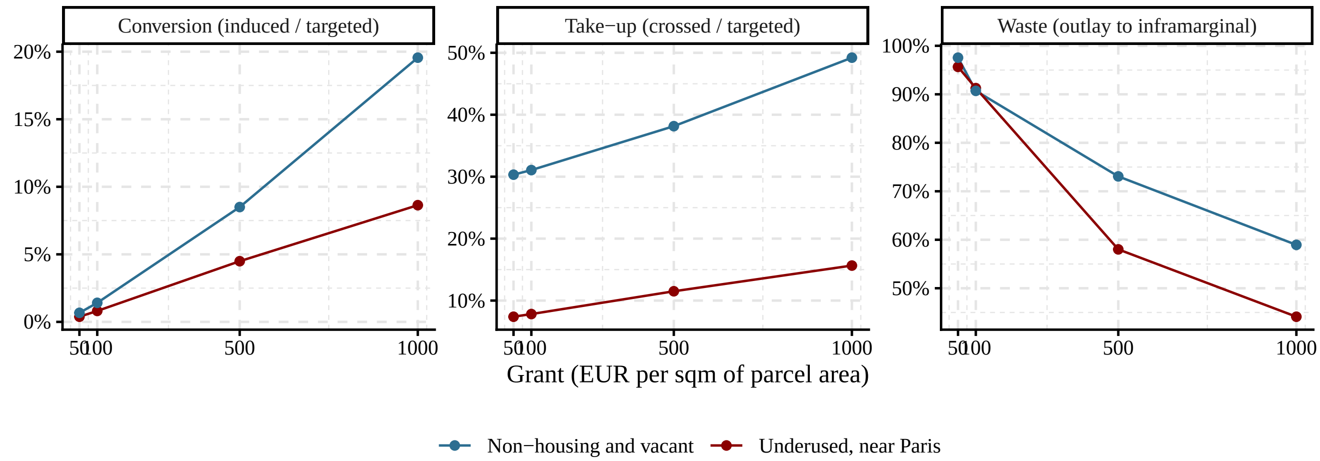

Counterfactual estimation under 2 counterfactual scenarios: subsidize non-housing and vacant parcels (blue); subsidize parcels built 50% or below the density cap located up to 1km around Paris (red). Conversion is the share that redevelop and would not have done so absent the policy. Take-up is the share of targeted parcels that redevelop. Outlay is the transfer to parcels that would have redeveloped absent the policy.

Figure 9: Take-up, conversion, and the inframarginal share of the outlay, by subsidy level

The comparison between the two schemes inverts a natural intuition. The non-residential scheme has by far the higher take-up, because vacant and non-housing parcels redevelop more readily to begin with, yet at low subsidies its conversion rate is no better and its inframarginal share is worse. Targeting the parcels most likely to redevelop is precisely the wrong rule: it maximises the transfer and minimises the causal effect. At high subsidies the ranking reverses, since the lower standing value of non-residential parcels means a given transfer moves $R_i^\star$ further. By then the subsidy is comparable to the construction cost itself.

6.2.3 The reachable set

The minimum inducing subsidy $\tilde{s}_i$ explains why the targeted schemes fail. Figure 10 reports the composition of the parcel stock across horizons. At twenty years, 13.4% of parcels redevelop under the baseline rent and are therefore beyond the reach of any policy that could only accelerate them, 57.1% are reachable, in the sense that some finite subsidy would tip them, 27.3% are blocked by their standing value, and 2.2% by the density cap. The reachable share is monotonically decreasing with time: as rent increases, the share of *redeveloping anayway* parcels is increasing.

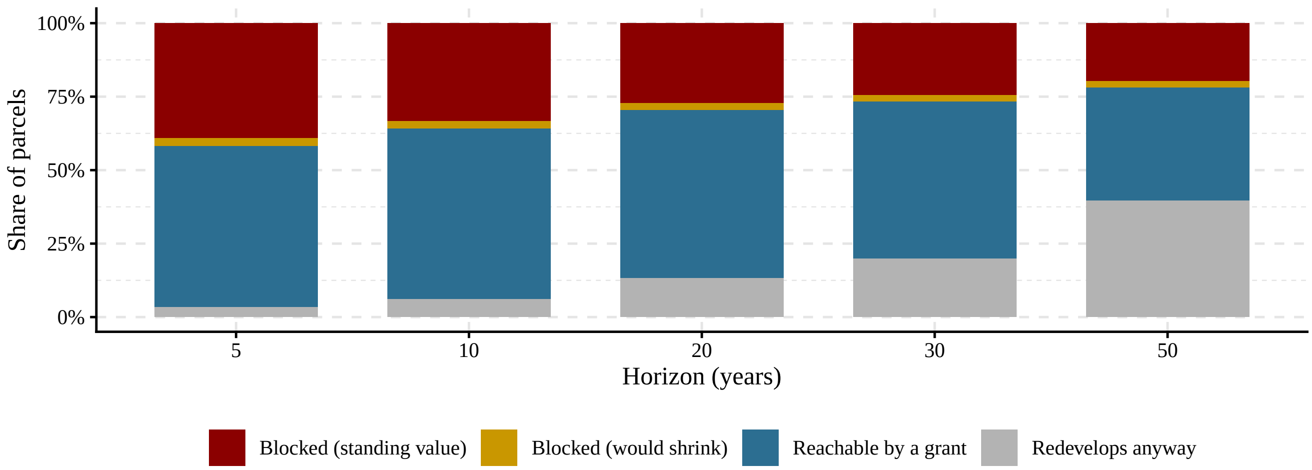

A parcel is reachable if some finite per-square-metre subsidy would bring its threshold below the prevailing rent; it is blocked by standing value if no reduction in the fixed cost suffices, and blocked by the density cap if redevelopment would replace the existing structure with a smaller one.

Figure 10: Composition of the parcel stock by horizon

Among reachable parcels the required subsidy is large. The median is 1,839 €/m^2 of parcel area, with an interquartile range of 924 to 3,038 (Figure 11). For comparison, the construction cost is $c_0 = 1{,}186$€ per square metre of floorspace: the fixed cost that must be offset to induce the median parcel is of the same order as the cost of building on it. The subsidies considered in the targeted schemes lie in the left tail of this distribution. A 50 €/m^2 transfer reaches 1.2% of the reachable set.

Table 10 traces the frontier obtained by ranking reachable parcels by $\tilde{s}_i$ and treating the cheapest first. This rule is infeasible — it conditions on F_i, which the planner does not observe — but it bounds what any targeting rule can achieve. Comparing it with the observable rules is instructive: at the highest

Table 9: Comparison between instrument classes at a 20-year horizon

Family	Scenario	ΔQ (GE)	ΔQ (PE)	ΔR	Projects	Net floor	Outlay
Fiscal	No fixed cost	104.1%	148.1%	-22.7%	193,331	282.59	–
	Underused, near Paris	2.3%	5.2%	-0.8%	-2,215	6.35	6,515
	Non-housing and vacant	2.3%	5.0%	-0.7%	-3,327	6.13	10,719
Regulatory	Density cap lifted	21.6%	61.3%	-6.5%	-35,478	58.74	–
	Below-median to median	9.9%	21.1%	-3.1%	-19,638	26.79	–

Counterfactual estimation under 5 counterfactual scenarios: Integral fixed cost removal; subsidize (1,000€/sqm) parcel built 50% or below the density cap located up to 1km around Paris; subsidize (1,000€/sqm) non-housing and vacant parcels; set $\overline{h}_i \to \infty$; raise the below-median density cap to median density. Floorspace is expressed in million square meters. Outlay (transfer to parcels that would have redeveloped absent the policy within 20 years) is expressed in million euros. All deviations are relative to the baseline path solved under the same A_t draw.

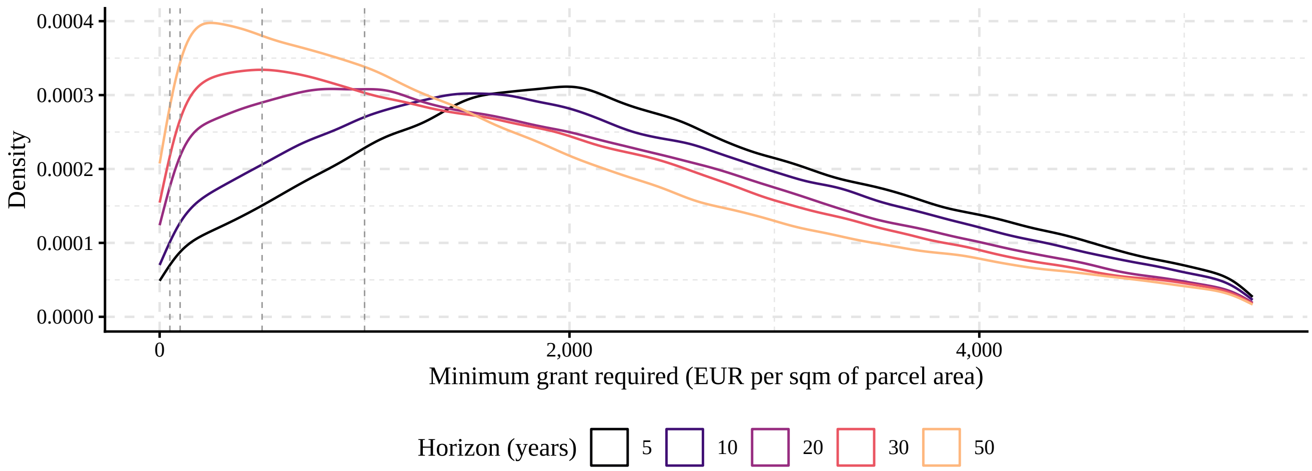

A parcel is reachable if some finite per-square-metre subsidy would bring its threshold below the prevailing rent. Dashed vertical lines mark the subsidy levels used in the targeted counterfactuals.

Figure 11: Distribution of the minimum inducing subsidy among reachable parcels, by horizon

subsidy considered, the underused scheme spends 105times as much per induced redevelopment as the frontier, and the non-residential scheme 543times as much.

Table 10: Reachable parcel set and optimal targetting at a 20-year horizon

	Reachable set					1bn EUR budget	
Horizon	Reachable	Share	Median $s_{\min}$	$\leq$ 500/sqm	$\leq c_0$	Induced (1bn)	Cost/induced (k)
5	297,469	54.7%	2,309	5.5%	18.7%	7,981	123
10	315,016	57.9%	2,104	8.0%	24.3%	10,697	93
20	310,701	57.1%	1,839	12.7%	32.7%	14,403	68
30	289,833	53.3%	1,714	15.4%	36.1%	15,884	63
50	208,923	38.4%	1,482	19.0%	41.7%	17,826	56

A parcel is reachable if some finite per-square-metre subsidy would bring its threshold below the prevailing rent. $s_{\min}$ is the smallest such subsidy. $c_0 = 1,186$ is the average construction cost per square meter of housing. The last two columns rank reachable parcels by $s_{\min}$ and subsidize them until a 1 billion euro budget is exhausted. The cost per per parcel is expressed in 1,000s of euros. All equilibria are solved under the same A_t draw.

6.2.4 Relaxing the density cap

The sizeable share of parcels that would never redevelop, even with a maximum subsidy on fixed cost, points to a non-fiscal instrument. Raising $\overline{h}_i$ acts on both margins at once: parcels that redevelop build denser, and the higher payoff lowers $R_i^\star$ so that additional parcels cross. It carries no public cost.

No uniform multiplier on $\overline{h}_i$ reproduces the supply obtained by eliminating the fixed cost: an unconstrained density cap ($\overline{h}_i \to \infty$) closes only 41% of the gap. The reason is that the cap is binding for a large share of parcels, but not all (it binds for 76.6% of parcels at the twenty-year rent). Multiplying it leaves 24% of thresholds unchanged, whereas removing F_i lowers every threshold at once.

Inverting the exercise gives the comparison in interpretable units. Removing the density cap entirely delivers the same supply at the twenty-year horizon as a uniform subsidy of 1,343 €/m^2 (against a median inducing subsidy of 1,839 €/m^2). Relative to the fixed cost, the density cap is a second-order friction.

Figure 13 shows that a relaxation of the density cap does very little to the composition of the stock. Lifting the cap entirely removes two-thirds of the parcels blocked by the density constraint, from 12,032 to 3,681 — some 1.5% of the stock. The parcels that remain blocked are not zoning-constrained: for 93 % of them the unconstrained optimal building is smaller than the standing structure even with no cap at all, so what holds them back is a prevailing rent too low to justify replacement. The other three categories are unchanged. In particular, the mass blocked by continuation value does not move. Raising the permitted envelope does increase the payoff to redevelopment and therefore lowers the threshold $R_i^\star$, but by far too little to overcome a continuation value that already exceeds what any permitted building would earn. The partial harmonisation of local plans is indistinguishable from the baseline.

6.2.5 Discussion

Three results emerge. First, fixed costs are the binding friction: removing them raises supply by 104.1% and more than doubles the number of redevelopment projects, while relaxing the density cap, the natural alternative instrument, is worth an order of magnitude less. Second, the friction is too large for

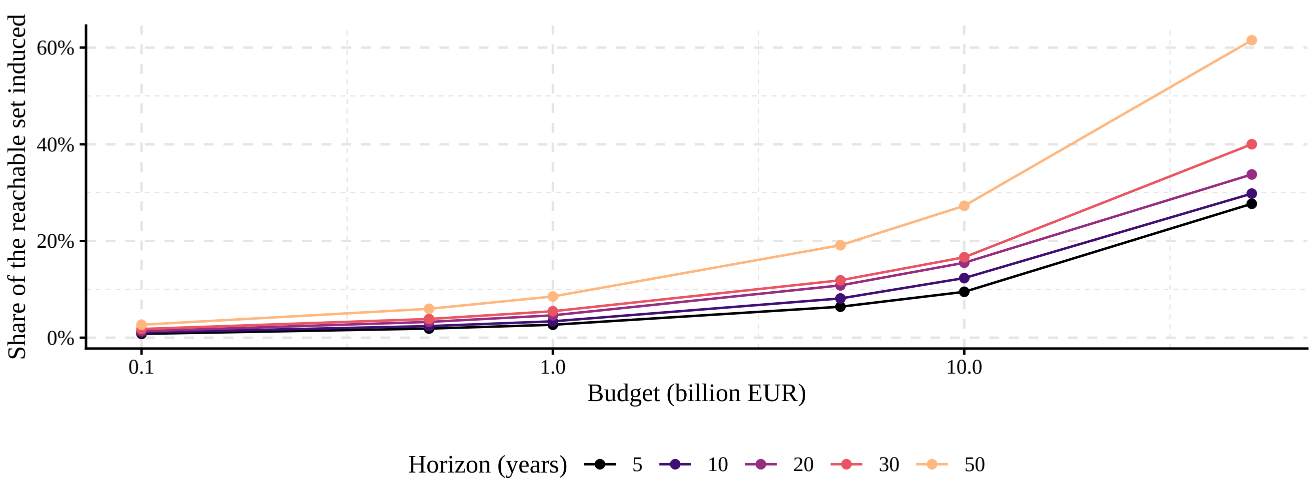

Note: The subsidy rule conditions on the unobserved fixed cost and therefore bounds feasible targeting.

Figure 12: Share of the reachable set induced to redevelop, by budget and horizon

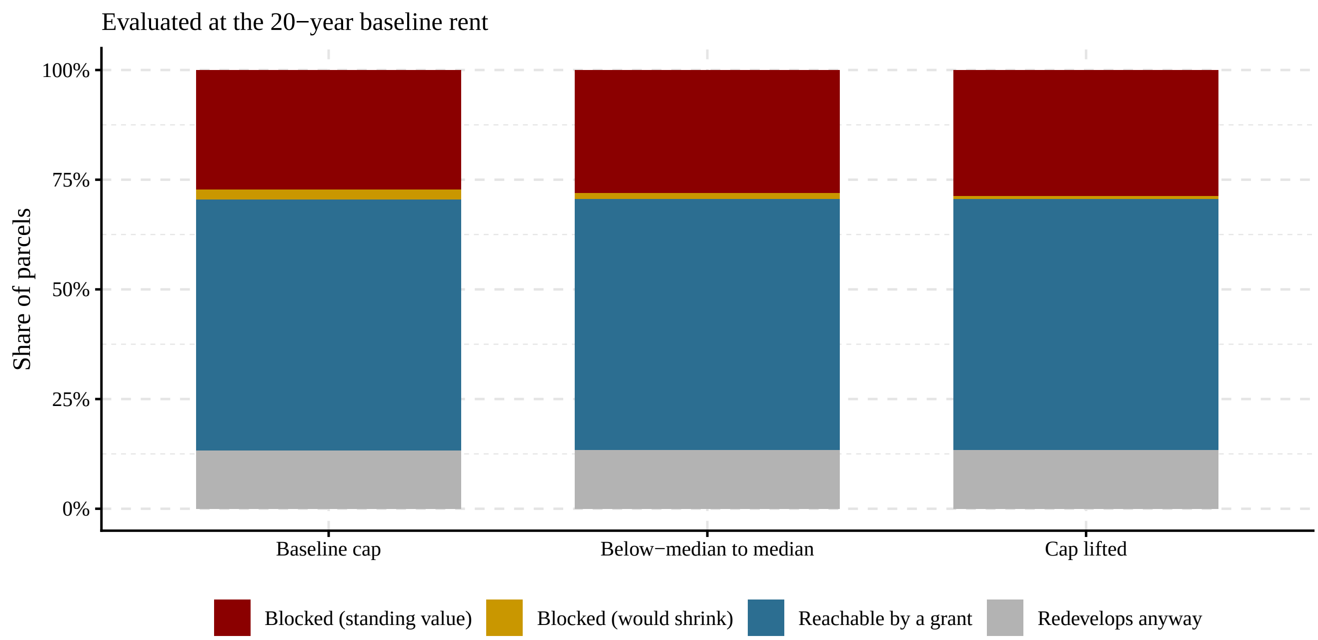

A parcel is reachable if some finite per-square-metre subsidy would bring its threshold below the prevailing rent. The two counterfactual scenarios, all evaluated at a common A_t draw, are: putting all below-median municipal FAR to the Paris area median FAR and totally removing the FAR limit.

Figure 13: Composition of the parcel stock under three density regimes, evaluated at the twenty-year baseline rent

marginal subsidies, since the median parcel requires a transfer comparable to the construction cost of the building that would replace it, so subsidies at the scale of existing programmes reach only the left tail of the distribution. Third, targeting on observables is expensive: the two rules considered spend an order of magnitude more per induced redevelopment than the infeasible frontier, because the characteristics a planner observes are uncorrelated with proximity to the redevelopment threshold.

The last point suggests that the design problem is one of screening rather than of budget. A rule that conditions on observable characteristics selects parcels that are more likely to redevelop, not parcels that are more likely to be induced to redevelop. An instrument that is self-targeting, one whose value accrues only on redevelopment, would avoid the inframarginal transfer entirely, and is a natural direction for future work.

Finally, a last set of counterfactual would deal with the option value parameters. On the one hand, the policy leader could reduce the interest rate ρ by offering preferred-rate loan (a scheme that exists for housing refurbishment or access to property at the household level). On the other, she could reduce uncertainty σ. Both parameters reduce the option value multiplier and increase the net present value of redevelopment projects.

7. Extensions and additional counterfactual results

7.1. Tax pass-through and local fiscal revenues

Taxes are levied at several stages of the redevelopment process. Construction projects involve three main taxes. First, the building permit applicant pays the *taxe d'aménagement*, whose base is a function of the project floorspace and whose rates are set at the municipal and departmental levels. Second, transactions are taxed under the *droits de mutation*, especially when they involve buildings older than five years and all land parcels.[9] The base is the transaction price and the rate is also locally decided. Finally, construction work is taxed under the value added tax, with a nationwide reduced rate of 5.5%. I extend the model by including these taxes at the different project stages, simulate the variation in housing supply, and conclude on their role. The bargaining structure of Section 3.3 makes the incidence question tractable, since a tax on the transaction reduces the surplus and is shared in the same proportion as a subsidy, which gives a pass-through of θ to the land price and $1 - \theta$ to the developer.[10]

7.2. Rent control and social housing

Lebret et al. (2026) include a parameter capturing the mandatory inclusion of a public good on the site, modelled as a share of total floorspace. I adapt their approach to model the corresponding requirement in French construction projects. The institutional context in the Paris area is well suited to the exercise. Since 2001, municipalities must target a specific share of social housing units (Loi SRU), which varies by municipality but is usually between 20 and 30 percent of the total housing stock. In Paris, rents are controlled, and local representatives vote yearly to cap them. Under rent control an increase in R_t does not trigger new construction for parcels at the threshold, which breaks the mapping from the aggregate rent to the local rent that (3) assumes.

7.3. Reducing the option value

A last set of counterfactuals concerns the parameters of the option value itself. A policy maker could reduce the discount rate ρ by offering preferential loans, a scheme that already exists for refurbishment and for first-time buyers, or reduce uncertainty σ by committing to a stable regulatory environment.

[9] Bérard and Trannoy (2018) estimate the role of an increase in the real estate transfer tax on transactions. They show that a moderate increase in the rate (0.7pp) implied a drop of 6% of transaction volume over the short run.

[10] The real estate transfer tax may already be capitalised in the land price. In the counterfactual I assume the pass-through remains the same, which yields a lower bound.

Both lower the multiplier $\beta/(\beta-1)$ and raise the net present value of redevelopment. Unlike the instruments of Section 6, they act on the value of waiting rather than on the payoff to acting.

8. Conclusion

In this effort, I measure the cost of putting an already-occupied parcel back into housing use, and ask what that cost implies for supply policy in a dense city. Redevelopment has long been described as lengthy and expensive, but its cost has been treated either as a regulatory tax or as an under-studied structural residual. I quantify the technical component directly, and find substantial and systematic heterogeneity across parcel types: non-housing-to-housing transitions supply more housing per square metre of land but face materially higher upfront costs.

Embedding these measured costs in an option-value model of redevelopment, in which developers face price uncertainty, construction takes time, and land is acquired through bilateral bargaining rather than frictionless bidding, I recover the parcel-level fixed cost as a determined residual and decompose it. The construction delay and the developer's bargaining rent account for 16.8% and 0.0% of the gap between the frictionless redevelopment value and the observed land price, leaving the fixed cost as the remainder. Within that remainder, land-use regulation accounts for 21% of the cross-sectional variance, demolition for 15%, and 37% is attributable to neither. The barrier to urban densification is therefore not reducible to zoning, and not reducible to engineering either.

Three conclusions follow from the counterfactuals. First, fixed costs are the binding friction: removing them raises supply by 104.1%, while removing the density cap entirely closes only 41% of that gap. Where the cap does bind it truncates the intensive margin, but it is not what keeps parcels from redeveloping at all. Second, the friction is too large for marginal subsidies: the median reachable parcel requires a transfer of the same order as the construction cost of the building that would replace it, so programmes at the scale of the *Fonds Friches* reach only the left tail of the distribution. Third, targeting on observables is expensive, by one to two orders of magnitude per induced redevelopment relative to an infeasible frontier conditioning on F_i itself.

References

BAR-ILAN, A., AND W. C. STRANGE (1996): "Investment Lags", *The American Economic Review*, 86, 610–622p.

BAUM-SNOW, N. (2023): "Constraints on City and Neighborhood Growth: The Central Role of Housing Supply", *The Journal of Economic Perspectives*, 37, 53–74p.

——— , AND L. HAN (2024): "The Microgeography of Housing Supply", *Journal of Political Economy*, 132, 1897–1946p, Publisher: The University of Chicago Press.

BROOKS, L., AND B. LUTZ (2016): "From Today's City to Tomorrow's City: An Empirical Investigation of Urban Land Assembly", *American Economic Journal: Economic Policy*, 8, 69–105p.

BÉRARD, G., AND A. TRANNOY (2018): "The impact of the 2014 increase in the real estate transfer taxes on the French housing market", *Economie et Statistique / Economics and Statistics*, 179–200p.

CAPOZZA, D., AND Y. LI (1994): "The Intensity and Timing of Investment: The Case of Land", *The American Economic Review*, 84, 889–904p, Publisher: American Economic Association.

CAPOZZA, D. R., AND R. W. HELSLEY (1990): "The stochastic city", *Journal of Urban Economics*, 28, 187–203p.

CHAPELLE, G., J. EYMÉOUD, AND C. WOLF (2023): "Land-use regulation and housing supply elasticity: evidence from France", THEMA Working Papers 2023-08, THEMA (THéorie Economique, Modélisation et Applications), Université de Cergy-Pontoise.

——— , J. PAUL-VENTURINE, AND C. URVOY (2025): "Mayors strike back: Evidence from upzoning in France", Number: halshs-05329475 Publisher: HAL.

——— , E. WASMER, A. LEMOINE, A. TRANNOY, AND J. PAUL-VENTURINE (2026): "What is France worth ? Residential land value estimates".

CHARREIRE, M., M.-N. LEFEBVRE, AND M. GIRARD (2025): "Effets de la concurrence sur la localisation des promoteurs et la production de logements : le cas de l'aire d'attraction de Lyon", *Working papers - Cahiers ESPI2R*.

COMBES, P.-P., G. DURANTON, AND L. GOBILLON (2021): "The Production Function for Housing: Evidence from France", *Journal of Public Economics*, 51p.

D'AMICO, L., E. L. GLAESER, J. GYOURKO, W. R. KERR, AND G. A. PONZETTO (2024): "Why Has Construction Productivity Stagnated? The Role of Land-Use Regulation".

DIAMOND, R., T. MCQUADE, AND F. QIAN (2019): "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco", *American Economic Review*, 109, 3365–3394p.

DIXIT, A. K., AND R. S. PINDYCK (1994): *Investment under uncertainty*, Princeton, New Jersey Chichester, West Sussex: Princeton University Press.

FAVILUKIS, J., E. S. PIKULINA, AND T. WU (2026): "House Prices, Quantities, and the Propagation of Shocks when Developers Are Imperfectly Competitive".

FROLKING, S., R. MAHTTA, T. MILLIMAN, T. ESCH, AND K. C. SETO (2024): "Global urban structural growth shows a profound shift from spreading out to building up", *Nature Cities*, 1, 555–566p.

GLAESER, E., AND J. GYOURKO (2018): "The Economic Implications of Housing Supply", *Journal of Economic Perspectives*, 32, 3–30p.

GLAESER, E. L., AND J. GYOURKO (2003): "The impact of building restrictions on housing affordability", *Economic Policy Review*, 21–39p, Publisher: Federal Reserve Bank of New York.

GLANCY, D., R. KURTZMAN, AND L. LOEWENSTEIN (2026): "JUE Insight: Shovel ready projects and commercial construction activity's long and variable lags", *Journal of Urban Economics*, 151, 103816p.

GOEYVAERTS, G., AND T. BOOGAERTS (2021): "Hold out in land assembly for residential development".

GROVENSTEIN, R. A., J. B. KAU, AND H. J. MUNNEKE (2011): "Development Value: A Real Options Approach Using Empirical Data", *The Journal of Real Estate Finance and Economics*, 43, 321–335p.

GYOURKO, J. (2009): "Understanding Commercial Real Estate: Just How Different from Housing Is It?",Technical Report w14708, National Bureau of Economic Research, Cambridge, MA.

——— , AND R. MOLLOY (2015): "Regulation and Housing Supply", in *Handbook of Regional and Urban*

Economics Volume 5: Elsevier, 1289–1337p.

LANGE, R.-J., AND C. N. TEULINGS (2024): "Irreversible investment under predictable growth: Why land stays vacant when housing demand is booming", *Journal of Economic Theory*, 215, 105776p.

LEBRET, D., C. H. LIU, AND M. VALENTIN (2026): "Carrot & Stick Policies and the Housing Affordability Puzzle: Redevelopment with Durable Assets in a Spatial Equilibrium".

LEFEBVRE, M.-N., AND I. DALY (2026): "Who Builds What When Subsidies Flow? Evidence from Place-Based Housing Policies in Small and Mid-Sized French Cities *".

MAJD, S., AND R. S. PINDYCK (1987): "Time to build, option value, and investment decisions", *Journal of Financial Economics*, 18, 7–27p.

MCDONALD, R., AND D. SIEGEL (1986): "The Value of Waiting to Invest*", *The Quarterly Journal of Economics*, 101, 707–727p.

MICELI, T. J., AND C. F. SIRMANS (2007): "The holdout problem, urban sprawl, and eminent domain", *Journal of Housing Economics*, 16, 309–319p.

MURPHY, A. (2018): "A Dynamic Model of Housing Supply", *American Economic Journal: Economic Policy*, 10, 243–267p.

PENG, E. (2023): "The Dynamics of Urban Development: Evidence from Land Use Reform in New York".

QUAN, D. C., AND S. TITMAN (1999): "Do Real Estate Prices and Stock Prices Move Together? An International Analysis", *Real Estate Economics*, 27, 183–207p, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1540-6229.00771.

QUINTERO, L. (2023): "Fewer Players, Fewer Homes: Concentration and the New Dynamics of Housing Supply".

REDDING, S. J. (2024): "Quantitative Urban Economics", Series: Working Paper Series Type: Working Paper.

ROLLET, V. (2025): "Zoning and the Dynamics of Urban Redevelopment".

SALIN, M., AND E. DE L'ESTOILE (2024): "Who Takes the Land? Quantifying the Use of Built-Up Land by French Economic Sectors to Assess Their Vulnerability to the 'No Net Land Take' Policy".

WOMACK, K. S. (2015): "Real Options and Urban Land Values: A Review of the Literature", *Journal of Real Estate Literature*, 23, 53–63p, Publisher: Routledge _eprint: https://doi.org/10.1080/10835547.2015.12090398.

APPENDIX

A. Data and cleaning procedure

A.1. Data presentation

Units of analysis The main unit of analysis is the cadastral parcel, a continuous land area with homogenous property rights. 99% of parcel delineations have been digitized by the French treasury. I focus the analysis on the Paris area, defined as the city of Paris (75), and the three conterminous *départements*: Hauts-de-Seine (92), Seine-Saint-Denis (93), and Val-de-Marne (94), accounting for 143 municipalities, including the 20 Parisian *arrondissements*. There are 679,513 parcels in 2024. Parcels can appear, disappear, split, or merge between years. As a result, 8% of parcels change definition across time. I keep track of those changes throughout the analysis. Parcel size is heavily skewed, with a median parcel of 313 square meters, an average parcel of 928 square meters. The largest parcels are non-housing parcels (airport infrastructure, agricultural land, penitentiaries, military infrastructure, factories).
Parcels can be easily aggregated at the municipal level as they are fully contained within a municipality. They can be aggregated within cadastral sections (N=5,503). There is no clear rule to define a cadastral section, but they are usually time-consistent aggregations of parcels. One can think of a section as the union of one or several US blocks. The median section contains 95 parcels (first quartile is 47 and third quartile is 158).

***Fichiers Fonciers*: fiscal approach** Every privately (as opposed to publicly) owned built units are taxed in France. Tax is collected yearly based on the unit destination (housing, tertiary, secondary, etc.), location (part of the tax rate is decided locally), and size (except for industrial units, which are taxed in a different way, hence no size is reported for most of them). A unit is part of a building. For instance, a single-family house is hosting one unit. A condominium with 5 flats hosts 5 units. A mixed-use building, with one restaurant on the floorspace and two flats above hosts 3 units (one tertiary and two housing). Taxed housing annexes, such as parking spaces or basement storage, are counted separately. Each building is mapped to a parcel. A parcel can host one or more fiscal building but a fiscal building cannot be located on more than one parcel. This leads to a number of "false positive" when mapping units to parcels. Indeed, the link between the unit and the parcel is done using the unit postal address. If, for some reason, the mailbox is located on another parcel (most of the time, an extremely small one, contiguous to the one hosting the actual physical building), then this latter parcel will host the units. If we were to use this parcel information to compute floorspace-area ratio, we would dramatically overestimate it. The same logic applies for "false" empty parcels. If the address declared is common to two parcels (same building front door for two contiguous parcels), then one of the two parcels is declared to be empty as all units are recorded on the same parcel. Starting in 2024, most buildings in the *Fichiers Fonciers* can be matched with the *BD Topo* (see below), hence allowing the mapping of building spanning over several parcels. See A.2 for more details on how parcels occupation is defined and how I deal with false positives after 2024.

***BD Topo*: physical approach** The BD TOPO is a 3D database produced by the French national geography institute (*IGN*). It provides the height and footprint of each physical building in France. 98% of buildings have an ID, the *ID-RNB*, which is common to various database. Although the dataset is issued yearly, the data collection can be lagged across time. According to the user manual, each building is referenced every 4 to 5 years. To measure the envelope (height times parcel coverage) precisely, I use the latest vintage before the permit was issued and the earliest vintage after the construction ends. Three cases are possible: the redeveloped building appears in the BD TOPO and I can compare the before and after envelopes; the redevelopment is undergoing in the BD TOPO and the measurement exhibits a torn down building; the latest vintage available is before the construction starting date and it is not possible to precisely ensure that the redevelopment implied a tear down. In this latter case, I rely on the fiscal data only and consider as torn downs buildings those where taxed units change identifiers between two given years.

***SITADEL*: building permits** Any project creating a housing unit must be accepted by the municipality where the project is taking place. When a project is accepted, a permit is issued and the permit data are reported in SITADEL database (with a lag up to 6 months). I use the data provided as part of *Base Nationale des Bâtiments* project, where a cleaning of the parcel identifier has been done. The "raw" database is not retreated and some parcel identifiers cannot be retrieved. Although a large number of information can be filled, most of them are not reported, especially regarding the transition in use (transition from office to housing units for instance). I rely on the information contained in the Fichiers Fonciers to infer the destination pre- and post-construction. The start and end construction dates are only filled for 60% of projects. Among the safest information, I particularly focus on the total surface and count of units created. I keep projects creating at least one housing unit, where this unit's floorspace must be 10 square meters or larger.

A.2. Define 2025 parcel occupation

Since 2024, in a joint effort, the *IGN* and the *Fichiers Fonciers* are matching units to their physical buildings. Around 75% of units are matched to a physical building as of 2025. I rely on spatial join for the remaining 25%. To obtain, for each physical parcel its "real" occupation, I apply the following algorithm:

1. Spatially match each parcel to the BD Topo building. I obtain a *idpar* $\leftrightarrow$ *ID-RNB* match.
2. Attribute the parcel composition to the BD Topo building based on the previous match, weighting for the coverage
3. Aggregate at the building level
4. Ventilate the building composition to all parcels it overlaps, weighting for the coverage

Each parcel and each BD Topo building has a polygon shapefile. In a first step, I match the parcels and the buildings and compute the overlapping area. Using this match, I attribute the parcels' composition to the BD Topo building (step 2). If the building is overlapping several parcels, or if a parcel hosts several buildings, I weight the contribution of each parcel to the buildings. Finally (step 4), I ventilate the building composition to all parcels the building lies upon, weighting for the coverage. Figure XXX helps visualizing the procedure.
I use this match in the model calibration for the current housing and non-housing stock at the parcel level. I describe the parcel occupation (Figure A1) in 2025. The decision rule is the following. Let i be the destination and $n = \sum_i n_i$ the number of units. If $n_i/n \geq 0.90$, then parcel destination is i. Otherwise, the parcel is said to be of mixed use.

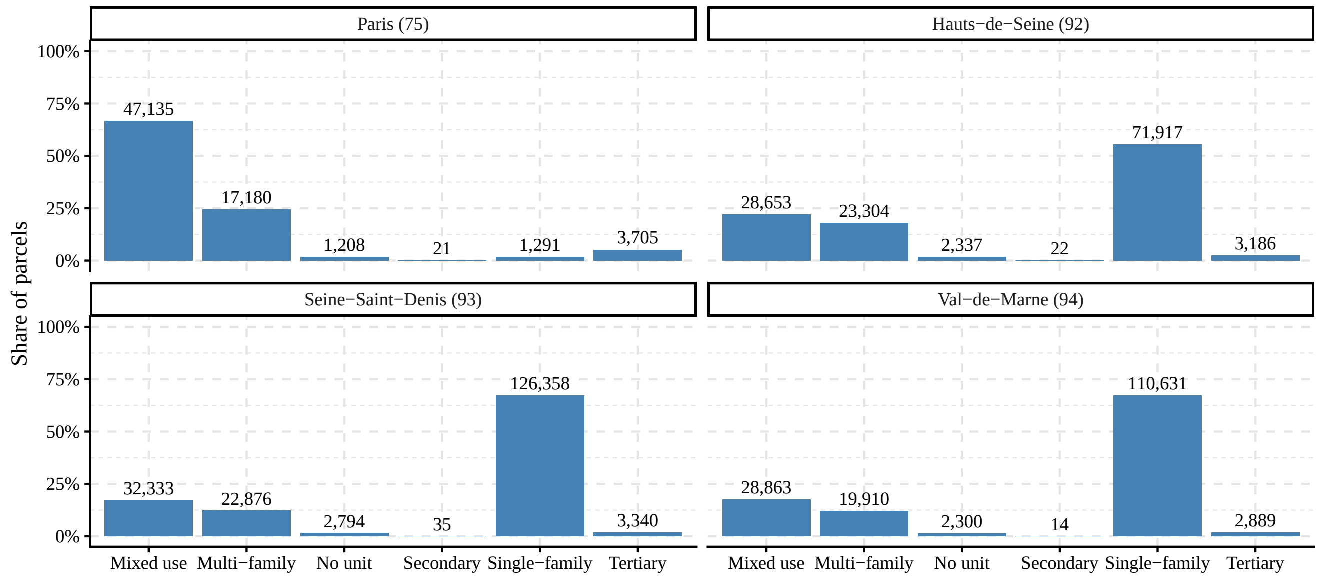

Data: Fichiers Fonciers (2025). Share of parcels by destination (x-axis), across Paris area *départements*. Count of parcels above each column. A parcel has destination i if $n_i/n > 0.9$ where n_i is the count of unit with destination i and n is the total count of units on the parcel.

Figure A1: Parcels destination across *départements*

A.3. Define land-use regulations and housing supply potential

Land-use regulations are set by the municipal council. Each council sets a very precise set of rules, often at the street or at the parcel level, along a large number of dimensions. To the best of my knowledge, there is no precise and exhaustive open-source dataset summarizing those regulations. Chapelle et al. (2025) rely on proprietary data and manual recoding.

It is possible to summarize the land-use regulations along three main dimensions: zoning, the building-coverage ratio (BCR), and the maximum height. In the model estimation, I focus on parcels zoned for housing and remove parcels specifically zoned for activity, natural zones, and other types of non-housing zoning, from the analysis.

To overcome the absence of precise data on BCR and maximum height, I infer the land-use regulations in place using the current building envelope. For each cadastral section, I measure the 90th percentile of building height, building-coverage ratio, floors, and housing floorspace. I then infer the maximum floorspace a developer could supply as:

$$\bar{h}_i = \overline{\text{BCR}}_i \times T_i \times \left\lfloor \frac{\bar{H}_i}{h_{\text{floor}}} \right\rfloor \tag{25}$$

where I assume a standard floor-to-ceiling height of $h_{\text{floor}} = 2.5$ meters. In absence of precise, georeferenced data on land-use regulations, I infer the maximum envelope using the 90$^{\text{th}}$ percentile envelope in each parcel section.

B. Calibration appendix

This appendix develops the calibration summarised in Section 4. Subsection B.1 reconstructs the latent demand process and its parameters. Subsection B.2 sets the financial block and derives the option value. Subsection B.3 estimates the construction technology. Subsections B.4 and B.5 calibrate the two frictions. Subsection B.6 measures demolition costs. Subsection B.7 documents the parcel-level variables.

B.1. Demand dynamics

Identifying the demand shifter A_t The demand shifter A_t is latent. Market clearing gives $R_t = A_t Q_t^{-\alpha}$, so $A_t = R_t Q_t^{\alpha}$. Identification requires an estimate of α and observations of the rental price and the aggregate stock over time. I build an annual series for the Paris area from 2014 to 2023, using rents from the Observatoires des Loyers and the housing stock from the Fichiers Fonciers. Given $\alpha = 0.35$ (Chapelle et al., 2023), I compute $\hat{A}_t = R_t Q_t^{\alpha}$ for each year, take logs, fit a linear time trend for μ and take the residual variance for σ^2. The construction lag does not enter this step. It shifts the timing of deliveries, and Q_t is measured on the standing stock, which is the delivered quantity by construction.

Expected growth rate of demand μ μ is the annualised drift of $\log \hat{A}_t$, estimated by regressing $\Delta \log \hat{A}_t$ on a constant. I obtain $\hat{\mu} = 0.01$ (0.00). The estimate aggregates population growth, real income growth and preference shifts over the sample period. As an external check I compare it with average annual population growth in the Paris region from the Census (INSEE RP), which bounds μ from below since population is one driver of demand among several. I find $\mu^{\text{RP}} = NA$. Note that μ is a property of the fundamental process alone. It governs the aggregate simulation of A_t and is absent from the individual stopping problem, which is driven by the perceived rent drift δ (Section 3.1).

Volatility of demand σ σ is the standard deviation of the annual increments $\Delta \log \hat{A}_t$, taken from the residuals of the regression above, giving $\hat{\sigma} = 0.04$. It is estimated on the fundamental process but applied to the perceived rent process in (4). Under the atomistic assumption the agent treats Q_t as fixed, so $\text{Var}(d \log R_t) = \text{Var}(d \log A_t)$ and the two volatilities coincide by construction. This parameter is the sole driver of the option multiplier $\beta/(\beta - 1)$ once ρ and δ are set. A higher σ raises β, raises the redevelopment threshold and raises the value of waiting.

Inverse demand elasticity α α governs how strongly a supply expansion depresses rents. I take $\alpha = 0.35$ from Chapelle et al. (2023), who estimate the inverse demand elasticity for the Paris region by instrumental variables. It enters the calibration in two places only, the reconstruction of $\hat{A}_t$ and the aggregate elasticity check of Section 5.1. It enters neither β nor any parcel-level threshold $R_i^{\star}$, since atomistic agents do not internalise the supply response. It does govern the strength of the general equilibrium level feedback in the transition path, where a larger α means a given wave of deliveries depresses the market-clearing rent more and slows subsequent redevelopment.

B.2. Rental price and financial parameters

Rental price growth rate δ δ is the expected growth rate of rents from the individual agent's perspective, treated as a primitive of the developer's price expectations (Section 3.1). I pin it down through the net residential yield rather than from realised rent growth, setting $\delta = \rho - \text{net\_cap}$ with $\text{net\_cap} = 0.045$ the external net residential capitalisation rate for Île-de-France, which yields $\delta = 0.02$. The yield-based construction is the appropriate one. The realised rent series is not a geometric Brownian motion, being the product of the diffusion A_t and the step function $Q_t^{-\alpha}$, so the trend of observed rents confounds the perceived drift with the supply path. The net yield is a market-clearing object that reveals what investors expect rents to do going forward, which is the object entering the agent's problem. It is also the object that prices the lag, since the revenue wedge $\lambda_u^R = e^{-(\rho-\delta)\ell_u}$ depends on the spread $\rho - \delta$ and not on either rate separately.

Discount rate ρ ρ is the real annual rate applied to future cash flows. I set $\rho = 0.065$, the real risk-free rate in France over the sample period (average 10-year OAT yield, Banque de France) augmented by a modest developer risk premium. Idiosyncratic parcel-level uncertainty is diversifiable, but the premium accounts for illiquidity and for the undiversifiable component of development risk. The discount rate enters in a second place, through the cost wedge $\lambda_u^C = (1 - e^{-\rho\ell_u})/(\rho\ell_u)$, which discounts a construction spend spread evenly over the build.

Option value β β is the positive root of the fundamental quadratic $\frac{1}{2}\sigma^2\beta(\beta-1) + \delta\beta - \rho = 0$ associated with the perceived rent process. With $\hat{\sigma} = 0.04$, $\delta = 0.02$ and $\rho = 0.065$ I obtain $\beta = 2.98$ and $\beta/(\beta-1) = 1.50 > 1$. The redevelopment threshold therefore exceeds the naive Marshallian trigger by that factor. Neither μ nor α enters the quadratic, so the multiplier is invariant to the realised aggregate supply path. By Proposition 1 it is invariant to the construction lag, and by Proposition 2 to the bargaining weight. The entire option-value block is thus untouched by the two frictions.

B.3. Construction technology

Cost convexity ζ ζ governs the curvature of variable costs. A developer building h^{new} square metres on parcel i faces $c_i(h^{\text{new}})^{(\zeta+1)/\zeta}$, convex whenever $(\zeta+1)/\zeta > 1$. Combes et al. (2021) report a capital share in housing production of 0.54 in a comparable context, implying $\zeta = 1.17$, which I adopt as the baseline. Section C, in progress, estimates ζ directly with a focus on multi-family units.

Average construction cost c_0 The parcel-specific scale is $c_i = c_0 T_i^{-1/\zeta}$, so that the variable cost of one square metre at unit density is c_i. I strip out the fixed cost and log-linearise (5),

$$\log C_i^{\text{var}} = \log c_0 - \frac{1}{\zeta}\log T_i + \left(1 + \frac{1}{\zeta}\right)\log h_i + e_i$$

and estimate it by OLS on the PRLN sample, extracting the intercept. I find $\hat{c}_0 = 1{,}186$€. Reported costs are contemporaneous with the works and are therefore already expressed on the construction schedule, so I do not apply the wedge λ_u^C a second time when recovering c_0.

B.4. Construction lag

The lag ℓ_u is the time between irreversible commitment and the date at which the completed building starts earning rent. It is indexed by the use of origin, since Section 2 shows that permit-to-completion duration differs sharply across origins. The adopted values are $\ell_{\text{hous}} = 2.0$, $\ell_{\text{mixed}} = 2.2$, $\ell_{\text{nonhous}} = 2.2$ and $\ell_{\text{vacant}} = 1.8$ years.
Table B1 reports the implied wedges. Two features stand out. The ratio κ_u is much closer to one than the raw revenue discount λ_u^R, because spreading the construction spend over the build offsets most of the delay on revenue, so what survives is the wedge on the net cash flow. And the implied movements are modest. For housing parcels the exercise payoff falls by 11.3%, the unconstrained optimal density by 3.0% and the redevelopment threshold rises by 5.7%. The lag is a real friction but a second-order one at the level of the individual project. Its first-order consequence is aggregate rather than parcel-level and operates through the delivery lag of Section 3.6.

B.5. Bargaining weight

The weight θ splits the redevelopment surplus between the landowner and the developer. Section 2 documents that land acquisition is locally concentrated, which motivates departing from the competitive benchmark $\theta = 1$. I proceed in three steps: a floor implied by the cross-section, a mapping from local concentration, and a validation against observed developer margins.

A floor implied by the cross-section A fixed cost cannot be negative, so (19) bounds θ from below parcel by parcel at $\theta_i = (p_i^L - V_i^w)/(\pi_i - V_i^w)$, an object computable without knowledge of F_i. Table B2 reports its distribution by use of origin. Any θ below the 95th percentile of that distribution implies that more than 5% of parcels carry a negative fixed cost, which the model cannot rationalise. The floor is $\theta \geq 0.38$.

A mapping from local concentration Within the floor I calibrate θ from the concentration of land acquisition. Table 4 reports a median municipal Herfindahl index corresponding to an effective number of developers of about seventeen. I adopt the transparent mapping

$$\theta_c = 1 - \mathrm{HHI}_c \tag{26}$$

so that a monopsonist captures the entire surplus and atomistic competition for land restores $\theta = 1$. Equation (26) is a calibration device and not a derivation, and I treat it as one point in the robustness band of Table C11 rather than as a structural claim. Two properties recommend it. It is monotone in the object the data measure, and it delivers a value inside the floor established above, which a mapping with more curvature would not.

Table B1: Construction lags and the implied wedges

Use of origin	ℓ_u	λ_u^C	λ_u^R	κ_u	Effect on $\Delta\pi$	Δh^*	ΔR^*
Housing	2.00	0.9377	0.9139	0.9746	-11.3%	-3.0%	5.7%
Mixed	2.25	0.9303	0.9037	0.9714	-12.6%	-3.3%	6.4%
Non-housing	2.25	0.9303	0.9037	0.9714	-12.6%	-3.3%	6.4%
Vacant	1.75	0.9452	0.9243	0.9778	-10.0%	-2.6%	5.0%

Data: Sitadel, model calibration. ℓ_u is the construction lag in years, by use of origin. λ_u^C is the present-value factor applied to the construction spend. λ_u^R is the present-value factor applied to rental income. $\kappa_u = \lambda_u^R/\lambda_u^C$ is the effective discount on the net cash flow. The last three columns report percentage changes relative to instantaneous construction ($\ell_u = 0$), holding the fixed cost and the continuation value fixed. $\Delta\pi$ is the change in the exercise payoff, Δh^* the change in the unconstrained optimal density, ΔR^* the change in the redevelopment threshold.

Table B2: Lower bound on the bargaining weight implied by $F_i \geq 0$

Use of origin	N	θ_i Median	p75	p95
Housing	344,803	-0.261	0.068	0.370
Mixed	47,898	-0.319	0.135	0.389
Non-housing	6,873	0.201	0.321	0.508
Vacant	1,600	0.372	0.464	0.624

Data: Model estimation. θ_i is the smallest landowner bargaining weight compatible with a non-negative fixed cost on parcel i, obtained by setting $F_i = 0$. Columns report the median, the 75th and the 95th percentile of its distribution within each use of origin. Adopting a weight below the reported 95th percentile implies that more than 5% of parcels in that category have a negative fixed cost. N is the number of parcels with $\pi_i > V_i^w$, for which θ_i is defined and increasing in the observed land price.

The adopted value is $\theta = 1.00$, one minus the median municipal Herfindahl index over permits with an identified developer. At that value developers retain 0.0% of the aggregate redevelopment surplus at the 2023 rent. I hold θ constant across parcels in the baseline. A municipality-specific θ_c is a direct extension of the same mapping and I report it as a robustness exercise, because the municipal index is estimated on few transactions in the thinner markets and inherits their sampling error.

Validation against developer margins The calibration has an independent counterpart in the accounts of the development industry. The share retained by the developer, $(1-\theta)S_i$, is the object that appears as the operating margin on a residential development operation. Two caveats bound the exercise. The developer's disagreement point is normalised to zero, so any positive outside option enters the price equation exactly as an addition to F_i and is absorbed into the residual. And θ is the weight in a bilateral negotiation over a single parcel. It is silent about multi-parcel assembly, which Section 2 shows carries a price premium of its own, and that premium is likewise absorbed into F_i.

B.6. Direct measurement of demolition costs

Data I measure development costs from the *Prix de revient des logements neufs* survey. It draws building permits at random from Sitadel, stratified by project size. Eligible developers, that is projects with at least one new unit and at least 75% of floorspace dedicated to housing, are asked about project size and project cost. The survey separates construction cost from demolition cost. I use both. The microdata report the permit identifier, which I match to Sitadel to recover the parcels covered and the project characteristics, in particular demolished floorspace by destination. The match yields 701 projects with demolition costs covering 1,247 parcels in 68 *départements*, of which 65 projects are in the Paris area. Reporting project characteristics is not mandatory in Sitadel, so the estimation sample is the 341 projects with complete characteristics.

Are demolition costs fixed? The model treats demolition as a component of F_i, so the first requirement is that it does not scale with the project built. I regress demolition cost on project size, measured by the count of new units and by new floorspace, with year and *région* fixed effects. Authorised floorspace carries no significant effect. The coefficient on authorised dwellings is significant only at the 10% level. Table C4 reports the results. Costs are fixed in the relevant sense: they depend on the ex ante composition of the parcel, not on the density chosen.

Levels by use of origin Demolition costs are heterogeneous across uses of origin. Projects on former secondary and tertiary parcels are the more expensive at face value, averaging 110,000€ and 148,000€ per project. Demolishing housing averages 48,000€. The ranking reverses once demolished floorspace is held constant. Housing demolition then costs up to three times the secondary rate, 150€ against 50€ per square metre, while retail, warehouse and tertiary demolition sit around 100€. Figure B2 plots both, per square metre in the left panel and in absolute terms in the right panel.

Hedonic decomposition I recover the per-square-metre cost of demolition by destination from a linear decomposition of total project cost. For project i,

$$\text{Demolition Cost}_i = \beta_0 + \gamma\,\text{Area}_i + \sum_k \beta_k \text{Surf}^{\text{dem}}_{ik} + X_{i,m(i)} + \varepsilon_i \tag{27}$$

where k indexes destinations (housing, tertiary, retail, artisan, warehouse, secondary). Each β_k is the marginal cost per square metre of demolishing floorspace of type k and β_0 is the mobilisation cost incurred irrespective of the quantity removed. X collects parcel-level controls (built density in a buffer, building age) and municipal controls (population density, average income). The sample is small ($N =$

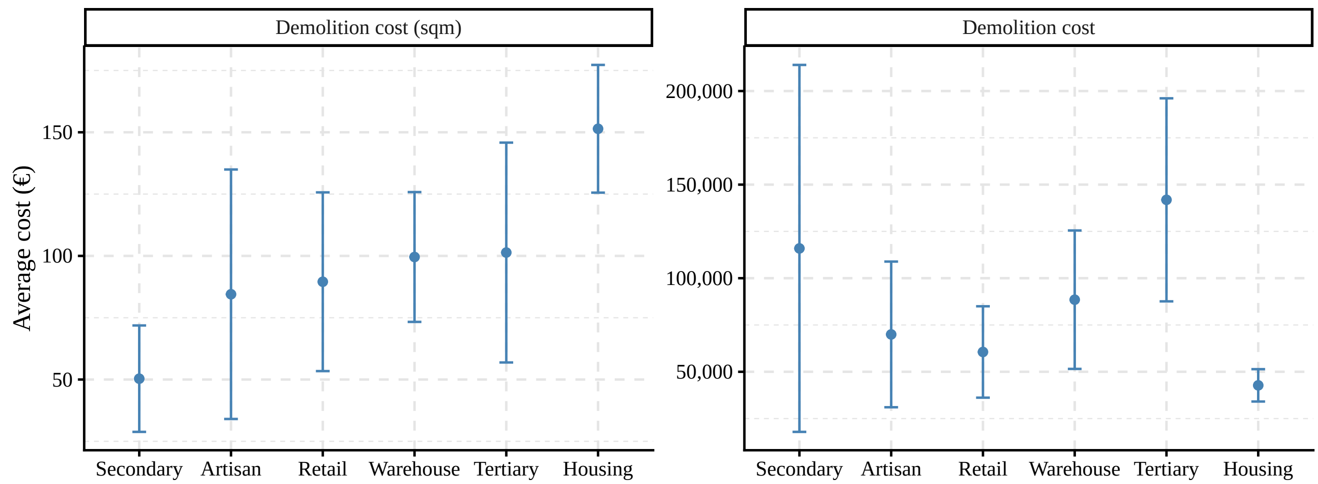

Data: PRLN. Average demolition costs by parcel's use of origin. Use of origin is the main floorspace occupation of the parcel. 95% confidence intervals.

Figure B2: Average demolition cost by use of origin

341) and cost is heavily right-skewed, so I report heteroskedasticity-robust analytical standard errors in parentheses alongside non-parametric bootstrap standard errors in brackets (1,000 resamples). Table C6 reports the estimates.

Robustness Columns (1) to (5) address the two threats specific to this design, heteroskedasticity driven by project scale and the leverage of a few very large projects. Column (1) is OLS in levels. It targets the estimand directly and predicts levels well, but weights every project equally and is therefore sensitive to outliers. Column (2) reweighs by inverse demolished floorspace, damping the influence of large projects at the cost of a lower R^2 and of down-weighting the observations that identify the rarer destinations. Columns (3) and (4) add year $\times$ region fixed effects to the OLS and WLS specifications, absorbing spatial and temporal differences in cost levels at the cost of thin cells for the rarer origins. Column (5) re-estimates the OLS after trimming observations flagged by Cook's distance ($N = 321$).
The estimates are stable across the OLS-type columns (1), (3) and (5) and cluster in the 35-56€ range: housing 45.7-52.0€, tertiary 36.8-45.7€, retail 37.4-44.9€, artisan 38.1-55.6€, warehouse 36.7-42.3€ and secondary 35.2-35.5€. All destination coefficients are significant at the 1% level under both inference schemes and the bootstrap never overturns significance. The parcel-area coefficient is indistinguishable from zero throughout, confirming that cost scales with the structures removed and not with the land. The mobilisation intercept is sizeable and precisely estimated at 25,269€ in column (1). The WLS columns shift the residential and tertiary slopes upward by roughly a quarter, reflecting the reweighing away from large low-unit-cost projects. That the OLS slopes barely move when fixed effects are added, column (1) against column (3), shows that regional and temporal variation shifts the level of cost and not the per-square-metre structure of interest.

Prediction I retain column (1) as the prediction equation. It estimates the target quantity directly, uses the full sample, delivers the best out-of-sample fit among the specifications considered and yields an interpretable intercept. I apply the coefficients to the destination-specific demolished floorspace of every parcel in the simulation sample, taken from the Fichiers Fonciers, to obtain a predicted demolition cost $\widehat{DC}_i$. In the tradition of Glaeser and Gyourko (2003) the residual fixed cost is $F_i = RC_i + FC_i$ with FC_i the construction fixed cost proper, of which $\widehat{DC}_i$ is the measured part. Section 5.6 uses $\widehat{DC}_i$ to bound the share of F_i attributable to physical demolition and reports the remainder as unexplained. Further iterations will add controls for built density around the project, municipal income and distance to the core city.

B.7. Parcel-level variables

Land price The DVF reports transactions involving parcels. French property rights link buildings to the parcel, so a parcel cannot be bought without the structures on it and reciprocally. A developer buying a parcel therefore buys land together with the pre-existing structure, which is exactly the object p_i^L of the model, since the bargained price compensates the landowner for the standing value V_i^w she gives up. I run a hedonic regression of log parcel price on observable characteristics (floorspace, unit destinations, height, footprint, number of units) with area and time fixed effects, and predict the price of every parcel in the sample.[11]

Fixed costs F_i Given the predicted land price $\hat{p}_i^L$, the wedges (λ_i^C, κ_i) and the weight θ, I recover F_i from (18), evaluating π_i at (3.2) or (3.2) depending on whether the density cap binds. The residual also absorbs costs the model does not name, including land assembly premia (Brooks and Lutz, 2016; Goeyvaerts and Boogaerts, 2021) and the developer's outside option. Given the data and the model I cannot disentangle these further, and Section 5.6 treats the unexplained share accordingly.

Land use regulation $\bar{h}_i$ $\bar{h}_i$ is the maximum floorspace that can legally be built on parcel i. The PLU constrains three dimensions. Zoning determines whether residential construction is permitted at all, and I restrict the sample to zones where it is. Maximum height $\bar{H}_i$ and the building coverage ratio $\overline{\text{BCR}}_i$ translate into a floorspace limit,

$$\bar{h}_i = \overline{\text{BCR}}_i \times T_i \times \left\lfloor \frac{\bar{H}_i}{h_{\text{floor}}} \right\rfloor \tag{28}$$

with a standard floor-to-ceiling height of $h_{\text{floor}} = 2.5$ metres. Absent precise georeferenced regulatory data, I infer the envelope from the 90th percentile envelope within each cadastral section.

Rental price R_t and location quality ϕ_i R_t is the median rental price index for the Paris metropolitan area in euros per square metre per year, from the Observatoires des Loyers. I set ϕ_i to the ratio of the observed median rent in the municipality containing parcel i to the metropolitan median in 2023,

$$\hat{\phi}_i = \frac{\bar{R}_{m(i),2023}}{\bar{R}_{2023}}$$

where $m(i)$ is the municipality of parcel i. This matches location-specific rent levels exactly in the cross-section and loads all time variation onto the common factor. The distribution of $\hat{\phi}_i$ runs from 0.77 in peripheral municipalities to 1.57 in inner Paris, capturing the rent gradient of the metropolitan area.

Land area T_i Land area in square metres comes from the cadastre (Fichiers Fonciers). It is a precisely measured administrative variable requiring no imputation. It enters through the regulatory constraint (28) and through the variable cost.

Non-housing revenue v_i^C For commercial and industrial parcels, v_i^C is the annual flow from the current use, entering the continuation value (13). I take it from the *valeurs locatives des locaux professionnels* of the Direction Générale des Finances Publiques, which assigns a fiscal rental value to each parcel as a function of floorspace weighted by building type and destination, and of a local coefficient. I deflate by the CPI to obtain real values consistent with the rest of the model and set $v_i^C = 0$ for vacant parcels,

[11] See Chapelle et al. (2026) for a discussion of land value estimation in France.

consistent with (14). Figure C7 plots imputed commercial rent against housing rent per square metre at the municipal level. The two correlate, with commercial rents relatively higher in Paris. I further assume $\delta^C = \delta$, that non-housing and housing rents follow the same trend. This is a strong assumption but one supported by the empirical literature (Quan and Titman, 1999; Gyourko, 2009).

Existing housing stock $h_i^{\mathbf{old}}$ For parcels in housing use, h_i^{old} is total residential floorspace from the Fichiers Fonciers. It enters the continuation value (12) and determines the densification gain $\bar{h}_i - h_i^{\text{old}}$, a key determinant of the threshold $R_i^\star$.

C. Elasticity of construction costs and project size

I estimate the variable cost of construction on project-level data from the PRLN survey. I assume a Cobb-Douglas housing production function with land and capital as inputs (Combes et al., 2021; D'Amico et al., 2024), without imposing constant returns to scale:

$$h_i = A_i \cdot T_i^a \cdot K_i^b$$

where h_i is total floorspace produced, T_i is parcel area, K_i is capital input, A_i is total factor productivity, and (a, b) are the land and capital elasticities of output. Inverting for the required capital and multiplying by the rental rate of capital r yields the variable cost of producing h_i on a parcel of fixed size T_i:

$$C_i^{\text{var}}(h_i) = r \cdot A_i^{-1/b} \cdot T_i^{-a/b} \cdot h_i^{1/b}. \tag{29}$$

Define the convexity parameter $\zeta \equiv 1/b$, the land gradient $\eta \equiv a/b$, and the baseline cost parameter $c_0 \equiv r \cdot A^{-\zeta}$, where A is average productivity. Equation (29) becomes

$$C_i^{\text{var}}(h_i) = \underbrace{c_0\, T_i^{-\eta}}_{c_i}\, h_i^{\zeta}, \tag{30}$$

which is exactly the variable cost component entering the developer's problem in (7). Total development cost adds the scale-invariant fixed cost F_i (demolition, remediation, permitting), so that $C_i = F_i + C_i^{\text{var}}(h_i)$. Variable cost decreases in parcel area (land economies of scale, governed by $-\eta$) and is convex in project scale (engineering limits, coordination costs, governed by ζ). Convexity arises from decreasing returns to the variable input ($b < 1$, i.e. $\zeta > 1$) and is separate from the land gradient.
Taking logarithms of (30) yields the estimating equation

$$\ln C_i^{\text{var}} = \beta_0 + \beta_T \ln T_i + \beta_h \ln h_i + \gamma X_i + \mu_{mt} + \epsilon_i, \tag{31}$$

where X_i collects project-type indicators (new construction, demolition-reconstruction, within-building transformation), parcel characteristics (slope, flood-zone exposure, soil pollution), and neighbourhood density, μ_{mt} are municipality-by-time fixed effects absorbing local input prices and construction productivity, and standard errors are clustered at the municipality level.
The two slopes map directly onto the technology:

$$\beta_h = \frac{1}{b} = \zeta, \qquad \beta_T = -\frac{a}{b} = -\eta, \tag{32}$$

so the convexity parameter is recovered directly as $\hat{\zeta} = \beta_h$, the land gradient as $\hat{\eta} = -\beta_T$, and the underlying elasticities as $\hat{b} = 1/\beta_h$ and $\hat{a} = -\beta_T/\beta_h$. Convexity ($\beta_h > 1$) and land economies of scale ($\beta_T < 0$) are identified separately from the two coefficients, and returns to scale $a + b$ are left unrestricted. The intercept identifies $\ln c_0$, so ζ, η, and c_0 that enter (7) are all estimated from the same regression.
The parameters of interest are technological: ζ and η govern the shape of the cost function $C_i^{\text{var}}(h_i)$, i.e. the mapping from floorspace and parcel area to variable cost. Consistent estimation requires that, conditional on controls and fixed effects, the cost shifter ϵ_i be uncorrelated with $\ln T_i$ and $\ln h_i$.
Two concerns arise, both from developer optimisation. First, density h_i is chosen given expected profitability, so a positive cost shock ($\epsilon_i > 0$) may induce a smaller project, correlating h_i with ϵ_i. This simultaneity moves the developer along the cost curve rather than rotating it: it affects which points on $C_i^{\text{var}}(\cdot)$ are observed, not the curvature itself. Because the structural model treats $c_i h_i^{\zeta}$ as a technology and does not endow parcels with cost-shock-correlated heterogeneity beyond $c_0 T_i^{-\eta}$, the OLS estimate of β_h is the object the model requires, and the municipality-by-time fixed effects plus parcel controls absorb the local cost heterogeneity that would otherwise contaminate it. Second, parcel size T_i is itself

partly chosen (assembly, subdivision). The rich parcel controls and fixed effects mitigate, but do not eliminate, correlation between T_i and unobserved cost attributes.[12]

PRLN costs are self-reported and subject to reporting error. Because this error is in the dependent variable, it inflates standard errors but does not bias the slopes under classical assumptions. To guard against non-classical outliers, I trim the top and bottom 1% of the cost distribution. I additionally verify robustness to: (i) the level-form non-linear least squares estimator with an explicit additive F_i, (ii) a translog specification adding $(\ln h_i)^2$ and $\ln h_i \ln T_i$, whose extra terms test the Cobb-Douglas restriction, and (iii) interacting $\ln h_i$ with project type, which allows the convexity to differ between new construction and demolition-reconstruction.

Table xxx reports the results. The estimates imply $\hat{\beta}_h \approx 1.05$ and $\hat{\beta}_T \approx -0.05$, hence $\hat{\zeta} \approx 1.05$, $\hat{\eta} \approx 0.05$, a capital elasticity $\hat{b} \approx 0.95$, and $\hat{c}_0 = 1,064$ euros per square metre: construction costs are mildly convex in project scale with weak land economies of scale.

Work in progress.

[12] A potential robustness against the simultaneity in h_i would be to instrument density with exogenous shifters of optimal scale that plausibly do not enter the cost technology, such as the maximum allowable floorspace under local land-use regulations $\bar{h}_i$, or a Bartik-style demand shock leaving out the construction sector.

Tables

Table C3: Housing supply intensity by project type and previous use

		FAR		UAR (×10)		Units		Floorspace (m^2)	
Category	N	Mean	Median	Mean	Median	Mean	Median	Mean	Median
By: project type									
Demolition	9113	1.12	0.57	0.17	0.05	20.33	2.0	1224.83	225.0
New units	2853	0.82	0.48	0.09	0.04	8.85	1.0	616.34	150.0
No demolition	5246	0.86	0.45	0.14	0.05	11.05	1.0	603.17	152.0
By: previous destination									
House	7218	0.79	0.46	0.10	0.04	7.59	1.0	521.56	162.0
Mixed use	2875	1.48	0.97	0.26	0.17	29.30	12.0	1646.23	673.0
Multi-unit housing	3057	1.01	0.55	0.16	0.05	18.91	1.0	1093.86	198.0
New units	3084	0.81	0.48	0.09	0.04	9.03	1.0	626.02	150.0
Non-housing	978	1.60	1.16	0.30	0.22	44.82	15.5	2364.45	800.5

Data: Fichiers Fonciers, BD TOPO, Sitadel. FAR and UAR are floorspace-to-area and units-to-area ratio, respectively. Demolition are projects with a change in the building envelope; new units are projects on bare land; no demolition are projects without change in the building envelope. Previous use is defined as the destination of the parcel before the projects.

Table C4: Correlation between project size and demolition costs

	Demolition cost (1)
Authorised floorspace	31.3*
	(16.3)
Number of dwellings	-1,081.5
	(1,005.1)
Observations	341
R^2	0.330
Year fixed effects	✓
Region fixed effects	✓

Data: PRLN, Sitadel. OLS regression with year and region fixed effects of the observed demolition costs (in euro) on project size (count of new dwelling and new floorspace). Heteroskedasticity-robust standard errors reported in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table C5: Land-price premium in land assembly (robustness)

Gap (months)	Model	Naive assembly	(s.e.)	Sequential assembly	(s.e.)	Obs.
0	(1)	.	(.)	89.437	(76.76)	9,482
	(2)	.	(.)	92.917	(76.68)	9,482
	(3)	.	(.)	137.331*	(70.19)	9,482
	(4)	.	(.)	152.425**	(67.78)	9,482
	(5)	.	(.)	125.871*	(70.64)	9,482
1	(1)	176.296*	(90.75)	66.890	(89.73)	9,482
	(2)	174.899*	(90.90)	71.619	(89.51)	9,482
	(3)	149.197*	(84.48)	134.205	(81.95)	9,482
	(4)	156.142**	(78.06)	151.444*	(80.09)	9,482
	(5)	131.926	(86.57)	124.238	(81.73)	9,482
3	(1)	235.787**	(98.82)	-3.618	(114.67)	9,482
	(2)	234.827**	(98.99)	2.601	(114.54)	9,482
	(3)	219.760**	(90.05)	83.677	(98.70)	9,482
	(4)	229.435**	(88.40)	102.256	(96.90)	9,482
	(5)	204.484**	(87.81)	73.244	(101.76)	9,482
6	(1)	140.922	(85.84)	29.607	(143.11)	9,482
	(2)	140.796	(85.76)	37.221	(143.01)	9,482
	(3)	128.151	(79.73)	148.335	(122.95)	9,482
	(4)	131.222*	(76.66)	177.871	(120.51)	9,482
	(5)	106.172	(81.08)	149.989	(121.82)	9,482

Data: Sitadel, DVF. OLS regression (1) robustness for different windows to define Naive *vs.* Sequential assembly. All models include year and municipal fixed effects. Model numbering matches Table 3. Municipal-level clustered standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table C6: Hedonic decomposition of demolition costs

	(1) OLS	(2) WLS	(3) OLS	(4) WLS	(5) OLS
Parcel area	0.43	0.35	0.34	0.41	−0.43
	(0.53)	(0.30)	(0.50)	(0.29)	(0.55)
	[0.80]	[0.66]	[0.82]	[0.74]	[0.59]
Housing	45.66***	68.55***	45.57***	70.84***	51.99***
	(9.73)	(13.83)	(9.10)	(14.24)	(5.93)
	[10.39]	[14.10]	[10.25]	[14.79]	[6.16]
Tertiary	36.76***	55.24***	37.54***	55.86***	45.72***
	(13.82)	(12.69)	(11.76)	(12.87)	(3.97)
	[15.33]	[12.66]	[13.91]	[12.42]	[4.72]
Retail	37.35**	46.88***	41.91***	52.13***	44.91***
	(14.49)	(11.74)	(10.64)	(11.21)	(3.84)
	[15.48]	[12.59]	[11.99]	[11.86]	[5.01]
Artisan	41.12***	52.63***	38.08***	53.07***	55.63***
	(6.22)	(12.92)	(6.16)	(14.45)	(13.80)
	[11.89]	[15.91]	[12.47]	[17.47]	[16.27]
Warehouse	37.62***	40.46***	36.68***	39.84***	42.25***
	(6.13)	(6.76)	(6.50)	(6.99)	(5.55)
	[8.88]	[7.94]	[9.74]	[8.13]	[9.78]
Secondary	35.54***	31.76***	35.18***	31.98***	35.50***
	(9.34)	(9.25)	(7.14)	(8.81)	(9.51)
	[13.89]	[9.88]	[10.77]	[9.57]	[20.72]
Constant	25,268.59***	14,469.20***			19,864.92***
	(4,001.72)	(2,869.75)			(2,802.99)
	[4,113.69]	[3,076.48]			[2,783.10]
Year × Region FE	No	No	Yes	Yes	No
Cook's trim	No	No	No	No	Yes
Observations	341	341	341	341	321
R^2	0.387	0.177	0.437	0.208	0.545

Data: PRLN, Sitadel. Column 1, 3, 5: OLS regressions. Column 2, 4: WLS regressions by the inverse of total demolished floorspace. Cook's trim drops the 20 most influential projects. Columns 3 and 4 include *région* and year fixed effects. Heteroskedasticity-robust standard errors reported in parentheses. Bootstrapped (1,000 iterations) standard errors reported in brackets. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table C7: Extensive *vs.* intensive margin parcel predicted redevelopment

		Share of parcels		
Use	Dept.	N	Interior share	$\Pr(\pi < 0)$
Housing	75	10,140	98.9%	0.19%
	92	95,953	26.6%	0.02%
	93	157,056	17.4%	0.01%
	94	136,295	16.2%	0.02%
Mixed	75	55,189	99.6%	0.05%
	92	28,575	66.8%	0.12%
	93	26,182	54.0%	0.32%
	94	24,582	48.4%	0.07%
Non-housing	75	1,472	86.0%	9.71%
	92	2,209	63.2%	3.67%
	93	2,613	48.2%	2.95%
	94	2,094	41.5%	1.77%
Vacant	75	42	85.7%	11.90%
	92	541	28.8%	1.11%
	93	672	22.8%	0.15%
	94	359	21.2%	0.56%

Data: Model estimation. Under the baseline estimation of the model, evaluated with 2023 rent, the table reports the share of parcels that are hitting the density cap $\bar{h}_i$ and the share of parcels with a negative fixed cost.

Table C8: Fixed cost linear decomposition

	log(Fixed cost per sqm)					
		Disaggregated		Predicted cost		+ FE
	(1)	(2)	(3)	(4)	(5)	(6)
notationCE	0.046	0.045	0.019	0.051	0.020	-0.117**
	(0.080)	(0.077)	(0.068)	(0.078)	(0.068)	(0.056)
notatione2Cr-BE	0.271***	0.251***	0.224***	0.265***	0.229***	0.171
	(0.080)	(0.076)	(0.080)	(0.080)	(0.081)	(0.146)
notatione4AM	0.282***	0.278***	0.244***	0.272***	0.238***	0.108
	(0.100)	(0.097)	(0.084)	(0.101)	(0.085)	(0.104)
notatione4AP	0.058	0.060	0.047	0.064	0.050	-0.083
	(0.114)	(0.106)	(0.093)	(0.110)	(0.095)	(0.060)
notatione4GS	0.322***	0.304***	0.249***	0.313***	0.250***	0.004
	(0.082)	(0.079)	(0.073)	(0.081)	(0.073)	(0.063)
notatione4SC-AH	0.304***	0.286***	0.232**	0.289***	0.231**	-0.008
	(0.084)	(0.092)	(0.099)	(0.082)	(0.092)	(0.058)
notatione5CG	0.086	0.090	0.078	0.090	0.077	-0.017
	(0.077)	(0.075)	(0.062)	(0.075)	(0.061)	(0.052)
notatione5MC	0.104	0.095	0.071	0.107	0.076	-0.068
	(0.087)	(0.083)	(0.071)	(0.085)	(0.071)	(0.057)
notatione6-7MGC	0.017	0.016	0.024	0.022	0.027	-0.084
	(0.076)	(0.074)	(0.067)	(0.075)	(0.066)	(0.056)
notatione6CSO	0.004	0.010	0.011	0.011	0.012	-0.089
	(0.079)	(0.077)	(0.067)	(0.077)	(0.066)	(0.058)
notatione6SB-A	0.064	0.057	0.054	0.066	0.059	-0.069
	(0.078)	(0.075)	(0.064)	(0.077)	(0.064)	(0.054)
notatione6SM	0.129	0.121	0.105	0.131	0.109	-0.027
	(0.084)	(0.081)	(0.070)	(0.083)	(0.070)	(0.063)
notatione7C	0.118	0.106	0.056	0.134	0.068	-0.193
	(0.274)	(0.272)	(0.284)	(0.274)	(0.286)	(0.186)
notatione7CCh-MP	-0.082	-0.078	-0.073	-0.073	-0.070	-0.082
	(0.110)	(0.109)	(0.109)	(0.110)	(0.109)	(0.064)
notatione7G	-0.001	-0.001	-0.002	0.005	0.001	-0.122**
	(0.077)	(0.075)	(0.067)	(0.076)	(0.066)	(0.056)
notatione7MC	0.064	0.059	0.096	0.071	0.104	0.006
	(0.097)	(0.093)	(0.097)	(0.094)	(0.097)	(0.056)
notatione7MP	-0.008	0.0001	-0.002	0.002	0.001	-0.031
	(0.091)	(0.088)	(0.081)	(0.090)	(0.082)	(0.062)
notatione7MS	-0.042	-0.037	-0.042	-0.034	-0.041	-0.139**
	(0.077)	(0.075)	(0.068)	(0.076)	(0.067)	(0.057)
notationFv	0.155	0.133	0.046	0.155	0.054	-0.057
	(0.128)	(0.125)	(0.117)	(0.126)	(0.117)	(0.127)
notationFx	0.028	0.025	0.099	0.034	0.105	-0.050
	(0.177)	(0.177)	(0.155)	(0.176)	(0.155)	(0.193)
notationFy	0.152*	0.137*	0.100	0.153*	0.106	-0.094
	(0.084)	(0.081)	(0.071)	(0.082)	(0.071)	(0.060)
notationFz	0.077	0.068	0.046	0.081	0.052	-0.111*
	(0.086)	(0.083)	(0.072)	(0.085)	(0.072)	(0.059)
notationg1AR	-0.004	0.0008	0.004	0.004	0.006	-0.145**
	(0.078)	(0.076)	(0.069)	(0.077)	(0.069)	(0.058)
notationg1BS	-0.005	0.002	-0.003	0.003	-0.002	-0.138**
	(0.081)	(0.079)	(0.074)	(0.080)	(0.073)	(0.058)
notationg1CB	0.061	0.048	0.013	0.066	0.020	-0.058
	(0.093)	(0.090)	(0.086)	(0.092)	(0.086)	(0.066)
notationg1MH	-0.015	-0.0003	-0.020	-0.005	-0.022	-0.139**
	(0.077)	(0.075)	(0.066)	(0.076)	(0.065)	(0.062)
notationg1SF	0.0008	0.008	-0.008	0.009	-0.008	-0.153**
	(0.085)	(0.083)	(0.076)	(0.084)	(0.075)	(0.071)
notationg1SO	-0.240***	-0.223***	-0.157**	-0.231***	-0.159**	-0.234***
	(0.081)	(0.078)	(0.071)	(0.079)	(0.070)	(0.057)
notationHydro	0.259**	0.274**	0.212**	0.328***	0.237**	0.001
	(0.125)	(0.131)	(0.106)	(0.125)	(0.102)	(0.076)
notationLP	-0.029	-0.040	-0.061	-0.026	-0.055	-0.183***
	(0.088)	(0.086)	(0.079)	(0.087)	(0.079)	(0.061)
notationp-IVMM	-0.075	-0.074	-0.065	-0.069	-0.062	-0.215***
	(0.104)	(0.103)	(0.095)	(0.103)	(0.094)	(0.068)
notationX	-0.013	-0.010	-0.009	-0.008	-0.010	0.018
	(0.110)	(0.109)	(0.110)	(0.107)	(0.109)	(0.093)

Table C9: Shapley shares across the parameter band

	Share of total variance (%), [min, max]	
Component	With controls	Without controls
Demolition	[8.8, 10.4]	[11.3, 13.8]
Parcel characteristics	[14.2, 17.6]	[19.7, 26.2]
Land-use regulation	[11.8, 17.2]	[17.5, 24.3]
Municipal (location)	[19.4, 26.8]	–
Unexplained	[28.0, 45.9]	[35.7, 51.5]

Data: Model estimation. Each cell reports the minimum and the maximum share of total variance across the combinations of $\rho \in \{0.060, 0.065, 0.070\}$, $\zeta \in \{1.00, 1.17\}$ and $c_0 \in \{1064, 1186, 1300\}$. In every cell $\delta = \rho - \text{net\_cap}$, the construction wedges λ_u^C and κ_u and the continuation value V_i^w are recomputed at the cell's parameters, and F_i is re-recovered from (18). The partition is coarser than in Table 8: the four parcel blocks are collapsed into one, and the demolition block is the single predicted cost of Table C6 rather than the disaggregated surfaces, which lowers its share by roughly one percentage point relative to the main table.

Table C10: Shapley shares across first-stage price specifications

	Share of total variance (%), [min, max]	
Component	With controls	Without controls
Demolition	[10.3, 13.6]	[13.7, 17.6]
Parcel characteristics	[16.8, 17.5]	[24.6, 25.8]
Land-use regulation	[14.5, 16.3]	[20.7, 23.3]
Municipal (location)	[24.4, 26.2]	–
Unexplained	[29.5, 30.7]	[37.0, 37.5]

Data: Model estimation. Structural parameters are held at their calibrated values and only the hedonic prediction of the parcel price entering (18) is varied, across the baseline specification, a parsimonious one, a log specification, a saturated one, and the bootstrap mean of the hedonic prediction. Each cell reports the minimum and the maximum share of total variance across those five specifications. The partition matches Table C9.

Table C11: Sensitivity of the second stage to the construction lag and the bargaining weight

	θ	Median F_i/T_i	$\Pr(F_i < 0)$	Median R^*	Share ripe at R_0	$\varepsilon^{\text{model}}$
	0.80	4,522	0.2%	370	3.0%	0.85
	0.90	4,413	0.2%	367	2.4%	0.89
	0.95	4,355	0.2%	366	2.2%	0.53
$\ell_h = 0.0$	1.00	4,280	0.2%	365	2.0%	0.54
	0.80	4,177	0.2%	372	3.2%	1.07
	0.90	4,071	0.2%	369	2.5%	0.63
	0.95	4,013	0.2%	368	2.3%	0.74
$\ell_h = 1.0$	1.00	3,932	0.2%	367	2.1%	0.73
	0.80	3,847	0.2%	375	3.4%	0.80
	0.90	3,741	0.2%	372	2.7%	0.67
	0.95	3,680	0.2%	370	2.4%	0.70
$\ell_h = 2.0$	**1.00**	**3,600**	**0.2%**	**369**	**2.2%**	**0.63**
	0.80	3,535	0.3%	377	3.7%	1.78
	0.90	3,426	0.2%	374	2.9%	1.38
	0.95	3,365	0.2%	373	2.6%	0.68
$\ell_h = 3.0$	1.00	3,277	0.2%	372	2.3%	0.66

Data: Model estimation. Each cell re-runs the second stage of Section 4.2 at the stated pair, holding every structural parameter of Table 6 fixed. Rows scale the whole lag map of Table B1 by a common factor, and are labelled by the resulting housing lag ℓ_h, so that $\ell_h = 0$ is instantaneous construction. Columns vary the landowner bargaining weight θ, with $\theta = 1$ the competitive benchmark in which the landowner captures the whole surplus. Median F_i/T_i is the median recovered fixed cost per square metre of parcel area, in real 2023 euros. $\Pr(F_i < 0)$ is the share of parcels with a negative recovered fixed cost. Median R^* is the median redevelopment threshold in euros per square metre per year. Share ripe at R_0 is the share of parcels whose threshold lies below the 2023 rent. $\varepsilon^{\text{model}}$ is the supply-curve elasticity of (22), evaluated at the 2023 rent by bumping R_t one percent holding A_t fixed. The shaded cell is the adopted calibration.

Figures

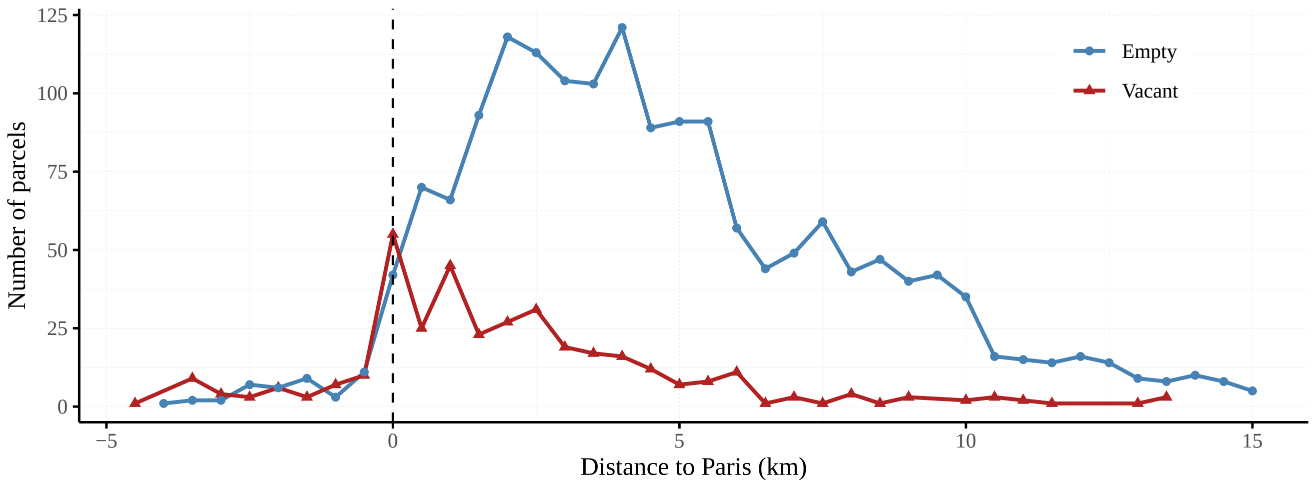

Data: Fichiers Fonciers, Sitadel. Empty parcels are parcels without physical building. Vacant parcels are parcel with no taxed units and past economic activity. Data as of 2024.

Figure C3: Empty and vacant parcels with distance to Paris

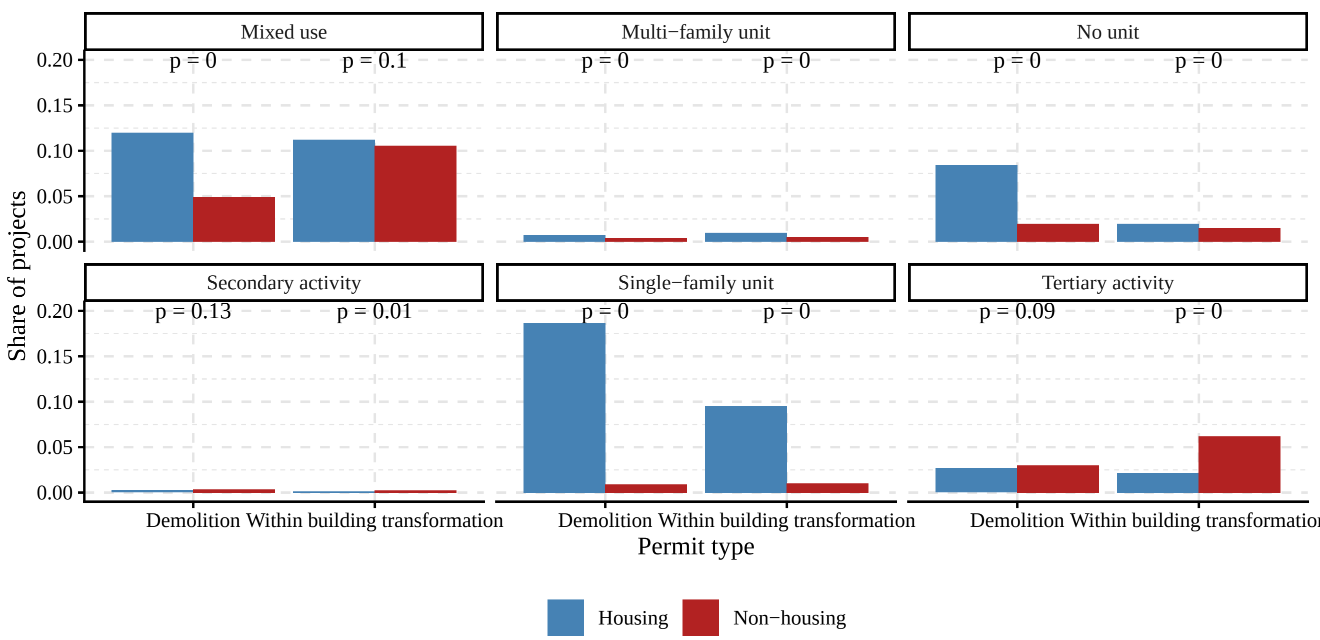

Data: Fichiers Fonciers, Sirene (2014-2024). Permits type depends on whether the permit involves a demolition or not. The use of origin is measured using the parcel destination prior to the permit. p-values of equality of means between each distribution (within-origin between-type) are printed.

Figure C4: Share of projects by type and use of origin

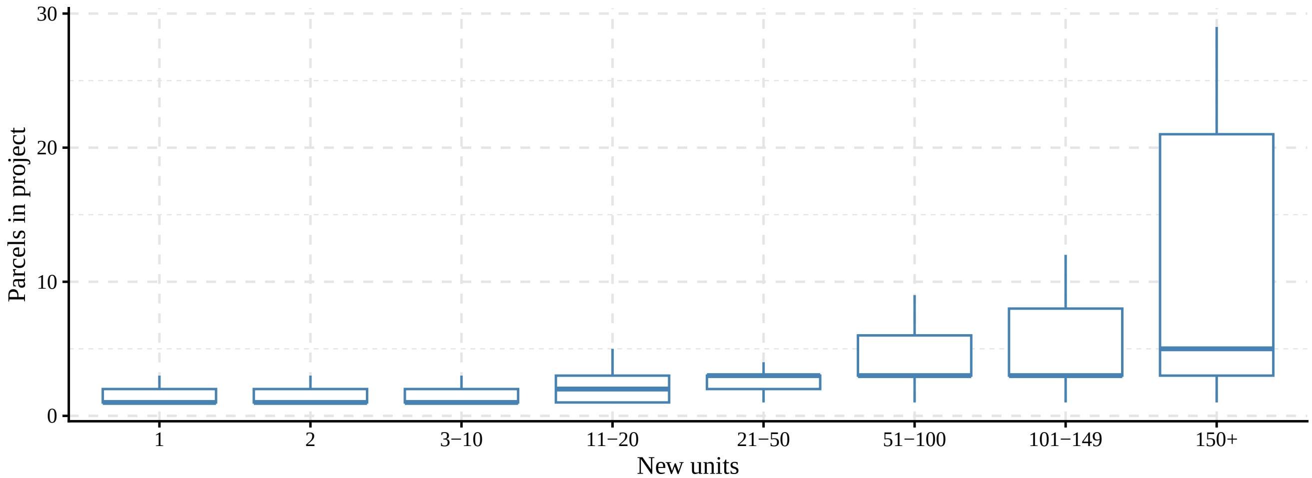

Data: Sitadel. Count of parcels involved in the project, by project size bins (measured by the number of new units). Boxplot with the 1st and the 99th percentile. The box represents 1 interquartile range and the solid horizontal line is the median count.

Figure C5: Count of parcels across project size

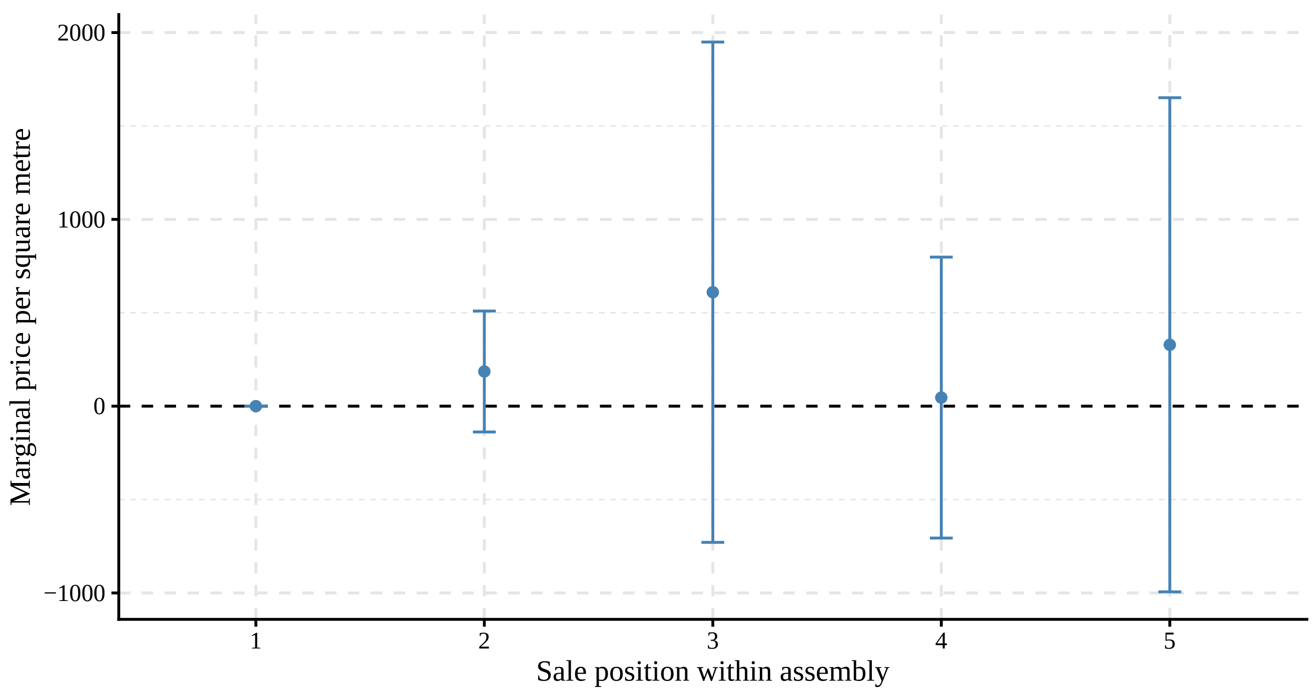

Data: Sitadel and DVF. Regression of parcel price per square meter on the position in the assembly program, and a set of controls.

Figure C6: Price premium by position in the assembly

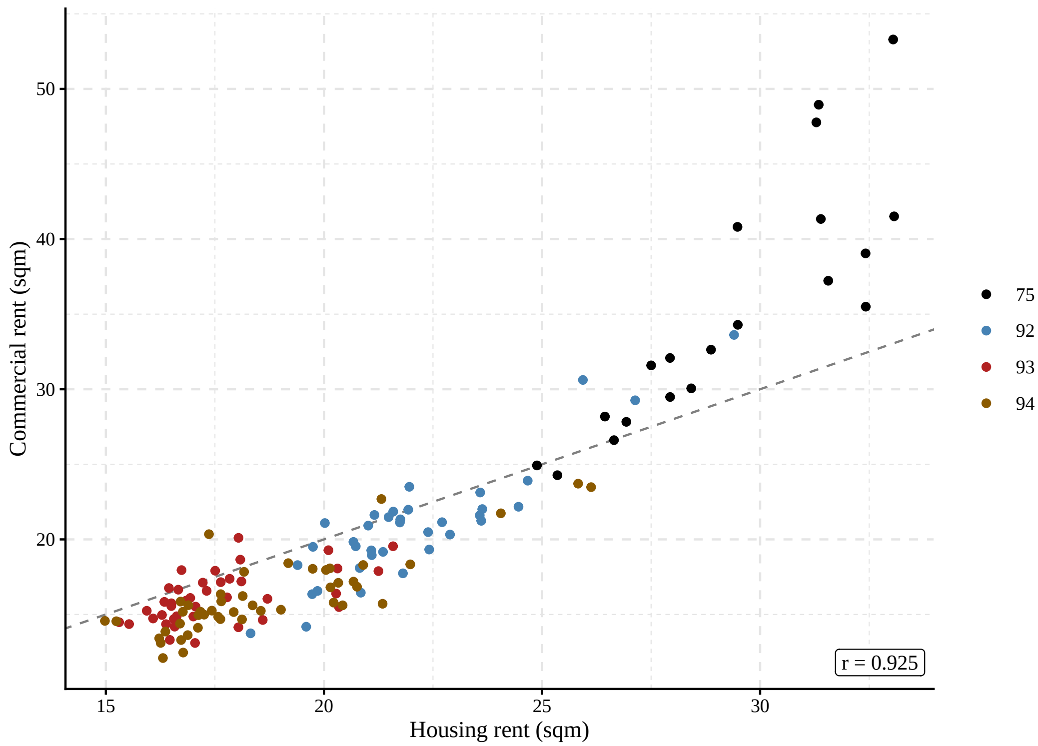

Data: Fichiers Fonciers, Observatoire des Loyers. Median housing and imputed commercial rents per square meter. Each dot is a municipality. Dashed gray line the 45° degree line.

Figure C7: Correlation between housing and commercial rents (2023)

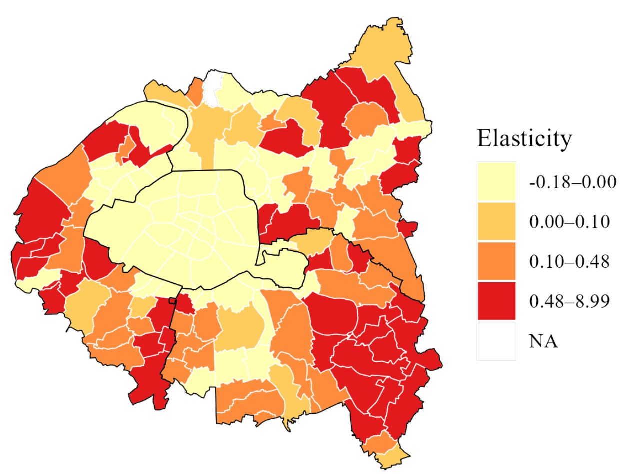

Data: Model estimation. Housing supply elasticity evaluated for a 1% increase of 2023 aggregate rent level.

Figure C8: Municipal-level housing supply elasticities

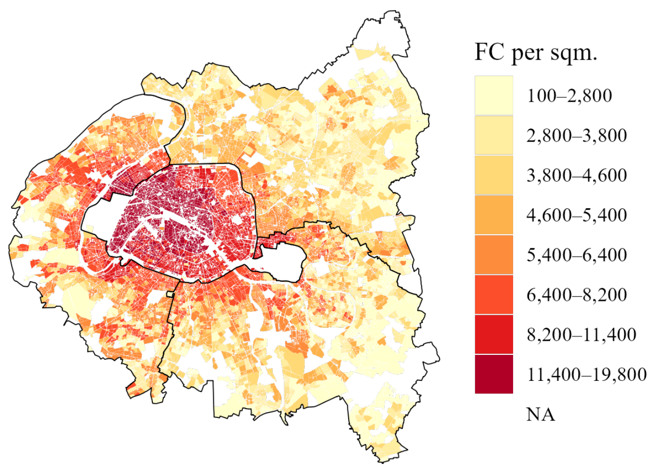

Data: Model estimation. Mean recovered fixed cost per square metre of parcel area, in real 2023 euros, averaged within cadastral section. Sections are shaded by octile of the distribution. Vacant parcels and parcels with a non-positive residual are excluded.

Figure C9: Fixed cost per square metre, by cadastral section