Commuting, Air Quality and Welfare

Pascale Champalaune * Pol Cosentino †

This version: 11th March, 2026

Abstract

We study the welfare effects of public transport infrastructure investments in the Paris region, highlighting the role of local air quality improvements typically omitted by standard quantitative urban models (QUMs). We first provide reduced-form evidence that tramway expansions improve local air quality. We then develop a QUM with endogenous worker and firm location choice, transport mode choice, and local air pollution affecting amenities and productivity. Applying the model to the *Grand Paris Express* metro project, we find welfare gains of around 1.5%. We show that omitting the air quality channel leads to a severe underestimation of welfare gains and spatial skill sorting, given the substantial heterogeneity in pollution exposure and its valuation across skill groups.

JEL Classification: D62, R13, R41, R42, Q51, Q52, Q53

Keywords: Air pollution, Public transport, Quantitative urban model, Welfare, Commuting, Infrastructure investment, Spatial sorting

*TSE, INRAE. Champalaune acknowledges financial support from a French government subsidy managed by the Agence Nationale de la Recherche under the framework of the "Investissements d'avenir" programme reference ANR-17-EURE-001, and from Cepremap.

†LEDa, Dauphine-PSL

Access to some confidential data, that part of this work is based on, has been made possible within a secure environment offered by CASD – Centre d'accès sécurisé aux données (Ref. 10.34724/CASD). We thank Clément Bosquet, Dominique Bureau, Olivier Chanel, Pierre-Philippe Combes, Gabrielle Fack, Laurent Gobillon, Andrea Herrera Bórquez, Hélène Ollivier and Sefi Roth, conference participants in Bordeaux (EEA) and University Lyon 2 (AFET), workshop participants in Paris (CREST) and Barcelona (IEB) and the members of the scientific committee of SGP (Société des Grands Projets) for their feedback. All remaining errors are ours.

1 Introduction

Local air pollution represents a major urban negative externality. This congestion force affects not only residential choices, *via* a capitalisation into housing prices (Chay and Greenstone, 2005), but also worker productivity (Graff Zivin and Neidell, 2012; Chang et al., 2019). Within cities, air pollution is largely dependent on local road traffic (Karagulian et al., 2015; Tessum et al., 2022).[1] Therefore, recent policies aimed at improving urban air quality have focused on reducing the number of polluting vehicles in city centres,[2] sometimes with the unintended effect of displacing pollution to other neighbourhoods (Bou Sleiman, 2023). We argue that public transport (PT) infrastructure represents a strong lever for reducing local air pollution, by incentivising commuters to shift away from private car use. Yet existing evaluations of PT infrastructure have omitted the formal quantification of air quality effects (Severen, 2026), while the comparatively scant literature linking PT to air quality has not yet considered general equilibrium effects at the city scale (Gendron-Carrier et al., 2022).

In this context, this paper investigates the effects of large-scale PT infrastructure on air quality and welfare, through commuting behaviour and firm and worker location choices, focusing on the Paris region. We document that major roads significantly affect local air pollution up to 3.5 km away, meaning that car traffic reductions could benefit about 80% of the city's population. We provide reduced-form evidence that the launch of new tramway lines, between 2008 and 2018, significantly increased the share of PT users in neighbourhoods in which a new tramway stop opened, as opposed to those in which it had not yet opened. This was accompanied by a substantial improvement in local air quality, and a rise in housing prices in these same neighbourhoods.

Motivated by these facts, we develop and calibrate a quantitative urban model (QUM) *à la* Ahlfeldt et al. (2015) to assess general equilibrium effects. The model incorporates worker heterogeneity (distinguishing between high- and low-skill workers, both with non-homothetic preferences) and transport mode choice between (polluting) cars and (non-polluting) PT. Locations in this city differ in terms of amenities, productivity, housing supply, and commuting costs. Hence, workers choose where to live and where to work, based on amenities, wages, housing costs and commuting costs. The model also allows for migration to the city, as large transport infrastructure was shown to attract new inhabitants (Heblich et al., 2020; Morten and Oliveira, 2024; Tsivanidis, 2026).

A key ingredient of our model relative to standard QUM is the incorporation of a tractable air pollution-traffic module. We model neighbourhood-level air pollution as generated by the number of workers commuting through the neighbourhood by car. Accordingly, air pollution is produced continuously along the trajectories that workers follow from their residence to their workplace. In practice, for each origin-destination pair, we thus construct a list of neighbourhoods that are crossed by the least-cost path by car, and collapse this by crossed neighbourhood to get an estimation of

[1] Appendix Figure A.1 overlays the current Paris road network onto 2018 concentration in fine particulate matter ($PM_{2.5}$), our pollutant of interest. In addition to the distance gradient in air pollution from the city centre to the outer suburbs, higher air pollution around major roads is clearly visible.

[2] See, for instance, work on German low-emission zones by Wolff (2014); Gehrsitz (2017); Pestel and Wozny (2021) and Gruhl et al. (2025), or on the London congestion charge by Leape (2006); Green et al. (2020) and Herzog (2024).

local car traffic, which we then match to local air pollution data. We find that when $1,000$ extra car commuters cross a neighbourhood, this increases annual average air pollution by 1.3%, which is similar to previous findings in other European cities (like Milan and London in Gibson and Carnovale, 2015; Green et al., 2020). Local air pollution feeds back into the neighbourhood amenity level and into local firm productivity, thus shaping both the residential choices of workers and the location decisions of firms. Importantly, because amenities are skill-dependent, our model is able to capture the spatial sorting of workers by skill driven by local air pollution, an empirically persistent phenomenon (Heblich et al., 2021).

Using neighbourhood-level information on housing prices and commuting behaviour by skill and transport mode, we estimate several structural parameters in our framework. First, our estimates of the (semi-)elasticity of commuting flows with respect to commuting time are consistent with existing literature (e.g., Ahlfeldt et al., 2015; Brinkman and Lin, 2024; Koster, 2024), but bring their own contribution by being disaggregated down to the skill×transport-mode level, which was not considered in earlier studies. We find that car commuters are much more sensitive to increases in commuting time than PT users, which we explain by the fact that driving does not allow one to legally perform other types of activities, while time spent in PT can. Second, we rationalise the within-city spatial equilibrium in 2008 and 2018 to recover location productivity and amenities and to provide estimates of the sensitivity of amenities to air pollution based on within-city variation, and by skill level. In order to tackle endogeneity concerns related to reverse causality and sorting, we exploit the fact that particulates emitted a few kilometres away may influence local air pollution, as they are carried by the wind. Given that in Paris, prevailing winds blow from the West to the East, we instrument the change in local air pollution with the change in the number of car commuters crossing neighbourhoods to the West of the local neighbourhood. Distant upwind traffic does not affect local amenities or productivity through noise, but still contributes to air pollution, as particulates disperse over distances of several kilometres (Deryugina et al., 2019). We find that local air pollution is perceived as a stronger disamenity by higher-skilled workers, with an estimated elasticity that is three times larger than that of lower-skilled workers. We also find strong negative effects on local productivity, consistent with recent literature (Chang et al., 2019; Holub and Thies, 2023; Champalaune, 2025).

We take the spatial equilibrium observed in the Paris region in 2018 as the baseline equilibrium to conduct counterfactual analyses.[3] First, we provide a prospective evaluation of the opening of the *Grand Paris Express* (GPE) metro lines, the largest public transport infrastructure project in

[3] Commuting behaviour likely has changed since 2018 due to Covid-19, but it still represents the bulk of car traffic and overall trips made by households. According to a survey run by Institut Paris Région in 2023, the urban planning agency of the Paris region, more than 50% of daily trips taken by 25-65 year-olds are work commutes. This figure accounts for all 25-65 year-olds at the denominator, regardless of employment status. Moreover, this number pertains to the number of trips, not the distance travelled. Daily trips that are not related to work are much shorter, such that home-work commuting should represent an even larger share of distance or time spent travelling outside the home. This is what we aim to capture here.

Europe, which will effectively double the length of the Paris metro network by adding 200 km of new lines. Our counterfactual analysis of GPE suggests that the welfare gains to be expected are around 1.5%. We find that the expected gains in travel time generate a substantial switch from using the car to using PT, in particular close to new metro stations, with moderate spillovers by distance. These changes trigger a strong improvement in air quality, since car traffic decreases locally. We find evidence of spillovers in terms of air quality improvements throughout the entire metropolitan area, and a significant decrease in CO_2^e emissions, driven by the avoidance of approximately 2.7 million kilometres of car travel per day. We estimate that omitting the air quality channel leads to a severe underestimation of spatial skill sorting and of welfare gains, with the latter biased downwards by about 40% for both skill groups. Despite the much stronger valuation of air quality by high-skilled workers, low-skilled workers end up being similarly affected owing to their greater initial exposure to road pollution. Hence, we show that it is crucial to account not only for induced changes in car traffic and air pollution when evaluating the welfare effects of PT infrastructure, but also for how air pollution exposure and its valuation differ across skill groups.

Second, our estimation of the effects of the creation and extension of tramway lines suggests significantly lower welfare gains than the GPE metro, around 0.6%. They are however evenly distributed across skill groups, unlike the GPE project, where high-skilled workers gain twice as much as low-skilled workers. This is largely due to the location of the new infrastructure, in outer suburbs, where a large fraction of low-skilled workers are located and to the fact that as the gains in travel time are more limited, fewer high-skilled workers are attracted to neighbourhoods with new tram stops. This limits housing price increases and the eviction of lower-skilled workers, who benefit from the new infrastructure.

Third, we find that combining both the GPE metro project and the new tramway lines into a single new network expansion entails 14% larger welfare gains and air quality improvements than if only the GPE metro were built. The welfare gains are also more evenly distributed across skill levels than under GPE only. All in all, these results show the importance of accounting for changes in local air pollution in the evaluation of PT infrastructure investments, due to their impact on both welfare and skill sorting.

Finally, as a complementary exercise, we simulate a ban on private cars from the entire Paris municipality, as in Ubeda (2021). While the latter performs a similar exercise without endogenous air quality responses and finds welfare losses for all workers, we find that accounting for the resulting improvement in air quality reverses this conclusion. The policy generates positive welfare effects despite the inflow of new residents and the associated rise in housing prices. The air quality gains are strongest in the city centre, where cars are banned, but also spread to the rest of the city.

This paper connects to the literature in urban and environmental economics. First, we add to the small literature on the air quality effects of PT investments. Initial work focuses on large infrastructure in developing contexts: Chen and Whalley (2012) examine the effect of a new subway in Taipei, Goel and Gupta (2017) the expansion of the metro system in Delhi, Li et al. (2019) Beijing, and Xie et al. (2024) China as a whole. The first find a negative effect on carbon monoxide concentration. Goel and Gupta (2017) find no effect on fine particulate matter ($PM_{2.5}$), our pollutant of interest, which they attribute to data limitations, while Li et al. (2019) do. Gendron-Carrier et al. (2022) take a broader perspective by studying 23 global cities that newly get a metro system between 2000 and 2017, finding effects limited to the most polluted cities. We show that both a large metro expansion like GPE and smaller-scale investments (tramway lines) improve local air quality in a city where PT is already well-developed and its take-up is already high at 45%. This has the important implication that PT infrastructure can serve as a policy lever for locally improving air quality, particularly where the acceptability of congestion pricing or low-emission zones is low (Morton et al., 2021; Frondel et al., 2025). In the French context, earlier work on PT infrastructure has investigated employment subcentre formation (Garcia-López et al., 2017), city-level employment growth (Mayer and Trevien, 2017), and spatial sorting (Ubeda, 2021). We provide the first evidence of its effect on air pollution in the Paris region, including general equilibrium effects.

Second, we contribute to the growing literature on within-city quantitative spatial models. Since the seminal paper of Ahlfeldt et al. (2015), a large body of work has used QUMs to study the impact of new road infrastructure (Fretz et al., 2022; Brinkman and Lin, 2024; Bagagli, 2025), neighbourhood effects (Redding and Sturm, 2024; Weiwu, 2024), congestion (Allen and Arkolakis, 2022; Herzog, 2024), consumption access (Miyauchi et al., 2021; Cosentino, 2025), and, more closely related to our work, upgrades in urban PT systems (Heblich et al., 2020; Severen, 2023; Warnes, 2024; Tsivanidis, 2026), and natural amenity preservation (Koster, 2024). Our paper contributes to this literature by newly endogenising local air pollution. We do so in a tractable way, by incorporating both a polluting and a non-polluting transport mode, and by modelling the negative externality that air pollution exercises on local amenities and productivity. This allows us to show that evaluations of PT infrastructure that omit air quality effects substantially underestimate welfare gains, with the largest bias for low-skilled workers.[4]

The remainder of this paper is organised as follows. Section 2 begins by laying out information about (our treatment of) the data sources. In Section 3, we first exploit the combination of these data to provide descriptive evidence on skill-biased sorting patterns by proximity to both PT and road infrastructure. Section 4 then provides reduced-form evidence of the impact of recent investments in PT infrastructure on transport mode choice and local air quality. Then, we lay out our structural model in Section 5. We take it to the data in Section 6, by estimating original model parameters, and calibrating the most common ones. Finally, we perform and discuss counterfactual policy analyses in Section 7. Section 8 concludes.

[4] Bahlali and Petit (2024) also model pollution emissions within a spatial equilibrium framework, using a dispersion model to convert emissions into concentrations. Their model, however, abstracts from agglomeration forces and worker heterogeneity, and remains theoretical rather than taken to data.

2 Data

We start this section by motivating the definition of our study area (*city definition*) and the observation levels used in the reduced-form section and in the structural model, respectively. We then turn to the description of our data sources.

2.1 Delimiting the city and its neighbourhoods

City definition Our study area is the Paris region, which we define as the Paris urban core, as delineated by the French Statistical Institute (INSEE). This city definition is based on two criteria: continuity in the built residential environment, and commuting flows within the area.[5] This means that we focus on the central municipality, Paris, and its inner suburbs, totalling 10 million inhabitants forming an "urban core", and exclude outer suburbs from the analysis (2.5 million inhabitants).[6] Relating this definition to the one used in the United States, this means that we focus on the equivalent of "central counties" of the Paris Metropolitan Statistical Area (MSA). Let us note that more than 80% of the population of the Paris MSA live in our zone of interest. Finally, we use the urban unit delineation from 2010, which was established based on data from the 2008 Census, the latter being the first year of our study period.

Observation level Reduced-form evidence uses IRIS (Îlots de Regroupement pour l'Information Statistique) as the unit of observation, the finest neighbourhood definition available in French censuses and roughly equivalent to US census block groups. IRIS are delineated by INSEE to form homogeneous blocks of approximately 2,000 inhabitants, distinguishing between residential, commercial, and other land uses (such as parks or harbours). All municipalities above 10,000 inhabitants and most municipalities between 5,000 and 10,000 inhabitants are divided into IRIS. The sample comprises 4,284 IRIS in the Paris urban unit, with an average population of 2,531 inhabitants (in 2008) over 0.90 km^2 (90 hectares, of which 62 developable), yielding a density of approximately 12,000 inhabitants per km^2.

While IRIS allow for very local reduced-form results, we adopt a slightly coarser partition for the structural analysis. As detailed in Section 5, our model allows workers to cross intermediate neighbourhoods when commuting by car, which requires expanding the origin-destination matrix to include all neighbourhoods traversed between each origin-destination pair. With 4,284 IRIS, this generates a 4,284^3 matrix (over 78 billion cells), which is computationally intensive. More importantly, using such a fine-grained definition yields extremely sparse matrices: only 8% of households are surveyed annually, so many low-volume flows would be unobserved in the data, and replaced

[5] More specifically, an urban core (called "urban unit" by INSEE) is a group of contiguous municipalities, without enclaves, made up of an urban centre with more than 10,000 jobs (here, the municipality of Paris), and other municipalities (inner suburbs, or *petite couronne*), where at least 40% of the employed resident population works in the urban centre or in municipalities in its vicinity.

[6] While the total inflow of workers living in this area to the urban core may be non-negligible, each municipality-to-municipality flow is very small. This would render very sparse matrices of commuting flows, and even sparser vectors of car commuters crossing each municipality, with a lot of missing values, hence our focus on the urban core.

by zeroes. More generally, Dingel and Tintelnot (2026) demonstrate that high-dimensional settings, with many small neighbourhoods, pose challenges for calibration based on observed shares to recover fundamentals. Partly re-aggregating neighbourhoods provides a solution to circumvent these issues. We therefore group IRIS using the first 7 digits of their 9-digit identifier, yielding 735 neighbourhoods in our final sample. These are approximately twice as large (1.82 km^2) and twice as populated (5,061 inhabitants on average in 2008) as the IRIS neighbourhood definition, with a density of about 10,000 inhabitants/km^2. This remains highly granular compared to related work. Appendix Figure A.2 provides a map overlaying the two neighbourhood definitions, with IRIS in purple and IRIS7 in black. The absence of purple lines within some black neighbourhoods means that in these cases, the two definitions match.

2.2 Data sources

All sources to raw data and links to access them are provided in Appendix Table C.1, and we describe here the treatment applied to them.

Commuting flows We use the restricted-access individual-level information of the supplementary files of the French census (*Recensement de la Population*) to build commuting flows by skill and by mode of transport. The census allows us to compute these flows from the IRIS of residence to the municipality of work. In order not to create imbalances in our commuting flow matrices, we proceed to break the flow down to the neighbourhood of work. To do so, we complement the census data by merging it with employer-employee data (DADS), which we aggregate to the firm establishment level by counting the number of employees by skill level. Then, in order to locate the establishment, we merge this information with the SIRENE registry, which contains the precise geographic location of the establishment. This allows us to break the workplace (destination) down to the IRIS level. We apply the same factor to each flow regardless of the mode of transport, since this information is not contained in DADS, but we are able to make the distinction between skill levels. Hence, we make the assumption that, within a workplace municipality, the distribution of transport modes is homogenous across IRIS, for a given skill level. This seems reasonable, as jobs that are similar in skill are concentrated in a given (cluster of) neighbourhood(s), and therefore have similar road and PT accessibility. At the end of this procedure, we end up with all commuting flows in 2008 and 2018, from IRIS of residence to IRIS of workplace, broken down by skill level and transport mode. We provide more details on the precise method and the specificities of French census data, including the fact that it follows a rolling collection system, in Appendix C.2.

Road transport network In order to be able to compute travel times between neighbourhoods, we create a road transport network. We make use of the "road section" files provided by the National Geographic Institute (IGN) in BD-TOPO. While they provide reliable and precise information on the location of road segments and their connections (nodes and edges), we do not directly use the mean speed variable, which has a lot of missing or unrealistic values. Instead, we build on the elements used in the Metric tool developed by INSEE. Metric assigns one value of mean speed in peak hours to each road or street category. Each category is itself based on a typology used by Open Street Map (OSM). As such, we proceed to mimic the types of segment used by OSM using the variables available in BD-TOPO (namely, its nature, level of importance, type of administrator, whether it is located in an urban area, and whether it is located in the municipality of Paris). We provide the full categorisation, with the corresponding assigned speeds, in Appendix Table B.1. We also use this dataset to compute distance to a major road, major roads being defined as motorways and primary roads from the table.

Public transport network We build a PT network based on several datasets. For the baseline quantitative model, reduced-form evidence and the tram counterfactual, we use information from IDFM, the Paris transport authority. We retrieve the location and opening dates of all tram stops using open data from IDFM as well. For the *Grand Paris Express* counterfactual, we use information from both IDFM and *Société du Grand Paris* (SGP), in charge of implementing the project.

Housing prices Housing price data availability is central to the analysis, as housing prices are necessary to derive measures of the amenity and productivity fundamentals of our model. An exhaustive geocoded register of all transactions, called DVF, is available from 2010. We take all transactions from 2010 to 2012 to approximate 2008 prices, all transactions from 2017 to 2019 to measure 2018 prices, and intersect them with IRIS contours. We measure IRIS-level housing price as a fixed effect from a regression of (log) total transaction amount on total floor area, the total lot size, and a control for the year of the transaction.

Air pollution Fine particulate matter ($PM_{2.5}$) concentration data comes from Airparif, the association in charge of monitoring air quality in the Paris region. They combine information from their monitors with a transport model to derive measures of $PM_{2.5}$ at a very fine spatial resolution of 50m$\times$50m for the whole study period, which we aggregate to the IRIS level (reduced-form) and to the neighbourhood/IRIS7 level (structural analysis).

Developable land We need data on developable land, so as to approximate a "floorspace" quantity as used in the structural model and as a control in the reduced-form section. This is especially important in the case of IRIS neighbourhoods, as they are designed to be homogenous in terms of land use, and some might seem like they could offer significant housing supply, when they are actually covered by forests. We use CORINE Land Cover, and define as developable land all land that is already artificial, thus most likely impervious and/or already built-up, and agricultural areas. This definition deducts all land covered by protected forests, rivers or wetland from the total surface area of the neighbourhood.

Local wages The estimation of some of the model's deep parameters relies on the use of local data on wages. To this end, we use again employer-employee data (DADS), which provides information on individual wages at the municipality level. Our measure of log wages is a fixed-effect at the municipality $\times$ skill level, accounting for (squared) age differences.

Conversion to CO_2^e emissions In order to translate the change in the number of kilometres travelled to a change in CO_2^e emissions in Section 7.6, we first use data from the Ministry for the Environment on the characteristics of the car fleet of the Paris region for 2021 (the year closest to our baseline) to get the share of vehicles by fuel type. Then, we multiply the change in kilometres driven by car by emission factors by fuel type by kilometre provided specifically for the French context by the French Environment Agency (ADEME), called “Base Empreinte”, to get an estimate of the change in CO_2^e emissions.

3 Background and descriptives

3.1 Public transport infrastructure in Paris

The Paris region has an extensive PT network, and both existing and planned infrastructure reflect a sustained commitment by local authorities to the development of a dense system relying on a range of technologies. Appendix Figure A.3 shows a map of the network at the date of 31st of December, 2024. The total number of metro stations (purple on the map) is a staggering 321, spread over 16 lines, all located in the centre or the very close suburbs. There are 416 suburban train stations (in pink), placed on 14 lines which extend back to the outer suburbs and even outside of the administrative region of Île-de-France. Finally, following the inner ring road and in the suburbs, there are 225 tramway stops (in orange), which allow for connections with both the city centre's metro system and the suburban train network.

Indeed, while Paris was one of the first cities to open a metro line, in 1900, it also built a large-scale suburban train network, RER, for *Réseau Express Métropolitain* in the 1960s. Previous work has shown that the opening of RER had a positive effect on employment in connected municipalities (Mayer and Trevien, 2017), thus favouring the formation of employment subcentres (Garcia-López et al., 2017). No investment of this scale has been implemented again, apart from the awaited *Grand Paris Express* (GPE). This project is a part-circular, part-radial addition to the existing (radial) metro network that is still under construction, and is not expected to be fully launched before the mid-2030s. We provide further information on GPE when we reach our prospective analysis of its effects in Section 6.

In the meantime, between the 1990s and the 2030s, the Paris region has rather focused on developing tramway infrastructure, a type of equipment that is both less heavy, hence less costly, but also of lower capacity. In this sense, it provides a potential complement to larger-scale projects like suburban trains, by connecting these to less dense neighbourhoods in which trains would not be cost-effective. This is exactly what happened in the case of the recently opened line T12, which is included in our sample below: the project was originated in the 2000s as an extension of an existing RER line, which would link Versailles, in the west, to Melun, in the south-east, but it was

Figure 1: Public transport infrastructure, sorting and car use

(a) Sorting by distance to PT

(b) Car use by distance to PT

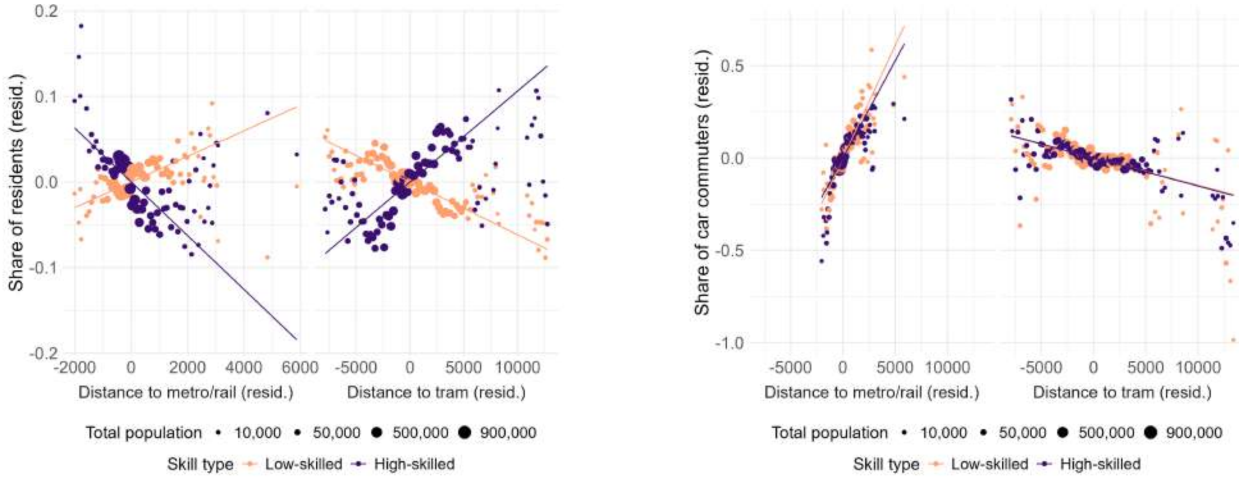

Notes: Panel 1 plots the share of residents by skill level, as a function of distance to the nearest metro/suburban train or tramway station, controlling for distance to the CBD. Panel 1b plots the share of resident commuters that are car users, again as a function of distance to the nearest metro/suburban train or tramway station, controlling for distance to the CBD. Low-skilled workers are in orange, and high-skilled workers are in purple. Bins are defined by 100-metre intervals of (residual) distance to the station. Point size is proportional to total population in the bin.

abandoned in 2006 because it was deemed too costly compared to the expected benefits. When it was revived in 2008, it took the form of a tramway line, which opened in 2022.

3.2 Public transport accessibility and car use

We thus begin by considering descriptive evidence on PT accessibility in the Paris region, and whether it relates to car usage. We make a distinction between access to heavy rail infrastructure (metro and suburban rail) and access to trams, as they differ in many aspects, including travel time, reliability, type of neighbourhoods served (in terms of density, but also skill composition) and date of construction (as described in the previous subsection). Figure 1 suggests that residential sorting associated with PT accessibility (Panel 1a) is disconnected from car reliance (Panel 1b) within the city. First, Panel (1a) examines the relationship between neighbourhood skill composition and proximity to PT, controlling for distance to the CBD. The results are strikingly asymmetric across the two sub-modes. Proximity to metro and suburban rail stations is strongly associated with higher shares of high-skilled residents: as distance to heavy rail increases, the share of high-skilled residents declines very sharply, while the share of low-skilled residents rises. This pattern is consistent with a sorting equilibrium in which high-skilled households outbid low-skilled households for housing near high-quality infrastructure, even controlling for centrality. The relationship is exactly reversed in the case of trams: as distance to the nearest station rises, the share of high-skilled (resp., low-skill) workers strongly increases (resp., decreases). Doing the same exercise with buses in Appendix Figure A.4 renders a flat relationship, regardless of skill. Related to this, again conditional on

Figure 2: Roads, car traffic and $PM_{2.5}$ exposure

(a) Major road proximity by skill (b) Major road proximity and $PM_{2.5}$

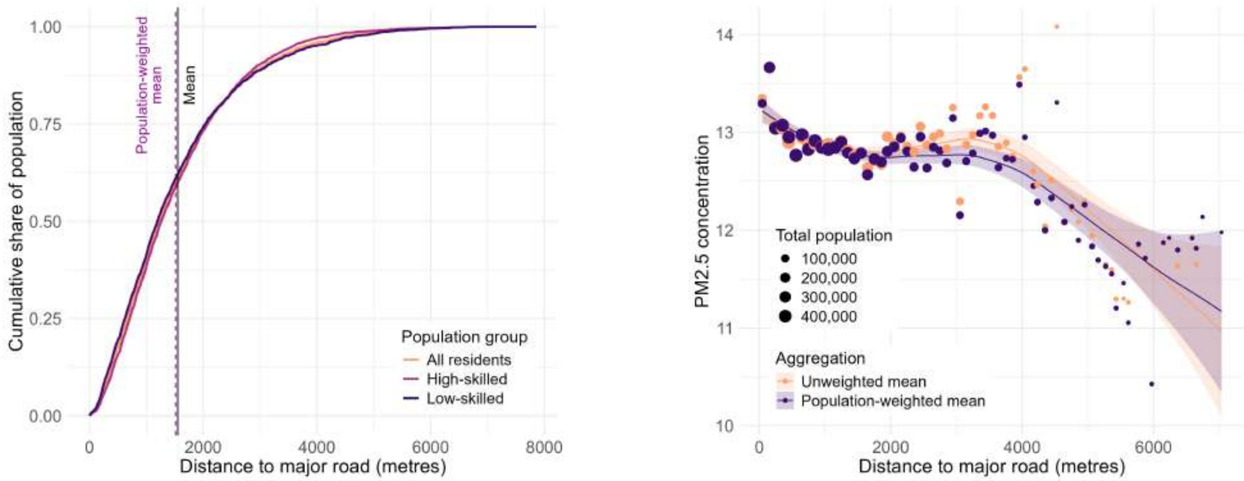

Notes: Panel (2a) shows cumulative neighbourhood population, as a function of distance to a major road (motorway, trunk road or backbone boulevard). The orange line corresponds to entire population, the purple line corresponds to low-skilled workers, and the pink line corresponds to high-skilled workers. The purple dashed vertical line shows the population-weighted mean distance to a major road, and the black plain line shows the unweighted mean distance to a major road. Panel (2b) shows $PM_{2.5}$ concentration against distance to a major road, with orange points showing raw $PM_{2.5}$ concentration, and purple points population-weighted $PM_{2.5}$ concentration. Observations are binned by 100-metre intervals. Point size is proportional to total population in the bin.

distance to the CBD, car usage appears to increase very sharply with distance to a metro or RER station (see Panel 1b), while there is a slight negative correlation in the case of distance to a tram station. This suggests that the latter does not provide as strong a substitute for the individual car as heavy rail does, although in neighbourhoods closest to a tram station, the share of car users remains far below that in neighbourhoods far from metro and RER stations.

All the graphs of Figure 1 thus point to a hierarchy between PT types, whereby heavy rail is much more highly valued than (slower) trams that often serve already disadvantaged areas, and buses that are more spatially spread out and used for shorter trips. This hierarchy is reflected in skill sorting patterns, which are strongest for metro, weaker for trams, and almost absent for buses. They also suggest that low-skilled workers face a disadvantage, as they tend to live in neighbourhoods that are further away from the types of PT infrastructure (metro and suburban trains) that are most strongly associated with lower car use. As a result, lower-skilled workers likely incur higher private transport costs and rely more on cars, not necessarily due to different preferences, but because of housing market constraints.

3.3 Road traffic and exposure to air pollution

Building on the previous section, we evaluate whether disparities in PT proximity and car usage translate into differences in road traffic and air pollution exposure in Figure 2. Panel (2a) shows the cumulative population distribution by distance to a major road.[7] It first exhibits an important fact in our context: people live close to major roads. Indeed, 75% of individuals in the Paris region live less than 2 km away from a very heavy-traffic, major road, and the mean distance to such roads is 1.5 km. Panel (2b) plots local air pollution against distance to major roads, in 100-metre bins, where each bin's size is proportional to population. It shows that such a distance of 2 km remains small considering the induced pollution burden, since average $PM_{2.5}$ concentration plateaus at 13 $\mu g/m^3$ on average up to 4 km away from a major road, as particulates can travel.[8] Finally, Panel 2 suggests that skill-based differences are less pronounced than in the previous subsection, though low-skill workers are still somewhat more concentrated near major roads (up to 3 km). Additional evidence in Appendix Figure A.5 suggests that there is no correlation between the share of high-skill workers and total traffic in a neighbourhood, while neighbourhoods with higher low-skill shares experience modestly higher traffic. Taken together with the evidence presented in Section 3.2, this suggests that by inducing skill-based spatial sorting, PT infrastructure, in particular heavy rail, acts as one of the vectors of disparities in exposure to car traffic and local air pollution.

4 Reduced-form evidence: New tramway lines

Having documented the relationship between PT access, car traffic, and pollution exposure, we turn to causal reduced-form evidence using the staggered opening of new tramway lines. We show evidence that tramway developments during the 2010s increased PT use and improved local air quality in neighbourhoods newly connected to the network.

4.1 Empirical strategy

We investigate the effects of tramway lines that were entirely or partly launched between 2010 and 2017, that is, both creations and extensions of existing lines. Due to limitations in the quality of our data, which we exposed above in Section 2.2, we use a long-difference specification, in which we compare outcomes in 2018 to those in 2008. We estimate the effect around new tram stops: do residents use PT more after the opening, and does local air pollution react?

Our approach thus consists in estimating the following equation:

$$\Delta_{2018-2008}\ Y_i = \alpha_1 + \alpha_2(\text{new tram stop})_i + \phi_{f(i)} + \varepsilon_i \tag{1}$$

where $\Delta_{2018-2008} Y_i$ is the change, between 2008 and 2018, in a) $PM_{2.5}$ concentration, b) the share of residents of neighbourhood i that commute using PT, or c) the share of residents of i that commute by car. The treatment variable $(\text{new tram stop})_i$ takes the value 1 if a new tram stop opens near neighbourhood i between 2010 and 2017, and the value 0 if a new tram stop opens in the vicinity of neighbourhood i later, during the 2019-2024 period. We use these "not-yet-treated" neighbour-

[7] This includes the infamous ring road (*boulevard périphérique*), but also motorways and boulevards.

[8] The WHO air quality guidelines state that annual average $PM_{2.5}$ concentration should not exceed 5 $\mu g/m^3$. The current EU limit is set at 25 $\mu g/m^3$, and it is set to decrease to 10 $\mu g/m^3$ in 2030.

hoods as they are the ones that are most comparable to our neighbourhoods of interest. In some specifications, we include fare zone fixed effects $\phi_{f(i)}$. As abovementioned, the fare zoning system was abandoned in 2015, making this a potential confounder of the effect we aim to quantify. As shown in Table B.2 in Appendix, the subscription fee to travel from the suburbs to the city centre increased quite sharply with distance before September 2015. For tramways launched between 2010 and August 2015 included, the incentive to become a new PT user was smaller if one lived in a fare zone $f(i)$ that was further away, as the cost was higher. Then, as the zoning was removed, the monetary cost of taking PT became uniform.[9] The fixed effects capture this phenomenon.

Control group: Why the not-yet-treated At first glance, already-treated neighbourhoods (in purple on the map in Figure 3) are good candidates. But some of them are treated as late as 2007, the year right before our pre-treatment year, and they would themselves be following a trend as a reaction to their recent treatment. The option of building "doughnuts" around treated neighbourhoods to select those further away as controls is not convincing either, due to several potential sources of contamination. First, residents of these doughnut neighbourhoods might be living too far from the tram station to reach it by foot, but they might still use the bus to reach it. Using such a control group would lead to an underestimation of the true effect. Second, even absent such an effect, if the road alternative to using the tram goes through an adjacent, control neighbourhood instead of a treated neighbourhood, then having commuters switch from using their car to using the tram would imply a decrease in air pollution in the control neighbourhood, again underestimating the effect. Third, $PM_{2.5}$ in one neighbourhood is correlated to that in adjoining neighbourhoods, and as we can see on the map in Figure 3, neighbourhoods located around black (treated) ones are particularly small. A decrease in local pollution in treated neighbourhoods would likely spill over these control IRIS, and thus render underestimated estimates. Finally, if some road lanes are removed to make room for the new tram tracks, part of the car traffic may be redirected to adjacent neighbourhoods, this time posing a risk for an overestimation of the effects.

As such, we opt for using only not-yet-treated units as a control group. These are shown in orange in the map in Figure 3. Most are located in the suburbs, rather than very close to the city centre like our (black) treated neighbourhoods. One advantage of this is that they are less likely to suffer from contamination issues, which we discussed above in the case of "doughnuts". However, this might make these IRIS less comparable to the treated ones. We comment on balancing tests below, after describing our set of treated and control IRIS, and control for distance to the CBD in our $PM_{2.5}$ regressions. Unlike in never-treated IRIS, tram lines and stops end up being constructed in not-yet-treated IRIS, which makes them the more robust control group: as shown in balancing tests below, they do have similar population and resident characteristics in terms of skill level. Their future exposure to the treatment also minimises concerns about both unobserved heterogeneity and selection bias more effectively than never-treated IRIS.

[9] The cost became uniform only for subscribers, and the price of tickets still increased step-wise as a function of distance to the city centre until January 2025. In this paper, we choose to focus on those who use PT to commute to work, and are thus frequent users, who are mostly subscribers, as opposed to ticket holders.

Figure 3: Map of IRIS by treatment status

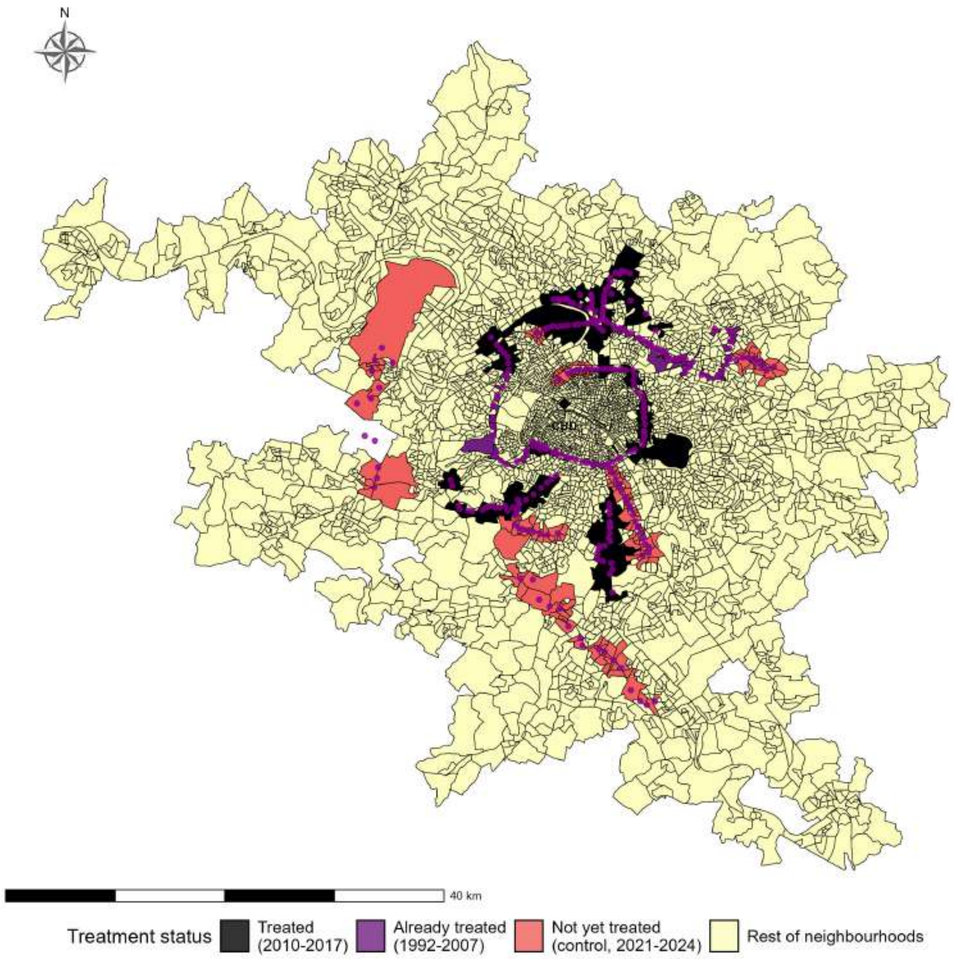

Notes: Authors' own elaboration based on IGN BD-TOPO (for IRIS contours), IDFM open data (for tramway stop information) and INSEE open data (for delineating the Paris urban unit). Purple dots show all tram stations in use as of 2024. Black neighbourhoods are part of the treatment group, and orange neighbourhoods are part of the control group.

Set of treated/control IRIS We define as treated all IRIS whose closest boundary is within at most 500 metres of a new tram stop. We choose 500 metres as the crow flies so as to approximate a reasonable walking distance from the boundary of the IRIS to the new tram stop. Not-yet-treated IRIS, our control group, are selected following the same rule. Our main sample thus comprises 284 treated IRIS and 210 not-yet-treated (control) IRIS, out of the 3,727 IRIS in the city.

Balancing tests and exogeneity Table B.4 shows pre-treatment (2008) descriptive statistics by treatment status, for IRIS that have a new tram station (the "strictly-defined" subsample) and for the main sample, taking all IRIS within 500 metres of a new tram stop. The top panel provides statistics for the main sample, while the bottom panel provides statistics for the restricted sample. For each variable, we show the p-value of the difference between the treated and control means,

Table 1: Number of tram stop (and line) openings by period and year (1992-2024)

Before 2008 *Already treated* (excluded)			Study period (2008-2018) *Treated*			After 2008 *Not-yet-treated* (control)		
Year	Tram stops	Line(s)	Year	Tram stops	Line(s)	Year	Tram stops	Line(s)
1992	12	1	2010	4	1	2019	9	2
1993	12	1	2013	43	4	2020	9	1
1997	13	1	2014	37	2	2021	21	1
2004	5	1	2015	36	2	2022	12	1
2007	28	2	2016	2	1	2023	13	1
			2017	7	1	2024	16	1

Note: Authors' own elaboration from IDFM data.

as well as the standardised mean difference (SMD). Focusing on the main sample, it appears that treated IRIS are similar to control IRIS in terms of housing prices, population, number of workers, and number of both high-skill and low-skill workers. They also have a similar initial number of car commuters, and rather similar developable land surface areas. Treated IRIS are significantly more polluted than control IRIS at baseline, though this is only a 2% difference (17.18 $\mu g/m^3$ against 16.80 $\mu g/m^3$, p-value 0.038). More importantly, the share of residents that use their car to go to work is highly significantly lower in treated IRIS than in not-yet-treated IRIS, respectively at 28% and 37%. This 9-point difference is mirrored in the difference in the share of PT users, which is respectively at 48% and 39%. This is consistent with the significant difference in distance to the CBD, with treated IRIS being on average 3.39 km closer than control IRIS.

This has consequences with regard to the unbiased estimation of the treatment effects. In treated neighbourhoods, where the share of PT users is already higher than in control neighbourhoods, it is likely that the outcome trends would have differed even in the absence of the treatment. Indeed, we expect the evolution of the share of PT users to be concave: in rather central neighbourhoods where the bus likely already runs more often, and hence where PT usage is already high, there are less individuals that may be convinced to quit using their car to using PT. There is also an incompressible share of individuals whose job demands that they use a car to commute. Hence, the share of PT users might have evolved differently between treated and not-yet-treated had the tram not been built. To mitigate such concerns, we control for the initial share of car/PT resident commuters, the initial number of workers[10], and/or the log distance to the CBD, in some regressions.

No treatment date manipulation It is reasonable to assume that planners prioritise building closer to the city centre, where the associated benefits are likely greater, though construction may be more expensive, before investing in more remote parts of the city. This could have been cause for concern with regard to the exogeneity of the timing of the treatment, had we exploited the staggered nature of the investments for identification. But, in our setting, we compare the change in outcomes

[10]Using population instead does not change the results, due to their very high correlation.

from 2008 to 2018 between neighbourhoods treated between 2010 and 2017 to neighbourhoods treated between 2019 and 2024. Pooling different treatment years allows us to mitigate the timing concerns. In particular, we do not have to worry about endogeneity in treatment timing within the 2010-2017 treatment period in this long-difference strategy. With that said, the key assumption here is that there were no non-random factors that acted such that a tram stop opened before 2018 rather than after 2018. Table 1 gives the number of tram stop openings by year and treatment cohort. Very few tram stops opened during the 2016-2020 period, around our observation date, while 43 opened at once in 2013, 37 in 2014, and 21 in 2021.

One concern could be that treatment would occur after 2018 rather than before because of other local characteristics, in particular to local resident commuting behaviours. Given the limited number of tram lines that opened during our period of interest,[11] we look at how the construction process unfolded for all seven lines that are part of our control group. Among these seven lines, four were open entirely during the control period (T9, T10, T12, T13), and three were extensions of existing lines (T1 – only one stop, T3b and T4). The T9 line, connecting Paris to Orly, was built to replace a bus route that was operating over capacity on a daily basis and already ran on a dedicated right-of-way. In 2016, it was targeted to open in 2020,[12] but due to the Covid-19 pandemic, and opened in 2021. This is innocuous, as our cut-off is in 2018. The case of line T10 is analogous, as it was always set to be launched in 2023.[13] Lines T12 and T13 are similar to each other, as they are both "tram-trains", which have the ability to operate on both the heavy rail network and the urban tramway tracks, and were built for the most part on existing unused railway tracks. Line T12 was planned to open in 2019,[14] after our cut-off date, and opened in 2022. Line T13, which, when the names of the stops were decided on in 2019, after our 2018 cut-off, was planned to open only in 2021 (which it did in 2022).[15]

Likelihood of anticipation Neighbourhoods treated early (2019–2020) may not be ideal controls. Because they are next in line for treatment, they may already be experiencing differential trends driven by anticipatory responses, such as a housing price increase, that would bias our estimates towards zero. Coincidentally, these neighbourhoods have new stations that are extensions of tramway lines which first opened during our treatment period, raising the concern that they may be affected by spillovers from the already-operating sections of the same line. We therefore exclude them from the control group. After this exclusion, our controls were not treated until at least three years after our measurement year of 2018, making potential anticipatory effects negligible. For completeness, we re-include these neighbourhoods as controls in Appendix Table B.6 and find estimates very similar to our main results.

[11] Out of the 14 existing lines, seven have had stops open only in 2010-2017, four have had stops open only in 2019-2024, and three have had stops open in both 2010-2017 and 2019-2024. We exclude the latter from the analysis.

[12] See this article by local newspaper *Le Parisien* (in French), which mentions that line T9 would be "built on exclusive lanes by 2020": Tram T9.

[13] The information is given in this May 2017 document disseminated by local authorities (in French): Tram 10. The document also mentions that road traffic was not disrupted during the construction period.

[14] See this 2014 pre-project summary by the Paris region transport authority: Tram T12.

[15] See this 2019 press release by the Paris transport authority (in French): Tram 13 Express.

4.2 Effect of new tramway stops and lines

Air quality We estimate the effect of opening a tram stop in a given neighbourhood in Table 2. For all outcomes, namely the changes in $PM_{2.5}$ concentration, the share of PT users and the share of car users, Columns (1)-(4) give results without fare zone fixed effects, while Columns (5)-(8) provide the exact same results, introducing fare zone fixed effects. These account for the fact that PT fares increasing step-wise in distance to the city centre were removed in 2015. We find a consistently positive effect of the opening of a tram station on air quality, although the point estimates differ quantitatively between specifications. Our preferred specifications, in Columns (6)-(8), suggest that compared to not-yet-treated neighbourhoods, $PM_{2.5}$ concentration decreased by approximately 0.6 $\mu g/m^3$ in treated neighbourhoods, a 3.5% decrease compared to the 2008 mean. This represents 16% of the average decrease of 3.76 $\mu g/m^3$ in $PM_{2.5}$ observed between these two dates.

Housing prices Looking at the effect on housing prices, the fourth panel of Table 2 shows a highly significant increase in housing prices in treated neighbourhoods, compared to not-yet-treated ones. The point estimate remains largely similar as we extend the set of controls, though it is slightly smaller once the effect is estimated within fare zones in Columns (5) to (8). Housing prices rose by 5 to 8% more close to a new tram station than close to not-yet-opened tram stations. If we divide this point estimate by the one we obtain for $\log(PM_{2.5})$, around .29 to .40, to approximate an elasticity, we obtain a rough estimate of 17-20%. This is slightly lower than previous estimates for total suspended particulates and PM_{10}, which are coarser than $PM_{2.5}$, in the United States (Chay and Greenstone, 2005; Bayer et al., 2009) or China (Chen et al., 2022), but larger than estimates for gases like nitrogen dioxide (Amini et al., 2022).

Transport modes In addition, treated neighbourhoods saw their share of PT users increase more than not-yet-treated, but significantly so only when controlling for the initial share of PT users. This is consistent with the fact that treated neighbourhoods were already more PT-intensive prior to the treatment, as 45% of its commuting residents used PT, as compared to 34% in not-yet-treated neighbourhoods. Such a high share of PT usage constrained the potential for its adoption. Holding the initial share of PT users constant, Columns (7) and (8) suggest that the share of PT users increased by 5 percentage points more in treated neighbourhoods than in not-yet-treated neighbourhoods. With that said, we fail to detect an effect on the change in the share of car users: the point estimate itself is negative, but the standard errors are comparatively large.

4.3 Discussion and robustness

Mediation analysis In order to disentangle the housing price effect of improved air quality from that of improved accessibility, we turn to a mediation analysis. Appendix Table B.7 shows the corresponding results. The results suggest that while housing prices increased by 8.2% in treated neighbourhoods, they would have increased by 6.8% had air quality nor varied: hence, we estimate that about 17% of the total housing price effect is attributable to the variation in $PM_{2.5}$.

Table 2: DiD results: Treated IRIS vs not-yet-treated IRIS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				$\Delta_{2018\text{-}2008}$ $PM_{2.5}$				
Treated	-0.75***	-0.73***	-0.37**	-0.39**	-0.76***	-0.75***	-0.60***	-0.61***
	(0.21)	(0.21)	(0.19)	(0.17)	(0.14)	(0.14)	(0.19)	(0.19)
(log) workers in 2008		-0.35**		-0.24		-0.15		-0.12
		(0.16)		(0.15)		(0.16)		(0.15)
(log) distance to CBD			0.87**	0.80**			0.51	0.48
			(0.36)	(0.34)			(0.33)	(0.32)
2008 mean $PM_{2.5}$	16.94	16.94	16.94	16.94	16.94	16.94	16.94	16.94
Mean outcome	-3.76	-3.76	-3.76	-3.76	-3.76	-3.76	-3.76	-3.76
R^2	0.105	0.120	0.154	0.161	0.485	0.487	0.493	0.495
				$\Delta_{2018\text{-}2008}$ log housing price				
Treated	0.08***	0.08***	0.08***	0.07***	0.05**	0.05**	0.05**	0.05**
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
(log) workers in 2008		0.05		0.05		0.02		0.04
		(0.04)		(0.04)		(0.04)		(0.04)
(log) housing price in 2008			0.02	0.005			-0.07	-0.08*
			(0.04)	(0.03)			(0.05)	(0.04)
2008 mean housing price	8.12	8.12	8.12	8.12	8.12	8.12	8.12	8.12
Mean outcome	0.033	0.033	0.033	0.033	0.033	0.033	0.033	0.033
R^2	0.048	0.055	0.049	0.055	0.124	0.125	0.138	0.144
				$\Delta_{2018\text{-}2008}$ Share resident PT commuters				
Treated	0.01	0.01	0.06**	0.06***	0.01	0.01	0.05**	0.05**
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
(log) workers in 2008		0.03		0.06*		0.03		0.05*
		(0.04)		(0.03)		(0.04)		(0.03)
2008 mean share PT	0.403	0.403	0.403	0.403	0.403	0.403	0.403	0.403
Control 2008 share PT			Yes	Yes			Yes	Yes
R^2	0.001	0.005	0.197	0.213	0.007	0.012	0.210	0.222
				$\Delta_{2018\text{-}2008}$ Share resident car commuters				
Treated	-0.004	-0.002	-0.03	-0.03	0.01	0.01	-0.02	-0.02
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
(log) workers in 2008		-0.04*		-0.02		-0.04*		-0.003
		(0.02)		(0.02)		(0.02)		(0.02)
2008 mean share car	0.286	0.286	0.286	0.286	0.286	0.286	0.286	0.286
Control 2008 share car			Yes	Yes			Yes	Yes
R^2	0.001	0.010	0.169	0.172	0.032	0.040	0.219	0.219
Observations	328	328	328	328	328	328	328	328
Fare zone FE					Yes	Yes	Yes	Yes

Notes: Standard errors clustered at the tram stop level in parentheses. Signif. codes: ***: 0.01, **: 0.05, *: 0.1. Log housing price is a neighbourhood-year fixed effect from a transaction-level regression of log housing price per square metre on floor area, lot size and a fixed effect for quarter of transaction.

No significant reaction of car usage The fact that the effect on car commuters does not exactly mirror the effect on PT commuters can be explained by the fact that as a tram stop opens, the number of PT users increases, but not all these new users are former car commuters. These new PT users may also be individuals that start using PT instead of walking or riding a motorbike because PT supply is now more convenient.[16] Indeed, tramways offer a much faster commute than buses, especially since their timeliness is not as dependent on traffic conditions. It may also be that these new PT users are also people that were not commuting before, and the new PT fostered a better access to jobs. Carantino and Haramboure (2020) do not find a significant effect of new light rail transit and tramway infrastructure on employment in deprived French "priority" neighbourhoods, but their identification strategy excludes cities with existing metro or tram networks (which rules Paris out), when we argue that much of the value of tramways lies in their connectivity to a dense existing network. On the other hand, also in the French context, Briant et al. (2015) find that place-based interventions only have a significant effect on local employment in neighbourhoods that are well-connected to the rest of the city, suggesting that PT does play a role in fostering access to jobs. Garcia-López et al. (2017) and Mayer and Trevien (2017) confirm that the Parisian suburban rail network created jobs in the suburbs.[17]

The role of residential mobility We do not interpret the effects on commuter shares by mode as coming solely from residents that were already living in the neighbourhood prior to the tram stop opening, since it is also very likely that the latter triggered an inflow of new residents that were attracted by the newly constructed infrastructure. Unfortunately, data limitations preclude the reliable identification of resident inflows and outflows.[18] As an alternative, we proceed to a re-estimation of the results in the two bottom panels of Table 2, but we change the definition of the outcome. Instead of computing the change in the share of PT or car users using the contemporaneous total number of commuters at the denominator, we hold it constant to 2008 levels. If the estimates were substantially larger, this would be a hint that an inflow of PT users external to the neighbourhood plays a role in the effects we see in Table 2. The results of this exercise, shown in Appendix Table B.8, are akin to the main ones. While only the total number of commuters, as opposed to the exact neighbourhood composition, is held constant, this does not hint at a prominent role of preferential residential mobility towards treated neighbourhoods.[19]

[16] Motorbikes and bicycles were grouped as one mode of transport until 2017 in the census, which prevents us from estimating separate effects on motorbike and bicycle usage.

[17] The fact that our estimates on transport mode shares are less precise might also be attributable to limited data quality. Indeed, the sample of housing units surveyed is made representative across municipalities, but it might not be as representative at the *within*-municipality (IRIS) level, thus introducing measurement error in the outcome, and hampering precision.

[18] Section 2 discusses the specificities of the census data, which has two other features that prevent us from properly identifying moves. The absence of an individual identifier precludes creating a panel, so we would need to rely on the variable on whether the individual has moved in the past year, and where from, at the municipality level. We would thus have to focus on the year preceding treatment. Given that the annual collection rate is about 8%, we are unable to extract significant information.

[19] Our control group is made up of not-yet-treated neighbourhoods, and tramway lines are planned several years

5 A quantitative urban model with heterogeneous workers

By definition, the above reduced-form results pertain only to treated neighbourhoods, and thus preclude any conclusions regarding city-wide effects. In this section, we develop a model that is suitable to evaluate the general equilibrium effects of investments in transport infrastructure, accounting for air pollution as a channel. We extend the model *à la* Ahlfeldt et al. (2015), first by differentiating between the individual car and PT as two distinct transport options, each with its own disutility associated with travel time. Second, we also allow workers to differ by skill. Third, we endogenise air pollution, which is generated by car use, and allow it to affect both local amenities and local productivity.

5.1 Model set-up

We study an open city made up of N discrete locations, in which a mass of workers both live and work. Due to its openness, the city is embedded in a greater economy of M locations offering a constant level of utility. We make a distinction between two skill groups, indexed by $g \in \{H, L\}$: type-H workers are high-skilled, and type-L workers and low-skilled. These workers can also choose between two transport modes: the (petrol-powered) car, which generates air pollution, and PT, which does not.[20]

5.1.1 Workers

Worker problem Each worker o of type g chooses the pair of locations (n, i) where they want to locate, with n the neighbourhood where they will live in, and i the neighbourhood they will work in. Simultaneously, they choose their transport mode m, with $m \in \{car, PT\}$, to commute from n to i. They derive the following Cobb-Douglas indirect utility:

$$U_{nim,g}(o) = \frac{B_{n,g} w_{i,g}}{d_{nim,g} P_n^{\beta_g} Q_n^{1-\beta_g}} z_{nim,g}(o), \tag{2}$$

where $B_{n,g}$ are type-specific amenities enjoyed at residence n, and $w_{i,g}$ is the type-specific wage worker o earns in workplace i. Turning to the denominator, $d_{nim,g}$ is the utility cost of having to commute from residence n to workplace i, using transport mode m, P_n is the price of final consumption good, taken as the numéraire, such that $P_n = 1$ in all n, and Q_n is the housing rent paid in neighbourhood n, and $(1 - \beta_g)$ is the share of income devoted to housing, which differs across worker types.

ahead. Although construction work does not start very long ahead (1 or 2 years), surveys are run among local authorities and inhabitants to assess the acceptability of the project earlier (around 5 years prior). More generally, the opening of a new tramway station may be anticipated by potential movers, who may choose to move to a not-yet-treated neighbourhoods, rather than to a treated neighbourhood, during our treatment period. Note that this does not affect our main results, since the PT infrastructure needs to be open for individuals to actually use it.

[20] We therefore consider that emissions from PT are negligible. In the Paris context, this assumption is appropriate: metros and suburban trains (which cross the city centre) and tramways are electric, and the entire bus fleet is fuelled by natural gas or electricity, a type of combustion that does not emit particulates. This leaves tyre wear, which we view as negligible compared to thermal vehicle emissions, especially at the road segment or neighbourhood level.

The final component, $z_{nim,g}(o)$, is a type- and mode-specific idiosyncratic shock which is revealed to the worker prior to them choosing in which pair to locate. This shock is drawn from the following Fréchet distribution:

$$F(z_{nim,g}(o)) = e^{-T_{nm,g}\, E_{im,g}\, z_{nim,g}^{-\epsilon_g}} \qquad z_{nim,g} > 0,\ \epsilon_g > 1 \tag{3}$$

The location-type-mode-specific scale parameter $T_{nm,g}$ controls the average utility derived by a worker of type g by living in n and commuting by transport mode m. $E_{im,g}$ does the same for the average utility derived by working in i and commuting by transport mode m. The shape parameter ϵ_g governs the dispersion of the shocks, hence the degree of substitutability between alternative locations: the larger ϵ_g, the lower the variance of the distribution, i.e., the larger the substitutability between locations. For instance, with a high ϵ_g, a very small degradation in local amenities in a neighbourhood n will result in a large outflow of residents away from n to a shockless neighbourhood n'. This is due to the fact that individuals are less willing to tolerate the negative change, since their preference for n is not that different from their preference for $n' \neq n$.

Assuming that these idiosyncratic preferences follow a (type-II) extreme value distribution, here a Fréchet, is what allows us to generate a finite elasticity of substitution between modes of transport, such that all workers need not necessarily choose the fastest mode to travel from n to i. An alternative modelling choice would have been to use a nested logit model, as in Tsivanidis (2026). However, this involves taking a range of non-innocuous modelling decisions on the structure of the worker's choice process, and is also less tractable. Indeed, as shown in the next paragraph, the Fréchet shock option allows us to derive simple, tractable, gravity equations of commuting flows for each workplace-residence-mode triad, which we can take to the data in a straightforward way.[21] In this sense, we extend previous works by Allen and Arkolakis (2022); Koster (2024) by allowing for heterogeneity across worker skill levels.

Deriving probabilities Using standard properties of the Fréchet distribution (McFadden, 1974), we derive the probability that a worker of type g chooses to live in n and work in i using a transport mode m as follows:

$$\begin{aligned}\lambda_{nim,g} = \frac{L_{nim,g}}{L_{N,g}} &= \frac{T_{nm,g}E_{im,g}(B_{n,g}w_{i,g})^{\epsilon_g}(d_{nim,g}Q_n^{1-\beta_g})^{-\epsilon_g}}{\sum_{k\in N}\sum_{l\in N}\sum_{m'\in\{car,PT\}} T_{km',g}E_{lm',g}(B_{k,g}w_{l,g})^{\epsilon_g}(d_{klm',g}Q_k^{1-\beta_g})^{-\epsilon_g}} \\ &= \frac{\Phi_{nim,g}}{\sum_{k\in N}\sum_{l\in N}\sum_{m'\in\{car,PT\}}\Phi_{klm',g}}\end{aligned} \tag{4}$$

where $L_{nim,g}$, at the numerator, is the number of type-g workers that reside in n and commute to work in i using transport mode m. At the denominator, $L_{N,g}$ denotes the total number of workers of type g that reside in the city (which has N neighbourhoods). Hence, the probability

[21] In the context of choosing trading routes to ship a good, Appendix D.2 in Allen and Arkolakis (2022) shows that if the shape parameter of the Fréchet distribution of the first choice is equal to that of the second one, then the nested logit model is isomorphic to the single-step model we use here. They also show that when these elasticities are different, they lose tractability.

of commuting between locations n and i depends on residence characteristics $B_{n,g}$, Q_n and $T_{nm,g}$, workplace characteristics $w_{i,g}$ and $E_{im,g}$, and the bilateral commuting cost d_{nim} at the numerator, but also, at the denominator, on the characteristics of all other possible residences and workplaces, and corresponding bilateral commuting costs (Redding, 2024).

By taking the sum of this probability over all workplaces i, we obtain the probability that a worker of type g chooses to commute by mode m and live in residence n:

$$\lambda^R_{nm,g} = \frac{R_{nm,g}}{L_{N,g}} = \frac{\sum_{i\in N} \Phi_{nim,g}}{\sum_{k\in N}\sum_{l\in N}\sum_{m'\in\{car,PT\}} \Phi_{klm',g}}, \tag{5}$$

Similarly, by summing across all available neighbourhoods of residence n, we get the probability that a worker of type g chooses to commute by mode m to work in i:

$$\lambda^L_{im,g} = \frac{L_{im,g}}{L_{N,g}} = \frac{\sum_{n\in N} \Phi_{nim,g}}{\sum_{k\in N}\sum_{l\in N}\sum_{m'\in\{car,PT\}} \Phi_{klm',g}}. \tag{6}$$

Finally, we assume perfect population mobility within the economy, which ensures that each worker in the city derives the same expected utility at equilibrium. We denote it $\overline{U}_g$, and write it as:

$$\overline{U}_g = \delta_g \Big[\sum_{k\in N}\sum_{l\in N}\sum_{m'\in\{c,p\}} \Phi_{klm',g}\Big]^{1/\epsilon_g}, \tag{7}$$

with $\delta_g \equiv \Gamma\left(\frac{(\epsilon_g-1)}{\epsilon_g}\right)$, and $\Gamma(.)$ the gamma function.

The share of type-specific workers choosing the Paris region to live and work in depends on the attractiveness of the city relative to the rest of the economy:

$$\frac{L_{N,g}}{L_{M,g}} = \left(\frac{\overline{U}_g}{\mathbb{U}_g}\right)^{\phi} \tag{8}$$

with $\mathbb{U}_g$ the (type-specific) constant utility in the wider economy, and ϕ the migration elasticity. $L_{N,g}$ is the number of type-specific workers living and working in the Paris region, and $L_{M,g}$ is the total number of type-specific workers in the economy.

The commuting market clearing condition implies that:

$$L_{im,g} = \sum_{n\in N} R_{nim,g} = \sum_{n\in N} \lambda_{nim|nm,g} R_{nm,g}. \tag{9}$$

This means that the number $L_{im,g}$ of type-g workers that work in i and commute by m is equal to the sum, over all residences n, of the number of type-g residents who live in n, work in n and commute by m. The latter is equal to the total number of type-g residents who live in n and commute by n, $R_{nm,g}$, multiplied by the conditional probability that a type-g worker who lives in n and commutes by m lives in neighbourhood i.

Endogenous, type-specific local amenities As written in equation (2), workers enjoy amenities $B_{n,g}$, which we allow to be both type-specific and endogenous, according to the equation below:

$$B_{n,g} = b_{n,g} \left(\frac{R_n}{K_n}\right)^{\eta^R} e^{\,\zeta_g^R\, \Xi_n}, \tag{10}$$

where $b_{n,g}$ is an exogenous component capturing exogenous first-nature geography factors (such as a scenic view). $\frac{R_n}{K_n}$ is the residential density within neighbourhood n, with R_n the total number of residents and K_n the (developable) land surface. Recall that R_n is itself endogenous, and the amenity level $B_{n,g}$ reacts to it through elasticity η^R.

Ξ_n is local air pollution, which we expect to act as a disamenity, i.e., $\zeta_g^R < 0$.[22] The sensitivity of amenities to air pollution is dependent on worker type g, allowing for the possibility for high-skill workers to value clean air more strongly than lower-skill workers. Finally, air pollution is also itself endogenised, as we model below.

5.1.2 Local air pollution

We model local air pollution in a given location j (i.e., a neighbourhood of residence n or a workplace i) as follows:

$$\Xi_j = \psi_j\, e^{\,\theta^F\, F_j} \tag{11}$$

where ψ_j is a component that captures all sources of air pollution that we abstain from modelling here, and thus leave exogenous.[23] Most importantly, local air pollution is generated by $F_j \equiv \sum_n \sum_i \lambda_{ni\,car}\, \mathbf{1}\big[j \in \text{path}(n \to i)\big]$, the cumulative flows of cars that cross a given neighbourhood j. In other words, F_j is the number of workers who use mode of transport ($m = \text{car}$) and pass through location j on their way from their residence n, to their workplace i. By making air pollution a product of forces external to the residents or workers of a given neighbourhood j, we capture how worker commuting decisions affect the exposure of residents in neighbourhoods that they do not live or work in, but only commute through.

Note that we do not make a distinction between high-skill and low-skill commuters, and thus assume that they drive cars that have similar levels of $PM_{2.5}$ emissions, both primary (e.g., through tyre wear) and secondary (e.g., through NO_x emissions that generate $PM_{2.5}$).

5.1.3 Firms

Firm problem The final good Y_i is produced by a representative firm using labour of both types $L_{i,H}$ and $L_{i,L}$, as well as commercial housing H_i^L. We assume it has Cobb-Douglas technology under constant returns to scale:

$$Y_i = A_i \left(\frac{L_i}{\alpha}\right)^{\alpha} \left(\frac{H_i^L}{1-\alpha}\right)^{1-\alpha}, \tag{12}$$

[22] Local air pollution enters the expression in an exponential, such that, once we linearise the expression by taking logs, it is technically possible for neighbourhoods to have zero pollution. This will be important when we simulate a scenario where cars are entirely banned from Paris, which implies that there is zero traffic-related pollution.

[23] These can include any polluting economic activity, such as the presence of a coal plant or, more fitting in the case of Paris, a waste incinerator, which we assume exogenous here.

A_i represents the productivity at workplace i, which, similarly to amenities, we will characterise below. First, we model total workforce L_i used in production as a constant elasticity of substitution (CES) function between both types of workers, high-skilled H and low-skilled L, where ρ governs the substitution between the two types. We thus write it $L_i = (\sum_g a_{i,g} L_{i,g}^{\rho})^{1/\rho}$.[24] $a_{i,g}$ represents the skill intensity of type g in location i, with $\sum_g a_{i,g} = 1$, and α is the share of labour in the production function. Finally, H_i^L is the commercial floorspace used for production.

First-order conditions The two first-order conditions of the firm's profit function, derived with respect to type-specific labour and housing, are the following. We express them as the inverse labour demand and the inverse demand for commercial floorspace:

$$w_{i,g} = \alpha^{1-\alpha}\ a_{i,g} L_{i,g}^{\rho-1} A_i L_i^{\alpha-\rho} \left(\frac{H_i^L}{1-\alpha} \right)^{1-\alpha}, \tag{13}$$

$$Q_i = (1-\alpha)^{\alpha} A_i \left(\frac{L_i}{\alpha} \right)^{\alpha} H_i^{-\alpha}. \tag{14}$$

Endogenous local productivity We model productivity in workplace i as follows:

$$A_i = a_i \left(\frac{L_i}{K_i} \right)^{\eta^L} e^{\,\zeta^L\, \Xi_i} \tag{15}$$

where a_i is an exogenous component that captures first-nature factors making the location more productive. The density of workers in the neighbourhood, $\left(\frac{L_i}{K_i}\right)$, is assumed to have a positive effect on the productivity level, through agglomeration economies, governed by elasticity η^L. This reflects how density, that is proximity to other workers, can foster productivity at the local scale, through better skill matching, or knowledge spillovers (Ahlfeldt et al., 2015; Duranton and Kerr, 2018).

Like we do for amenities, we also introduce local air pollution Ξ_i as a factor influencing local productivity. Indeed, a growing literature shows that the productivity of both high-skill and low-skill workers, and hence total factor productivity more broadly, is negatively affected by local air pollution (Chang et al., 2019; Dechezleprêtre et al., 2019; Holub and Thies, 2023; Champalaune, 2025). We thus assume ζ_L to be negative.

5.1.4 Housing

Following Combes et al. (2021), developers produce both residential housing and commercial housing using land K and some machinery capital M, with a Cobb-Douglas functional form:

$$H_i^S = K^{\mu} M_i^{1-\mu} \tag{16}$$

[24] We also have $L_{i,g} = \sum_m L_{im,g}$, meaning that the total number of workers of type g working in location i is equal to the sum of workers of type g commuting to i using mode m, over all modes of transport.

This housing is then rented at price Q_i. Maximising profit, we obtain:

$$H_i = k_i\, Q_i^{\frac{(1-\mu)}{\mu}} \tag{17}$$

with $k_i = (1-\mu)^{\frac{(1-\mu)}{\mu}} K_i$. Housing supply H_i increases with land availability k_i and rent Q_i. $\frac{1-\mu}{\mu}$ is the price elasticity of housing supply.

5.2 Model inversion

5.2.1 Wages

Using the probability to commute from n to i by mode m, given by equation (4), and the probability to reside in neighbourhood n and commuting by m, given by equation (5), we are able to write down the probability that a worker commutes from n to i, conditionally on residing in neighbourhood n:

$$\lambda^R_{nim|nm,g} = \frac{\lambda_{nim,g}}{\lambda^R_{nm,g}} = \frac{E_{im,g}(w_{i,g}/d_{nim,g})^{\epsilon_g}}{\sum_{l\in N} E_{lm,g}(w_{l,g}/d_{nlm,g})^{\epsilon_g}} \tag{18}$$

Let us now introduce a functional form for commuting costs $d_{nim,g}$, which we assume to be equal to $d_{nim,g} = e^{\kappa_{m,g}\tau_{nim}}$. This is an iceberg cost, which generates a loss in the worker's net wage, and which is attributable to the cost of devoting time to travel to work. τ_{nim} is the time it takes to travel from residence n to workplace i using transport mode m. We assume that the elasticity of commuting costs with respect to travel time, $\kappa_{m,g}$, is mode-specific, since it is likely that one extra minute spent in a car bears a higher cost to a worker than an extra minute spent in PT. For instance, PT allows passengers to engage in other activities, such as using their phones or reading, which can lower the perceived marginal cost of travel time. In contrast, car travel does not (legally) give this opportunity, making each additional minute a direct loss, or a "pure disutility". $\kappa_{m,g}$ is also specific to worker type, since we expect higher-skilled workers to value time more highly than lower-skill workers.

We also define "transformed" wages as the product of the average utility derived from working in i, $E_{im,g}$, and the actual wage perceived from working in i, $w_{i,g}^{\epsilon_g}$:

$$\omega_{im,g} = E_{im,g} w_{i,g}^{\epsilon_g} \tag{19}$$

Combining the conditional commuting probability given in (18) with the labour market clearing condition given in (9), and using our notation in (19), we obtain:

$$\begin{aligned} L_{im,g} = \sum_{n\in N} \frac{\lambda_{nim,g}}{\lambda^R_{nm,g}} &= \sum_{n\in N} \frac{E_{im,g}(w_{i,g}/d_{nim,g})^{\epsilon_g}}{\sum_{l\in N} E_{lm,g}(w_{l,g}/d_{nlm,g})^{\epsilon_g}} R_{nm,g} \\ &= \sum_{n\in N} \frac{(\omega_{im,g}/e^{\upsilon_{mg}\ \tau_{nim}})}{\sum_{l\in N}(\omega_{lm,g}/e^{\upsilon_{mg}\ \tau_{nlm}})} R_{nm,g} \end{aligned} \tag{20}$$

where we denote $\upsilon_{m,g} = -\kappa_m \epsilon_g$. There exists a unique vector of transformed wages $\omega_{im,g}$ that solves the labour market clearing condition, given the observed vectors of residents $R_{nm,g}$, workplace employment $L_{im,g}$, and the commuting cost parametrisation. As such, given the recovered transformed wages for PT $\omega_{i\ PT,\ g}$, a normalisation of the scale parameter $E_{i\ PT,\ g}$ for PT, and the Fréchet shape parameter ϵ_g, it is possible to recover wages for car users $w_{i\ car,\ g}$ and the car scale parameter $E_{i\ car,\ g}$.

5.2.2 Skill intensity

Once type-specific wages have been retrieved, it is possible to recover the intensity in low-skilled workers $a_{i,L}$ within each workplace i using the labour demand curve (13):

$$\frac{1-a_{i,L}}{a_{i,L}} = \frac{w_{i,H}}{w_{i,L}}\Big(\frac{L_{i,L}}{L_{i,H}}\Big)^{\rho-1} \tag{21}$$

where $w_{i,H}$ is the wage of high-skilled workers in i, and $w_{i,L}$ is the wage of low-skilled workers in i. $L_{i,L}$ (resp., $L_{i,H}$) is the total number of low-skilled (resp. high-skilled) workers in workplace i.

5.2.3 Productivity

Assuming free entry for firms, and using both labour demand (13) and firm demand for housing (14), we can deduce the productivity of location i:

$$A_i = W_i^{\alpha} Q_i^{1-\alpha} \tag{22}$$

This is an increasing function of both aggregate wages W_i[25] and rents Q_i in workplace i. Intuitively, a higher productivity must induce higher labour and floorspace costs for the zero profit condition to continue to hold.

5.2.4 Amenities

We started by recovering type-specific transformed wages $\omega_{im,g}$ from Section 5.2.1. Combining type-specific expected utilities in equation (7) with the residential choice probability in equation (5), we can recover the type-specific transformed amenities:

$$\Omega_{nm,g} = \frac{\lambda^R_{nm,g} Q^{(1-\beta_g)\epsilon_g}}{\sum_{i\in N}(\omega_{im,g}/e^{\nu_{mg}\tau_{nim}})} \tag{23}$$

Given the recovered transformed amenities for PT $\Omega_{n\ PT,g}$, a normalisation of the scale parameter $T_{n\ PT,g}$ for PT to 1, and a calibration of the Fréchet shape parameter ϵ_g, it is possible to recover the average utility at the place of residence for car users, that is, the car scale parameter $T_{n\ car,g}$, as well as the vector of amenities $B_{n,g}$.

[25] Aggregate wage cost for the representative firm equals $W_i = \Big(\sum_g a_{i,g}^{\frac{1}{1-\rho}} w_{i,g}^{\frac{\rho}{\rho-1}}\Big)^{\frac{\rho-1}{\rho}}$. It is a weighted average of type-specific wages, obtained through $W_i L_i = \sum_g w_{i,g} L_{i,g}$ and type-specific inverse labour demand (13).

5.2.5 Housing and density of development

Finally, we write the market-clearing condition for housing as:

$$H_n^S = H_n^R + H_n^L, \tag{24}$$

H_n^R is the residential housing demand in n, and H_n^L the commercial housing demand in n. This allows us to recover the demand for housing addressed by residents in n:

$$H_n^R = \sum_{m'} \sum_{g} (1 - \beta_g) \sum_{i \in N} \lambda^R_{nim|nm,g} \frac{w_{i,g}}{Q_n} R_{nm,g} \tag{25}$$

where we recall that $(1 - \beta_g)$ is the share of housing in worker expenditure, which depends on worker type g. Since $\lambda^R_{ni|n,g}$ is the conditional probability for a type-g worker to work in workplace i given that they live in neighbourhood n, $\sum_{i \in N} \lambda^R_{ni|n,g} w_{i,g}$ is the expected wage of residents of type g living in neighbourhood n (and earning their wage at workplace i).

We derive the firm's demand for housing in workplace i from the first-order condition given in equation (14). This writes:

$$H_i^L = \left((1 - \alpha)\frac{A_i}{Q_i}\right)^{1/\alpha} L_i \tag{26}$$

The firm's demand for housing depends positively on its productivity A_i and the number of workers that it employs L_i, but decreases with the housing price Q_i, and the labour share α.

5.2.6 Static spatial equilibrium definition

Given the calibration of economic parameters, and a parametrisation of commuting cost $\{d_{nim,g}\}$, a static spatial equilibrium is defined with a set of observed vectors $\{R_{nm,g}, L_{im,g}, Q_n\}$, and a set of unobserved vectors $\{w_{i,g}, E_{im,g}, a_{i,L}, A_i, T_{nm,g}, B_{n,g}, H_n^S\}$ such that: i) commuting markets clear (20); ii) the first-order conditions of the firm's program (13, 14) and skill-intensity (21) are satisfied; iii) the firm's zero profit condition is satisfied (22); iv) population mobility is satisfied (23); and v) the housing market clears (24).

6 From theory to data: Quantification

This section takes the model to the data. First, we estimate a large fraction of the model's parameters. Next, we calibrate the remaining parameters, that have already been examined extensively by the existing literature.

6.1 Estimation

We now lay out our estimation strategy and corresponding results for the commuting cost elasticities $\nu_{\text{car},H}$, $\nu_{\text{car},L}$, $\nu_{\text{PT},H}$, and $\nu_{\text{PT},L}$ and the Fréchet parameters ϵ_H and ϵ_L. We make the decision to calibrate agglomeration parameters η^N and η^R, as they have been estimated extensively by the literature, so as to focus on the proper estimation of the type-specific local air pollution disamenity parameters ζ_H^R and ζ_L^R. Finally, we also estimate the elasticity of local air pollution to commuting θ^F.

6.1.1 Commuting time semi-elasticity

Given the functional form for commuting costs $d_{nim} = e^{\kappa_{m,g}\tau_{nim}}$, we are able to estimate the semi-elasticity of commuting flows with respect to travel times ($\nu_{mg} = -\epsilon_g \kappa_{m,g}$). Indeed, linearising the commuting probability in equation (4) rewrites:

$$\begin{aligned} ln\ \lambda_{nim,g} = &\underbrace{ln(E_{im,g}\ w_{i,g}^{\epsilon_g})}_{\zeta_{im,g}} + \underbrace{ln(T_{nm,g}\ B_{n,g}^{\epsilon_g}\ Q_n^{(\beta_g-1)\epsilon_g})}_{\vartheta_{nm,g}} \\ &- \underbrace{\epsilon_g \kappa_{m,g}}_{\nu_{mg}}\ \tau_{nim} - \underbrace{ln\Big(\sum_{k\in N}\sum_{l\in N}\sum_{m'\in\{car,PT\}} \Phi_{klm',g}\Big)}_{\xi_{nim,g}} \end{aligned} \tag{27}$$

This is a gravity equation, in the sense that the strength of the commuting flow $\lambda_{nim,g}$ between any pair of residence and workplace (n, i) increases with the size of each of these two neighbourhoods, $\zeta_{im,g}$ and $\vartheta_{nm,g}$.[26] It also decreases with the bilateral frictions between them (here embodied by the travel time τ_{nim} between the two locations). On the left-hand side, the empirical equivalent of $\lambda_{nim,g}$ is simply the share of workers of each type that commute from residence n to workplace i using transport mode m. We estimate each equation separately by transport mode m × type g, such that $\zeta_{im,g}$ boils down to ζ_i, a destination fixed-effect capturing workplace characteristics, and ϑ_n is an origin fixed-effect capturing residence characteristics. $\xi_{nim,g}$ is a residual term. As is standard in the literature (Ahlfeldt et al., 2015; Redding, 2024), we estimate this equation using Poisson pseudo maximum likelihood (PPML), which accommodates the presence of zero values in the dependent variable, unlike the OLS log specification (Silva and Tenreyro, 2006).

Our estimations of equation (27) are shown in Table (3). The coefficients are highly significant and have the same sign across all types and transport modes. The widely accepted range for this semi-elasticity is between -0.07 and -0.2 (see, e.g., Ahlfeldt et al., 2015; Brinkman and Lin, 2024). These estimates are weighted averages of ours, so it is reassuring that the latter lie in between -0.05 for PT and -0.12 for cars. As expected, the elasticity of commuting flows with respect to travel time is indeed larger in absolute value for car users than for PT users. Although the literature on this aspect is very limited, we do find that this is consistent with Koster (2024), who compares rail and road transport usage in the UK context. We explain it by the fact that PT utility costs are likely less directly linked to time than car usage costs. In the case of car users, commuting costs, be they monetary or non-monetary, increase substantially with travel time: longer trips imply a higher fuel

[26] More specifically, the size of each workplace i is embedded into the average utility of working in i and the wage it offers, and the size of each residence n is embedded into the average utility derived from living in n, its rents and its amenities.

Table 3: Estimation of $\nu_{m,g}$

	HS car (1)	LS car (2)	HS public (3)	LS public (4)
Commuting time by car	-0.1042***	-0.1224***		
	(0.0006)	(0.0008)		
Commuting time by PT			-0.0529***	-0.0582***
			(0.0002)	(0.0003)
Origin FE	Yes	Yes	Yes	Yes
Destination FE	Yes	Yes	Yes	Yes
Observations	544,643	539,484	544,644	545,382

Notes: Signif. Codes: ***: 0.01, **: 0.05, *: 0.1. Heteroscedasticity-robust standard errors in parentheses. Estimation by Poisson Pseudo-Maximum Likelihood (PPML) to accommodate zeros on the left-hand side.

consumption, a higher depreciation of the vehicle, but also more time lost, since it is not possible to perform other tasks from behind the wheel. On the other hand, the monetary cost for PT users does not rise as sharply with each extra minute of travel time, since the structure of fares is, in our context, completely flat.[27] The non-monetary cost also increases less quickly with time, since the perceived burden of one extra minute spent in PT is less cumbersome, as one is able to engage in other activities.

Regarding heterogeneity by skill, though statistically significant, differences across types of workers are not economically significant. Semi-elasticities of commuting cost are larger in absolute value for lower-skill workers, which appear to be slightly more sensitive to travel time. Recall that the estimated semi-elasticity ν_{mg} is equal to the product of the type-specific Fréchet shape parameter ϵ_g and the mode-specific cost elasticity κ_m. We thus formulate the hypothesis that $\|\nu_{m,H}\| < \|\nu_{m,L}\|$ is attributable to the fact that $\epsilon_H < \epsilon_L$. In words, this means we assume that higher-skill workers have a larger dispersion in the utility they derive from each location (Redding and Sturm, 2024).

6.1.2 Fréchet parameters

Using municipality-level data on the dispersion of 2018 wages across workplaces, we are able to derive estimates for Fréchet parameters (following Ahlfeldt et al., 2015). Indeed, the labour market clearing condition in (20) implies that transformed wages $\omega_{im,g}$ can be deduced from information on workplace employment, residence employment and commuting times, independently from the value of the Fréchet parameters ϵ_g. This means that as equation (19) suggests, the Fréchet shape parameters ϵ_g only convert wages into transformed wages, and thus only scale the variance of wages. We call the latter $\sigma^2_{w_{i,g}}$. Thus, we calibrate the values of ϵ_g, for $g \in \{H,\ L\}$, such that we minimise

[27] Our baseline model is calibrated on 2018 data, after the removal of the fare zoning that occurred in 2015.

the squared difference between our observed wages and estimated wages:

$$\epsilon_g = \text{argmin} \left(\sigma^2_{\hat{\omega}_{i,g}} - \sigma^2_{w_{i,g}} \right)^2 \tag{28}$$

where σ^2 denotes the variance, $\hat{\omega}_{i,g}$ are transformed log wages which we recover from the model, and $w_{i,g}$ are observed log wages in 2018. Doing so, we get values equal to $\epsilon_L = 9.3$, and $\epsilon_H = 6.4$. This confirms the hypothesis made right above that $\epsilon_H < \epsilon_L$.

6.1.3 Air pollution as a disamenity

We log-linearise the expressions for amenities (10) and productivity (15), and go on with the estimation of local pollution disamenity parameters. We expect omitted variables to generate bias in the estimation of the effect of air pollution on contemporary amenity and productivity levels, including worker and firm sorting mechanisms. Hence, we opt for an estimation in first-difference, which allows us to get rid of the influence of time-invariant confounding factors.

$$\begin{aligned} \Delta \ln\ B_{n,g} &= \eta^R\, \Delta \ln\ \big(R_n\big) + \zeta^R_g\, \Delta \Xi_n + \Delta \ln\ b_{n,g} \\ \Delta \ln\ A_n &= \eta^L\, \Delta \ln\ \big(L_n\big) + \zeta^L\, \Delta \Xi_n + \Delta \ln\ a_n \end{aligned} \tag{29}$$

where η^R is the residential elasticity of amenities, and ζ^R_g is the disamenity effect of local air pollution, as perceived by workers of type $g \in \{H,\ L\}$, and ζ^L is the production disamenity effect of local air pollution. We view $\Delta \ln\ a_n$ and $\Delta \ln\ b_{n,g}$ as error terms.

Even taking the first-difference, we fear that there may be time-varying elements that affect both the change in local amenities $\Delta \ln\ B_{n,g}$ or local productivity $\Delta \ln\ A_n$ and the change in local $PM_{2.5}$ concentration $\Delta \Xi_n$. For instance, at the instigation of successive Paris mayors Bertrand Delanoë and Anne Hidalgo, some streets were pedestrianised over the 2008-2018 period which affected local air quality (Bou Sleiman, 2023). Targeted neighbourhoods were not chosen randomly, but because of their proximity to parks, museums, commercial streets or other amenities. Other city-level actions implemented by French cities, including Paris, were also shown to have differential effects by neighbourhood income level (Champalaune, 2026). Hence, we develop an instrument for the local change in $PM_{2.5}$ concentration. We turn the fact that a) prevailing winds blow from the West to the East of the city, and that b) descriptive statistics from the reduced-form exercise show that the share of car commuters decreased between 2008 and 2018, to our advantage, and use the variation in upwind car commuters as an instrument. More specifically, we use the sum of all variations in the number of car commuters crossing all neighbourhoods $j \neq i$ that are located West of neighbourhood i, weighted by proximity. Formally, the IV is defined as:

$$\Delta \text{ upwind car commuters}_i = \sum_j \mathbb{1}\{\text{longitude}_j < \text{longitude}_i\} \frac{1}{\text{distance}_{ij}} \Delta \text{ car commuters}_j \tag{30}$$

where $\mathbb{1}\{\text{longitude}_j < \text{longitude}_i\}$ is a dummy equal to 1 if neighbourhood j is located west of neighbourhood i. As such, we exploit the fact that decreases in car traffic upwind of a neighbourhood

Table 4: Effect of local air pollution on local amenity and productivity, by skill

	$\Delta\ln B_{n,H}$			$\Delta\ln B_{n,L}$			$\Delta\ln A$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ΔPM$_{2.5}$	0.003***	-0.090***	-0.090***	0.003***	-0.027	-0.026**	-0.076***
	(0.0010)	(0.025)	(0.016)	(0.0008)	(0.018)	(0.013)	(0.020)
Initial PM$_{2.5}$		-0.051***	-0.049***		-0.021***	-0.018***	-0.028***
		(0.011)	(0.007)		(0.007)	(0.005)	(0.005)
Δln residents			0.165***			0.104***	
			(0.007)			(0.005)	
Δln jobs							0.113***
							(0.002)
IV for ΔPM$_{2.5}$		Yes	Yes		Yes	Yes	Yes
Extra controls			Yes			Yes	Yes
Observations	732	732	732	733	732	732	732
R^2	0.009	0.166	0.577	0.016	0.133	0.462	0.916
F-stat, ΔPM$_{2.5}$		82.63	89.90		82.63	89.90	80.14

Notes: Signif. Codes: ***: 0.01, **: 0.05, *: 0.1. Heteroscedasticity-robust standard errors in parentheses. ln$\Delta\ln B_{n,H}$ (resp., $\Delta\ln B_{n,L}$) denotes the change in amenities as perceived by high-skill (resp., low-skill) workers in neighbourhood of residence n between 2008 and 2018. ΔA denotes the change in productivity in workplace i between 2008 and 2018. ΔPM$_{2.5}$ denotes the change in PM$_{2.5}$ between 2008 and 2018. In Columns (2)-(3) and (5)-(7), ΔPM$_{2.5}$ is instrumented using Δupwind car commuters, defined in equation (30). All regressions control for neighbourhood area K. Columns (3), (6) and (7) control for log distance to CBD.

influence its change in local air pollution, without directly affecting its local amenity or productivity level. The first-stage regression results in Table B.9 provide formal evidence of the relevance of the instrument, with the lowest F-statistics at 72.

Table 4 shows the results of the OLS estimation of equations (29), for amenity levels as perceived by high-skilled workers in Columns (1)-(3), by low-skilled workers in Columns (4)-(6) and productivity in Column (7). OLS specifications render (biased) slightly positive estimates, due to omitted variables such as localised PT investments and housing market dynamics. Instrumenting for the change in local PM$_{2.5}$ concentration, we find that the latter is perceived as a consumption disamenity among both types of workers, though it is a stronger disamenity for high-skill workers. This is consistent with a more pronounced preference of the high-skilled for clean air, stemming from plausibly better awareness about its effect on health outcomes. Given that local air quality improved during the period, we interpret these coefficients as follows: a 1μg/m^3 decrease in PM$_{2.5}$ (from a 15.8 mean in 2008) generates a perceived improvement of amenities by 9% for high-skilled workers, and 2.6% for low-skill workers. A similar change generates an increase in 7.6% in local productivity.

While we have included the change in the number of residents R and the number of jobs L in some specifications, we are not able to instrument for these (potentially endogenous) variables simultaneously with local air pollution. We thus resort to calibrating them as standard values in the literature, with $\eta_L = 0.07$ and $\eta_N = 0.10$ (Ahlfeldt et al., 2015; Heblich et al., 2020). Nonetheless, we are confident that the inclusion of R and L in the estimations of agglomeration forces in Columns (3)

Table 5: Estimation of θ^F

	$\Delta ln\ \Xi_j$	
ΔF_j	0.013***	0.028***
	(0.002)	(0.006)
Geometric mean correction	No	Yes
Observations	672	672
R^2	0.078	0.061

Notes: Authors' estimations based on Airparif $PM_{2.5}$ concentration data and INSEE Census data. Heteroscedasticity-robust standard errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1. Column (1) defines ΔF_j as the change in the flow of car commuters between 2008 and 2018, in thousands. Column (2) defines ΔF_j as the change in the flow of car commuters between 2008 and 2018, standardised by the geometric mean.

and (6) of Table 4 does not strongly contaminate our estimates of the disamenity effects of local air pollution ζ_H^R, ζ_L^R and ζ^L (through a "bad control" problem), as there is no difference in point estimates between Columns (2) and (3), and Columns (5) and (6).

6.1.4 Local air pollution from car traffic

First, we (log-)linearise the expression for local air pollution given in equation (11), and take the first-difference, which gives:

$$\Delta ln\ \Xi_j = \psi_0 + \theta^F\ \Delta F_j + \Delta ln\ \psi_j \tag{31}$$

Then, we estimate equation (31) using the same data sources as in the reduced-form section. In particular, we still suspect that there is endogeneity in the relationship between the number of car commuters driving through a given neighbourhood and local air pollution, due to, e.g., the non-random placement of major roads in already polluted neighbourhoods. We do not completely circumvent this issue, but instead of regressing current (2018) $PM_{2.5}$ concentration in j on the current (2018) flow of commuters that cross j, we also resort to using first-differences, and regress the *change* in $PM_{2.5}$ level on the *change* in the number of commuters driving through neighbourhood j, between 2008 and 2018. This accounts for the presence of any time-invariant, unobserved factors that may confound the relationship, such as the presence of polluting facilities or major roads. The corresponding results are shown in Table 5. An increase of 1,000 car commuters driving through a neighbourhood, relative to the broader trend across all neighbourhoods, is associated with a 1.3% increase in annual $PM_{2.5}$ concentration in the neighbourhood. This is close to estimates obtained for Milan (Gibson and Carnovale, 2015) or London (Green et al., 2020). In Column (2), we normalise both of these variables by their geometric mean, so as to circumvent any issues related to zeros or changes in the overall distribution of car commuting flows within neighbourhoods. This estimate is therefore ready to be employed in other work.

Finally, our measure of car traffic relies on commuting flows, which do not capture freight or leisure travel. We view this as a minor limitation. Freight is irrelevant to our counterfactual exercises, since goods do not transit by public transport. Leisure and shopping trips, while numerous, tend to be much shorter than commutes. According to a 2023 survey by Institut Paris Région, over 50% of daily trips made by working-age adults are home-to-work commutes, a share that is expectedly even larger when measured in distance.[28] Commuting thus plausibly accounts for the bulk of car-generated pollution exposure in neighbourhoods.

6.2 Calibration of common parameters

We rely on existing literature to calibrate the remaining parameters, as their values are rather consensual and their estimation is not central to this study. In particular, as above-mentioned, we calibrate agglomeration parameters η^R and η^L to 0.10 and 0.07, respectively, since they have been reliably estimated in a significant fraction of existing papers (Ahlfeldt et al., 2015; Heblich et al., 2020; Allen and Arkolakis, 2022; Takeda and Yamagishi, 2024). We take a labour share equal to $\alpha = 0.75$, following Cette et al. (2019) and Gutiérrez and Piton (2020). Regarding the share of income devoted to housing, $(1-\beta_g)$, we deduce values from Combes et al. (2019) for the city of Paris, who find that for homeowners, this share is equal to 0.344, and to 0.369 for renters. While this does not give a distinction between high- and low-skill households, we use them as an indication to set $1 - \beta_H = 0.35$, and $1 - \beta_L = 0.40$. We use the housing supply elasticity with respect to machinery capital derived by Combes et al. (2021) specifically for the context of Paris, with $1 - \mu = 0.54$. The elasticity of substitution between high- and low-skilled labour is taken to be equal to $\rho = 0.3$ from Card (2009), also used by Tsivanidis (2026). Finally, we take a value $\phi = 3$ for the migration elasticity, as is standard in the literature (Monte et al., 2018; Bryan and Morten, 2019), and as was confirmed in the French context (Bilal, 2023).

7 Counterfactual analyses

We start by an exercise where we simulate the effect of banning all cars from the Paris municipality. We then come to our prospective counterfactual analyses of PT infrastructure investments, which will generate improvements in travel time: the construction of additional tramway lines, and the creation of *Grand Paris Express* metro around the city centre.

7.1 Exact-hat procedure

We use the exact hat-algebra method popularised by Dekle et al. (2008) to conduct several experiments. This technique works as follows. Let us choose a variable x in the initial equilibrium. After a change in model fundamentals, this variable now takes value x'. The hat-form version of variable x will be defined as the ratio between the new variable and the original one, such that $\widehat{x} \equiv \frac{x'}{x}$. The idea behind this method is thus to write the entire original system of equilibrium equations in *changes*, instead of computing levels to then compare. We provide all expressions in transformed, "hat"-form, in Appendix D.

[28] The main results of the survey are given here, in French: https://www.institutparisregion.fr/mobilite-et-transports/deplacements/enquete-regionale-sur-la-mobilite-des-franciliens/

7.2 Projected impact of banning cars from Paris

We start with a simple exercise, where we set a strict policy that bans all cars from entering the city centre of Paris, or formally, the Paris municipality. Since November 2024, it is already the case that all motor vehicles are banned, regardless of their emission levels, from the most central districts. The policy we simulate thus amounts to generalising this measure to the entire municipality, which is on the agenda of the new candidate for mayoralty, endorsed by the current mayor Anne Hidalgo.[29] This policy is therefore a realistic scenario, at the very least in terms of interest from policy-makers. Ubeda (2021) performs a similar analysis without including endogenous variations in local air quality, and finds a welfare loss of 2.49% for low-skill workers, and 2.95% for high-skill workers. Our results are shown in Appendix Table B.10. Column (1) replicates Ubeda (2021), whose setting is that of a closed city without local air quality responses. Then, as we open the city in Column (2), welfare gains diminish due to the rise in housing prices triggered by the inflow of new residents. Our preferred, complete estimation of the effects is shown in Column (3): as opposed to Ubeda (2021), we find positive welfare effects of the implementation of the policy, despite allowing for in-migration, due to the significant improvement in air quality. The latter is especially strong in the city centre, where cars are banned, but the effect also spreads out to the rest of the city, as the cost of increased car travel time deters about 20% of former car users of commuting with their individual vehicle.

7.3 Projected impact of new tramway lines

As a continuation of the reduced-form evidence provided in Section 4 on the opening of tramway lines between 2008 and 2018, we proceed to simulate the effect of the opening of new such lines, from 2018 on. Between 2018 and 2024, 7 lines were built or extended, thus creating 78 new stations, and 68 (mostly extensions of existing lines) are in the pipeline. All additions from 2018 on represent about 150 km of new lines. Figure A.6 in Appendix displays the average gain in travel time from each origin neighbourhood to all other neighbourhoods following tramway station openings. It is clearly visible on this map that the objective behind the construction of this infrastructure (in black) is to provide connections with the existing network (in purple), so as to generate alternative routes to help decongest it, but also to connect isolated areas that only have bus transit.[30] Some municipalities gain up to 15 minutes of average commuting time to all possible destinations in the city.[31]

[29] The policy bears the name of "*Zone à trafic limité*" (ZTL), and is thus more stringent than the "*Zone à faibles émissions*" (low-emission zone) that covers the Paris municipality and 76 neighbouring municipalities, and only excludes the most polluting vehicles from entering its perimeter. There are very few exceptions to the ban: vehicles belonging to the PT fleet, health services, taxi cabs, and those that live and/or work in the targeted area. All information on ZTL provided by the municipality is available on the Paris municipality website (in French, accessed 22/09/2025). The former candidate for mayor Rémi Féraud mentioned the extension of ZTL to the whole municipality in an interview to local newspaper *La Parisien* in November 2024 (see here, in French).

[30] Indeed, before the implementation of these lines, suburb-to-suburb transport options are very limited, and such trips often imply having to go through the central municipality of Paris to leave it again, instead or circumventing it.

[31] The average gain in travel time is less substantial than in the case of GPE, which we study below and extends further out. This is also due to the fact that these new lines have an average commercial speed of 45km/h, as opposed to 60km/h for GPE. 45 km/h is nonetheless high as compared to the speed of regular trams, because most of these new lines are "tram-trains", which can operate on both urban tracks and mainline railway tracks.

Table 6: Projected effects of new tramway lines

	Counterfactual scenario					
	(1)	(2)	(3)	(4)	(5)	(6)
Welfare $\Delta\%$						
High-skilled	0.54	0.24	0.32	0.23	0.28	0.28
Low-skilled	0.28	0.15	0.19	0.14	0.17	0.16
λ_{Car} $\Delta\%$						
High-skilled	-1.82	-1.1	-1.03	-1.25	-1.14	-1.12
Low-skilled	-1.88	-1.42	-1.57	-1.7	-1.65	-1.62
λ_{Public} $\Delta\%$						
High-skilled	0.66	1.39	1.69	1.4	1.55	1.57
Low-skilled	0.59	1.06	1.26	1.06	1.18	1.16
(mean) Pollution $\Delta\%$						
Whole area	-0.08	-0.05	-0.05	-0.06	-0.05	-0.05
Paris	-0.1	-0.06	-0.06	-0.07	-0.06	-0.06
Outside Paris	-0.08	-0.05	-0.05	-0.06	-0.05	-0.05
(mean) Rent $\Delta\%$						
Whole area	0.05	0.3	0.41	0.3	0.39	0.34
Total population $\Delta\%$						
Whole	0.0	0.63	0.82	0.59	0.71	0.72
High-skilled	0.0	0.72	0.96	0.69	0.83	0.85
Low-skilled	0.0	0.47	0.58	0.41	0.51	0.5
Parameters						
Migration elasticity	0.0	3.0	3.0	3.0	3.0	3.0
η^L	0.0	0.0	0.07	0.07	0.07	0.07
η^R	0.0	0.0	0.1	0.1	0.1	0.1
ζ_L^R	-0.025	-0.025	-0.025	0.0	0.0	-0.025
ζ_H^R	-0.09	-0.09	-0.09	0.0	0.0	-0.09
ζ^L	-0.075	-0.075	-0.075	0.0	-0.075	0.0

Notes: Each column reports a separate counterfactual effects of tramway lines built since 2018 and planned to this date. Counterfactuals take the 2018 equilibrium as the baseline, and consider different structural parameters. Column (1): closed city, no agglomeration effects, local air pollution effects. Column (2): open city, no agglomeration effects, local air pollution effects. Column (3): open city, agglomeration effects and local air pollution effects. Column (4): open city, agglomeration effects, no local air pollution effects. Column (5): open city, agglomeration effects, productivity effect of local air pollution, no amenity effect of local air pollution. Column (6): open city, agglomeration effects, disamenity effects of local air pollution, no productivity effect of local air pollution.

City-level effects Table 6 provides our estimations of the change in welfare for both high-skilled and low-skilled workers, and the changes in the fraction of workers using each transport mode λ_m, in air pollution, in rents and in total population, all attributable to new tramway lines. Column (3) provides the full effects, including all channels. As expected given the small scale of the infrastructure, we find welfare effects that are smaller than for previously studied projects (e.g., Tsivanidis, 2026). We also find regressive effects, in the sense that high-skilled workers benefit 68% more than low-skill workers from the new infrastructure (0.32% *versus* 0.19%). As variations in travel time are small, the incentive to use PT is not sufficient for high-skilled workers to react a lot, as they are less sensitive to travel time than low-skill workers on average. There is still an increase in PT take-up by the high-skilled, but it is low. On average, at the city level, air quality does not improve much, if at all (-0.05%).

Spatial distribution of economic activity and air pollution Appendix Figure A.7 shows neighbourhood-level changes in number of commuters, local air quality, employment and housing prices. The results here are more mixed than for GPE, and the averages we just discussed mask substantial heterogeneity in the effects of the new tramway infrastructure over space. Indeed, Panel (a) shows a decrease in the number of car commuters in the northern two-thirds of the city, including the Paris municipality. But in the southern part, the number of car commuters actually increases. This translates into a rise in local air pollution in this zone, as opposed to the rest of the city where it decreases (Panel (b)). This can be explained by looking at the evidence in Panel (c). Indeed, the latter a sizeable rise in employment in the south-west, all along and to the south-west of the new T12 tramway line, which connects Evry (in the south) to Massy-Palaiseau (about halfway between the Saclay hub point and the Orly airport). The new tramway stations draw a significant number of jobs along the new line, but the latter remains unconnected to the neighbourhoods located to its south. Hence, their residents, in which there are no new stations or sufficiently fast buses, have to keep on commuting by car, and this for a longer distance as jobs are being displaced close to the tramway line.

7.4 Background information on *Grand Paris Express*

The *Grand Paris Express* (GPE) project, as we quickly mentioned in Section 3.1, is an extensive project of the Paris region, which is, to this date, still under construction. It will not be delivered before the mid-2030s for the most part. Given the very large scale and cost of this project, we offer to provide a prospective analysis of the effects of this investment.

The project involves the construction of four new automated metro lines around Paris, outside the city centre, together with the extension of two existing lines, adding up to 200 km of new tracks. To this day (or 31/12/2024), only one extension, of line 14, has been opened, and the rest of the new stations are expected to be put into service starting in late 2026, with a full completion date in 2031. The construction of the metro is part of a broader initiative aimed at further developing the Paris area.[32] For instance, the south-western part of the project, metro line 18, aims to better connect the Saclay research and development hub to the rest of the suburbs, and, by extension, to the city centre. The authorities at the origin of this initiative, including the French government and the Île-de-France administrative region, expect a number of other benefits, including better access to employment for disadvantaged neighbourhoods. 53 “priority neighbourhoods” will be newly connected to the rest of the PT network *via* these new lines, and a reduction of the use of

[32] This is the reason why the project was also met with criticism from some parties, who argued that it would come to reinforce the already very strong power of the capital region, as compared to other large French cities.

Figure 4: Gain in travel time due to *Grand Paris Express* at the origin-neighbourhood level

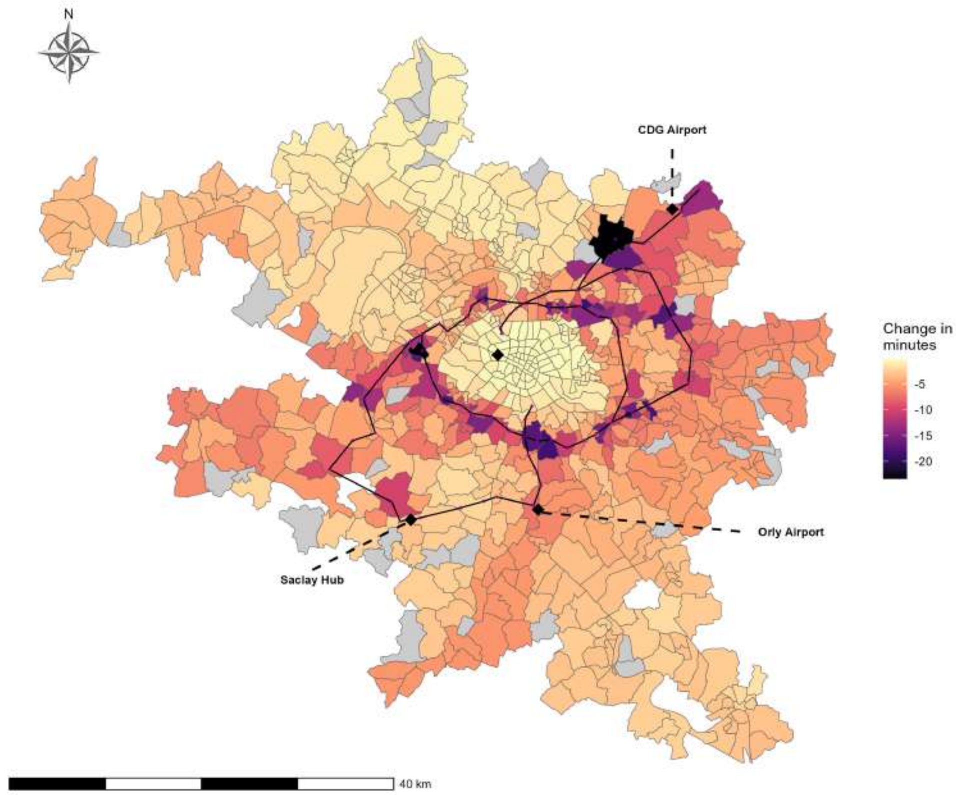

Notes: Authors' own computation based on data from IDFM and *Société du Grand Paris* for building the public transport networks, and INSEE open data for delineating neighbourhoods and Paris the urban unit.

cars in favour of PT.[33]

Indeed, of high relevance of this study, the construction of these new metro lines will provide substantial gains in travel time for PT users, as shown in Figure 4. The largest gain, in the north-east, corresponds to Gonesse, a municipality located in-between the Bourget secondary airport and the Paris Charles-de-Gaulle (CDG) airport, which is the main one in Paris and the second busiest in Europe, and where the poverty rate reaches 25% (i.e., more than three times the 2024 national average of 7.3%).[34] On the other hand, slightly darker municipalities in the south-west belong to the abovementioned Saclay hub. They are much more dynamic economically, and host a variety of higher education institutions and research labs.[35] The estimated gain in travel time is of 15 minutes to reach any of the neighbourhoods of the city.

Figure 5: Variation in air pollution, number of residents and employment around GPE metro stations

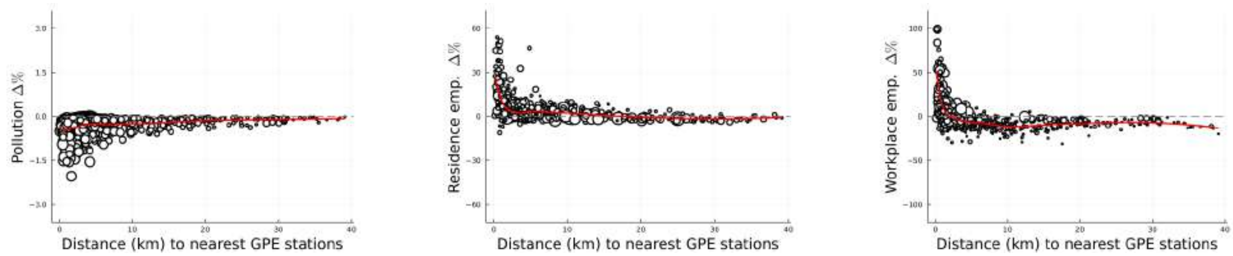

(a) Pollution (b) Residents (c) Employment

Notes: Counterfactual effects of *Grand Paris Express*, taking the 2018 equilibrium as the baseline. Changes in percentages. All channels are enabled: open city, agglomeration effects and local air pollution effects on productivity and amenity levels.

7.5 Projected impact of *Grand Paris Express*

Effects around new metro stations Figure 5 provides an overview of the effects of the introduction of GPE, around the newly created stations. Panel (a) shows that there will be a decrease in local air pollution in the close vicinity of the station, which will fade out as distance to station increases. Both better connectivity and cleaner air will attract new residents to neighbourhoods close to the station, but the effect is very localised, as seen in Panel (b). The same applies to workplace employment, displayed in Panel (c), which also shows an increase in the number of jobs that is confined to the neighbourhoods that are very close to the new station. In terms of magnitude, we notice that the relative increase in employment is much larger than the relative increase in residents, suggesting a very strong reaction of firms so as to locate close to the new infrastructure. The attractiveness of these locations in terms of accessibility is reinforced by the fact that they are less polluted, which generates extra productivity gains for firms who locate there.

Spatial distribution of economic activity and air pollution Running this counterfactual exercise renders the results shown in Figure 6. As expected, Panel (a) shows that GPE causes a sharp drop in the number of car commuters that cross the neighbourhoods in which a new metro station opens. This is due to a decrease in car usage by the residents of these neighbourhoods themselves, but also to a decrease in car usage by residents outside the neighbourhood, who used to cross it by car but no longer do. Consequently, the first map shows a decrease in the number of car commuters that go through a large fraction of neighbourhoods, even in some that are further away from the new metro station. This is especially true in the easternmost part of the city, where GPE metro line 16 will allow to build a faster connection between three existing regional express train

[33] On a website dedicated to the project, the entity in charge lists a number of expected benefits for residents, see (in French): Which benefits from Grand Paris Express?, accessed 26/02/2025.

[34] Source: INSEE.

[35] For a full map of the project with all connections to other lines, see this Grand Paris Express Map by clicking on "Explorer la carte".

Figure 6: Variation in main variables due to *Grand Paris Express*

(a) Car commuters through

(b) Local air pollution

(c) Employment

(d) Housing prices

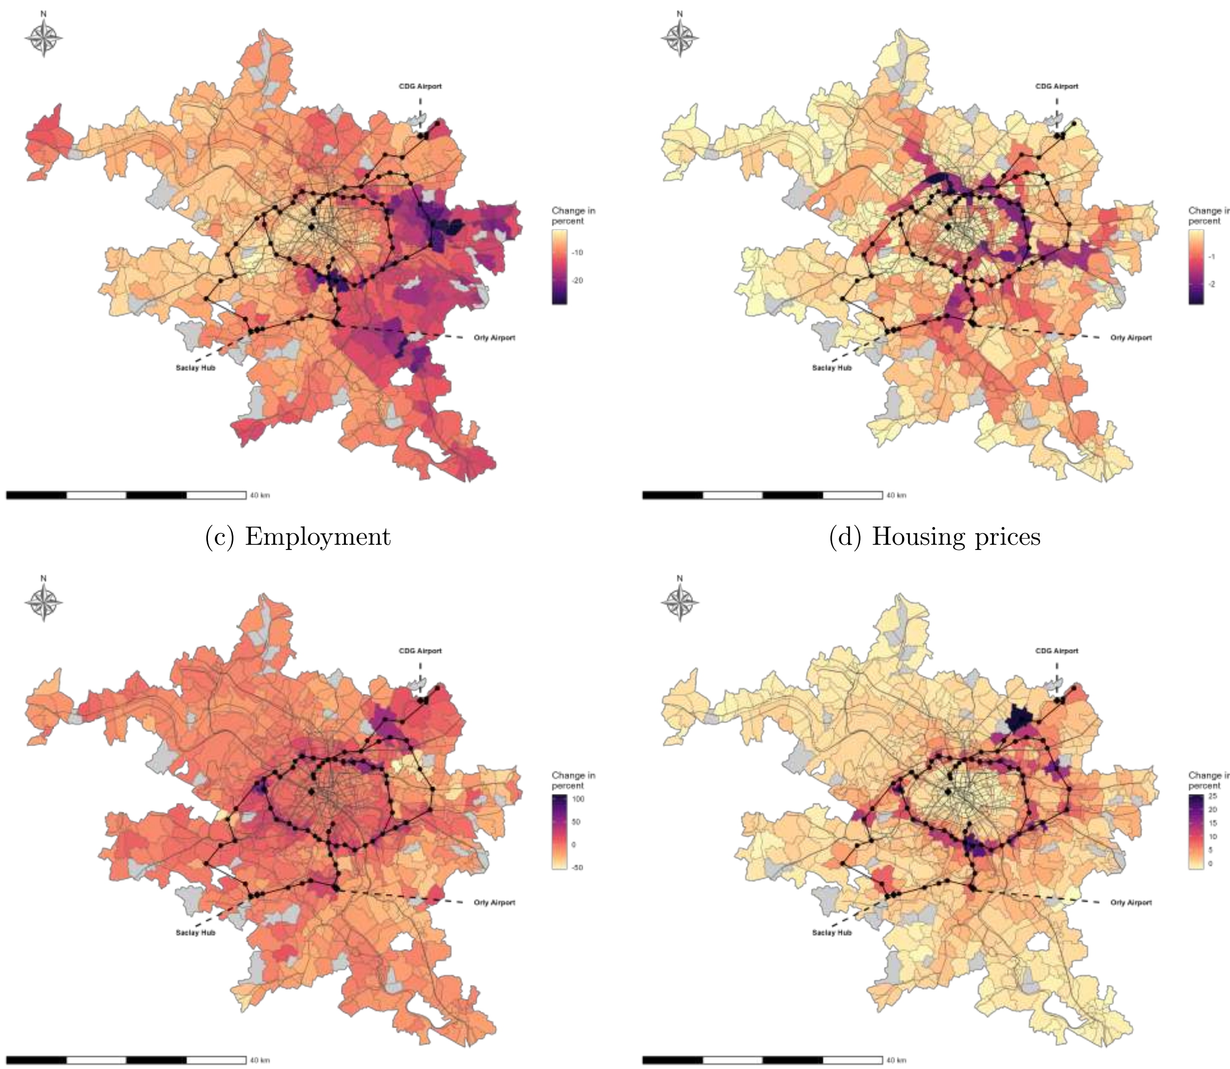

Notes: Variation in the number of crossing car commuters, local air pollution, employment, and housing prices in all neighbourhoods. Counterfactual simulation of the effects of GPE where all channels are enabled: open city, agglomeration forces, and reaction of amenities and productivity to air pollution. These results are aggregated to the city level in Column (3) of Table 7.

lines (RER A, B, D), as well as two other suburban rail services, and where the number of cars crossing drops by up to 30%. On the other hand, most neighbourhoods in the western part of the region see almost no change in the number of car users that commute through them (around -2%). This is to be expected, since this zone is left largely untreated by GPE. Given that the western part of the new infrastructure is connected to the three suburban rail lines that serve the west (RER A, C and E), we still see a small improvement in travel time in Figure 4, but given that it is less than 5 minutes, it does not translate into a strong reaction in terms of commuting behaviour.

Panel (b) shows the effects on local air quality. The results do not perfectly mirror those in Panel (a), although pollution is generated solely by commuting in the model, because the results are shown in proportion to baseline levels. The decrease in the number of car commuters is approximately the same in the entire eastern half-ring of the GPE circular line (line 15), around 15%. However, the baseline number of car commuters driving through the north-eastern portion is higher than those who drive through the south-eastern portion, such that there is a substantially larger improvement in air quality in the north-eastern part than in the south-eastern part. On the other hand, the western part of the region remains largely unaffected by GPE, such that the change in local air quality is more limited, although we see some spillovers. As the east and north-east areas are comparatively both more polluted, and more economically disadvantaged to begin with, these results suggest that the GPE project will attenuate current inequalities in exposure to $PM_{2.5}$.

Let us turn to Panels (c) and (d) on employment and housing prices. Employment moves towards new metro stations, it leaves some areas even more residential than before. This is especially the case in the east, where differences in changes in employment are very contrasted: there are large gains close to the CDG and Le Bourget airports (industrial areas, with logistics platforms), but decreases in employment right at the south of these airports. As commuting to work is made much faster, we see a reinforcement of the separation between residence and workplace (Heblich et al., 2020). This is less the case in the south, were the effects seem to be more spatially spread out. Improved commuting time, improved air quality, and an inflow of new residents and firms will also prop up housing prices, as shown in Panel (d).

City-level changes Table 7 provide our estimations of the effects of GPE, mirroring Table 6. As expected given the much larger scale of the project, we find welfare effects that are about twice the size of that of new tramway lines. With that said, GPE has “regressive” distributional effects, in the sense that higher-skilled workers benefit from a 54% larger welfare gain (1.79%) than low-skilled workers (1.16%).

Our estimate of the overall welfare improvement of welfare of 1.5% on average is slightly lower, than previous estimates on similar large-scale PT infrastructure (e.g., Tsivanidis, 2026, on Bogotá’s TransMilenio bus system).[36] We find that the share of PT users within the group of high-skilled (resp., low-skilled) workers increases by 9.7% (resp., 7.4%), from an initial share of 38% (bottom panel of Table B.3). It goes hand in hand with large decreases of 6.3% and 9% in the share of workers using their car to commute to work, from an initial share of 40%. This translates into an improvement in local air quality, albeit moderate: we estimate that $PM_{2.5}$ concentration decreases by 0.3%. Interestingly, although the new metro lines are built in the suburbs, air pollution decreases by approximately the same amount in the city centre as outside the Paris municipality. This is due to the fact that a significant fraction of car commuters within Paris actually came from much further out in the suburbs, and have now switched to using GPE, which is well connected to the centre’s existing metro network.

[36] These are the average welfare gains, obtained by using initial 2018 population share of high-skilled $s_H = 0.631$ and share of low-skilled $s_L = 0.369$.

Table 7: Projected effects of *Grand Paris Express*

	Counterfactual scenario					
	(1)	(2)	(3)	(4)	(5)	(6)
Welfare $\Delta\%$						
High-skilled	3.07	1.35	1.79	1.25	1.53	1.56
Low-skilled	1.62	0.89	1.16	0.81	1.0	0.99
λ_{Car} $\Delta\%$						
High-skilled	-9.8	-6.0	-6.3	-7.5	-6.9	-6.8
Low-skilled	-10.3	-7.9	-9.0	-9.7	-9.5	-9.3
λ_{Public} $\Delta\%$						
High-skilled	3.6	7.8	9.7	7.9	8.9	9.0
Low-skilled	3.2	6.0	7.4	6.2	6.9	6.9
(mean) Pollution $\Delta\%$						
Whole area	-0.45	-0.3	-0.3	-0.35	-0.33	-0.33
Paris	-0.48	-0.28	-0.29	-0.37	-0.33	-0.3
Outside Paris	-0.45	-0.3	-0.3	-0.35	-0.32	-0.33
(mean) Rent $\Delta\%$						
Whole area	0.2	1.57	2.21	1.51	2.05	1.73
Total population $\Delta\%$						
Whole	0.0	3.6	4.74	3.29	4.05	4.11
High-skilled	0.0	4.12	5.46	3.79	4.65	4.76
Low-skilled	0.0	2.71	3.51	2.44	3.03	3.01
Parameters						
Migration elasticity	0.0	3.0	3.0	3.0	3.0	3.0
η^L	0.0	0.0	0.07	0.07	0.07	0.07
η^R	0.0	0.0	0.1	0.1	0.1	0.1
ζ_L^R	-0.025	-0.025	-0.025	0.0	0.0	-0.025
ζ_H^R	-0.09	-0.09	-0.09	0.0	0.0	-0.09
ζ^L	-0.075	-0.075	-0.075	0.0	-0.075	0.0

Notes: Each column reports a separate counterfactual effects of the planned GPE metro system. Counterfactuals take the 2018 equilibrium as the baseline, and consider different structural parameters.
Column (1): closed city, no agglomeration effects, local air pollution effects. Column (2): open city, no agglomeration effects, local air pollution effects. Column (3): open city, agglomeration effects and local air pollution effects. Column (4): open city, agglomeration effects, no local air pollution effects. Column (5): open city, agglomeration effects, productivity effect of local air pollution, no amenity effect of local air pollution. Column (6): open city, agglomeration effects, disamenity effects of local air pollution, no productivity effect of local air pollution.

Channels: Agglomeration Column (1) shuts down agglomeration forces by setting $\eta^L = eta^R = 0$, and closes the city by setting the migration elasticity to zero. Column (2) opens the city again, but leaves agglomeration parameters at a zero value. This allows us to quantify the role of in-migration. As commuting time by PT decreases and new jobs are being created, new residents are attracted to the city. We can see that most of these new residents are attracted by the GPE infrastructure, since the share of PT users doubles between Column (1) and Column (2). This

Table 8: Summary of benefits and monetary cost of counterfactual policies

Counterfactual policy	ΔWelfare (%)	ΔPM$_{2.5}$ (%)	ΔCO$_e^2$ emissions (kilotonnes)	ΔDaily km (millions)	Cost per km (million €)	Total cost (million €)
No cars in centre	2.45	-1.46	-329.3	-10.20		
GPE	1.54	-0.30	-86.6	-2.682	100-200	42,000
Tramways	0.27	-0.05	-14.7	-0.455	13-22	2,664
GPE + Tramways	1.69	-0.33	-94.6	-2.929	113-222	44,000

Notes: First four columns obtained from counterfactual simulations. Average welfare gains obtained by taking a share of high-skilled $s_H = 0.631$ and a share of low-skilled $s_L = 0.369$, based on 2018 initial population. CO$_e^2$ emission factors by fuel type are from ADEME, and information on the car fleet is from the Ministry for Ecology. We multiply daily km avoided by corresponding emission factors, and by 218 working days to obtain the yearly CO$_2^e$ emissions avoided. Cost information from IDFM and a Senate report on *Grand Paris Express*, available here.

generates an increase in housing prices, but also lowers the air quality gains, since part of these new inhabitants still make use of their car to commute. These two effects both induce a loss in welfare gains, which even halve for high-skilled workers, as the inflow of like-skilled new residents is larger (+4.1%, as compared to +2.7% for low-skilled workers). But then, allowing for gains from agglomeration by going from Column (2) to Column (3), part of the welfare loss from the in-migration is restored. This is despite a new increase in rents, since agglomeration forces attract even more new workers.

Channels: Local air quality We now turn to the second half of the columns in Table 7. Column (4) switches off the feedback of air quality on both perceived amenity levels and productivity. If feedback on air quality is not taken into account, welfare gains are $(1.79/1.25 - 1 =)$ 43% smaller for high-skilled workers, and $(1.16/0.81 - 1 =)$ 43% smaller for low-skilled workers. Indeed, the latter live in more polluted areas due to heavier car traffic, but they also have a lower valuation of air quality, which renders their benefit from the air quality improvement induced by new metro lines similar to that of the high-skilled. Column (5) switches only the effect on productivity back on, and Column (6) switches only the effect on amenities back on. The results of these last two columns are strikingly similar, which suggests that the amenity channel and the productivity channel play quantitatively comparable roles in explaining the overall effects.

7.6 Summary: local pollution, global pollution and construction costs

We summarise estimated welfare gains, changes in both ambient and global (CO$_e^2$) air pollution and information on construction costs in Table 8. We estimate the change in CO$_e^2$ emissions by multiplying the change in the total number of km commuted from each origin municipality with emission factors that we build at the municipality level. We construct the emissions factors by combining information on the characteristics of the car fleet in each origin municipality, and emission factors by fuel type. We estimate that GPE allows to avoid emitting about 87 kT of CO$_e^2$ per year, that is about 0.7% of the total annual emissions from road traffic of the Paris region. Completely banning cars from the Paris city centre would achieve about twice this number, which does seem large given the low acceptability of such a strict policy, although changing the entire fleet to electric cars would be equivalent, leaving aside non-exhaust emissions.

Finally, let us point out that the planned extensions of the tramway network are viewed as complementary to GPE by local authorities. As such, we provide a last row to Table 8, in which we run counterfactual simulations, using a new PT network that includes both new trams and GPE. The full results of this exercise are provided in Appendix Table B.11. Gains in commuting time being larger than with GPE only, the average welfare effects are also 9.7% higher, at 1.69%. They are of 1.97% for the high-skilled, and 1.27% for the low-skilled, which is slightly more redistributive than in the GPE-only scenario of Table 7, due to the fact that tramway lines have similar effects across skill levels (Table 6). This higher rise in welfare is partly attributable to a stronger improvement in air quality than in the GPE-only scenario (10% larger), driven by a larger take-up of PT, since the network is denser and provides more connections.

8 Conclusion

This paper investigates the welfare effects of PT infrastructure in the Paris region by incorporating a feedback that the existing literature typically abstracts away from: as cars cross neighbourhoods, they generate local air pollution, which in turn shapes residential amenities and local productivity. Our central finding is that this omission is quantitatively large. Ignoring the air quality channel causes welfare gains from PT investments to be underestimated by approximately 40% for both skill groups, suggesting that existing evaluations that abstract from this channel materially understate welfare gains.

First, we provide reduced-form evidence that new tramway lines generate a modal switch from car to PT, improve local air quality, and raise local housing prices, consistent with amenities gain for nearby residents, and with job creation in the vicinity of new stations. Second, we develop and calibrate a quantitative urban model that integrates the air quality channel to account for general equilibrium effects. Our prospective analysis of the Grand Paris Express projects aggregate welfare gains of about 1.5%. However, distributional results show these aggregate gains mask marked heterogeneity by skill. As lower-skilled neighbourhoods get better connected by the new infrastructure, their residents are displaced and replaced higher-skilled ones, who in the end derive higher gains. These findings are most directly relevant to dense cities where PT usage is already substantial.

Further research may apply this framework to other contexts, especially to more highly-polluted and car-reliant cities, so as to provide a comparison with ours. Finally, we believe this framework can also be appropriately extended to include health outcomes more explicitly, so as to quantify the potential associated gains.

References

Ahlfeldt, G. M., S. J. Redding, D. M. Sturm, and N. Wolf (2015). The economics of density: Evidence from the Berlin Wall. *Econometrica 83*(6), 2127–2189.

Allen, T. and C. Arkolakis (2022). The welfare effects of transportation infrastructure improvements. *Review of Economic Studies 89*(6), 2911–2957.

Amini, A., K. Nafari, and R. Singh (2022). Effect of air pollution on house prices: Evidence from sanctions on Iran. *Regional Science and Urban Economics 93*, 103720.

Bagagli, S. (2025). The (Express) Way to Segregation: Evidence from Chicago. Job Market Paper.

Bahlali, M. and Q. Petit (2024). An equilibrium model of city with atmospheric pollution dispersion. *Journal of Mathematical Economics 111*, 102943.

Bayer, P., N. Keohane, and C. Timmins (2009). Migration and hedonic valuation: The case of air quality. *Journal of Environmental Economics and Management 58*(1), 1–14.

Bilal, A. (2023). The Geography of Unemployment. *Quarterly Journal of Economics 138*(3), 1507–1576.

Bou Sleiman, L. (2023). Displacing Congestion: Evidence from Paris. CREST Working Paper No. 2021-03.

Briant, A., M. Lafourcade, and B. Schmutz (2015). Can tax breaks beat geography? Lessons from the French enterprise zone experience. *American Economic Journal: Economic Policy 7*(2), 88–124.

Brinkman, J. and J. Lin (2024). Freeway Revolts! The Quality of Life Effects of Highways. *Review of Economics and Statistics 106*(5), 1268–1284.

Bryan, G. and M. Morten (2019). The aggregate productivity effects of internal migration: Evidence from Indonesia. *Journal of Political Economy 127*(5), 2229–2268.

Carantino, B. and A. Haramboure (2020). A Streetcar Named Opportunity: Can Light Rail Transit Foster Social Integration? Working Paper.

Card, D. (2009). Immigration and inequality. *American Economic Review 99*(2), 1–21.

Cette, G., L. Koehl, and T. Philippon (2019). The Labor Share in the Long Term: A Decline? *Economics and Statistics* (510-511-512), 35–51.

Champalaune, P. (2025). Wages, Agglomeration and Air Pollution. Working Paper.

Champalaune, P. (2026). Cleaner Air for Whom? Air Quality Policy and Persistent Exposure Disparities. Working Paper.

Chang, T. Y., J. Graff Zivin, T. Gross, and M. Neidell (2019). The effect of pollution on worker productivity: Evidence from call center workers in China. *American Economic Journal: Applied Economics 11*(1), 151–72.

Chay, K. Y. and M. Greenstone (2005). Does air quality matter? Evidence from the housing market. *Journal of Political Economy 113*(2), 376–424.

Chen, S., P. Oliva, and P. Zhang (2022). The effect of air pollution on migration: Evidence from China. *Journal of Development Economics 156*, 102833.

Chen, Y. and A. Whalley (2012). Green infrastructure: The effects of urban rail transit on air quality. *American Economic Journal: Economic Policy 4*(1), 58–97.

Combes, P.-P., G. Duranton, and L. Gobillon (2019). The Costs of Agglomeration: House and Land Prices in French Cities. *Review of Economic Studies 86*(4), 1556–1589.

Combes, P.-P., G. Duranton, and L. Gobillon (2021). The production function for housing: Evidence from France. *Journal of Political Economy 129*(10), 2766–2816.

Cosentino, P. (2025). Trade, Commuting, and City Structure. Working Paper.

Dechezleprêtre, A., N. Rivers, and B. Stadler (2019). The economic cost of air pollution: Evidence from Europe. OECD Economics Department Working Paper, No. 1584.

Dekle, R., J. Eaton, and S. Kortum (2008). Global rebalancing with gravity: Measuring the burden of adjustment. *IMF Staff Papers 55*(3).

Deryugina, T., G. Heutel, N. H. Miller, D. Molitor, and J. Reif (2019). The mortality and medical costs of air pollution: Evidence from changes in wind direction. *American Economic Review 109*(12), 4178–4219.

Dingel, J. I. and F. Tintelnot (2026). Spatial Economics for Granular Settings. *Econometrica*.

Duranton, G. and W. R. Kerr (2018). The logic of agglomeration. In *New Oxford Handbook of Economic Geography*. Oxford University Press, Oxford.

Fretz, S., R. Parchet, and F. Robert-Nicoud (2022). Highways, market access and spatial sorting. *The Economic Journal 132*(643), 1011–1036.

Frondel, M., V. Helmers, and S. Sommer (2025). Fostering the acceptance of congestion charges: Experimental evidence for Europe. Ruhr Economic Papers No. 1148.

Garcia-López, M.-À., C. Hémet, and E. Viladecans-Marsal (2017). Next train to the polycentric city: The effect of railroads on subcenter formation. *Regional Science and Urban Economics 67*, 50–63.

Gehrsitz, M. (2017). The effect of low emission zones on air pollution and infant health. *Journal of Environmental Economics and Management 83*, 121–144.

Gendron-Carrier, N., M. Gonzalez-Navarro, S. Polloni, and M. A. Turner (2022). Subways and urban air pollution. *American Economic Journal: Applied Economics 14*(1), 164–196.

Gibson, M. and M. Carnovale (2015). The effects of road pricing on driver behavior and air pollution. *Journal of Urban Economics 89*, 62–73.

Goel, D. and S. Gupta (2017). The effect of metro expansions on air pollution in Delhi. *The World Bank Economic Review 31*(1), 271–294.

Graff Zivin, J. and M. Neidell (2012). The impact of pollution on worker productivity. *American Economic Review 102*(7), 3652–73.

Green, C. P., J. S. Heywood, and M. N. Paniagua (2020). Did the London congestion charge reduce pollution? *Regional Science and Urban Economics 84*, 103573.

Gruhl, H., N. Volkhausen, N. Pestel, and N. aus dem Moore (2025). Air pollution and the housing market: Evidence from Germany's Low Emission Zones. *Journal of Environmental Economics and Management 132*, 103161.

Gutiérrez, G. and S. Piton (2020). Revisiting the Global Decline of the (Non-housing) Labor Share. *American Economic Review: Insights 2*(3), 321–338.

Heblich, S., S. J. Redding, and D. M. Sturm (2020). The Making of the Modern Metropolis: Evidence from London. *Quarterly Journal of Economics 135*(4), 2059–2133.

Heblich, S., A. Trew, and Y. Zylberberg (2021). East-side story: Historical pollution and persistent neighborhood sorting. *Journal of Political Economy 129*(5), 1508–1552.

Herzog, I. (2024). The city-wide effects of tolling downtown drivers: Evidence from London's congestion charge. *Journal of Urban Economics 144*, 103714.

Holub, F. and B. Thies (2023). Air Quality, High-Skilled Worker Productivity and Adaptation: Evidence from GitHub. Working Paper.

Karagulian, F., C. A. Belis, C. F. C. Dora, A. M. Prüss-Ustün, S. Bonjour, H. Adair-Rohani, and M. Amann (2015). Contributions to cities' ambient particulate matter (PM): A systematic review of local source contributions at global level. *Atmospheric Environment 120*, 475–483.

Koster, H. R. (2024). Welfare Effects of Greenbelt Policy: Evidence from England. *Economic Journal 134*(657), 363–401.

Leape, J. (2006). The London congestion charge. *Journal of Economic Perspectives 20*(4), 157–176.

Li, S., Y. Liu, A.-O. Purevjav, and L. Yang (2019). Does subway expansion improve air quality? *Journal of Environmental Economics and Management 96*, 213–235.

Mayer, T. and C. Trevien (2017). The impact of urban public transportation: Evidence from the Paris region. *Journal of Urban Economics 102*, 1–21.

McFadden, D. (1974). The measurement of urban travel demand. *Journal of Public Economics 3*(4), 303–328.

Miyauchi, Y., K. Nakajima, and S. J. Redding (2021). Consumption access and agglomeration: Evidence from smartphone data.

Monte, F., S. J. Redding, and E. Rossi-Hansberg (2018). Commuting, migration, and local employment elasticities. *American Economic Review 108*(12), 3855–3890.

Morten, M. and J. Oliveira (2024). The effects of roads on trade and migration: Evidence from a planned capital city. *American Economic Journal: Applied Economics 16*(2), 389–421.

Morton, C., G. Mattioli, and J. Anable (2021). Public acceptability towards Low Emission Zones: The role of attitudes, norms, emotions, and trust. *Transportation Research Part A: Policy and Practice 150*, 256–270.

Pestel, N. and F. Wozny (2021). Health effects of low emission zones: Evidence from German hospitals. *Journal of Environmental Economics and Management 109*, 102512.

Redding, S. J. (2024). Quantitative Urban Economics. In *Handbook of Regional and Urban Economics*, Volume 6. Elsevier.

Redding, S. J. and D. M. Sturm (2024). Neighborhood effects: Evidence from wartime destruction in London. NBER Working Paper No. 32333.

Severen, C. (2023). Commuting, labor, and housing market effects of mass transportation: Welfare and identification. *Review of Economics and Statistics 105*(5), 1073–1091.

Severen, C. (2026). Evaluating Transportation Improvements within Cities using Quantitative Spatial Models. Working Paper.

Silva, J. S. and S. Tenreyro (2006). The log of gravity. *Review of Economics and Statistics 88*(4), 641–658.

Takeda, K. and A. Yamagishi (2024). The Economic Dynamics of City Structure: Evidence from Hiroshima's Recovery. *CEP Discussion Paper No. 1988*.

Tessum, M. W., S. C. Anenberg, Z. A. Chafe, D. K. Henze, G. Kleiman, I. Kheirbek, J. D. Marshall, and C. W. Tessum (2022). Sources of ambient $\mathrm{pm}_{2.5}$ exposure in 96 global cities. *Atmospheric Environment 286*, 119234.

Tsivanidis, N. (2026). Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio. *American Economic Review 116*(2).

Ubeda, M. (2021). Local amenities, commuting costs and income disparities within cities. GATE Working Paper No. 2038.

Warnes, P. E. (2024). Transport infrastructure improvements and spatial sorting: Evidence from Buenos Aires. Columbia University Job Market Paper.

Weiwu, L. (2024). Unequal access: Racial segregation and the distributional impacts of interstate highways in cities. MIT Job Market Paper.

Wolff, H. (2014). Keep your clunker in the suburb: Low-emission zones and adoption of green vehicles. *Economic Journal 124*(578), F481–F512.

Xie, L., T. Zou, J. Linn, and H. Yan (2024). Can building Subway systems improve air quality? New evidence from multiple cities and machine learning. *Environmental and Resource Economics 87*(4), 1009–1044.

A Appendix Figures

Figure A.1: $PM_{2.5}$ concentration and the Paris region road network

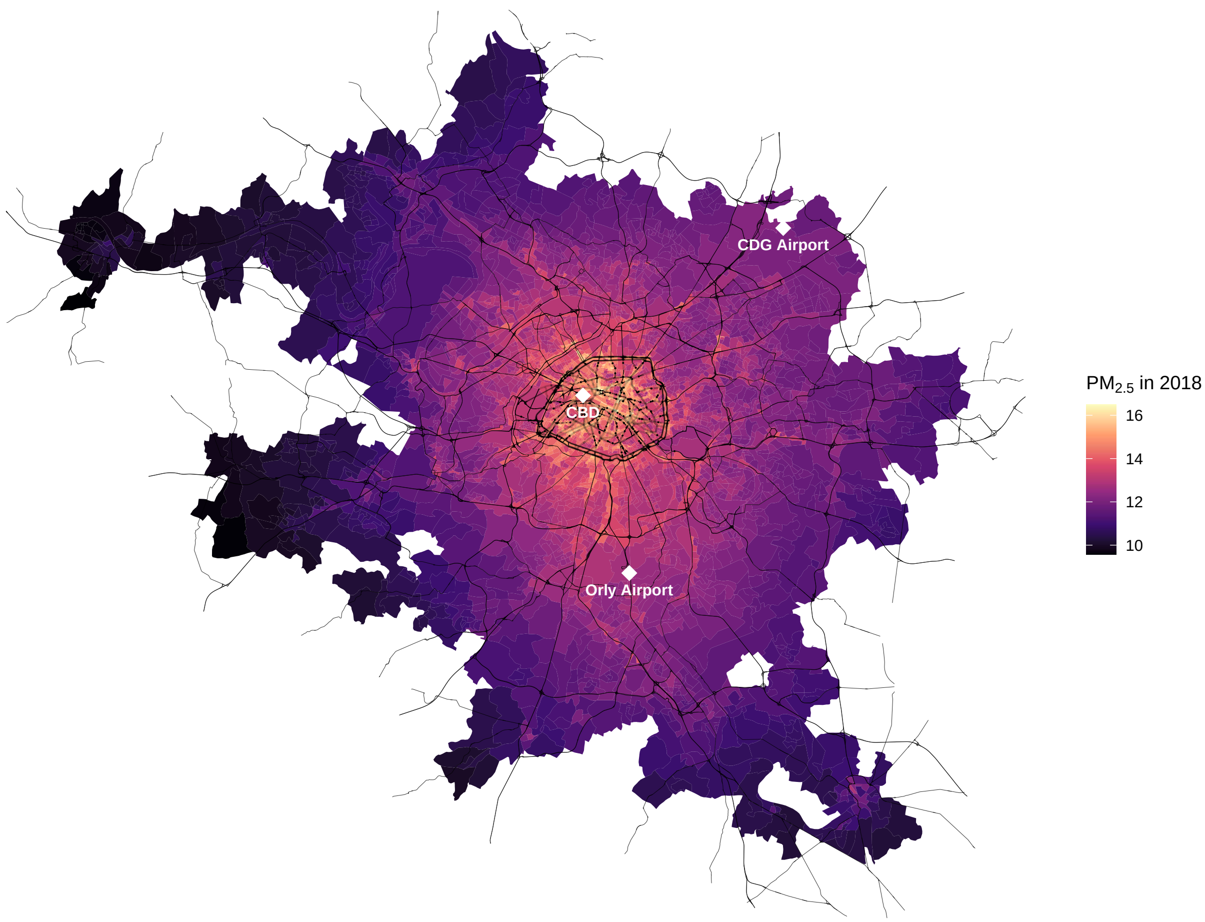

Notes: Authors' own elaboration based on Airparif ($PM_{2.5}$), IGN BD-TOPO (road network) and INSEE (to delineate the urban unit and neighbourhoods) data. Major roads (motorways, trunk roads and main boulevards in the city centre are shown in black. Neighbourhoods (IRIS), our observation level for the reduced-form analysis, are delineated by grey lines. White diamonds show the central business district (CBD) in the 8^{th} district of Paris, and the two main airports, Paris Charles-de-Gaulle (CDG) in the north, and Paris-Orly in the south.

Figure A.2: Neighbourhood definitions: IRIS and IRIS7

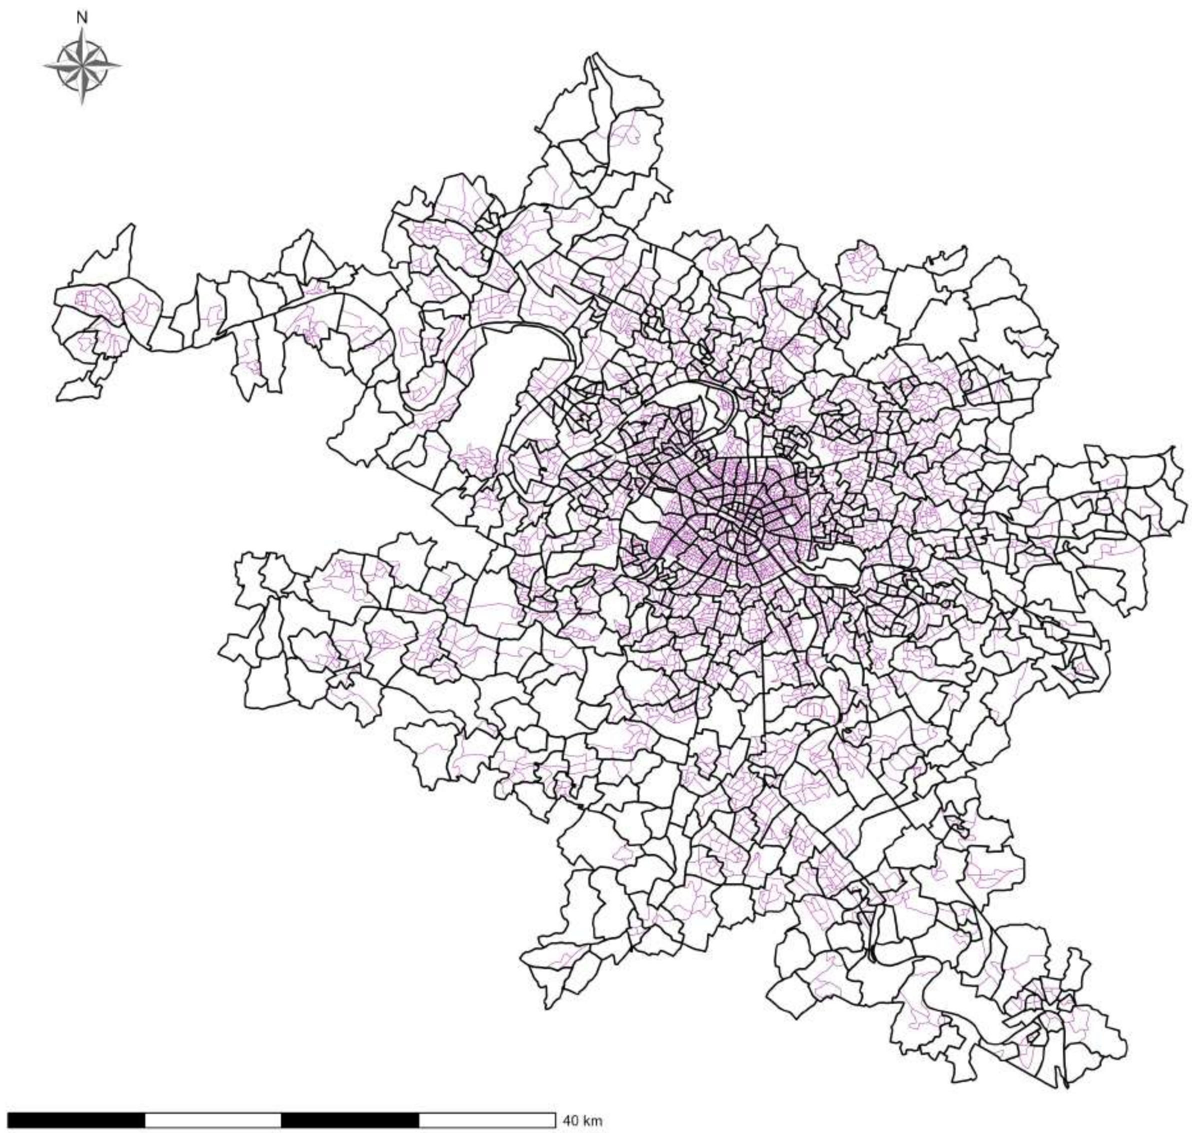

Notes: Authors' own elaboration based on IGN and INSEE data. Purple lines delineate IRIS neighbourhoods, the neighbourhood definition used in the reduced-form analysis. Black lines delineate IRIS7 neighbourhoods (defined using the first 7 digits of the IRIS identifier), the neighbourhood definition used in the counterfactual simulations.

Figure A.3: The Paris region public transport network (2024, excluding bus lines)

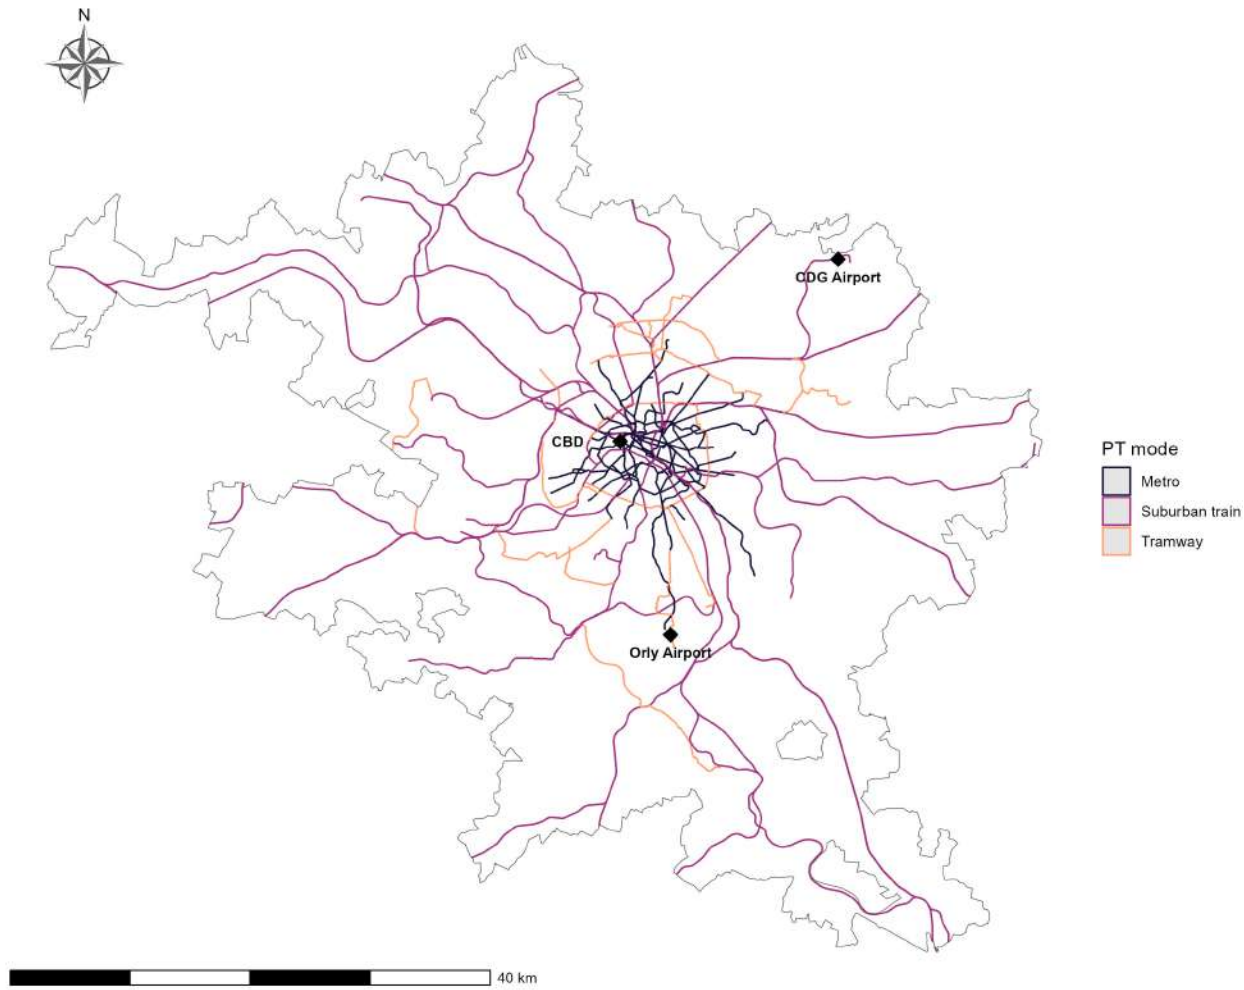

Notes: Authors' own elaboration based on IDFM (public transport network) and INSEE (urban unit contours) data. Black diamonds show the central business district (CBD) in the 8$^{\text{th}}$ district (*arrondissement*) of Paris, and the two main airports, Paris Charles-de-Gaulle (CDG) in the north, and Paris-Orly in the south. Metro lines (in purple) are concentrated in the city centre, suburban train lines (RER and Transilien, in pink) extend far into the outer suburbs, and tramway lines (orange) connect metro and suburban train lines together.

Figure A.4: Neighbourhood composition by distance to nearest rail, tram and bus stations

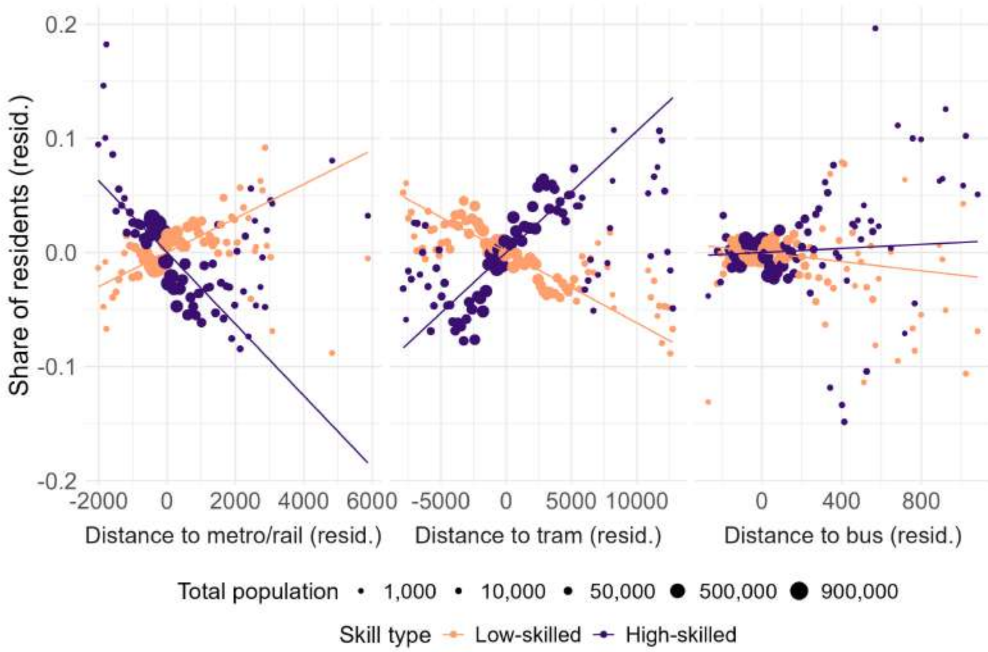

Notes: Authors' own elaboration based on IDFM (public transport stops), INSEE (neighbourhood composition) and IGN (IRIS contours) data.

Figure A.5: Car traffic as a function of worker shares by skill

(a) Raw relationship

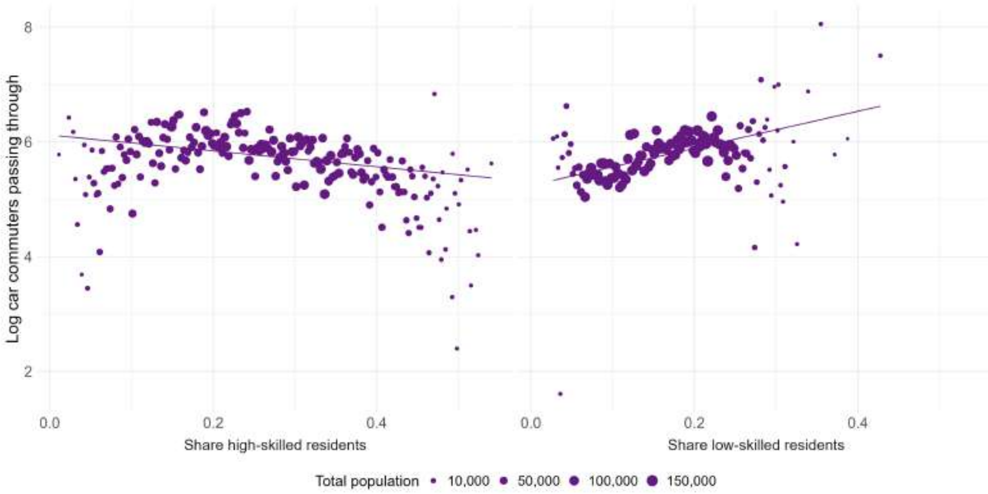

(b) Controlling for distance to CBD

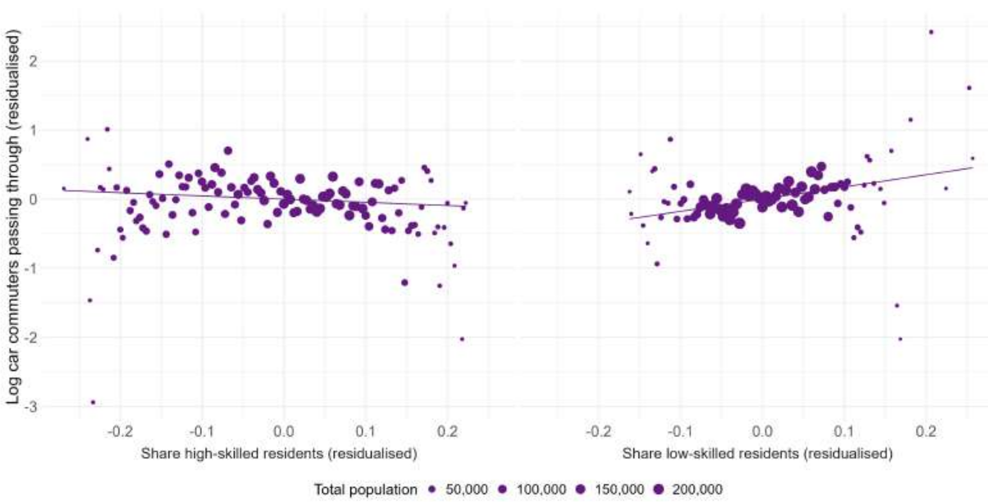

Notes: The number of car commuters passing through a neighbourhood is measured by combining census information on neighbourhood-to-neighbourhood commuting flows and the BD-TOPO road transport network. Panel (b) controls for distance to the CBD, defined as Paris's 8^{th} district. Bins are defined by intervals of 0.25 pp in the share of high- or low-skill workers. Bin size is proportional to total population in the bin.

Figure A.6: Gain in travel time due to new tramway line openings at the origin-neighbourhood level

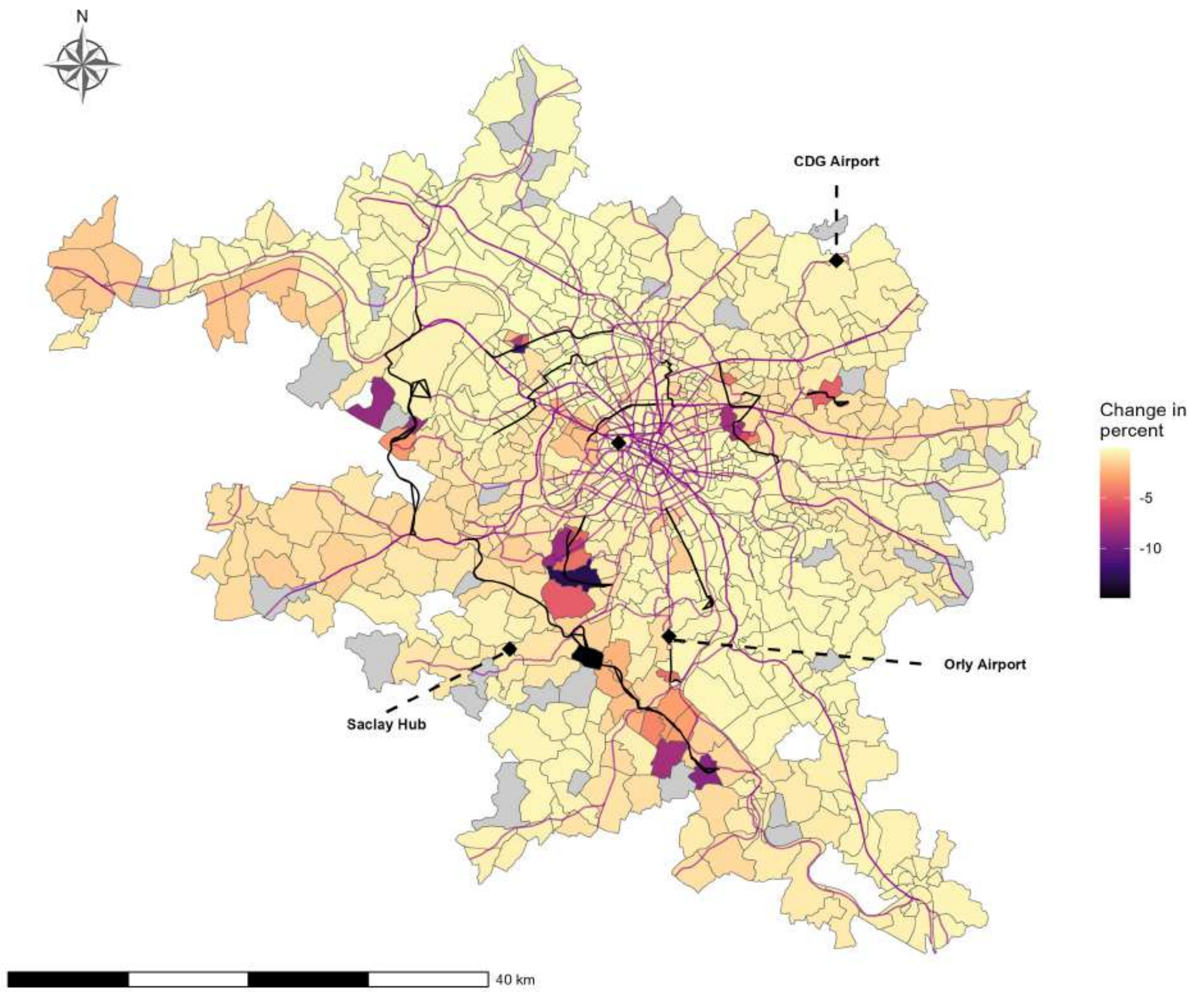

Notes: Authors' own elaboration based on IDFM (public transport network) and INSEE (urban unit contour) data. Existing network of tramway, metro and suburban train lines in purple. New tramway lines in black.

Figure A.7: Variation in main variables due to tramway line openings

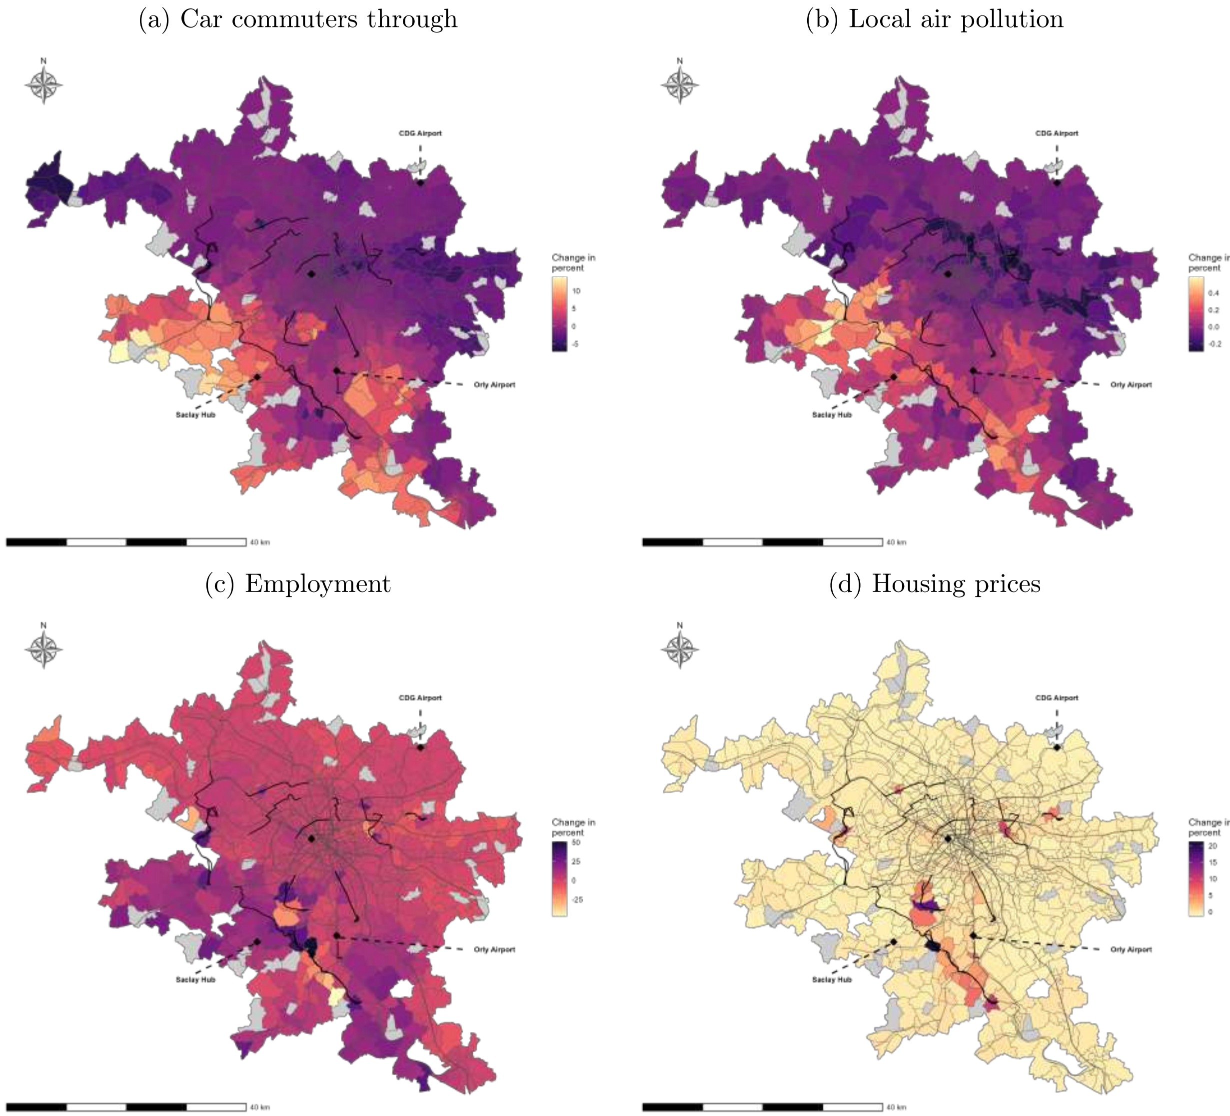

Notes: Variation in the number of crossing car commuters, local air pollution, employment, and housing prices in all neighbourhoods. Counterfactual simulation of the effects of new tramway line openings where all channels are enabled: open city, agglomeration forces, and reaction of amenities and productivity to air pollution. These results are aggregated to the city level in Column (3) of Table 6.

B Appendix Tables

Table B.1: Road typology - From BD-TOPO to the OSRM classification

OSRM/Metric road type	Mean speed (km/h)	BD-TOPO variable values
Motorway	90	NATURE $\in$ { ‘Route à 1 chaussée’, ‘Route à 2 chaussées’, ‘Type autoroutier’ } $\cap$ [CL_ADMIN = ‘Autoroute’ $\cap$ IMPORTANCE = 1]
	85	All above $\cap${Paris = 1}
Motorway link	45	NATURE = ‘Bretelle ’ $\cap$ [CL_ADMIN = ‘NA’ $\cap$ IMPORTANCE = 1]
	40	All above $\cap$ {Paris = 1 $\cup$ URBAIN = ‘Oui’}
Trunk	85	NATURE $\in$ { ‘Route à 1 chaussée’, ‘Route à 2 chaussées’, ‘Type autoroutier’ } $\cap$ { CL_ADMIN = ‘NA’ $\cap$ IMPORTANCE = 1 }
	65	All above $\cap$ {Paris = 1 $\cup$ URBAIN = ‘Oui’}
Trunk link	40	NATURE = “Bretelle” $\cap$ { [CL_ADMIN = ‘NA’ $\cap$ IMPORTANCE = 1 } $\cap$ { CL_ADMIN = “Autoroute” $\cap$ IMPORTANCE $\in$ {2, 3}] }
	30	All above $\cap$ {Paris = 1 $\cup$ URBAIN = ‘Oui’}
Primary	65	NATURE $\in$ { ‘Route à 1 chaussée’, ‘Route à 2 chaussées’, ‘Type autoroutier’ } $\cap$ [CL_ADMIN $\in$ { ‘Nationale’, ‘Autoroute’ } $\cap$ IMPORTANCE $\in$ {2, 3, 4}]
	55	All above $\cap$ {Paris = 1 $\cup$ URBAIN = ‘Oui’}
Primary link	30	NATURE = “Bretelle” $\cap$CL_ADMIN = ‘Nationale’
	25	All above $\cap$ {Paris = 1 $\cup$ URBAIN = ‘Oui’}
Secondary	55	NATURE $\in$ { ‘Route à 1 chaussée’, ‘Route à 2 chaussées’, ‘Type autoroutier’ } $\cap$ [CL_ADMIN = ‘NA’ $\cap$ IMPORTANCE = 2]
	40	All above $\cap$ {Paris = 1 $\cup$ URBAIN = ‘Oui’}
Secondary link	25	NATURE = “Bretelle” $\cap$ CL_ADMIN $\in$ [‘Départementale’, ’NA’] $\cap$ IMPORTANCE = 2
	20	All above $\cap$ {Paris = 1 $\cup$ URBAIN = ‘Oui’}
Tertiary	40	NATURE $\in$ { ‘Route à 1 chaussée’, ‘Route à 2 chaussées’, ‘Type autoroutier’ } $\cap$ [CL_ADMIN = ‘NA’ $\cap$ IMPORTANCE $\in$ {3, 4}]
Tertiary link	20	NATURE = “Bretelle” $\cap$ CL_ADMIN = ’NA’ $\cap$ IMPORTANCE $\in$ {3, 4, 5, 6}
Unclassified	25	NATURE $\in$ { ‘Route à 1 chaussée’, ‘Route à 2 chaussées’, ‘Type autoroutier’ } $\cap$ IMPORTANCE $\in$ {5, 6}
Living street	10	NATURE $\in$ { ‘Route empierrée’, ‘Chemin’ }

Notes: All segments labelled as inaccessible in BD-TOPO are excluded from the road network, i.e., we keep only those for which ACCES_VL is either ‘*Libre*’ (free) or ‘*A péage*’ (toll). We also exclude all segments that are likely impassable by car, or unlikely to be used to commute from home to work, namely those whose ‘nature’ is labelled as: ‘*Sentier*’ (pathway), ‘*Escalier*’ (stairs), ‘*Piste cyclable*’ (cycling path), or ‘*Bac ou liaison maritime*’ (ferry or sea connection).

Table B.2: Price and number of tram stations and IRIS by fare zone

Fare zone	RER stations	Tram stations	IRIS	Treated IRIS	Price in 2011	Price in 2018
1-2	34	40	41	84	62	75.20
3	38	92	97	157	80.30	75.20
4	69	42	42	43	98.10	75.20
5	150	4	4	0	109.90	75.20

Notes: Authors' calculations from IDFM open data.
The price is the price in current euros for a monthly Navigo subscription. For 2011, we show prices for access from zone 1 to the zone of interest of the row (or within zones 1 and 2 for the first row). Less expensive subscriptions were available so as to travel solely, e.g., from zone 2 to zone 5, or from zone 3 to zone 4, and so on.

Table B.3: Descriptive statistics: all IRIS in 2008 and 2018

Variable	2008 Mean	2008 SD	2008 Min	2008 Max	2008 Median	2018 Mean	2018 SD	2018 Min	2018 Max	2018 Median
All IRIS (reduced-form estimations)										
$PM_{2.5}$	16.66	2.45	12.63	25.41	16.02	12.88	1.21	9.65	16.66	12.80
Population	2516.29	791.22	11.00	7269.00	2380.00	2614.63	954.11	106.00	11643.00	2437.50
Workers	1145.65	408.62	9.89	3939.27	1078.33	1153.71	467.76	58.00	6737.00	1079.00
HS residents	336.63	241.10	0.00	1554.00	292.50	384.09	274.81	0.00	3313.00	345.00
LS residents	758.69	326.77	9.00	3163.00	722.50	710.72	347.67	18.25	4192.14	656.52
Potential car traffic	2812.08	3879.53	17.00	44353.00	1584.00	2949.37	4174.29	13.00	49301.00	1601.00
Estimated car traffic	770.77	847.44	0.00	7116.00	486.00	751.47	867.09	0.00	7126.00	440.00
Resident car commuters	91.22	171.57	4.00	4070.00	59.00	99.31	230.97	2.00	6040.00	55.00
Total resident commuters	245.95	277.94	5.00	5741.00	192.00	271.27	361.93	3.00	8047.00	196.00
Resident PT commuters	105.93	114.98	4.00	2455.00	71.00	121.69	145.00	1.00	2969.00	80.00
Share resident car commuters	0.36	0.22	0.01	1.00	0.35	0.35	0.22	0.01	1.00	0.32
Share resident PT commuters	0.40	0.19	0.02	1.00	0.39	0.42	0.19	0.02	1.00	0.41
Log housing price	8.24	0.44	5.73	9.25	8.19	8.28	0.50	6.49	9.26	8.22
All IRIS7 neighbourhoods (structural estimations)										
$PM_{2.5}$	16.10	2.12	12.79	23.33	15.61	12.60	1.13	9.81	15.71	12.54
Population	14465.75	12343.33	172.00	73860.00	10520.00	14997.85	12562.63	171.00	73416.00	10985.00
Workers	6596.59	5832.44	84.74	37422.40	4711.02	6627.09	5769.46	71.00	36192.00	4704.00
HS residents	1954.42	2524.98	15.00	19114.00	992.00	2221.90	2747.65	24.00	20046.00	1222.00
LS residents	4351.77	3631.58	24.00	20404.00	3117.00	4065.65	3332.88	25.43	19371.65	2978.04
Potential car traffic	15969.18	23487.53	57.00	196664.00	7894.00	16681.25	23981.82	50.00	205122.00	8526.00
Estimated car traffic	4370.95	4471.33	11.00	27887.00	2894.00	4247.90	4345.88	36.00	28514.00	2885.00
Resident car commuters	148.43	340.66	4.00	4070.00	69.00	175.88	473.96	3.00	6040.00	71.00
Total resident commuters	304.51	499.39	9.00	5741.00	193.00	354.83	666.73	3.00	8047.00	200.00
Resident PT commuters	106.31	150.68	4.00	2455.00	66.00	130.68	196.97	2.00	2969.00	80.00
Share resident car commuters	0.42	0.22	0.02	1.00	0.42	0.40	0.22	0.01	1.00	0.39
Share resident PT commuters	0.36	0.19	0.03	1.00	0.35	0.38	0.19	0.03	1.00	0.37

Table B.4: Pre-treatment statistics: Not-yet-treated IRIS and treated IRIS (2008)

	Mean (SD)		p-value	Std mean difference
	Control	Treated		
	Full sample: All IRIS within 500m of a (new) tram stop			
$PM_{2.5}$	16.81 (2.33)	17.04 (1.81)	0.229	0.107
Population	2486.21 (702.29)	2628.07 (811.34)	0.043	0.187
Working residents	1086.18 (361.14)	1148.68 (442.67)	0.095	0.155
HS residents	251.38 (196.27)	279.07 (264.90)	0.202	0.119
LS residents	789.89 (283.53)	823.84 (278.75)	0.185	0.121
Estimated car traffic	948.77 (1091.96)	867.43 (1093.05)	0.417	0.074
Total resident commuters	211.43 (174.13)	237.14 (199.45)	0.136	0.137
Resident car commuters	62.75 (60.51)	59.22 (82.00)	0.599	0.049
Resident PT commuters	87.76 (109.14)	115.54 (121.77)	0.009	0.240
% Resident car commuters	0.33 (0.24)	0.26 (0.19)	$<$0.001	0.323
% Resident PT commuters	0.34 (0.23)	0.45 (0.22)	$<$0.001	0.447
Log housing price	8.18 (0.41)	8.20 (0.40)	0.576	0.053
Distance to CBD (km)	11.95 (6.90)	8.95 (2.54)	$<$0.001	0.577
Total surface area (km^2)	0.66 (3.15)	0.34 (0.57)	0.091	0.143
Developable surface area (km^2)	0.41 (0.78)	0.32 (0.54)	0.141	0.131
Number of IRIS	210	284		
	Strict sample: Only IRIS that have a (new) tram stop			
$PM_{2.5}$	16.20 (1.74)	16.84 (1.73)	0.021	0.369
Population	2469.44 (765.39)	2663.83 (836.91)	0.133	0.242
Workers	1083.17 (404.65)	1163.60 (450.94)	0.245	0.188
HS residents	216.24 (158.50)	239.94 (244.88)	0.492	0.115
LS residents	820.70 (307.90)	884.59 (275.23)	0.163	0.219
Estimated car traffic	1472.89 (1263.68)	1421.77 (1457.93)	0.820	0.037
Total resident commuters	244.37 (207.25)	223.62 (185.31)	0.500	0.106
Resident car commuters	78.60 (73.21)	60.01 (51.73)	0.054	0.293
Resident PT commuters	86.02 (95.55)	106.88 (130.65)	0.270	0.182
Share resident car commuters	0.33 (0.21)	0.29 (0.21)	0.253	0.182
Share resident PT commuters	0.33 (0.22)	0.42 (0.23)	0.006	0.445
Log housing price	8.09 (0.40)	8.11 (0.41)	0.690	0.065
Distance to CBD (km)	13.46 (5.99)	9.42 (2.66)	$<$0.001	0.872
Total surface area (km^2)	2.23 (7.78)	0.71 (1.02)	0.047	0.274
Developable surface area (km^2)	0.86 (1.51)	0.68 (1.01)	0.366	0.136
Number of IRIS	63	108		

Table B.5: DiD results: Treated IRIS vs not-yet-treated IRIS, weighted by IRIS population

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				$\Delta_{2018\text{-}2008}$ log housing price				
Treated	0.08***	0.08***	0.08***	0.08***	0.05***	0.05***	0.05**	0.05**
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
(log) workers in 2008		0.04		0.03		0.02		0.03
		(0.03)		(0.03)		(0.02)		(0.03)
(log) housing price in 2008			0.03	0.02			-0.05	-0.06
			(0.04)	(0.04)			(0.04)	(0.04)
2008 mean $PM_{2.5}$	8.13	8.13	8.13	8.13	8.13	8.13	8.13	8.13
Mean outcome	0.031	0.035	0.035	0.035	0.035	0.035	0.035	0.035
R^2	0.050	0.056	0.055	0.059	0.139	0.140	0.147	0.150
				$\Delta_{2018\text{-}2008}$ Share resident PT commuters				
Treated	0.008	0.007	0.05**	0.05**	0.01	0.010	0.05**	0.05**
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
(log) workers in 2008		0.05**		0.07***		0.06**		0.07***
		(0.02)		(0.02)		(0.02)		(0.02)
2008 weighted mean share PT	0.368	0.368	0.368	0.368	0.368	0.368	0.368	0.368
Control 2008 share PT			Yes	Yes			Yes	Yes
R^2	0.001	0.013	0.174	0.201	0.008	0.024	0.201	0.210
				$\Delta_{2018\text{-}2008}$ Share resident car commuters				
Treated	-0.004	-0.002	-0.03	-0.03	0.01	0.01	-0.02	-0.02
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
(log) workers in 2008		-0.04*		-0.02		-0.04*		-0.003
		(0.02)		(0.02)		(0.02)		(0.02)
2008 weighted mean share car	0.316	0.316	0.316	0.316	0.316	0.316	0.316	0.316
Control 2008 share car			Yes	Yes			Yes	Yes
R^2	0.001	0.010	0.169	0.172	0.032	0.040	0.219	0.219
Observations	328	328	328	328	328	328	328	328
Fare zone FE					Yes	Yes	Yes	Yes

Notes: Standard errors clustered at the tram stop level in parentheses. Signif. codes: ***: 0.01, **: 0.05, *: 0.1. Log housing price is an IRIS-year fixed effect from a transaction-level regression of log housing price per square metre on floor area, lot size and a fixed effect for quarter of transaction.

Table B.6: Implicit DiD results: Treated IRIS vs not-yet-treated IRIS, including potentially contaminated controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				$\Delta_{2018\text{-}2008}$ $PM_{2.5}$				
Treated	-0.16	-0.13	-0.04	-0.04	-0.21	-0.21	-0.27**	-0.27**
	(0.31)	(0.30)	(0.20)	(0.19)	(0.16)	(0.16)	(0.12)	(0.13)
(log) workers in 2008		-0.70***		-0.20		-0.04		-0.002
		(0.16)		(0.14)		(0.15)		(0.13)
(log) distance to CBD			1.5***	1.4***			0.78***	0.78***
			(0.13)	(0.13)			(0.20)	(0.20)
R^2	0.003	0.042	0.359	0.362	0.572	0.572	0.610	0.610
				$\Delta_{2018\text{-}2008}$ log housing price				
Treated	0.06**	0.05*	0.05**	0.05**	0.03	0.03	0.02	0.02
	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
(log) workers in 2008		0.09***		0.06*		0.03		0.05*
		(0.03)		(0.03)		(0.03)		(0.03)
(log) housing price in 2008			0.08**	0.06*			-0.09**	-0.10**
			(0.03)	(0.03)			(0.04)	(0.04)
2008 mean $PM_{2.5}$	8.19	8.19	8.19	8.19	8.19	8.19	8.19	8.19
Mean outcome	0.056	0.056	0.056	0.056	0.056	0.056	0.056	0.056
R^2	0.022	0.05	0.050	0.062	0.205	0.207	0.225	0.233
				$\Delta_{2018\text{-}2008}$ Share resident PT commuters				
Treated	0.01	0.01	0.06**	0.06***	0.01	0.01	0.05**	0.05**
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
(log) workers in 2008		0.03		0.06*		0.03		0.05*
		(0.04)		(0.03)		(0.04)		(0.03)
Share PT in 2008			-0.38***	-0.39***			-0.42***	-0.42***
			(0.05)	(0.05)			(0.06)	(0.06)
R^2	0.001	0.005	0.197	0.213	0.007	0.012	0.210	0.222
				$\Delta_{2018\text{-}2008}$ Share resident car commuters				
Treated	0.003	0.004	-0.01	-0.01	0.02	0.02	-0.004	-0.004
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
(log) workers in 2008		-0.03*		-0.03**		-0.03*		-0.001
		(0.02)		(0.02)		(0.02)		(0.01)
Share car in 2008			-0.24***	-0.24***			-0.35***	-0.35***
			(0.03)	(0.03)			(0.04)	(0.04)
R^2	0.001	0.007	0.128	0.136	0.021	0.028	0.202	0.202
Observations	494	494	494	494	494	494	494	494
Fare zone FE					Yes	Yes	Yes	Yes

Notes: Standard errors clustered at the tram stop level in parentheses. Signif. codes: ***: 0.01, **: 0.05, *: 0.1. Log housing price is an IRIS-year fixed effect from a transaction-level regression of log housing price per square metre on floor area, lot size and a fixed effect for quarter of transaction.

Table B.7: Mediation analysis: $PM_{2.5}$ as a mediator of the treatment effect on housing prices

	Estimate	95% CI Lower	95% CI Upper	p-value
Average mediating effect (ACME)	0.014	0.002	0.027	0.036*
Average direct effect (ADE)	0.068	0.026	0.110	0.002**
Total treatment effect (TTE)	0.082	0.044	0.119	0.000***
Proportion mediated	0.171	0.013	0.439	0.036*

Notes: Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1. Sample size used: 301. Number of simulations: 1000. ADE is estimated from the complete model $\Delta \ln(\text{housing\_price}_i) = b_0 + b_1(\text{new tram stop}_i) + b_2 \Delta \text{PM}_{2.5} + \varepsilon_i$. TTE is estimated from the baseline model $\Delta \ln(\text{housing\_price}_i) = c_0 + c_1(\text{new tram stop}_i) + \epsilon_i$.

Table B.8: Implicit DiD results: Treated IRIS vs not-yet-treated IRIS, holding number of commuters constant

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
			$\Delta_{2018\text{-}2008}$ Share resident PT commuters					
Treated	0.03	0.03	0.07*	0.07*	0.02	0.02	0.06	0.06
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)
(log) workers in 2008		0.10		0.13		0.10		0.12
		(0.09)		(0.08)		(0.09)		(0.09)
Share resident PT commuters in 2008			-0.34***	-0.38***			-0.39***	-0.41***
			(0.11)	(0.11)			(0.13)	(0.13)
2008 mean share PT	0.403	0.403	0.403	0.403	0.403	0.403	0.403	0.403
R^2	0.002	0.011	0.033	0.048	0.005	0.014	0.040	0.053
			$\Delta_{2018\text{-}2008}$ Share resident car commuters					
Treated	-0.0003	0.000	-0.01	-0.01	0.01	0.01	-0.006	-0.008
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
(log) workers in 2008		-0.006		-0.007		0.009		0.04
		(0.03)		(0.03)		(0.03)		(0.03)
Share resident car commuters in 2008			-0.16***	-0.16***			-0.31***	-0.32***
			(0.04)	(0.04)			(0.06)	(0.06)
2008 mean share car	0.286	0.286	0.286	0.286	0.286	0.286	0.286	0.286
R^2	0.001	0.001	0.022	0.022	0.010	0.010	0.064	0.068
Observations	328	328	328	328	328	328	328	328
Fare zone FE					Yes	Yes	Yes	Yes

Notes: Standard errors clustered at the tram stop level in parentheses. Signif. codes: ***: 0.01, **: 0.05, *: 0.1.

Table B.9: First-stage results for $\Delta \text{PM}_{2.5}$

	$\Delta \text{PM}_{2.5}$			
	(1)	(2)	(3)	(4)
Δ Upwind car commuters	0.0006***	0.0006***	0.0006***	0.0005***
	(6.4×10^{-5})	(5.9×10^{-5})	(6.34×10^{-5})	(6.14×10^{-5})
Initial $\text{PM}_{2.5}$	-0.4671***	-0.4765***	-0.4632***	-0.4138***
	(0.0081)	(0.0071)	(0.0079)	(0.0067)
Log area	0.2544***	0.2564***	0.2341***	
	(0.0076)	(0.0067)	(0.0087)	
Δ log workers		0.1015***		
		(0.0104)		
Δ log residents			0.1732***	
			(0.0386)	
Log distance to CBD				0.2983***
				(0.0095)
F-test (1st stage)	81.36	72.36	84.47	59.41

Notes: Heteroskedasticity-robust standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1. Examples of first-stage regression results for estimations of the consumption and production disamenity effects of local air pollution.

Table B.10: Projected effects from banning cars from the Paris city centre

	Counterfactual scenario					
	(1)	(2)	(3)	(4)	(5)	(6)
Welfare $\Delta\%$						
High-skilled	4.31	2.0	2.87	-0.58	0.58	1.12
Low-skilled	2.17	1.28	1.79	-0.34	0.74	0.45
λ_{Car} $\Delta\%$						
High-skilled	-28.6	-24.5	-28.7	-20.2	-34.4	-20.3
Low-skilled	-19.0	-16.1	-22.1	-12.5	-30.6	-8.1
λ_{Public} $\Delta\%$						
High-skilled	10.4	17.3	22.6	5.0	15.0	12.1
Low-skilled	5.9	10.1	14.1	2.6	12.5	4.3
(mean) Pollution $\Delta\%$						
Whole area	-1.42	-1.27	-1.46	-1.11	-1.73	-1.05
Paris	-5.58	-5.58	-5.58	-5.58	-5.58	-5.58
Outside Paris	-0.91	-0.74	-0.96	-0.57	-1.26	-0.5
(mean) Rent $\Delta\%$						
Whole area	0.6	2.64	3.7	-0.53	4.56	0.3
Total population $\Delta\%$						
Whole	0.0	5.29	7.62	-1.48	1.94	2.65
High-skilled	0.0	6.11	8.88	-1.74	1.77	3.4
Low-skilled	0.0	3.88	5.46	-1.03	2.23	1.36
Parameters						
Migration elasticity	0.0	3.0	3.0	3.0	3.0	3.0
η^L	0.0	0.0	0.07	0.07	0.07	0.07
η^R	0.0	0.0	0.1	0.1	0.1	0.1
ζ_L^R	-0.025	-0.025	-0.025	0.0	0.0	-0.025
ζ_H^R	-0.09	-0.09	-0.09	0.0	0.0	-0.09
ζ^L	-0.075	-0.075	-0.075	0.0	-0.075	0.0

Notes: Each column reports a separate counterfactual effects of banning cars from the Paris city centre. Counterfactuals take the 2018 equilibrium as the baseline, and consider different structural parameters.
Column (1): closed city, no agglomeration effects, local air pollution effects. Column (2): open city, no agglomeration effects, local air pollution effects. Column (3): full effect, i.e., open city, agglomeration effects and local air pollution effects. Column (4): open city, agglomeration effects, no local air pollution effects. Column (5): open city, agglomeration effects, productivity effect of local air pollution, no amenity effect of local air pollution. Column (6): open city, agglomeration effects, disamenity effects of local air pollution, no productivity effect of local air pollution.

Table B.11: Projected effects from combining GPE and tramway expansions

	Counterfactual scenario					
	(1)	(2)	(3)	(4)	(5)	(6)
Welfare $\Delta\%$						
High-skilled	3.37	1.49	1.97	1.38	1.68	1.72
Low-skilled	1.79	0.99	1.27	0.89	1.1	1.1
λ_{Car} $\Delta\%$						
High-skilled	-10.71	-6.66	-6.82	-8.18	-7.5	-7.4
Low-skilled	-11.27	-8.65	-9.84	-10.6	-10.3	-10.13
λ_{Public} $\Delta\%$						
High-skilled	3.91	8.62	10.72	8.7	9.75	9.87
Low-skilled	3.53	6.63	8.16	6.86	7.62	7.53
(mean) Pollution $\Delta\%$						
Whole area	-0.49	-0.33	-0.33	-0.39	-0.35	-0.36
Paris	-0.53	-0.31	-0.31	-0.4	-0.36	-0.33
Outside Paris	-0.49	-0.33	-0.33	-0.38	-0.35	-0.36
(mean) Rent $\Delta\%$						
Whole area	0.24	1.75	2.46	1.7	2.29	1.94
Total population $\Delta\%$						
Whole	0.0	3.96	5.23	3.64	4.48	4.54
High-skilled	0.0	4.53	6.02	4.19	5.13	5.25
Low-skilled	0.0	2.99	3.87	2.7	3.35	3.32
Parameters						
Migration elasticity	0.0	3.0	3.0	3.0	3.0	3.0
η^L	0.0	0.0	0.07	0.07	0.07	0.07
η^R	0.0	0.0	0.1	0.1	0.1	0.1
ζ_L^R	-0.025	-0.025	-0.025	0.0	0.0	-0.025
ζ_H^R	-0.09	-0.09	-0.09	0.0	0.0	-0.09
ζ^L	-0.075	-0.075	-0.075	0.0	-0.075	0.0

Notes: Each column reports a separate counterfactual effects of tramway lines built since 2018 and the planned GPE metro system. Counterfactuals take the 2018 equilibrium as the baseline, and consider different structural parameters.

Column (1): closed city, no agglomeration effects, local air pollution effects. Column (2): open city, no agglomeration effects, local air pollution effects. Column (3): full effect, i.e., open city, agglomeration effects and local air pollution effects. Column (4): open city, agglomeration effects, no local air pollution effects. Column (5): open city, agglomeration effects, productivity effect of local air pollution, no amenity effect of local air pollution. Column (6): open city, agglomeration effects, disamenity effects of local air pollution, no productivity effect of local air pollution.

C Data Appendix

C.1 Data sources

Table C.1 provides a list of all sources of raw data used in the paper.

Table C.1: List of raw data sources

Purpose	Name of dataset	Producer	Access	Link
IRIS boundaries	Contours...IRIS	IGN	Open	Link
Skill×mode commuting flows	EAR (RP)	INSEE	CASD	Restricted
	Geo-SIRENE	INSEE	Open	Link
Road transport network	BD-TOPO v3.2	IGN	Open	Link
	Metric	INSEE	Réseau Quêtelet	Link
Public transport network	Registry of stops and stations	IDFM	Open	Link
	Planned transport lines in the Paris region	IDFM	Open	Link
	Planned stops in the Paris region	IDFM	Open	Link
	Location of GPE stations	SGP	Open	Link
Housing prices	DVF	Cerema	On demand	Open-data version
Air pollution	$PM_{2.5}$ modelled data	Airparif	On demand	Open-data version
Developable land	CORINE Land Cover 2018	Copernicus	Open	Link
Wages	DADS	INSEE	CASD	Restricted
	Geo-SIRENE	INSEE	Open	Link
CO_2^e emissions	Base Empreinte	ADEME	Open	Link
	Car fleet	MTE	Open	Link

C.2 Details on census data

We use French census data to derive commuting flows by skill and transport mode, at the neighbourhood level, which are at the basis of the estimation of the structural model.

French census data The census is exhaustive, and runs every 5 years, only in municipalities of less than 10,000 inhabitants. However, the vast majority of municipalities in the Paris region are above above this threshold, and are never surveyed exhaustively. Instead, in these municipalities, 8% of housing units are surveyed every year, such that about 40% of the population of these municipalities are present in each five-year census wave, that is, every five years. In the reduced-form section, we keep only workers surveyed in 2008 and 2018 exactly, so as to avoid making wrong conclusions about their commuting behaviour. Indeed, although we compare only these two pre- and post-treatment dates, combining years 2014 to 2019 to retrieve total population would clash with the fact that actual treatment dates (tramway stop openings) are in fact staggered from 2010 to 2017. Combining years of collection would prevent us from properly assigning treatment status. As such, we focus only on one pre-treatment year, 2008, and one post-treatment year, 2018. On the other hand, in the structural analysis, we aggregate years 2015 to 2019 and rely on survey weights provided by the French Statistical Institute so as to recover the total number of workers in 2018.

An additional challenge brought about by the survey collection method is that the collection rate is of about 8% per year at the *municipality* level, which does not guarantee a constant survey collection rate at the neighbourhood level, even if we take observations that are 10 years apart.[37] This is a problem especially for commuter counts, which are in absolute terms; if less people are surveyed in 2018 than in 2008 in a given neighbourhood, there is a mechanical drop in the number of commuters which we should not impute to a change in commuting behaviour. We tackle this issue by dividing the 2018 values by the change in survey rates from 2008 to 2018, thus converting all 2018 variables to the values they would have had the survey rate stayed constant. Finally, due to statistical secrecy concerns, we only exploit residence-IRIS-to-work-municipality flows for which we have at least 10 commuters after weighting. Our commuting flow data is thus "bottom-coded".

Sample selection In both the reduced-form and the structural portions of the paper, we select workers between 16 and 65 years of age, who live and work in the Paris urban unit, and get information on the IRIS they live in, the municipality they work in, their education level, their occupation and the mode of transport they use to commute. We define skill level using information on the last degree obtained: low-skill individuals are those who did not complete secondary education, and high-skill individuals are those who at least got secondary education. Turning to transport mode, we make the distinction between using a car, and using PT. We omit those who go by foot, cycle or ride a motorbike.[38]

Breaking commuting flows down to the neighbourhood-to-neighbourhood level While we do know the IRIS of residence, we only know the *municipality* of work. We tackle this issue by multiplying the flow of workers from each residence neighbourhood/IRIS to each work municipality by the share of employment that each potential work neighbourhood/IRIS represents in the work municipality. To do so, we make use of matched employer-employee data (DADS) managed by the French Statistical Institute, which contains information on occupation and an establishment identifier. Since establishment location is given at the municipality level, and we work at a within-municipality scale, we merge these data with the SIRENE registry, from which we know the exact geographic location of the universe of French establishments, using the establishment identifier. The match rate is 95%. We thus deduce the share of establishments and the share of employment that each IRIS of the work municipality represents, by skill level, and then apply it to our commuting data. We thus end up with commuting flows across all neighbourhoods, broken down by both skill and transport mode.

[37] Since the census is designed in such a way that one census wave corresponds to 5 years of collection, comparing municipalities $5n$ years apart should ensure internal validity.

[38] We do so especially because before 2016, the census makes no difference between cycling and riding a motorbike.

D Exact-hat expressions of relationships

We list here the different expressions for the change in each variable of interest, using exact-hat algebra.

The probability that a type-g worker commutes from n to i using transport mode m changes as follows:

$$\widehat{\lambda}_{nim,g} = \frac{\lambda'_{nim,g}}{\lambda_{nim,g}} = \frac{(\widehat{B}_{n,g}\widehat{w}_{i,g})^{\epsilon_g}(\widehat{d}_{nim}\widehat{Q}_n^{1-\beta_g})^{-\epsilon_g}}{\sum_{k\in N}\sum_{l\in N}\sum_{m'}(\widehat{B}_{k,g}\widehat{w}_{l,g})^{\epsilon_g}(\widehat{d}_{klm'}\widehat{Q}_k^{1-\beta_g})^{-\epsilon_g}\lambda_{klm',g}}. \tag{32}$$

The total amount of housing consumed by workers:

$$\widehat{H}^R_{nm,g}H^R_{nm,g} = (1-\beta_g)\left[\sum_{i\in N}\widehat{\lambda}_{nim|nm,g}\lambda_{nim|nm,g}\frac{\widehat{w}_{i,g}w_{i,g}}{\widehat{Q}_nQ_n}\right]\widehat{R}_{nm,g}R_{nm,g}. \tag{33}$$

The conditional probability for a type-g worker to choose workplace i, given that they live in neighbourhood n and use transport mode m:

$$\widehat{\lambda}_{nim|nm,g} = \frac{\left(\frac{\widehat{w}_{i,g}}{\widehat{d}_{inm}}\right)^{\epsilon_g}}{\sum_{l\in N}\left(\frac{\widehat{w}_{l,g}}{\widehat{d}_{lnm}}\right)^{\epsilon_g}\lambda_{nim|nm,g}}. \tag{34}$$

The amount of commercial floorspace evolves as:

$$\widehat{H}^L_iH^L_i = (1-\alpha)^{1/\alpha}\left(\frac{\widehat{A}_iA_i}{\widehat{Q}_iQ_i}\right)^{1/\alpha}\widehat{L}_iL_i. \tag{35}$$

The amount of residential floorspace evolves as:

$$\widehat{H}^S_nH^S_n = \widehat{Q}_n^{\left(\frac{1-\mu}{\mu}\right)}H^S_n. \tag{36}$$

We compute the variation in amenities for residents of n of type g as:

$$\widehat{B}_{n,g} = \left(\widehat{R}_n\right)^{\eta^R}e^{\zeta^R_g\,\Xi_n\,(\widehat{\Xi}_n-1)}. \tag{37}$$

Similarly, the variation in productivity at workplace i writes:

$$\widehat{A}_i = \left(\widehat{L}_i\right)^{\eta^L}e^{\zeta^L\,\Xi_n\,(\widehat{\Xi}_n-1)}. \tag{38}$$

As a consequence, production of the tradable good varies:

$$\widehat{Y}_i = \widehat{A}_i\left(\widehat{L}_i\right)^{\alpha}\left(\widehat{H}^L_i\right)^{1-\alpha}. \tag{39}$$

The variation in the number of car commuters that pass through a neighbourhood j to go to work generates the change in local air pollution:

$$\widehat{\Xi}_j = e^{\,\theta^F \; F_j \; (\widehat{F}_j - 1)}, \tag{40}$$

with $\widehat{F}_j = \frac{\sum_n \sum_i \lambda'_{ni\ car}\ \mathbf{1}'\left[j \in \text{path}(n \to i)\right]}{\sum_n \sum_i \lambda_{ni\ car}\ \mathbf{1}\left[j \in \text{path}(n \to i)\right]}$.

The change in welfare is the aggregation of the change in type-specific wage income, in type-specific amenities, in commuting costs, in housing prices, and in the probability to for a type-g worker to commute from k to l by mode m, moderated by the Fréchet taste parameter:

$$\widehat{\overline{U}}_g = \Big[\sum_{k \in N} \sum_{l \in N} \sum_{m'} (\widehat{B}_{k,g} \widehat{w}_{l,g})^{\epsilon_g} (\widehat{d}_{klm'} \widehat{Q}_k^{1-\beta_g})^{-\epsilon_g} \lambda_{klm',g}\Big]^{1/\epsilon_g}. \tag{41}$$

Finally, the change in the number of workers of each type g corresponds to the change in utility, moderated by the migration elasticity ϕ:

$$\widehat{L}_{N,g} = \left(\widehat{\overline{U}}_g\right)^{\phi}. \tag{42}$$