Bike-friendly cities: an opportunity for local businesses? Evidence from the city of Paris*

Federica Daniele[1], Mariona Segú[2], Pablo Warnes[3], David Bounie[4], and Youssouf Camara[4]

[1]OECD and Banca d'Italia

[2]CY Cergy Paris Université and IEB, Universitat de Barcelona

[3]Aalto University and Helsinky GSE

[4]CREST, CNRS, Télécom Paris, Institut Polytechnique de Paris

March 17, 2024

Abstract

Cities are increasingly interested in developing cycling infrastructure. Yet, little is known about the (potentially heterogeneous) economic impact of such investments. We evaluate the economic impact of a large-scale cycling infrastructure investment in Paris, the *Plan Vélo*. Using geolocated data covering nearly the universe of French card transactions we estimate a positive and statistically significant elasticity of local revenues to market access. We find a larger elasticity in clusters characterised by a concentration of smaller and younger establishments. The project increased the non-tradables consumption share of central/less densely populated districts at the disadvantage of peripheral/more densely populated ones.

Keywords: cities, cycling, infrastructure investment, local economic activity.

JEL Codes: D12, L81, L83, R2, R4

*Corresponding author: mariona.segu@cyu.fr.

The views expressed in this research are exclusively the authors', and they do not reflect the ones of the Organisation for the Economic Cooperation and Development, nor the ones of Banca d'Italia. This research has been conducted within the Research Chair "Digital Finance" under the aegis of the Risk Foundation, a joint initiative by Groupement des Cartes Bancaires CB, INSEE, Caisse des Dépôts et Consignations, Télécom Paris and University of Paris 2 Panthéon-Assas. We are grateful for the insightful comments received by: Gabriel Ahlfeldt, Raffaello Bronzini, Ana Moreno-Maldonado, Clara Santamaria and Marta Santamaria.

1 Introduction

The use of bikes as a commuting mode is increasingly common in cities, especially European ones. Local governments are ramping up their efforts to develop appropriate supporting infrastructures, such as the construction of large bike networks or the adoption of bike-sharing systems. While there is evidence on the positive impact of developing bike infrastructure on pollution reduction, the one on local economic activity is relatively under-studied. Anecdotal evidence suggests that shop-owners often share the conviction that the construction of bike lanes threatens their business since it may make harder for drivers to park their car, which often brings them to campaign against it.[1]

This article goes beyond anecdotal evidence in analysing for the first time the impact on economic activity of the development of a large-scale bike network during 2017-2019, the *Plan Vélo* in the city of Paris. We leverage the staggered development of the bike infrastructure and the changes induced on bilateral cycling transport costs to construct a time and space-varying firm-level market access metric. We then estimate the elasticity of local revenues to market access using geolocated data covering nearly the universe of French card transactions (provided by Groupement des Cartes Bancaires CB). With this estimated elasticity, we conduct a counterfactual exercise where we estimate the cumulative gain/loss for local businesses from the development of the new infrastructure.

Until recently, the norm in empirical studies evaluating the impact of newly developed transport infrastructure was to take the simple existence in a given location of new infrastructure as the treatment. These studies tended to disregard the network dimension of transport infrastructure benefits, ignoring that the magnitude of the gains depends on the characteristics of the destinations getting connected. In line with the more recent economic geography literature, we estimate the economic impact of the new transport infrastructure through a market access approach, where (firm-level) market access captures the demand each business can reach conditional on each consumer commuting according to its preferred route and means of transport.

We measure market access over a granularly-defined grid with each cell measuring measuring 180-by-180 meters spanning the city of Paris for every quarter during 2015-19. Market access is defined as the weighted sum of non-tradable demand across consumer locations, with weights inversely proportional to the bilateral commuting disutility. The commuting disutility depends on commuting costs through an estimated commuting elasticity, and it is also in

[1] Various case studies show that retailers systematically overestimate the share of customers that arrive by car. See the press article on Bloomberg CityLab "The Complete Business Case for Converting Street Parking Into Bike Lanes", accessed on March 16, 2024.

turn a weighted sum of bilateral commuting costs across transport modes, with the weights directly proportional to the estimated preference consumers place on individual transport modes.

We estimate the elasticity of business revenues at the grid cell level, as proxied by card transaction total revenues, to market access. The development of bike lanes can be endogenous to economic activity in a given location. To address this potential identification threat, we let the commuting elasticity be identified by the development of bike lanes in more distant parts of the network, which we assume to be exogenous with respect to the economic conditions in a given location. This is achieved by means of both controlling for local bike lane density and instrumenting market access with an alternative metric that relies only on changes occurring far away from a given location. We estimate an elasticity of local business revenues to market access between 4.10 and 5.17. More concretely and taking the half-value between these two coefficients, an average improvement in market access, corresponding to the development of 46 meters of bike lanes in a grid-cell, translated into an increase in grid-cell business revenues of 4.31%, amounting to 97,000 euros extra per quarter. Considering the average density of businesses per grid cell, this resulted into an average per business gain of 3,300 euros per quarter. Conversely, the elasticity on average revenues per transaction is found to be close to zero, thus implying that the network development did not have a statistically significant impact on the value of individual transactions, but instead increased the number of transactions for the businesses that were affected.

Business revenues in a given location could increase both because that location has become easier to commute to and thus it starts receiving more visits ("connectivity" channel), but also because it is part of a street where it has become more pleasant to shop, given for instance the presence of a new bike lane and the introduction of complementary pedestrianization measures ("amenity" channel). Thanks to the introduction of the local bike lane density control, our estimated elasticities correctly identify the first channel.

We implement a large battery of robustness checks to eschew endogeneity concerns. First, we address potential concerns driven by the so-called "centrality bias", namely the bias arising from the fact that certain places, owing to their central location, end up being systematically more than proportionally impacted by infrastructure development (Borusyak and Hull, 2020). We do so by removing central locations from our sample, where centrality is defined either in geographic terms or in a network connectivity sense. Second, we show the absence of pre-trends in our outcomes of interest, thus ruling out the endogenous placement of new bike lanes development. Similarly, we do not find evidence of systematic predictive power in the observables with respect to the timing of development, thus ruling out a potential bias driven by the endogenous timing of new bike lanes development. Third, we take advantage of the presence of a set of bike lanes that were part of the original plan but that by the end of 2019 had not been developed yet, to show that our results are robust when considering a smaller, more homogeneous sample of locations. Fourth, we check that our results do not appear to be driven by substitution between card and cash payments. Last, we implement a final large set of miscellanea robustness checks that might impair the interpretation of our main results.

Next, we investigate the existence of heterogeneous effects by dividing the city into clusters with similar business characteristics. We find that the positive and statistically significant elasticity of local business revenues to market access seems hinges on the presence of younger businesses and/or food-related industries, such as cafes, fast-food or bars. Thanks to this evidence, we are able to advance the hypothesis that younger and newer business locations getting access to the new infrastructure become more salient to the eyes of potential consumers and start receiving for this reasons more visits. Cycling arguably allows consumers to be aware of available shopping opportunities better than car ridership or public transport, since cyclists travel on surface and they can more easily stop to visit a store along their ride. Being pedestrianization often a by-product of the development of bike lanes, this mechanism is amplified by demand-side "footfall" externalities kicking in when consumers rely on active mobility as a transport mode for retail-related shopping trips (Koster et al., 2019).

We then explore the elasticity of further outcomes to market access. First, we do not detect a statistically significant positive elasticity of firm entry to market access, which would be consistent with a model of monopolistic competition and free entry. A possible explanation for this result is that commercial rents increased following the rise in demand brought about by the new infrastructure. Unfortunately, we are unable to test this hypothesis due to lack of data. Second, we test the impact of market access on house prices and again, do not detect any significant effect. We argue that any potential effect on prices should come via the amenity effect and less so from the connectivity effect. Since our setting focuses on the latter and accounts for the amenity effect through the local bike lane density measure, it is reasonable to observe a non-significant impact on house prices. Third, we find a negative impact of market access on the total number of cars transiting across a given grid cell. This can happen if car usage goes down once cycling becomes relatively more attractive. An alternative explanation is that bike lane deployment usually implies the removal of car lanes or the introduction of speed limits, which tend to disincentives car usage.

In the last part of the paper, we identify what are the locations that have gained and those that have lost in relative terms from the development of the new infrastructure. We find that the pre/post difference in total revenues generated exclusively by the variation in market access has been positive for 45% of cells in our sample and negative for the remaining others. The uneven impact of the infrastructure can be explained by analysing the characteristics of the locations that gained from the development of the bike network. Locations where market access improved tend to be 1) centrally located, and 2) characterised on average by lower purchasing power (or income). Our analysis highlights that an important consequence of the development of Plan Vélo was to increase the non-tradables consumption share of central/less densely populated districts at the disadvantage of peripheral/more densely populated ones.

Relation with the literature This paper contributes to several strands of the economic literature. First, we relate to the body of literature that looks at the link between transport infrastructure and economic activity. The first wave in this strand of literature explores this link by looking at the impact of getting access to a new infrastructure on outcomes such as employment (Duranton and Turner, 2012), population (Baum-Snow, 2007; Gonzalez-Navarro and Turner, 2018) and property prices (Gibbons and Machin, 2005; Billings, 2011). A more recent set of papers has adopted a market access approach with respect to the exploration of this link. The aim of this approach is to account for the fact that gains from transport infrastructure investment depend on whether locations get connected to more or less attractive places (Ahlfeldt et al., 2015; Heblich et al., 2020; Gorback, 2020; Tsivanidis, forthcoming; Warnes, 2024). With this paper, we contribute to this literature by providing the first empirical assessment of the economic impact of a cycling infrastructure.

This work is also related to a recent strand of the urban economics literature measuring the geography of consumption in the non-tradables sector by means of large-scale spatial datasets, such as the one used in this paper, i.e., card transactions data. Some examples are online review data (Davis et al., 2019), mobile phone data (Athey et al., 2018; Miyauchi et al., 2021) and card transactions data (Relihan, 2017; Allen et al., 2020; Agarwal et al., 2017; Diamond and Moretti, 2021; Bounie et al., 2023). We contribute to this literature by exploiting a high-frequency geolocalised card transaction-level dataset to measure local economic activity that covers the near totality of card transactions made by French citizens. This provides us with a level of representativeness that cannot be found in similar studies.

In evaluating the economic effects of a major green urban policy, we contribute to the strand of the environmental economics literature assessing the impact of pollution reduction policies in cities (see Currie and Walker (2019) for a summary of the literature). Some recent papers have assessed the impact of car usage restrictions in the center of cities on congestion (Sleiman et al., 2021; Tassinari, 2021) and on economic activity (Viard and Fu, 2015; Galdon-Sanchez et al., 2021). With this paper, we evaluate the consequences of a different pollution reduction policy: the development of a large-scale bike network.

The remainder of the paper is organised as follows. Section 2 describes the development of the new layer of bike lanes in the city of Paris in 2015. Section 3 presents the conceptual framework employed in the analysis. Section 4 describes the data and the construction of the market access measure. Section 5 details the empirical strategy used to identify the elasticity of local economic activity to market access. Section 6 discusses the results. Finally, Section 7 concludes.

2 The *Plan Vélo*

A sizeable expansion of the cycling network favoring an even more decisive shift towards active mobility was the priority of the administration that took office in 2015. The initiative was labeled *Plan Vélo* and it consisted of about 80 km of new bike lanes, for a total investment of 150 million euros (Mairie de Paris, 2015).[2] The new plan was organised around two main axes (North-South and West-East), and a series of large routes that were set to become the main arteries of the new network (Boulevard Voltaire, Haussmann, Avenue Friedland, the Quais de Seine and Rue de Rivoli).

The Plan Vélo got off to a slow start: only 4% of bike lanes were constructed from 2015 to 2017. In February 2017, an independent observatory, the *Observatoire du Plan Vélo de Paris*, was set up with the purpose of monitoring the advancement of the development of the Plan.[3]

Once it took off, however, the development of the Plan took place relatively quickly, with 57 km of bike lanes (71% of the original total length) developed between July 2017 and November 2019 - the last month included in our sample. Figure 1 represents different development stages of the plan. The transformation of Paris into an increasingly bike-friendly city is reflected into the swift increase in bike usage, which increased at the average monthly growth rate of 15% during 2018-2019 (Figure 2).

The construction of the network was coordinated at the central level, and it thus left

[2]The re-elected administration has launched in 2021 a Plan Vélo - stage II (2021-2026) with an increased budget of 250 million euros, thus keeping up with its ambition to transform Paris into a European capital of sustainable transport.

[3]See the press statement about the observatory launch here: https://parisenselle.fr/wp-content/uploads/2017/02/PeS_Observatoire-Plan-Velo_Presse_14022017.pdf.

little leeway to district mayors to steer the development towards their places of interest, thus relaxing potential concerns about the endogeneity of location and/or timing. The timing of development appeared rather to follow technical criteria that were independent of economic activity, such as, for example, the decision to develop first the areas located close by the two main axes (North-South and West-East), or the bike lanes that had a longer total extension. We argue that these two features of the development process are amenable from the point of view of our identification strategy, which relies on variation of development status and timing across different parts of the city.

3 Conceptual framework

We use the partial equilibrium model laid out in Gorback (2020) augmented with a modal choice block as the guiding framework of our analysis. Assume that there are $j \in J$ locations, each populated by R_j residents. A resident of location j maximises Cobb-Douglas utility by choosing how much to consume of a housing (h_j), tradable (c_j), and non-tradable (n_j^i) good, and in which location i to purchase the latter:

$$max_{n_j^i, c_j, h_j} \left(\frac{h_j}{\beta}\right)^{\beta} \left(\frac{c_j}{\alpha}\right)^{\alpha} \left(\frac{n_j^i}{1-\alpha-\beta}\right)^{1-\alpha-\beta} \frac{z_{ij}}{d_{ij}} \tag{1}$$

$$s.t. \quad I_j = q^j h_j + c_j + p^i n_j^i \tag{2}$$

where I_j is income for residents of location j, q^j is the price for the housing good in location j and p^i is the price for the non-tradable good sold in location i.

The choice on where to purchase the non-tradable good depends on an idiosyncratic preference term, z_{ij}, which is Fréchet distributed, $F(z_{ij} = e^{-E_i z_{ij}^{-\varepsilon}})$, with E_i being a destination-level amenity parameter and ε governing the substitutability between alternative consumption destinations. Secondly, it depends on the disutility of commuting from home j to the shopping destination i, d_{ij}.[4]

After having decided where to go shopping, consumers choose how to get there. The setup of the modal choice problem follows closely Tsivanidis (forthcoming). Consumers become car owners according to a Bernoulli probability, with expected value ρ. They can choose to

[4]While we acknowledge that shopping trips might have other locations as origin, such as the workplace location, our data do not allow us to differentiate revenues depending on trip origin, and thus we assume that shopping trips depart from the home location of consumers. Nonetheless, we do not see this as a major limitation of our analysis since non-commuting trips tend to be fairly concentrated around the home location of consumers (Miyauchi et al., 2021).

commute via a private transport mode (conditionally on owning a car), $\mathcal{M}_{Private} = \{\text{Car}\}$, or via a public one, $\mathcal{M}_{Public} = \{\text{Walking, Public Transport, Cycling}\}$. They have idiosyncratic preferences over their preferred mode of transport. The disutility of commuting via transport mode $m \in \{\mathcal{M}_{Private} \cup \mathcal{M}_{Public}\}$ is by $d_{ijm} = \exp(\kappa t_{ijm} - b_m + v_{ijm})$, where t_{ijm} is the time it takes to go from j to i via transport mode m, κ is the disutility of commuting elasticity to transport costs and v_{ijm} is a mode-specific idiosyncratic preference shock. Following McFadden (1974), preference shocks are drawn from a Generalized Extreme Value (GEV) distribution:

$$F(\mathbf{v}) = 1 - \exp\left(-\sum_k \left(\sum_{m\in\mathcal{M}_k} \exp((v_{ijm} - b_m)/\lambda_k)\right)^{\lambda_k}\right) \quad k \in \{Private, Public\}$$

The parameter λ_{Public}, or simply λ since by construction$\lambda_{Private} = 1$, allows for correlation within the public transport modes nest, with the correlation increasing as $\lambda \to 0$. Finally, b_m are mode-specific common preference shifters.

Expected disutility from commuting prior to the realization of the idiosyncratic commuting preference shocks is given by $d_{ij} = \exp(\kappa \bar{t}_{ij})$, where:

$$\begin{aligned} \bar{t}_{ij} &= -\frac{1}{\kappa} \ln\left[(1-\rho)\exp(-\kappa t_{ij,0}) + \rho \exp(-\kappa t_{ij,1})\right] \\ t_{ij,0} &= -\frac{\lambda}{\kappa} \ln\left(\sum_{k\in\mathcal{M}_k} \exp\left(\frac{b_k - \kappa t_{ij,k}}{\lambda}\right)\right) \\ t_{ij,1} &= -\frac{1}{\kappa} \ln(\exp(b_{car} - \kappa t_{ij,car}) + \exp(-\kappa t_{ij,0})) \end{aligned} \tag{3}$$

Using Equation 3 into the shopping location choice problem, the probability of purchasing the non-tradable good in location i is:

$$\rho_{ij} = \frac{E_i \exp(-\nu \bar{t}_{ij})}{\sum_s E_s \exp(-\nu \bar{t}_{sj})} \tag{4}$$

The parameter $\nu = \varepsilon\kappa$ identifies the semi-elasticity of commuting flows with respect to bilateral commuting costs and it is a combination of the disutility of commuting elasticity parameter, κ, and the commuting heterogeneity parameter, ε. The probability of commuting from i to j is then increasing in the destination-specific amenity parameter, E_i, and decreasing in the bilateral expected commuting cost, $\bar{t}_{ij}$.

The probability defined in Equation 4 can alternatively be interpreted as the share of residents living in location j that choose to purchase the non-tradable good produced in location i. Hence, in a context where $\bar{t}_{i,j}$ declines on average for all origin/destination pairs thanks to the development of new infrastructure, the expression in Equation 4 entails that a given shopping destination i experiences an increase in expected revenues accruing from consumers living in grid cell j only if $\bar{t}_{i,j}$ drops more than $\bar{t}_{i',j}$ with $i' \neq i$. Importantly, since consumers can shop in one and one only location, a generalised decline in $\bar{t}_{ij}$ does not entail an increase in aggregate expenditure, but rather an adjustment of expenditure towards locations that become easier to reach in relative terms.

By multiplying the expression in Equation 4 by location j's total number of residents, R_j, and their income spent on the non-tradable good, $(1-\alpha-\beta)I_j$, we obtain a proxy of firm-level market access:

$$MA_i = (1-\alpha-\beta)\sum_j \frac{E_i \exp(-\nu \bar{t}_{ij})}{\sum_s E_s \exp(-\nu \bar{t}_{sj})} \times R_j \times I_j \tag{5}$$

In what follows, we rely on Equation 5 to guide the construction of a market access indicator in our setting.

4 Data

In this section we describe the different data sources employed in this paper.

4.1 Geographical unit of analysis and time-frame

Our geographical unit of analysis are equally-sized squared grid cells (used interchangeably with cells from here onwards) covering all the territory of the city of Paris. We use the 9 arcseconds grid of the European Commission Global Human Settlement Layer project which divides the city into 2230 cells of approximately 180×180 meters (Schiavina et al., 2019). Of these 2230, we select 1787 cells with less than 75% of green surface, thus dropping those falling within Paris' two urban forests perimeter. We further drop from the sample cells where we do not consistently record economic activity in the transaction-level dataset during the 2015-2019 analysis period. This brings down the final sample to 1418 final cells. All our variables are computed at the grid cell level. In the case of data initially available for different geometries, we use weights to report the information at the grid cell level.[5]

The observation period we consider goes from January 2015 until November 2019. This time frame allows for two years of pre-period before the development of Plan Vélo started in mid-2017. We end our analysis in November 2019 because of 1) the disruptions to pub-

[5]Specifically, we construct weights based on the overlapping surface between polygons and grid cells.

lic transport caused by the national strike in December 2019, 2) the COVID-19 pandemic in March 2020. Both these events strongly impacted the mobility of Parisians and their consumption behavior (Bounie et al., 2023). The analysis is conducted at the quarterly one.

4.2 Card transaction data

Local economic activity is hard to measure at a very fine geographical scale. We use the best available dataset for the task presented in this paper, namely card transaction data. This type of data is becoming increasingly popular in the economics literature (Relihan, 2017; Allen et al., 2020; Miyauchi et al., 2021) since it offers the advantage of a high time frequency and geographical granularity.

Our dataset on card transactions comes from *Groupement des Cartes Bancaires CB* (CB), a consortium including the near totality of French banks created in 1984.[6] We have information at the merchant-month level for the period ranging from 2015 to 2019. Groupement des Cartes Bancaires CB collects each month the value and volume of transactions made via CB cards, i.e., cards issued by banks part of the CB network. As of 2019, there were 77 million cards in use in the CB system, and 1.8 million CB-affiliated French merchants (Cartes Bancaires, 2019). Figure 3 displays the evolution of quarterly nominal total value of transactions recorded on CB payment system during the period 2015-2019. The growth of nominal total value is related also to increasing card usage in retail payments and substitution from cash. In one of our robustness checks, we show however that increasing card usage and substitution from cash does not appear to threaten our analysis, and the generalisation of our findings to all payments, via both card and cash.

CB data contain the merchant business identification number (SIRET code) that allows us to match them to the French business registry (SIRENE). For each merchant, we thus have the date of creation, the sector of activity (NAF code) and the exact geographical location. For our analysis, we keep merchants located within the city of Paris and operating in the following sectors: retail commerce, restaurants, accommodation services, travel agencies, personal services, bakeries, sports clubs, cinemas and theaters.[7] Our final sample comprises 67,230 unique SIRET codes, accounting for 61% of total card activity. We measure the coverage of our dataset by calculating for each industry the ratio between the nominal total value of transactions recorded in the (full) CB dataset and nominal total value added ac-

[6] In 2020, *Groupement des Cartes Bancaires CB* had more than 100 members (including payment service providers, banks and e-money institutions).

[7] The sectors' NAF codes are: 47 (retail commerce), 56 (restaurants), 79 (travel agencies), 55 (accommodation), 96 (personal services), 1071 (bakeries and pastry shops), 9312Z (sports club), 5914Z (cinemas), 9001Z and 9004Z (theaters and shows).

cording to the national accounts. For the three largest non-tradables industries employed in this analysis (i.e., retail commerce, restaurants and accommodation), this ratio is well above 50%, which suggest that our dataset provides a good coverage of total economic activity.[8]

Our outcomes of interests are total revenues, transactions' volume and average revenues p/transaction. We collapse them at the grid cell/quarter level by taking averages and we subsequently log-transform them.

4.3 Measuring firm-level market access

The following steps describe the construction of our key independent variable, firm-level market access, MA_i, i.e., the size of potential demand in shopping destination i (Section 3).

Step 1: getting information on Plan Vélo — The first step consists of gathering information on Plan Vélo from the *Observatoire du Plan Vélo de Paris* (see Section 2), which has maintained since July 2017 a geolocalised repository keeping track of the daily development of the plan.

Step 2: getting mode-specific bilateral travel times — We then take each of the squared grid cells that form our basic unit of analysis and estimate the bilateral travel time in minutes between the centroid of each pair of cells by mode of transport. To do so, we combine three sources of data: i) the information gathered in step 1 on the development of the *Plan Vélo*; ii) the historical snapshot of the OpenStreetMap project (OSM) every quarter from 2015 to 2019[9]; iii) and General Transit Feed Specification (GTFS) files[10] for the city of Paris for each quarter of from 2015 to 2019, which were provided by the Parisian public transit operator, RATP group.

We calculate the bilateral travel time matrix (in minutes) between all centroids of the cells in our analysis for public transit by combining the OSM data and the GTFS files for the city of Paris and applying the *travel_time_matrix()* function of the *Rapid Realistic Routing with R5* package for R.[11] This packages uses Conveyal's R5 Routing Engine to calculate realistic travel times that allows us to incorporate multiple forms of public transit as well as walking within the same trip. For all public transit travel times, we fix the departure time at 17:00 hs, or the closest time available after 17:00 hs.

[8] Further details about the representativeness of the sectors can provided upon request.

[9] We take the information contained in OpenStreetMap on the first day of February, May, August, and November of every year from 2015 to 2019.

[10] GTFS files are an Open Standard system used to distribute relevant information about transit systems. They include all timetables for a public transit system, the location of all bus stops and metro stations, among other relevant information.

[11] See Pereira et al. (2021) for more details on this package.

In order to calculate the travel time matrices for the other modes of transit (walking, cycling, and driving), we use a different approach that gives us more flexibility when defining the travel network. We start by extracting the network of all *ways*[12] that are traversable by given a mode of travel (i.e, driving, walking or cycling) from the OSM data. We then assign travel speed along each of these ways for each mode. For the driving network, we assign the speed limit (in km/h) of each street segment, according to the information on OSM for that specific moment in time. For the walking network, we assign a fixed speed of 4.5 km/h for the entire network. For assigning speeds to the edges of the cycling network we rely on information from OSM to classify each edge into six categories. Table 1 summarizes the different categories used. We assign the fastest cycling speed to cycle tracks, which are protected bike lanes that are either off road or provide some form of physical barrier blocking car traffic.[13] The *Plan Vélo* provided this type of cycling infrastructure. For cycle tracks we assign a cycling speed of 16 km/h. For all other categories we adjust the speed downward by an adjustment factor that goes from half the speed (for cycle lanes[14] and residential roads) to 0.17 times the speed of cycle tracks for primary, trunk and motorway roads.[15]

Once we have defined the travel speeds for each edge in each type of network (driving, cycling and walking), we can calculate the time it takes to traverse each edge on the network[16]. We then apply Dijkstra's algorithm[17] on this network to find the minimum travel time between each pair of centroid.[18] We repeat this calculation for every ordered[19] pair of centroids to calculate the full travel time matrix for cycling, walking and driving.

Finally, for the cycling travel time matrix, as the new cycling infrastructure is built, we match the information from the Plan Vélo to the OSM network and transform the roads that overlap with the Plan Vélo into cycle tracks, irrespective of what the OSM data tells us. We do this because OpenStreetMap is a crowd sourced project that might have a small lag in updating the true conditions on the ground. Since we are only taking data from one

[12] A way is defined as any linear feature of a map, such as a road, a sidewalk, a river, etc.

[13] These are equivalent to what we refer to as "bike lanes" when discussing the *Plan Vélo* infrastructure.

[14] Cycle lanes, as opposed to tracks, usually lie within the roadway itself and do not provide any physical separation with traffic. They are usually marked with painted lines and signs on the pavement and may be shared with buses.

[15] These adjustments to the cycling speed follow a similar logic to Broach (2016), that provides generalized cost formulas for cycling, where the cost of cycling depends on the type of infrastructure and the density of traffic, among other factors.

[16] An edge in these networks is defined as an ordered pair of nodes, and a node is created every time two *ways* intersect. Within a city, most of these edges coincide with the intuitive notion of a city block.

[17] To do so, we treat the network as a directed graph, where the weights for each edge are defined by the time it takes to traverse said edge.

[18] We technically match the centroid to the closest node on the graph, so the travel times are to the closest nodes on the graph to the centroids of each cell.

[19] We allow for the travel time from a point A to a point B to be different from the travel time from point B to point A.

point in time to reflect the entire quarter, we find that we can be more precise by utilizing this additional source of information.

With these travel time matrices, we can define the bilateral travel time for each origin i/destination j pair, each mode m, and quarter: t_{ijm}. As the plan gets developed, bilateral travel times by bike on average decline, and more so in the proximity of the new Plan Vélo infrastructure. Figure 4 shows the evolution over time of commuting costs from the city hall to all possible destinations.

Step 3: conditional logit estimation — Next, we convert the mode-specific bilateral commuting costs into disutilities, d_{ij}, by means of a nested conditional logit model (Tsivanidis, forthcoming). Specifically, we rely on the commuting module of the 2018 Census (INSEE, 2018). We retain full-time workers residing and commuting within the city of Paris, for a total of 163.000 trips with non-missing information on the chosen commuting transport mode. Information on the place of residence/work in the survey is available at the district level, for a total of 20 districts, or *arrondissements*. We then aggregate the 2230-by-2230 commuting costs matrices into 20-by-20 matrices for the purpose of matching them with the information contained in the commuting survey. We choose simple bilateral averages, but test the robustness of the estimates to alternative aggregation routines.[20]

The output of the nested conditional logit estimation is displayed in Table 2. We estimate a disutility of commuting elasticity $\kappa = 0.003$: an increase in travel time via a specific transport mode by 10 minutes translates into a 3 percentage points lower probability of choosing to commute via that mode compared to walking (the base category). This elasticity is three times smaller than in Tsivanidis (forthcoming).[21] The smaller size of the municipality of Paris compared to Bogotá could explain the discrepancy: shorter distances reduces consumers responsiveness to differences in travel times across transport modes.[22] The inverse of correlation across mode-specific idiosyncratic preference shocks is 0.041, thus denoting a sizeable correlation across idiosyncratic mode-specific preference shocks. Both the cycling and car-specific common preference shifters are negative, thus implying a preference by consumers for walking, rather than cycling or driving, between two destinations holding travel time constant. Conversely, the common preference shifter for public transport is positive, thus denoting a preference for public transport compared to walking, once again assuming

[20]Specifically, we tried aggregating commuting costs using the median as well as bilateral population-weighted means.

[21]The conditional logit estimation relies on work-related trips. However, recent work by Miyauchi et al. (2021) finds the commuting elasticity estimated based on consumption-related trips to be higher than the one based on work-related trips.

[22]The modal shares are: 71% by public transport, 7% by bike, 16% by walking and 6% by car.

identical travel times.

The parameters in Table 2 are combined with information on the car ownership rate in the city of Paris $\rho = 0.37$ to obtain an estimate of expected bilateral commuting costs $\bar{t}_{ij}$ according to Equation 3.[23][24]

Step 4: estimating the semi-elasticity of commuting flows to commuting costs — Next, we use bilateral consumption flows to estimate ν in Equation 4. While these are not directly observed in our dataset, for the last year in our sample we developed an imputation procedure to calculate a proxy for them. Specifically, for the last year in our sample, we observe daily transactions indexed by the merchant and card identifier. We do not know where the cardholder lives, nor we have any demographic information on him/her. However, we can, based on each cardholder shopping history, impute a "most likely" residence location. We do so by taking the following steps:

1. we retain transactions occurring in the city of Paris on weekends and on weekdays after 18h;

2. we further retain transactions that are usually carried out in the proximity of one's residence, which we identify as those transactions occurring in merchants identified by one of the following sectoral codes: 1071, 1072, 4724 (bakeries), 4773 (pharmacies), 4711B-D (supermarkets, minimarkets), 4721, 4722, 4723, 4725, 4729 (food stores);

3. we finally keep cardholders for which the number of observed transactions is $N \geq 9$.

We end up with a sample of 3.2 million cards, about $1/8^{th}$ of the total number of cards present in the data, but amounting to nearly half (49%) of transactions total value. For this subset of cards, we calculate the modal shopping destination and we set it as "most likely" residence location, j. Next, we calculate bilateral consumption flows, x_{ij}, by summing across transactions carried out by cardholders with imputed residence location j towards merchants with (known) business location i, and estimate:

$$\ln x_{ij} = \alpha_i + \alpha_j + \nu \bar{t}_{ij} + e_{ij} \tag{6}$$

where $\bar{t}_{ij}$ are the just obtained expected bilateral commuting costs, and α_i and α_j are, respectively, business and residence location fixed effects. We repeat this estimation for four

[23] A normalization is implemented to ensure that $\bar{t}_{ii} = 0$. More specifically, $t_{ij,0}$ is rescaled by $-(\lambda/\kappa)\ln(\exp(b_{walking}/\lambda) + \exp(b_{cycling}/\lambda) + \exp(b_{pt}/\lambda)$ and $t_{ij,1}$ is rescaled by $-(1/\kappa)\ln(\exp(b_{car}) + 1)$.

[24] Paris car ownership rate is taken from https://www.lefigaro.fr/economie/le-scan-eco/dessous-chiffres/2017/10/12/29006-20171012ARTFIG00166--paris-la-voiture-est-deja-une-espece-en-voie-de-disp php.

different quarters of 2018, and consistently estimate $\widehat{\nu} = 0.1$, similar to 0.07 in Ahlfeldt et al. (2015).

Finally, we use $\widehat{\nu}$ to build an empirical counterpart to Equation 5. In order to isolate the variation in market access induced exclusively by the fall in travel time by bike, we use pre-plan population and median income, as well as pre-plan travel times by public transport, car and walking while constructing $\bar{t}_{ij,t}$.[25]

Figure 5 portrays the evolution of market access over time. A few considerations emerge from the inspection of this figure. First, market access in places situated in the proximity of a new bike lane rises if the latter connects with the rest of the network. For example, the construction of a bike lanes along the south-west Seine riverbank occurring between 2017q4 and 2018q4 did not trigger an expansion in market access due to their initial lack of connectedness with other parts of the network. Second, market access in given location increases only if the reduction in bilateral commuting costs exceeds the average reduction experienced by other locations. This is a direct consequence of the conceptual framework, according to which consumers shop in one and one location only. For example, the 15^{th} district located in the south-west quadrant of the city experienced a loss of market access during 2018q4 and 2019q4 in spite of the development of connected bike lanes. The loss in market access occurred because other places - most notably the central districts - experienced a more sizeable reduction in bilateral commuting costs.

4.4 Control variables

We assemble a rich dataset of time-varying grid cell level characteristics from multiple sources. We get annual socioeconomic and demographic characteristics from the *Institut national de la statistique et des études economiques* (INSEE) such as total population, population aged between 25 and 39 years old, foreign population, number of job seekers and working age population. Given the lag with which these data are released, these control variables enter our specification with a three-year lag period. These data are at the IRIS geographical level (the equivalent of census tracts in France). In order to report variable x to the grid cell level, we calculate for each IRIS unit i the share of surface overlapping with grid cell j, s_{ij}. Subsequently, we set $x_j = \sum_k s_{kj} x_k$, where k indexes all IRIS units overlapping with grid cell j. See Table 3 for the summary statistics of the different control variables employed in the analysis.

[25]The data refer to 2015 and come from French National Statistical Office (INSEE).

4.5 Other outcome variables

We finally consider a few additional alternative outcome variables, which might reasonably be also affected by the development of the new network. First, we get information on all newly created establishments at any given point in time and their address from the national business registry (SIRENE). We use the *OpenStreetMap* API to geolocalise the addresses and assign the new establishments to their respective grid cell. Since the number of newly created establishments in a given grid cell is for the majority of observations close to zero due to the highly geographically disaggregated nature of our analysis, we replace the number of new firms with a dummy taking value 1 if any entry takes place and 0 otherwise.

Further, we downloaded micro data on the universe of house transactions occurring in the city of Paris during our period of interest (*demandes de valeurs foncières*), containing information on the sale price and house characteristics. These data are geolocalised and we are thus able to directly assign each transaction to their respective grid cell. We implement a simple cleaning procedure aimed at the removal of outliers suggested by INSEE (see Cailly et al. (2019)). Subsequently run a hedonic regression of the log of the sale price on the number of bedrooms, the number of rooms, the number of squared meters, and a set of dummies identifying different housing types.

5 Empirical strategy

In this section, we detail the empirical strategy used to estimate the elasticity of local economic activity to market access. Evaluating the impact of transport infrastructure investments on local economic activity is often challenging given the typically non-random placement of transport infrastructure. A local government might decide that the new infrastructure should cross a declining neighborhood in order to revitalise it. Alternatively, it may choose to develop new infrastructure in an already booming neighborhood, thus supporting its expansion with adequate infrastructure. The endogeneity of transport infrastructure placement might also stem from the initiative of private interest groups, who can succeed at lobbying the local government into developing/not developing a new infrastructure in their neighborhood for their personal gain.

We address these identification concerns by letting our coefficient of interest be identified by variation in market access triggered by bike lane development in distant places (Hornbeck and Rotemberg, 2021). This identification strategy requires to include in our specification a local bike lane density measure. By controlling for the intensity of local bike lane development, the market access elasticity is identified by variation in market access stemming from bike infrastructure development occurring further away. The identifying assumption is that development in distant places is independent of local economic conditions measured in individual places. This strategy also allow us to focus on the *connectivity effect* of market access while controlling for any potential *amenity effect*, which should be captured by local bike lane density.

We measure the elasticity of local economic activity to market access, MA_{it}, by means of the following specification:

$$ln(Y_{it}) = \alpha_i + \alpha_{dt} + \beta ln(\text{MA}_{it}) + \gamma X_{it} + \delta LBLD_{it} + e_{it} \tag{7}$$

where Y_{it} can be 1) total revenues, 2) transactions' number, 3) average revenues p/transaction, all calculated across merchants located in grid cell i at time t. Equation 7 also includes grid cell fixed effects α_i and district×time fixed effects α_{dt}. Further, we control for a set of economic and demographic characteristics (X_{it}) varying at the grid cell and time level. Control variables are (log) population, ratio of foreigners, unemployment rate and (log) population aged between 25 and 39 years old.

We test two alternative measures of local bike lane density, $LBLD_{it}$. Like most infrastructure networks, Plan Vélo is articulated into a series of bike lanes (which we dub as "projects"). In our favorite specification, $LBLD_{it}$ corresponds to the total length of the bike lane or project crossing grid cell i as of time t. We take this as our favorite definition of local bike lane density since we deem likely that development in a given grid cell is influenced by development across cells belonging to the same project. As an alternative measure, we let $LBLD_{it}$ be equal to the total length of bike lanes situated in grid cell i and its neighboring cells as of time t. This second measure accounts for the fact that development in grid cell i and time t might be influenced by economic conditions not only in grid cell i but also in its most immediate neighbors.

Finally, we estimate the results using an alternative strategy where we instrument market access in a given location with measures relying on bilateral commuting costs with more distant locations. The advantage of this strategy is that we do not have to control for local bike lane density, which might be an endogenous control. We use two alternative instruments. First, we run a first stage in which MA is instrumented by a MA measure that relies on bilateral commuting costs with locations that are at least 1km away. Second, we instrument MA with a measure of MA that uses all origin-destination pairs except for pairs located along the same bike lane or "projects". Thus, variation in these instruments is shaped exclusively by bike lane developments occurring farther away in the network. For the first instrument, this is:

$$ln(\mathrm{MA}_{it}) = \alpha_i + \alpha_{dt} + \beta ln(\mathrm{MA}_{it}^{1Km}) + \gamma X_{it} + e_{it} \tag{8}$$

where MA_{it}^{1Km} is defined as:

$$\mathrm{MA}_{it}^{1Km} = \sum_{ij|dist_{ij}>1km} \frac{\exp(-\widehat{\nu}\bar{t}_{ij,t})}{\sum_s \exp(-\widehat{\nu}\bar{t}_{sj,t})} \mathrm{Population}_j \times \mathrm{Median\ income}_j. \tag{9}$$

We then run Equation 7 using the instrumented measure of MA and without including the variable of local bike lane density.

6 Results

6.1 Baseline results

Table 4 presents the estimates of β from Equation 7, for the two measures of local bike lane density and two instruments. We present the results in three panels, one for each outcome: the log of total revenues, the log of transactions volume and the log of average revenue per transaction. The coefficients of market access are positive and significant for the first two outcomes and across all specifications. In column 1, we report the estimated coefficients without including any local bike lane density control. In columns 2 and 3, we include, alternatively, our two measures of local bike lane density. Columns 4 and 5 report the IV results, which are quantitatively similar to the OLS and similarly significant. This suggests that the potential bias caused by the endogeneity is LBLD is less of a concern. Conversely, we do not detect a statistically significant elasticity with respect to average revenues p/transaction, regardless of the local bike lane density control included or the estimation strategy used.

The magnitude of the coefficients are between 4.10 and 5.17. To interpret them, we rely on variation in our market access (MA) measure. Variation in MA is low since it hinges on variation in cycling bilateral commuting costs only, where cycling is not the main means of transport. On average, grid cells that experienced a positive change in MA saw an increase of 0.0093, i.e., by approximately 1%, over the whole period. Hence, the average improvement in MA translates into a 4.81% increase in total revenues and a 4.29% increase in volume[26]. We provide a more detailed quantification of the impact of the infrastructure

[26] We multiply 0.0093 times 5.173 for the first one, and 0.0093 times 4.603 for the second one.

and its distributional consequences in Section 6.5.

We find that the positive impact of an improvement in market access on local economic activity materialises with a delay. In Table 5, the baseline measure is replaced with its lagged value. When we do so, we find the elasticity of total revenues to be positive and statistically significant across all lags. Similarly, the elasticity of transactions' volume is statistically significant across all lags, and it grows in magnitude as lags of market access further back in time are considered.

By the way that our empirical strategy is designed, our estimated coefficients measure primarily the impact on economic activity to improved access to transport infrastructure ("connectivity channel"). However, this may not be the only channel at work: the development of a new bike lane might in fact positively affect the revenues of local merchants also through an amenity channel, by for example improving the exterior looks of a place. These two channels are correlated but not entirely collinear. Consider the example of two places characterised by a main street where in both cases the city council decides to build a spacious bike lane. Both streets become safer and cleaner, because perhaps the city council also implements some speed restrictions for cars and car traffic goes down. As a result, by an *amenity* channel more people might choose to go shopping or eating out in those two places. However, consider now the case of the bike lane connecting with the pre-existing cycling network only in one of them. In this place, the business volume will rise even more, since not only this is a nicer place where to shop, but it is also easier to reach by a *connectivity* channel. In our setup, the local bike lane density control, $LBLD_{it}$, acts to soak up the amenity channel, thus allowing MA_{it} to identify the connectivity one.

6.2 Robustness

In this section, we run a series of robustness tests to address some of the potential issues that might challenge the causal interpretation of the estimated coefficients.

Centrality bias — We test the robustness of our results to the exclusion of highly-connected districts, which might mechanically benefit more from transport infrastructure development (Borusyak and Hull, 2020).[27] In column 2 of Table 6, we follow Chandra and Thompson (2000) and exclude central districts, specifically *arrondissements* 1 to 4. In column 3, we remove *transport hubs*, where connectivity is defined from the standpoint of the public transport network. Using information on the public transport network (metro, tramway, and

[27] Unfortunately, after some examination, we concluded that the centrality bias test proposed by Borusyak and Hull (2020) was not computationally feasible in our context due to the very granular scope of our analysis.

suburban trains),[28] we define as *transport hubs* those grid cells located less than 500 meters away from stations featuring three or more public transport connections.[29] All coefficients in Table 6 remain positive and statistically significant, thus eschewing centrality bias as a concern in our setting.

Non-random bike lane development — We investigate whether places that experienced a larger increase in market access did not feature a statistically significantly different evolution of the outcome variables before any bike lane development took place. To do so, we run a pre-trends analysis by means of the following regression:

$$ln(Y_{it}) = \alpha_i + \alpha_{dt} + \sum_t \beta^t \Delta ln(MA_{i,15-19}) \times \tau_t + \gamma X_{it} + \delta LBLD_{it} + e_{it} \quad (10)$$

where we regress our outcome variables on the (log) change in market access that occurred during 2015q1-2019q4. The (log) change in market access is interacted with time dummies (τ_t) and a full set of time-specific coefficients, β^t, is estimated. According to the evidence shown in Figure 6, places that experienced higher market access improvements started featuring higher levels of economic activity only after development began, thus suggesting that potential non-random bike lane *placement* does not challenge our strategy.

Non-randomness might characterise not only the location of new bike lanes but also the *timing* of development. Thus, we test whether development timing appears to be as good as random based on observable characteristics (Deshpande and Li, 2019). We take the sample of cells left to be developed (and eventually developed) at time $t = t_0$ and regress the development date on the set of demographic variables employed as controls in the main specification:

$$\text{Development date}_i | (D_i^{t_0} = 0) = \alpha + \beta X_i^{t_0} + e_i \quad (11)$$

We run Equation 11 for different choices of $t = t_0$ and report the results of the estimation in Table 7. During the first months of construction it appears that places characterised by a lower population, a lower percentage of job seekers and more young people received access to the bike network faster. In the middle of the period (column 2) only the percentage of job seekers seems to matter. Finally, towards the end our sample period, during which

[28] Data come from Île-de-France Mobilités website.

[29] The stations excluded are: Charles-de-Gaulle Étoile, Châtelet les Halles, Cité, Denfert-Rochereau, Gare Montparnasse, Gare Saint-Lazare, Gare de Lyon, Gare de l'Est, Gare du Nord, Haussmann Saint-Lazare/Havre-Caumartin, Invalides, La Motte Picquet - Grenelle, Magenta, Opéra, Place d'Italie, Porte de Choisy, Porte de Vincennes, Porte des Lilas, République, Saint-Michel Notre-Dame, Strasbourg - Saint-Denis.

most of the construction occurred, population is again associated in a statistically significant way with the development date. We interpret the absence of systematic correlation between the covariates and the development date as supportive evidence of development timing not systematically correlated with local economic conditions.

Exploiting the unfinished Plan Vélo — In a further check, we restrict the sample to cells that were planned to feature some bike lane development according to the original Plan Vélo (see Figure 7). Since all included cells were initially considered to be "in need" of a bike lane, we can expect this subsample to be more homogeneous than the baseline one. We test this hypothesis by means of a balancing test (Table 9), which confirms that cells that received some bike lane development did not differ in a statistically significant way from cells that did not, except for the length of planned bike lanes. This last element suggests that the decision to develop first certain bike lanes was probably also driven by the need to start first with the longer ones, supporting the argument of the quasi-random development timing. We re-run the baseline specification on this subsample and display the results in Table 8. Columns 1 and 3 reproduce the baseline results from column 2 of Table 4, while columns 2 and 4 show the coefficients of the subsample. We observe that the coefficients remain statistically significant and that they are slightly larger than in the baseline estimation. For interpretation, we refer to the average increase in market access for this subsample, which is 0.0125. Hence the average improvement in market access implies an increase of 10.57% in total revenues and an increase of 11.56% in transactions' volume.

Other potentially confounding factors — In Table 10, we conduct a set of further robustness tests. In column 2, we control for lagged sectoral shares to make sure that our results are not driven by changes in local economic activity composition.[30] Second, in columns 3 and 4, we test the robustness of our results to a law passed in 2015 that allowed businesses located in some places to stay opened on Sundays.[31] Since our card transaction dataset is available at the monthly frequency, we cannot exclude transactions carried out on Sundays and directly control for potentially endogenous self-selection into this policy. Instead, we first remove grid cells that were at all affected by the Sunday Law (column 3) and second, we include an interaction term between the grid cell-specific share of surface

[30]Specifically, we include the lagged share of revenues for each grid cell and time in the following non-tradables subsectors: non-specialised retail stores (Code NAF 471), specialised food retail stores (Code NAF 472), specialised non-food retail stores (Code NAF 474-477), fast food restaurants/bars (Code NAF 561), restaurants (Code NAF 562), bars specialised in the sale of drinks (Code NAF 563).

[31]See the map of concerned places in Figure 8. Data come from APUR, Mairie de Paris and DRIEA IF/UD75.

concerned by the law and time dummies (column 4). Across all tests, the elasticity of economic activity to market access remains positive and statistically significant.

Test for card usage — Finally, we test if the increase in revenues in places with higher market access improvement is partially driven by an increase in cards usage and substitution away from cash payments. To do so, we check if market access increases the share of establishments using cards. First, we build a card usage intensity index by dividing the number of establishments present in the *Cartes Bancaires CB* dataset over the number of establishments that should be active in that same place and quarter according to the business registry. Then, we re-run Equation 7 placing this index on the left hand side of the equation. If the estimated coefficient were to be found positive and statistically significant, this would entail that a market access improvement is associated with an increase in the share of establishments recording card payments, which would weaken our assumption of card payments as representative of all payments - card and cash ones. The lack of statistical significance in the coefficients displayed in Table 11 confirms that this is not the case.

6.3 Heterogeneity

In this section, we explore the heterogeneity in the estimated impact of market access improvements on local economic activity. First, we build clusters of grid cells that display similar local establishments characteristics. First, we select the characteristics on which we run the clustering algorithm: a set of (dummy-based) sub-industry indicators as of 2015 (supermarkets and malls, specialised food retail stores, specialised non-food retail stores, fast food restaurants/bars, restaurants, bars); a size dummy taking value 1 if in 2015 average merchant size in a given grid cell is greater than the median value; an age dummy taking value 1 if in 2015 average merchant age in a given grid cell is greater than the median value calculated across all cells.[32] Next, we run a *k-means* clustering algorithm (Bock, 2007) for different values of the number of clusters, k, and we select $k = 5$ by means of an elbow test as shown in Figure 9.[33]

The characteristics of the five clusters in terms of the variables used for the clustering exercise are reported in Table 12. The degree of specialisation into stores that sell essential goods,

[32]We define a grid cell as specialised in a given industry if the share of revenues coming from that industry is greater than the share of revenues coming from that industry at the city level.

[33]The elbow method is a heuristic method widely used in data science to determine the optimal number of clusters in a dataset. It consists of plotting the sum of squared errors (SSE) calculated across the identified clusters for each selected number of clusters, and then pick the number of clusters k^* such that the average reduction in the SSE obtained by moving from k_{i-1} to k_i for $k_i < k^*$ can be considered substantially larger than the one obtained for $k_i > k^*$, i.e., by looking for the value of k corresponding to the elbow of the curve.

such as specialised food retail stores, is fairly homogeneous across clusters, while places tend to differ quite substantially in terms of their degree of specialisation into fast-food, restaurants or bars. Table 13 contains the estimated coefficients from a fully-interacted version of Equation 7. The evidence suggests a positive and statistically significant impact of a market access improvement on economic activity for grid cells specialised into smaller/younger establishments or fast-foods/cafes/bars. In contrast, clusters specialised in retail and older businesses do not seem to be as sensitive to changes in market access. A potential explanation is that the development of new cycling infrastructure people helps improving establishments' salience, particularly with respect to younger and thus less known businesses. Thanks to the new bike lanes, people commute more on the surface and with a transport mode that is easier to park than cars. Hence, they become more aware of shopping opportunities and better able to exploit them.

6.4 Other outcomes

In this section, we inspect how a set of additional outcomes are related to market access. First, we build a dummy indicator that takes value one in the case of positive firm entry. We could expect that a higher market access could encourage new firms to locate in these places. Our estimates of column 1 of Table 14 do not support this hypothesis. One possible explanation for the lack of firm entry can be commercial rents. As market access improves in a given place, the rental rate a perspective business must pay to enter the market also rises, thus offsetting the benefit accruing from a market expansion. Unfortunately, we do not have access to commercial rate to test this hypothesis.

Next, we investigate whether an improvement in market access has an impact on house prices. We construct a price index by running a hedonic regression of the (log) of house prices on individual properties characteristics, in addition to grid cell and time fixed effects (see Section 4 for more details on the dataset used). Again, our estimated coefficient is not statistically significant (column 2 of Table 14). Bike lane construction is usually accompanied by the introduction of speed limits and the re-making of footpaths to make more room for active mobility, which might indeed increase the value of a property (so-called "amenity" channel). In our setting, the coefficient that is suited for measuring this channel is the one on local bike lane density, which is indeed positive and strongly significant. An extra kilometer of bike lanes is associated with an increase in house prices by 6.9%. In contrast, the coefficient on market access captures primarily the impact of an improvement in accessibility through cycling. The absence of a statistically significant link with house prices underscores the weaker capitalisation of cycling infrastructure improvements in house prices in this context, in contrast with the one found in other types of infrastructure investment, such as public transport.

Third, we use as an outcome the log for the total number of cars transiting across a given grid cell during a given time. We use publicly available data on car traffic measured by multiple sensors distributed across the city of Paris to obtain a measure of car flow at the grid cell-quarter level.[34] We find a negative impact of market access on car traffic (column 3 of Table 14). Multiple mechanisms can explain this negative elasticity. First, car usage goes down because cycling has become relatively more attractive and a subset of former car users switches to cycling. Second, bike lane deployment usually implies the removal of car parking slots or the introduction of speed limits, which tend to disincentivise car usage.

6.5 Distributional consequences of Plan Vélo

As a consequence of the preference specification adopted (see Section 3), we are unable to identify the value generated by the project at the aggregate level. By the way it is defined, market access in a given place increases only if bilateral commuting costs from that place decline on average more than what they do from other places. This means, essentially, that market access (and, thus, potential demand) gains in a given place can take place only at the expenses of others.

Our framework enables us to identify what are the places that have gained in relative terms and those that lost since the development of the new infrastructure. We multiply the difference between 2019q4 and pre-Plan Vélo market access by 5.17 (the baseline elasticity - see column 2 of Table 4), thus obtaining the percentage point change in total revenues implied by the development of Plan Vélo only. We find this difference to be positive for 40% of cells (+1.9 p.p. on average across "winners") and negative for 60% (-0.9 p.p. on average across "losers").

The uneven impact of the infrastructure can be explained by analysing the characteristics of the places that gained from the development of the bike network. In Figure 11, we show how market access has changed during our period of analysis. The first thing to notice is that places where market access improved tend to be 1) centrally located, and 2) characterised

[34] Data is collected by *données de comptage routier* and comes from 3,342 sensors distributed around the city of Paris. The data counts the number of cars flowing on specific streets at given times. We impute the number of cars transiting in a given grid cell and month of the year following a similar methodology to the one employed for the control variables, with the only difference that this time the rescaling factor is the share of the street monitored by a specific sensor overlapping with different cells. In case there is no sensor mapping to a given grid cell, we assign the imputed value given by an average of traffic measured in neighboring cells. We collapse the monthly frequency of the original dataset to the quarterly one by taking quarter-specific averages.

on average by lower purchasing power (or income) (Figure 10).[35] Hence, Plan Vélo acted to redistribute demand for non-tradables away from more peripheral but residents-rich districts towards more central but less residents-rich ones. As an example, before the development of Plan Vélo residents of the 15^{th} or the 17^{th} district (situated respectively in the south-west and north-west quadrant of the city) might have preferred to go shopping locally since it was for them costly to reach the city centre. After the development of new bike lanes connecting these two districts with the city centre, a larger fraction of residents chooses to take a bike ride and go shopping in the city centre instead. The same reasoning holds for residents of the central districts, who may now find easier to go shopping in peripheral ones. However, central districts in Paris are mostly shopping locations and have low population density, so that the change in consumption habits is more likely to benefit businesses located in the city centre at the disadvantage of those located in peripheral districts.

This analysis supports the conclusion of the non-trivial distributional consequences of the new infrastructure, both from a quantitative point of view and from the point of view of the type of spatial reallocation involved.

7 Conclusion

Despite many existing narrative accounts, sound quantitative assessments of the consequences of bike infrastructure development for local economic activity are scarce. The development of bike infrastructure can affect local economic activity in multiple ways. By bringing down bilateral travel time especially between given origin-destination pairs, it can reshape the geography of consumption towards locations that become better accessible. Furthermore, by favoring a switch to active mobility, the development of bike infrastructure can benefit local businesses by making them more salient, easier to visit, on top of potential "footfall" effects that materialise if the construction of the bike lane contributes to making streets more pedestrian-friendly.

We conduct an empirical evaluation of the impact of the construction of a large-scale infrastructure project, the Plan Vélo that occurred between 2017 and 2019 in the city of Paris, on businesses operating in the non-tradables sector. We find robust evidence in favor of an increase in economic activity in the non-tradable sector, as proxied by the total value and volume of card transactions directed at merchants located in parts of the city subject to an increase in market access triggered by the development of the new infrastructure.

[35] Purchasing power in the city of Paris tends to be concentrated in more peripheral districts, especially in the western part of the city, owing primarily to the distribution of population.

Our analysis of the distributional impact of the new infrastructure highlights that an important consequence of the development of Plan Vélo was to redistribute the demand for non-tradables away from more peripheral/more densely populated districts towards more central/more scarcely populated ones. These distributional consequences should be taken into consideration by current policy-makers, especially in light of the fact that the second edition of Plan Vélo - renewing the local administration commitment to make the city of Paris one of European active mobility capital cities - is currently under construction.

References

Agarwal, S., J. B. Jensen, and F. Monte (2017): "Consumer mobility and the local structure of consumption industries," Tech. rep., National Bureau of Economic Research.

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

Allen, T., S. Fuchs, S. Ganapati, A. Graziano, R. Madera, and J. Montoriol-Garriga (2020): "Is Tourism good for Locals? Evidence from Barcelona," .

Athey, S., D. Blei, R. Donnelly, F. Ruiz, and T. Schmidt (2018): "Estimating heterogeneous consumer preferences for restaurants and travel time using mobile location data," in *AEA Papers and Proceedings*, vol. 108, 64–67.

Baum-Snow, N. (2007): "Did Highways Cause Suburbanization?" *The Quarterly Journal of Economics*, 122, 775—-805.

Billings, S. B. (2011): "Estimating the value of a new transit option," *Regional Science and Urban Economics*, 41, 525–536.

Bock, H.-H. (2007): "Clustering Methods: A History of k-Means Algorithms," *In: Brito, P., Cucumel, G., Bertrand, P., de Carvalho, F. (eds) Selected Contributions in Data Analysis and Classification. Studies in Classification, Data Analysis, and Knowledge Organization.*

Borusyak, K. and P. Hull (2020): "Non-random exposure to exogenous shocks: Theory and applications," Tech. rep., National Bureau of Economic Research.

Bounie, D., Y. Camara, and J. W. Galbraith (2023): "Consumers' Mobility, Expenditure and Online-Offline Substitution Response to COVID-19: Evidence from French Transaction Data," *European Economic Review*, 151.

Broach, J. P. (2016): "Travel mode choice framework incorporating realistic bike and walk routes," Ph.D. thesis, Portland State University.

Cailly, C., J.-F. Côte, A. David, J. Friggit, S. Gregoir, A. Nobre, F. Proost, S. Schoffit, N. Tauzin, and H. Thélot (2019): "Les indices Notaires-Insee des prix des logements anciens Méthodologie v4," .

Cartes Bancaires (2019): "Cartes Bancaires en chiffres," https://www.cartes-bancaires.com/a-propos/cb-en-chiffres/.

Chandra, A. and E. Thompson (2000): "Does public infrastructure affect economic activity?: Evidence from the rural interstate highway system," *Regional Science and Urban Economics*, 30, 457–490.

Currie, J. and R. Walker (2019): "What do economists have to say about the Clean Air Act 50 years after the establishment of the Environmental Protection Agency?" *Journal of Economic Perspectives*, 33, 3–26.

Davis, D. R., J. I. Dingel, J. Monras, and E. Morales (2019): "How segregated is urban consumption?" *Journal of Political Economy*, 127, 1684–1738.

Deshpande, M. and Y. Li (2019): "Who is screened out? application costs and the targeting of disability programs," *American Economic Journal: Economic Policy*, 11, 213–48.

Diamond, R. and E. Moretti (2021): "Where is Standard of Living the Highest? Local Prices and the Geography of Consumption," Tech. rep., Discussion paper.

Duranton, G. and M. A. Turner (2012): "Urban growth and transportation," *Review of Economic Studies*, 79, 1407–1440.

Galdon-Sanchez, J. E., R. Gil, F. Holub, and G. Uriz-Uharte (2021): "Benefits and Costs of Driving Restriction Policies: The Impact of Madrid Central on Congestion, Pollution and Consumer Spending," *mimeo*.

Gibbons, S. and S. Machin (2005): "Valuing rail access using transport innovations," *Journal of Urban Economics*, 57, 148–169.

Gonzalez-Navarro, M. and M. A. Turner (2018): "Subways and urban growth: Evidence from earth," *Journal of Urban Economics*, 108, 85–106.

Gorback, C. (2020): "Your uber has arrived: Ridesharing and the redistribution of economic activity," *Wharton working paper*.

Heblich, S., S. J. Redding, and D. M. Sturm (2020): "The making of the modern metropolis: evidence from London," *The Quarterly Journal of Economics*, 135, 2059–2133.

Hornbeck, R. and M. Rotemberg (2021): "Railroads, Market Access, and Aggregate Productivity Growth," *mimeo*.

INSEE (2018): "Mobilités professionnelles en 2018 : déplacements domicile - lieu de travail Recensement de la population - Base flux de mobilité," .

Koster, H. R., I. Pasidis, and J. van Ommeren (2019): "Shopping externalities and retail concentration: Evidence from Dutch shopping streets," *Journal of Urban Economics*, 114, 103194.

Mairie de Paris (2015): "Paris Capital du Vélo 2020," *report*.

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

Miyauchi, Y., K. Nakajima, and S. J. Redding (2021): "Consumption access and agglomeration: evidence from smartphone data," *CEP Discussion Papers dp1745, Centre for Economic Performance, LSE.*

Pereira, R. H. M., M. Saraiva, D. Herszenhut, C. K. V. Braga, and M. W. Conway (2021): "r5r: Rapid Realistic Routing on Multimodal Transport Networks with R5 in R," *Findings*.

Relihan, L. E. (2017): "Is online retail killing coffee shops?" *mimeo*.

Schiavina, M., S. Freire, and K. MacManus (2019): "Clustering Methods: A History of k-Means Algorithms," *GHS population grid multitemporal (1975, 1990, 2000, 2015) R2019A*.

Sleiman, L. B. et al. (2021): "Are car-free centers detrimental to the periphery? Evidence from the pedestrianization of the Parisian riverbank," Tech. rep.

Tassinari, F. (2021): "Low emission zones and traffic congestion: evidence from Madrid Central," *Unpublished*.

Tsivanidis, N. (forthcoming): "Evaluating the impact of urban transit infrastructure: Evidence from bogota's transmilenio," *American Economic Review*.

Viard, V. B. and S. Fu (2015): "The effect of Beijing's driving restrictions on pollution and economic activity," *Journal of Public Economics*, 125, 98–115.

Warnes, P. E. (2024): "Transport infrastructure improvements and spatial sorting: Evidence from Buenos Aires," *Working paper*.

Appendix

A Figures and Tables

Table 1: Cycling speeds by type of edge in the cycling network

Type of edge in cycling network	Adjustment to default speed
Cycle track	1
Cycle lane	0.50
Primary, trunk or motorway	0.17
Secondary or tertiary	0.25
Unclassified one way streets	0.50
Other (mainly residential) with no bike signs or infrastructure	0.50

Notes: For each *highway* type as defined by the OpenStreetMap classification, we assign a different cycling speed. The speed is calculated as a fraction of the maximum cycling speed, which is fixed at 16 km/h for cycle tracks. All other ways are adjusted by the adjustment number in this table, so, for example, the cycling speed on secondary and tertiary roads is $16 \times 0.25 = 4$ km/h. Back to Section 4.

Table 2: Estimated modal choice parameters

Description	Parameter	Value
Disutility of commuting elasticity to travel time	κ	.003
Inverse of correlation across mode-specific idiosyncratic preference shocks	λ	.041
Cycling preference shifter	$b_{cycling}$	-.065
Public transport preference shifter	b_{pt}	.029
Car preference shifter	b_{car}	-2.76

Note: The conditional logit estimation is implemented using the `nlogit` STATA module. The base category is *walking*. Back to Section 4.

Table 3: Descriptive statistics

All grid cells employed in the analysis	Mean	Std. Dev.	Min	Max
Total revenues (in000s €)	1,652	4,857	0	95,003
Transactions' volume	27,492	42,041	5	453,622
Avg. revenues p/transaction (€)	68	126	8	2,693
Merchants (#)	28	27	1	232
Population	1,478	773	0	4,216
Population 25-39	395	248	0	1,348
Jobseekers (%)	9	2	0	19
Foreigners (%)	15	6	0	81
Cars (#)	20,782	22,714	29	166,834
House price (€p/m2)	8,543	1,441	6,118	12,733
N	1,418			

Note: the percentage of jobseekers is with respect to working age population, the percentage of foreigners is with respect to total population. All variables correspond to quarter-specific averages of the underlying monthly values (constant during the year for socioeconomic and demographic characteristics). The data are quarterly averages referring to 2015. Source: INSEE and *Groupement des Cartes Bancaires CB.*

Table 4: Elasticity of local economic activity to market access: baseline evidence

Panel A:	Log total revenues				
	(1)	(2)	(3)	(4)	(5)
Log market access	4.107**	5.173**	4.500*	4.406**	4.125**
	(1.882)	(2.280)	(2.447)	(1.91)	(1.85)
Panel B:	Log transactions' volume				
Log market access	5.062**	4.603**	4.996**	4.754**	5.005**
	(2.041)	(2.160)	(2.298)	(2.05)	(2.00)
Panel C:	Log average revenues p/transaction				
Log market access	-0.933	0.577	-0.482	-0.339	-0.859
	(1.202)	(1.404)	(1.484)	(1.24)	(1.18)
N	27,097	27,097	27,097	27,097	27,097
Controls	X	X	X	X	X
Grid cell FE	X	X	X	X	X
District×Time FE	X	X	X	X	X
LBLD	None	Same project	Neighbors	-	-
Instrument				Exclude 1km	Exclude project pairs
FS F-stat				1400.93	110718.35
Estimation		OLS		2SLS	

Notes: coefficients from the estimation of Equation 7. Standard errors are clustered at the grid cell level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Groupement des Cartes Bancaires CB*. Back to Section 6.1.

Table 5: Elasticity of local economic activity to market access: lagged impact

Panel A:	Log total revenues			
	(1)	(2)	(3)	(4)
Log MA	5.173**			
	(2.280)			
First lag log MA		4.970**		
		(2.304)		
Second lag log MA			4.575**	
			(2.209)	
Third lag log MA				5.515**
				(2.311)
Panel B:	Log transactions' volume			
Log MA	4.603**			
	(2.160)			
First lag log MA		5.378**		
		(2.284)		
Second lag log MA			5.442**	
			(2.320)	
Third lag log MA				6.582***
				(2.502)
N	27,097	25,744	24,391	23,038

Notes: baseline estimation as in Table 4 column 2, estimating the elasticity to lagged market access. Standard errors are clustered at the grid cell level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Groupement des Cartes Bancaires CB*. Back to Section 6.1.

Table 6: Robustness tests: dealing with centrality bias

Panel A:	Log total revenues		
	(1)	(2)	(3)
Log market access	5.173**	5.225**	5.278**
	(2.280)	(2.353)	(2.284)
Panel B:	Log transactions' volume		
Log market access	4.603**	4.793**	4.770**
	(2.160)	(2.230)	(2.158)
N	27,097	25,197	26,817
Test	Baseline	Remove central districts	Remove transport hubs

Notes: baseline estimation as in Table 4 column 2 (col.1); excluding districts 1-4 (col.2); excluding grid cells located within 500 meters of metro/train hubs featuring at least 3 metro and/or train connections (col.3). Standard errors are clustered at the grid cell level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Groupement des Cartes Bancaires CB*. Back to Section 6.2.

Table 7: Robustness tests: testing random development timing

	Treatment date		
	2017q2	2018q1	2018q4
Log population	-3.294***	-0.612	-0.625*
	(1.028)	(0.583)	(0.362)
% Foreigners	16.64***	-1.879	-2.106
	(4.818)	(2.883)	(1.937)
% Job seekers	-19.50**	-17.68***	-0.706
	(8.822)	(5.196)	(3.791)
Log population 25-39 yrs old	2.673***	0.561	0.480
	(0.907)	(0.501)	(0.307)
N	271	201	146

Notes: the dependent variable is the date in which the cells in the still-to-be-developed sample as of 2017q2 (col.1), 2018q1 (col.2) and 2018q4 (col.3) are going to be treated. The covariates refer to 2017q2 (col.1), 2018q1 (col.2), 2018q4 (col.3). Back to Section 6.2.

Table 8: Robustness tests: keeping only Plan Vélo subsample

	Log total revenues		Log transactions' volume	
	(1)	(2)	(3)	(4)
Log market access	5.173**	8.129***	4.603**	8.897***
	(2.280)	(2.953)	(2.160)	(3.010)
N	27,097	9,220	27,097	9,220
Long Diff MA	.009	.013	.009	.013
Controls	X	X	X	X
Grid cell FE	X	X	X	X
District×Time FE	X	X	X	X
LBLD	Same project	Same project	Same project	Same project
Sample	Baseline	Plan Vélo	Baseline	Plan Vélo
Estimation	OLS			

Notes: baseline estimation as in Table 4 column 2 restricted to the subsample of grid cells intersected by the original Plan Vélo. Standard errors are clustered at the grid cell level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Groupement des Cartes Bancaires CB*. Back to Section 6.2.

Table 9: Balancing test of local characteristics in the Plan Vélo subsample between grid cells where development had taken place by the end of 2019 and those where it did not

	Developed		Not developed				
	Mean	Std Dev	Mean	Std Dev	Difference	t-stat	p-value
MA (in000s)	3539	969	3491	803	-47	0.57	0.57
Roads (m)	1117	342	1125	313	8	-0.25	0.80
Planned bike lanes (m)	190	109	173	89	-17	1.88	0.06
Population	1414	872	1393	751	-21	0.28	0.78
Foreigners (%)	16	5	15	4	-0	1.16	0.25
Jobseekers (%)	9	2	9	2	-0	0.73	0.47
Population 25-39	397	281	374	231	-23	0.96	0.34
Entrant firms (#)	0	1	1	1	0	-0.57	0.57
Car flow (#)	24368	23820	19518	19703	-4850	2.38	0.02
House price (p/m2)	8821	1635	9048	1717	227	-1.45	0.15
Value (in000s)	1582	3060	2588	7283	1006	-1.93	0.05
Volume	30645	47632	37053	54856	6409	-1.34	0.18
Avg. value p/transaction	69	152	69	69	1	-0.05	0.96
Avg. value p/merchant	48504	77339	61199	127410	12695	-1.29	0.20
Merchants (#)	31	31	36	28	5	-1.75	0.08
N	229	.	232	.	.	.	.

Notes: the data refer to 2015. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Groupement des Cartes Bancaires CB*. Back to Section 6.2.

Table 10: Robustness tests: miscellanea

Panel A:	Log total revenues			
	(1)	(2)	(3)	(4)
Log market access	5.173**	4.264*	6.814***	5.211**
	(2.280)	(2.206)	(2.575)	(2.280)
Panel B:	Log transactions' volume			
Log market access	4.603**	4.140*	5.970**	4.652**
	(2.160)	(2.307)	(2.418)	(2.155)
N	27,097	25,740	22,157	27,097
Test	Baseline	Sectoral shares	Remove affected by Sunday Law	Sunday Law trend

Notes: baseline estimation as in Table 4 column 2 (col.1); augmented to include lagged sectoral shares as controls (col.2); excluding cells affected by the Sunday Law (col.3); augmented to include an interaction term between the grid cell-specific share of surface concerned by the 2015 "Sunday Law" and time dummies (col.4); excluding grid cells located within 100 meters from the itinerary of tramway T3b (col.5); . Standard errors are clustered at the grid cell level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Groupement des Cartes Bancaires CB*. Back to Section 6.2.

Table 11: Robustness tests: elasticity of credit card usage intensity to market access

	Card usage intensity index			
	(1)	(2)	(3)	(4)
Log MA	0.270			
	(0.750)			
First lag log MA		0.238		
		(0.871)		
Second lag log MA			0.291	
			(0.916)	
Third lag log MA				0.372
				(0.989)
N	26,948	25,604	24,260	22,916

Notes: baseline estimation as in Table 4 column 2 applied to the ratio between the number of establishments reporting transactions in the *Groupement des Cartes Bancaires CB* dataset in a given quarter and grid cell, and the number of establishments active in that same quarter and grid cells according to the business registry (SIRENE). Source: SIRENE, *Observatoire du Plan Vélo de Paris*, INSEE and *Groupement des Cartes Bancaires CB*. Back to Section 6.2.

Table 12: Local merchant characteristics clusters: descriptive statistics

Cluster	Retail			Restaurants			Firm characteristics	
	Non spec.	Spec./food	Spec./other	Fast-food	Restaurants	Bars	Large	Old
1	30	23	10	78	8	18	18	28
2	18	63	21	90	12	65	75	63
3	6	61	89	47	19	43	74	83
4	39	28	32	1	7	8	64	80
5	29	25	64	43	60	23	71	34
All	28	35	30	56	15	27	49	52

Notes: col.1-6 contain the % of cells per each cluster specialised in 2015 in the corresponding activities. Col.7-8 contain the % of cells in each cluster such that in 2015 average merchant size (col.7) or age (col.8) was greater than the median value. Source: *Groupement des Cartes Bancaires CB*. Back to Section 6.3.

Table 13: Testing heterogeneous effects with respect to local merchant characteristics

	Log total revenues	Log transactions' volume
Log MA× Small and new businesses	7.189*	5.074
	(3.849)	(3.600)
Log MA× Spec. food stores/fast food/bars	8.011***	5.802*
	(2.703)	(3.282)
Log MA× Spec. retail + old businesses	1.689	3.651
	(3.816)	(4.086)
Log MA× Retail + old businesses	6.330	0.811
	(5.236)	(3.946)
Log MA× Spec. retail/restaurants	0.454	6.674
+ new businesses	(3.271)	(4.438)
N	24,777	24,777

Notes: baseline estimation as in Table 4 column 2, testing heterogeneous effects through the inclusion of interaction terms between market access and cluster-specific dummies. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Groupement des Cartes Bancaires CB*. Back to Section 6.3.

Table 14: Elasticity of other outcomes to market access

	(1) Business Entry	(2) Log of House Prices	(3) Car Volume
log MA	-0.390	0.742	-9.345***
	(0.744)	(0.800)	(3.246)
LBLD - same project (km)	-0.013	0.069***	-0.209***
	(0.011)	(0.007)	(0.051)
N	27,097	26,993	26,946

Notes: coefficients from the estimation of Equation 7, testing the elasticity of other outcomes to market access. Standard errors are clustered at the grid cell level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Groupement des Cartes Bancaires CB*. Back to Section 6.4.

Figure 1: Development of Plan Vélo

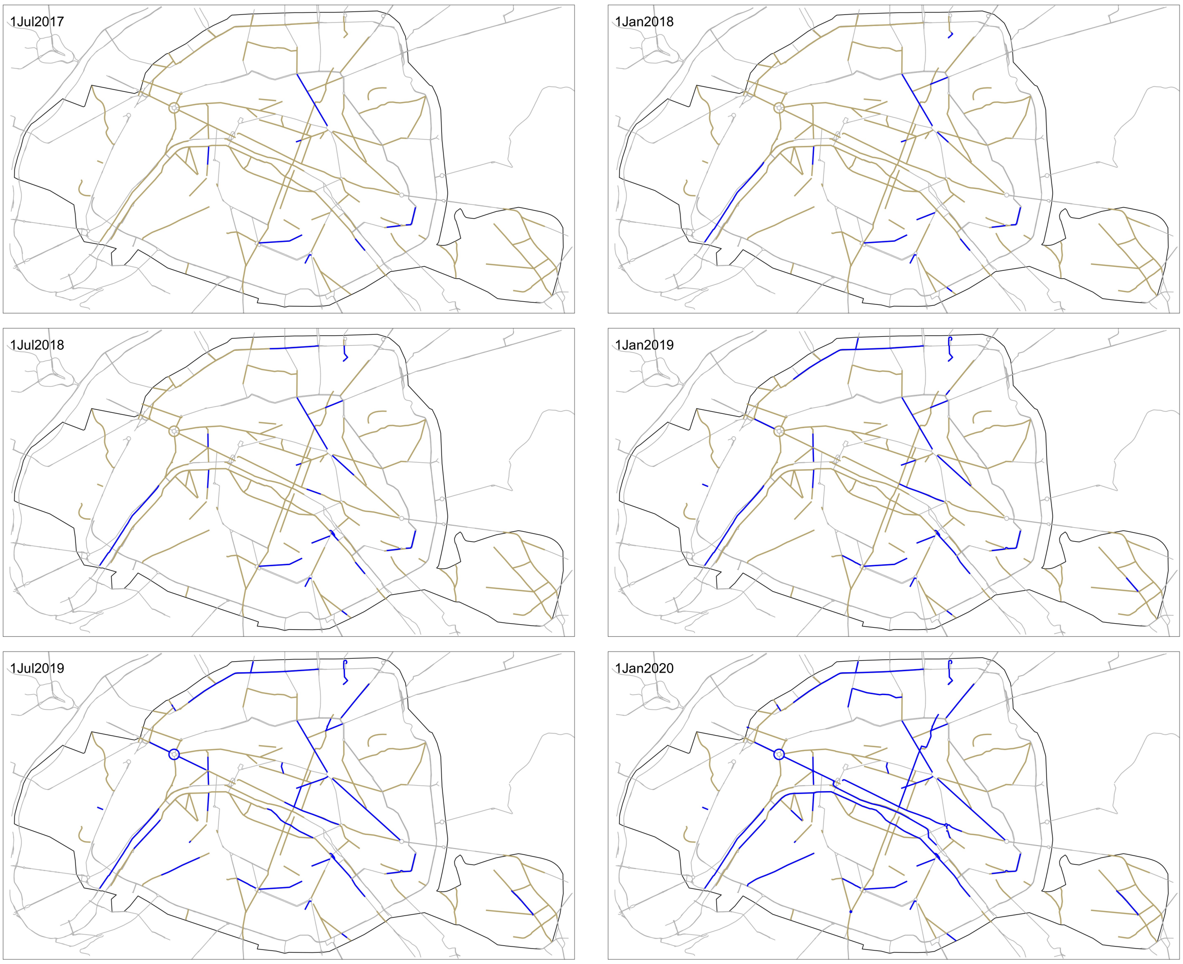

Notes: blue lines show the development of Plan Vélo; gold lines show the original Plan. Source: *Observatoire du Plan Vélo de Paris*.

Figure 2: Total number of bike trips recorded in Paris over time

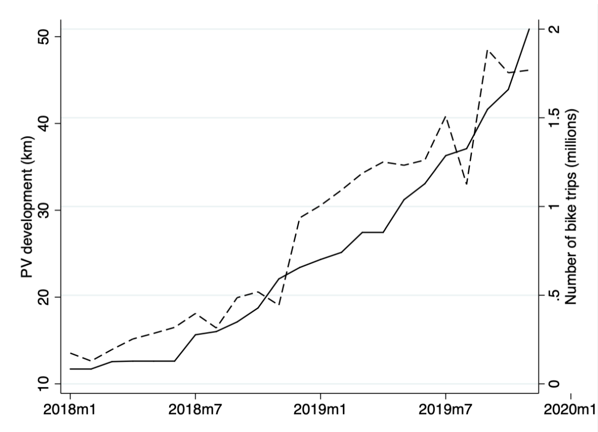

Notes: all bike trips are recorded by sensors distributed across the city. Source: *Comptage vélo - Données compteurs* dataset from https://www.data.gouv.fr/en/datasets/comptage-velo-historique-donnees-compteurs-et-sites-de-comptage/.

Figure 3: Total card transaction revenues taking place in Paris over time

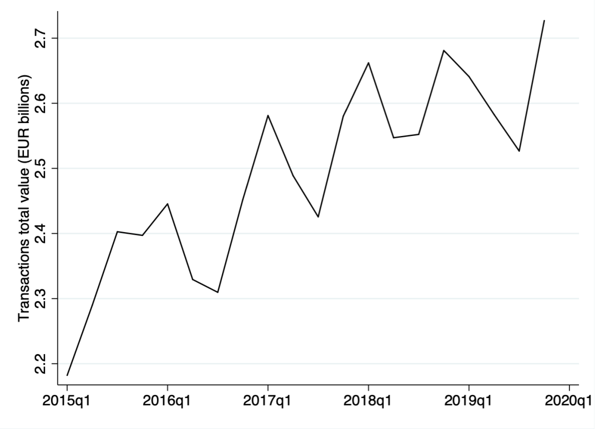

Source: *Groupement Cartes Bancaires CB*.

Figure 4: Bilateral travel time by bike at different points in time

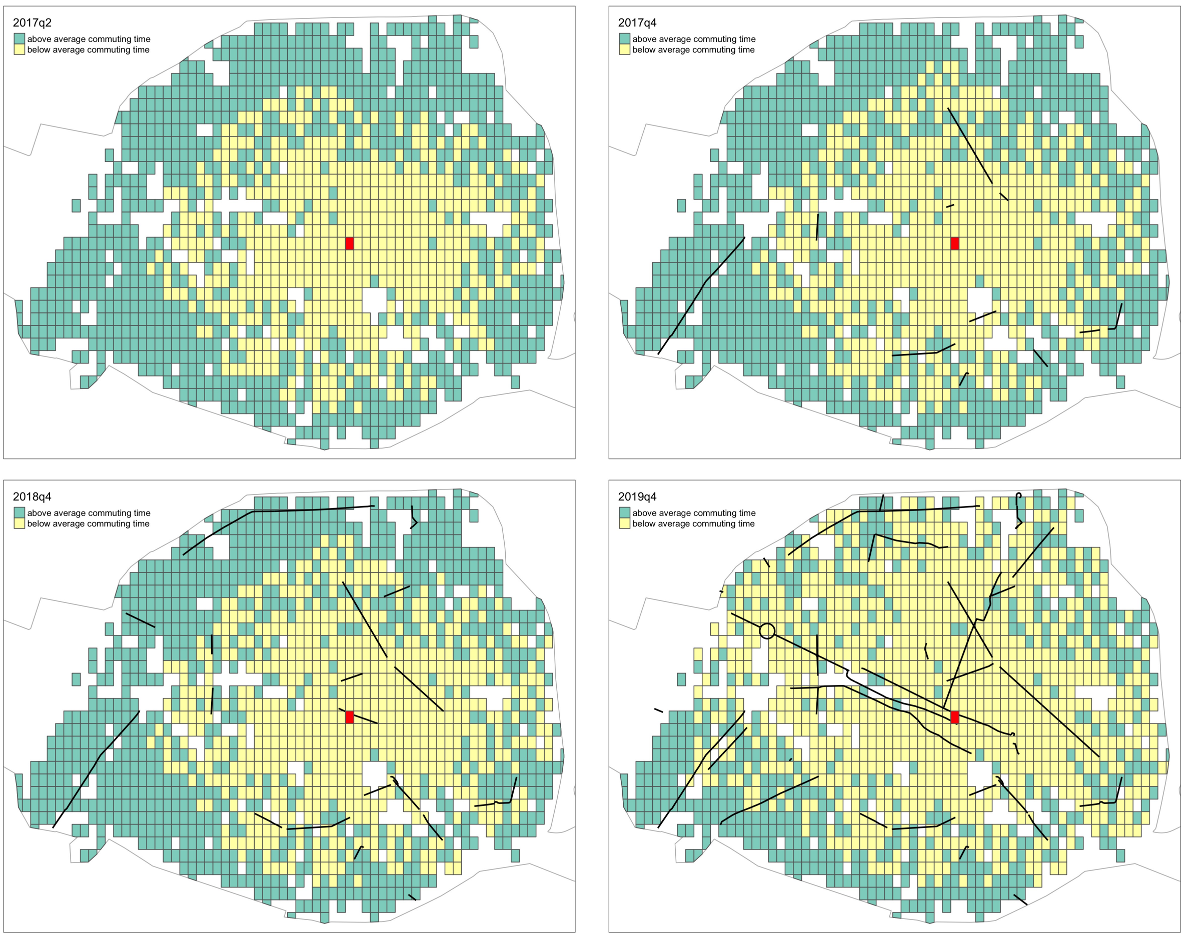

Notes: bilateral travel time by bike over time from Hôtel de Ville (red dot) to other parts of the city. Overlaid black lines capture Plan Vélo development over time. Back to Section 4.3.

Figure 5: Quartiles of market access at different points in time

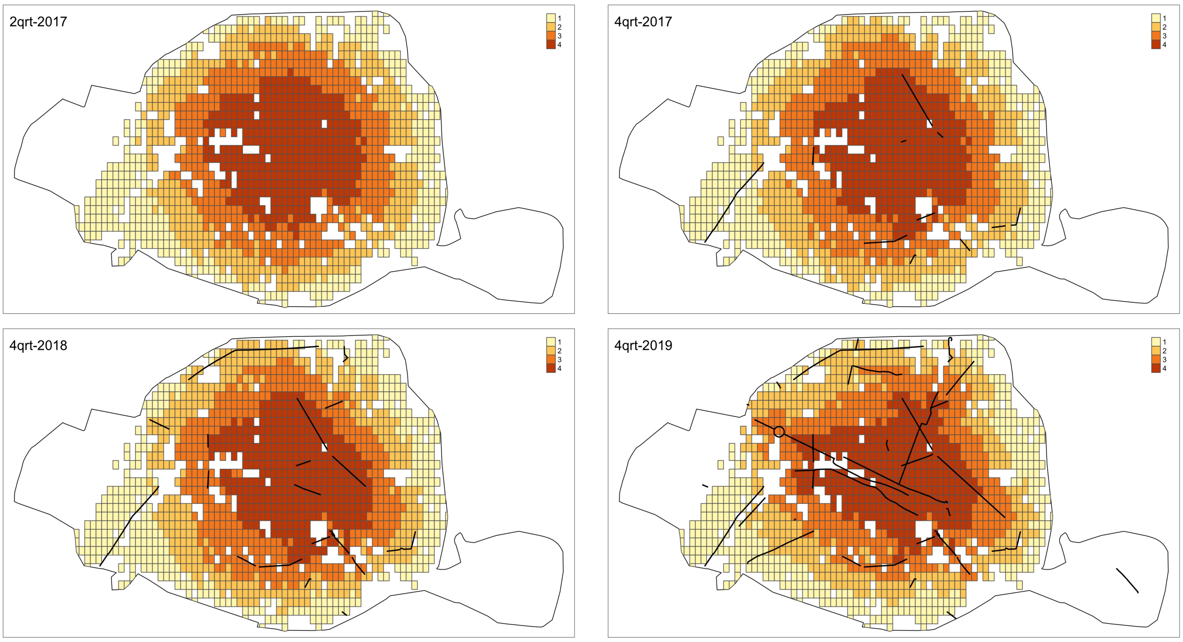

Notes: overlaid black lines capture Plan Vélo development over time. The quartiles refer to the cross-space and time market access distribution and are thus held constant. Back to Section 4.3.

Figure 6: Pre-trends analysis

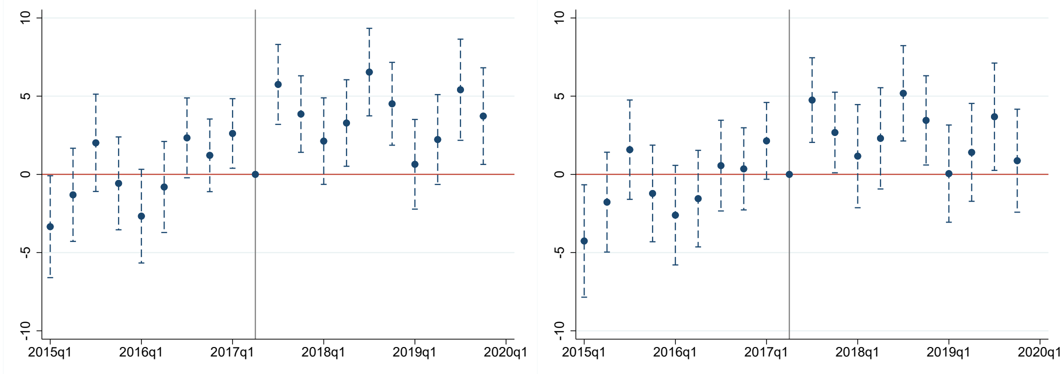

(a) Log total revenues

(b) Log transactions' volume

Notes: estimated β^t from Equation 10 on the y-axis. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Groupement des Cartes Bancaires CB*. Back to Section 6.2.

Figure 7: Grid cells crossed by 2015 Plan Vélo

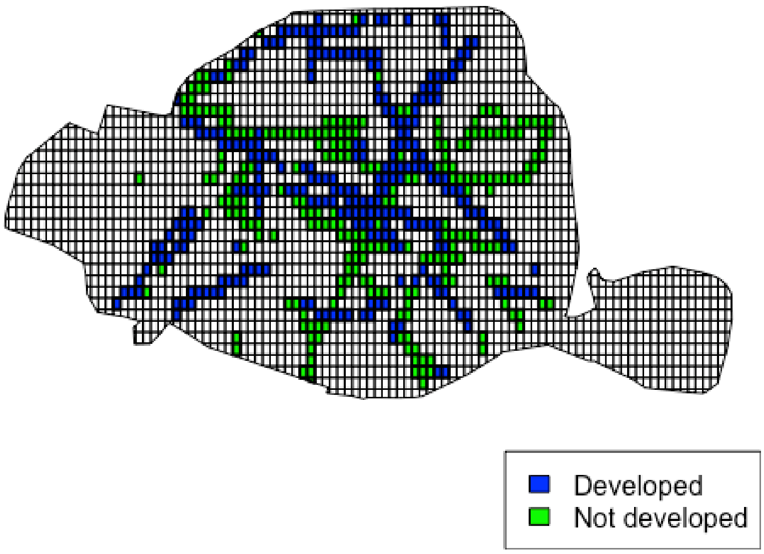

Notes: blue grid cells are those that were actually developed, green grid cells correspond to chunks of the original Plan Vélo that were not yet developed as of 2019q4. Source: *Observatoire du Plan Vélo de Paris*. Back to Section 6.2.

Figure 8: Places concerned by the Sunday Law

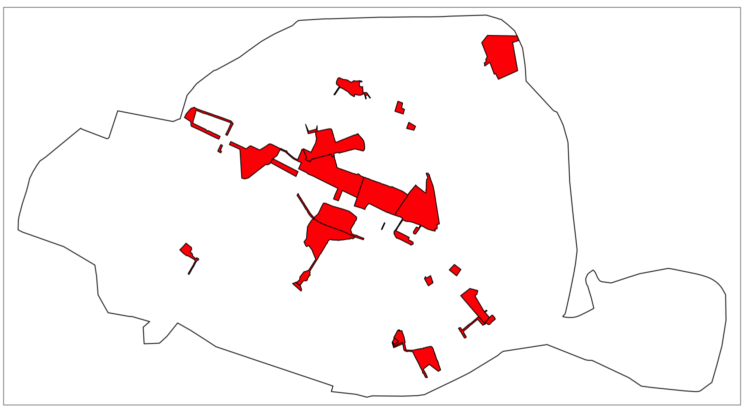

Notes: red zones include international tourism areas, tourism areas, commercial areas and train stations. Source: APUR, Mairie de Paris and DRIEA IF/UD75. Back to Section 6.2

Figure 9: Elbow test for the selection of the optimal number of clusters

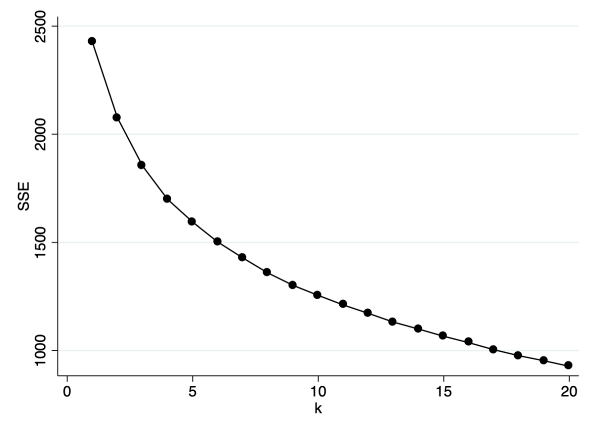

Notes: sum of squared errors on the vertical axis. Back to Section 6.3

Figure 10: Spatial distribution of purchasing power

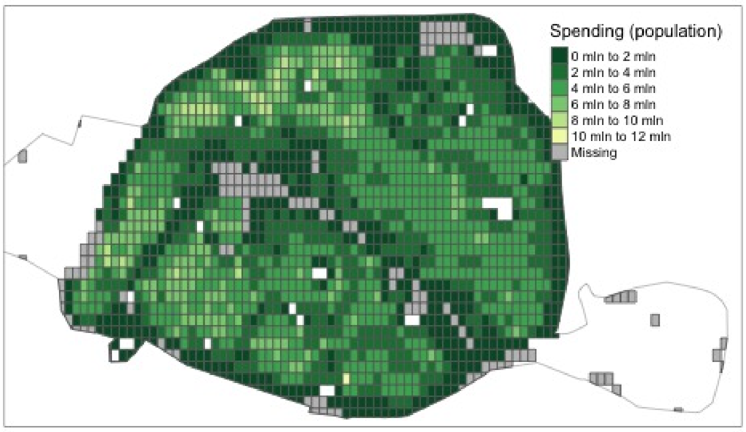

Notes: spending in a given grid cell j corresponds to Population$_j$ $\times$ Median income$_j$ in 2015. Source: INSEE. Back to Section 6.2.

Figure 11: Bike-induced change in market access (2015q1-2019q4)

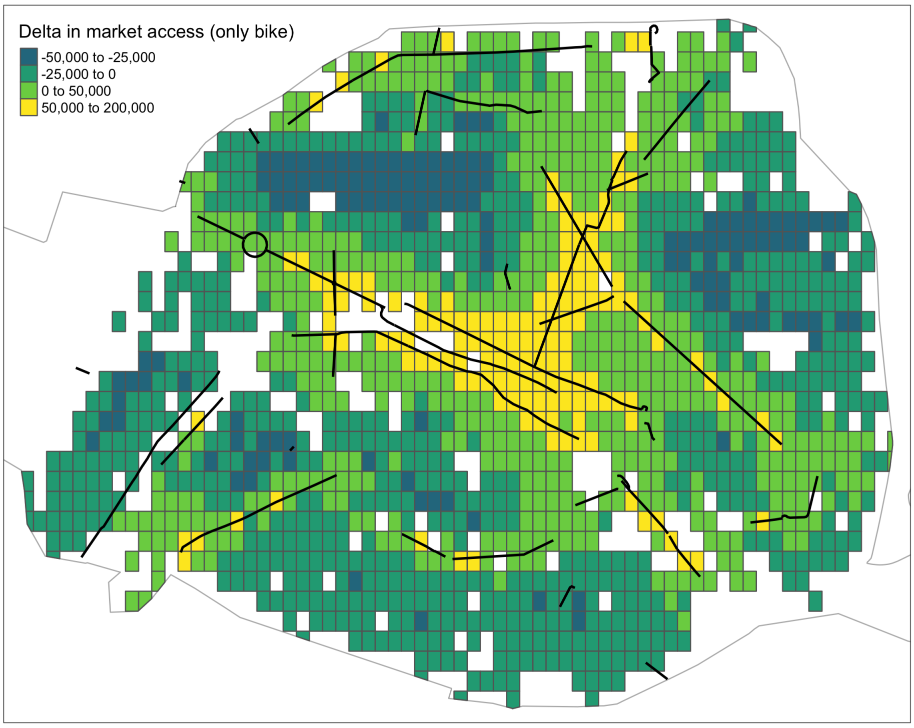

Notes: bike-induced change in market access (in absolute terms) following the development of Plan Vélo (black lines). Back to Section 4.3.