



THEMA

théorie économique,
modélisation et applications

THEMA Working Paper n°2022-09
CY Cergy Paris Université, France

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

Federica Daniele, Mariona Segu, Pablo Warnes, David Bounie,
Youssef Camara



April 2025

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

Federica Daniele¹, Mariona Segú², Pablo Warnes³, David Bounie⁴, and Youssouf Camara⁴

¹Banca d'Italia

²CY Cergy Paris Université and IEB, Universitat de Barcelona

³Aalto University and Helsinki GSE

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

July 31, 2024

Abstract

Cities are increasingly interested in developing cycling infrastructure to shift residents' mobility towards more sustainable modes of transport. Yet, little is known about the economic impact of such investments. We estimate the effects of a large-scale cycling infrastructure investment in Paris, the *Plan Vélo*, on non-tradable spending. Using geolocated data covering nearly the universe of French card transactions, we estimate a positive and statistically significant elasticity of non-tradable spending to increases in market access driven by new cycling infrastructure. We find a larger elasticity in clusters characterised by a concentration of smaller and younger establishments. The new bike infrastructure increased the non-tradable 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 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, Marion Leroutier, Ana Moreno-Maldonado, Clara Santamaria and Marta Santamaria.

1 Introduction

Road transport accounted for nearly 25% of the world’s total of carbon dioxide emissions in 2023 (IEA, 2023), the leading contributor to climate change. To curb emissions, many cities around the world are stepping up efforts to incentivise a shift towards more sustainable modes of transport, such as investing in the development of bike infrastructure. Nonetheless, a trade-off between social and private gains may arise, with special interest groups sometimes ending up opposing these investments. For instance, shop owners are often against the development of bike lanes because they could make it harder for cars to park nearby, thus harming their business.¹

This paper sheds light on this potential trade-off by providing, for the first time, robust evidence on the impact of cycling infrastructure on local economic activity. In particular, we study the effect of the development of a large-scale bike lane network on non-tradable spending as the primary outcome of interest, in addition to firm entry, housing prices and car traffic as secondary ones. We focus on the deployment of the *Plan Vélo* in the city of Paris, a major initiative aimed at promoting a transition towards active mobility that consisted in the construction of approximately 80 km of bike lanes between 2017 and 2020.² We leverage the staggered development of the infrastructure and the ensuing changes in bilateral travel costs by bike to construct a time- and space-varying measure of firm-level market access. We then estimate the elasticity of local revenues to market access using geolocated data covering nearly the universe of card transactions made by French residents provided by Groupement des Cartes Bancaires CB.³ With the estimated elasticity at hand, we assess the cumulative gain/loss for local businesses from the development of the new bike infrastructure and the spatial reallocation of spending that it entailed. To complement our results on the effects of this new cycling infrastructure on business revenues, we study its effects on housing prices, firm entry and car traffic.

We estimate the economic impact of the new bike transport infrastructure through a market access approach. Firm-level market access captures the potential demand that each business can reach conditional on consumers’ preferred route and mode of transport. We measure market access using a 180-by-180-meter grid covering the city of Paris for every

¹Various case studies show that retailers overestimate the share of customers that arrive by car and believe that bike lanes might lead to a loss in revenue. See the press article on Bloomberg CityLab “[The Complete Business Case for Converting Street Parking Into Bike Lanes](#)”, accessed on March 16, 2024.

²We end our analysis in November 2019 to avoid the disruptions to public transport caused by the national strike in December 2019 and the start of the COVID-19 pandemic.

³These data were made available thanks to a partnership with Groupement des Cartes Bancaires CB, and we exploit the card payments data in accordance with the EU General Data Protection Regulation, in application of Article 89. We use the abbreviation ‘CB’ hereafter to indicate the source of the card payments.

quarter between 2015 and 2019. Market access for a business location is defined as the weighted sum of demand for non-tradables across all possible origin locations, with weights that are inversely proportional to the bilateral travel disutility between the origin locations and the location of a business. In turn, these bilateral travel disutilities depend, via a travel elasticity, on a weighted sum of bilateral travel costs across transport modes, with weights proportional to consumers' preferences for each mode. The travel elasticity and consumers' preferences for each mode are obtained by feeding measured travel costs and commuting survey data into the preliminary estimation of a conditional logit model disciplining consumers' choice of how to go shopping. The nesting of a modal choice into our modelling of consumers' behaviour allows us to account for imperfect substitution across transport modes and thus for an increase in the number of consumers going shopping by bike potentially matched by a decline in the number of those going by car.

Having measured market access for each grid cell in Paris for every quarter between 2015 and 2019, we regress the total value of card transactions on these measures of market access to estimate the elasticity of business revenues to market access. Since our measure of market access captures the effect of the new cycling infrastructure on the transport network, this elasticity can be interpreted as the effect of the new bike lanes on business revenues.

One important concern, however, is that the development of bike lanes can be endogenous to economic activity in a given location, presenting a potential identification challenge. To address this concern, we instrument our measure of market access with an alternative measure that relies only on changes in the travel costs (due to the new cycling infrastructure) far away from a given location. The identifying assumption underlying this strategy is that the development of bike lanes in more distant parts of the city is uncorrelated with the economic conditions in any given location.

Our preferred estimation of the elasticity of business revenues to market access is 5.17. This means that the average improvement in market access, a 0.93% increase, equivalent to the development of 46 meters of bike lanes in a given grid cell, resulted in a 4.81% increase in business revenues, amounting to approximately 100,000 euros per quarter. Considering the average number of businesses per grid cell, this resulted in an average gain of 3,000 euros per business. We find that this effect is driven by an increase in the volume of transactions and not by an increase in size, suggesting that bike infrastructure induced more frequent but equally large transactions.

Business revenues in a given location can increase for two main reasons: improved connectivity and better amenities. First, easier commutes due to new bike lanes can attract

more visitors to the area (connectivity channel). Second, the area may become a more pleasant shopping destination if complementary pedestrianization measures are introduced alongside the bike lanes (amenity channel). We use a measure of the density of local bike lanes, measured as the total length of bike lanes in the immediate surroundings, to separate the effect of the amenity channel on business revenues from the effect of the connectivity channel. If bike lanes provide a direct amenity effect to businesses located near them, then, for a given increase in connectivity, being exposed to more local bike lanes (because of a higher density of bike lanes in the grid cell) should lead to increases in business revenues through changes in the amenity value of the area. Conversely, for a given density of local bike lanes, a larger increase in market access should affect business revenues only through the connectivity channel.

We implement a large battery of robustness checks to eschew endogeneity issues. First, we address the potential bias arising from the fact that certain places, owing to their central location, end up being more than proportionally impacted by infrastructure development (Borusyak and Hull, 2023). We do so by removing from our sample either central locations in a geographical sense, or locations qualifying as transport hubs according to the public transit network. Second, we show the absence of pre-trends in our outcomes of interest, addressing the potential bias arising from the endogenous placement of the new bike lanes. Similarly, we find no evidence that observables systematically predict the timing of bike lane development, thereby ruling out potential bias from endogeneity in the network development timing. Third, we show that our results survive when considering a smaller, more homogeneous sample of locations. This is achieved by comparing developed areas to those with bike lanes included in the original network plan that remained undeveloped by the end of 2019. Fourth, we show that household sorting is not driving our results as the demographic composition of neighbourhoods remains stable throughout our study period. This is consistent with household sorting being a medium-term process rather than a short-term one. Fifth, we check that our results are not driven by substitution between card and cash payments.

Next, we investigate the existence of heterogeneous effects by dividing the city into clusters with similar business characteristics. We find evidence that the positive effects of bike lanes on local business revenues (through changes in market access) are driven by the presence of small and young businesses as well as food-related industries, such as cafes, fast food restaurants or bars in the affected locations. This is consistent with bike lanes attracting more traffic to small and newer businesses by increasing their salience. Travelling by bike as opposed to car or public transit arguably allows consumers to be better aware of avail-

able shopping opportunities, since cyclists travel above ground and at lower speeds, which allows them to stop more easily to visit a store along their ride. Moreover, pedestrianisation infrastructure that often come along with the development of bike lanes might favour combined cycling-walking shopping trips, with “footfall” externalities adding up to the potential benefits for businesses in areas affected by the new cycling infrastructure (Koster et al., 2019).

Having described the effects of the *Plan Vélo* on business revenues, we then consider how this new infrastructure might have affected other important outcomes. First, we do not detect a statistically significant effect of market access changes (due to the new bike lanes) on firm entry. While a monopolistic competition model with free entry might be consistent with a positive relationship, a possible explanation for this null result is that commercial rents might have increased following the rise in local business profitability brought about by the new infrastructure. Unfortunately, we are unable to test this hypothesis in this paper due to a lack of data on commercial rents. Second, we test the impact of market access on residential housing prices. Again, we do not find any significant effect of the increase in market access due to the new infrastructure on residential housing prices. We argue that the potential effect of the new infrastructure on housing prices might accrue especially via an amenity channel, as opposed to a pure connectivity or market access channel. We detect a positive elasticity of housing prices to the local bike lane density control in line with this hypothesis and with existing work (Garcia-López et al., 2024). Third, we find a negative impact of market access on the total number of cars transiting through a given grid cell. This effect is consistent with a modal shift, where consumers reduce their travel by car and increase their travel by bike to locations that see an increase in connectivity by bike. An alternative explanation is that bike lane development occurred together with the removal of car lanes or the introduction of speed limits, which further discouraged car usage. However, since we are controlling for the density of new bike lanes in each grid cell, it seems unlikely that the possible removal of car lanes when introducing bike lanes can explain the entire magnitude of the effect.

In the last part of the paper, we identify which locations gained or lost in relative terms from the development of the new bike infrastructure. We find that the pre/post difference in business 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 benefited the most from the development of the new bike network. Locations where market access

improved are on average 1) more centrally located, and 2) characterised by lower total income, mostly due to very low residents density. Based on our analysis, an important consequence of the development of *Plan Vélo* was to increase non-tradable spending in central/less densely populated districts at the expense of peripheral/more densely populated ones. For instance, our model predicts that after the development of Plan Vélo the fraction of residents of the south- and north-west quadrants of the city that give up shopping locally and prefer going to the city centre goes up since it is now easier to get there by bike.

Related literature This paper contributes to different bodies of the economic literature. First, it is related to the research strand that looks at the link between transport infrastructure and economic activity. The first wave of papers in this literature investigate the impact of getting access to a new infrastructure on outcomes such as employment (Duranton and Turner, 2012; Mayer and Trevien, 2017), population (Baum-Snow, 2007; Gonzalez-Navarro and Turner, 2018) and property prices (Gibbons and Machin, 2005; Billings, 2011). A relatively more recent set of papers follows a market access approach, where the potential gains stem from getting connected to more attractive/richer places as opposed to merely getting connected to the new infrastructure (Ahlfeldt et al., 2015; Hebllich et al., 2020; Gorbach, 2020; Tsivanidis, forthcoming; Warnes, 2024). This paper contributes to this literature by providing the first empirical assessment of the economic impact of a large-scale new cycling infrastructure in a major European capital city.

This study is also related to a recent strand of the literature that focuses on the measurement of the geography of spending in the non-tradable sector by means of large-scale spatial datasets, such as online review data (Davis et al., 2019), mobile phone data (Athey et al., 2018; Miyauchi et al., 2021) and card transactions data (Relihan, 2022; Allen et al., 2020; Agarwal et al., 2017; Diamond and Moretti, 2021; Bounie et al., 2023). We contribute to this literature by leveraging a high-frequency, geolocalised card transaction level dataset to measure the effects of new cycling infrastructure on local economic activity covering the near totality of card transactions made by French residents (Landais et al., 2020).

Our work is especially related to Galdon-Sanchez et al. (2023), who evaluate the impact of car driving restrictions on local spending in the city of Madrid, which is another type of transport policy increasingly common among European cities. Galdon-Sanchez et al. (2023) find that areas directly affected by the driving restrictions (in the form of Low Emission Zones) witnessed a decline in local spending as proxied by card spending. If we consider that businesses in these areas saw a decrease in accessibility by car with no compensating increase in accessibility by other modes of transport, businesses in these areas likely decreased their

market access. In this sense, our findings are conceptually aligned with those in [Galdon-Sanchez et al. \(2023\)](#). In our case, areas undergoing an improvement in available transport infrastructure experience an increase in local spending while at the same time observing a decrease in the volume of car traffic. This difference in outcomes between these two policies highlights the potential benefits of providing alternative modes of transport that are more sustainable, as opposed to command-and-control policies for reducing car traffic in cities. Methodologically, a distinguishing feature of our paper relative to [Galdon-Sanchez et al. \(2023\)](#), is our use of a market access-based approach. This approach allows us to identify the impact of the new cycling infrastructure on the entire city, and not just the areas near the new bike lanes. This methodological choice adds realism to the quantification exercise, given that transport infrastructure investments, even localised ones, rarely have an isolated impact, but rather interact with the pre-existing network, which amplifies and propagates their effects.

This research evaluates the economic effects of the development of a large-scale cycling network, a policy that had the stated objective of significantly greening the transportation sector in the city of Paris. Hence, we contribute to that body of the environmental economics literature that focuses on the impact of pollution reduction policies in cities, such as car usage restrictions on congestion ([Bou Sleiman, 2024](#); [Tassinari, 2024](#)) and on economic activity ([Viard and Fu, 2015](#); [Galdon-Sanchez et al., 2023](#)). [Currie and Walker \(2019\)](#) provide a summary of this body of research. Closer to our setting, we relate to the recent literature evaluating the consequences of cycling infrastructure on bike use, car congestion, pollution and housing prices ([Thorne, 2022](#); [Bernard, 2023](#); [Garcia-López et al., 2024](#)).

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. 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*

The administration that took office in 2015 in Paris prioritised a significant expansion of the cycling network, favouring a substantial transition towards active mobility. The initiative was labelled *Plan Vélo* and it consisted of about 80 km of new bike lanes, for a total investment

of 150 million euros to be developed between 2015 and 2020 (Mairie de Paris, 2015).⁴ 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 to monitor the advancement of the development of the plan.⁵

Once it took off, the plan unfolded relatively quickly, with 57 km of bike lanes (71% of the original total length) developed between July 2017 and November 2019 - the last month in our study period. Figure 1 represents different development stages of the plan. The transformation of Paris into an increasingly bike-friendly city is mirrored by the swift increase in bike usage, which grew at the average monthly rate of 15% during 2018-2019 (Figure A1).

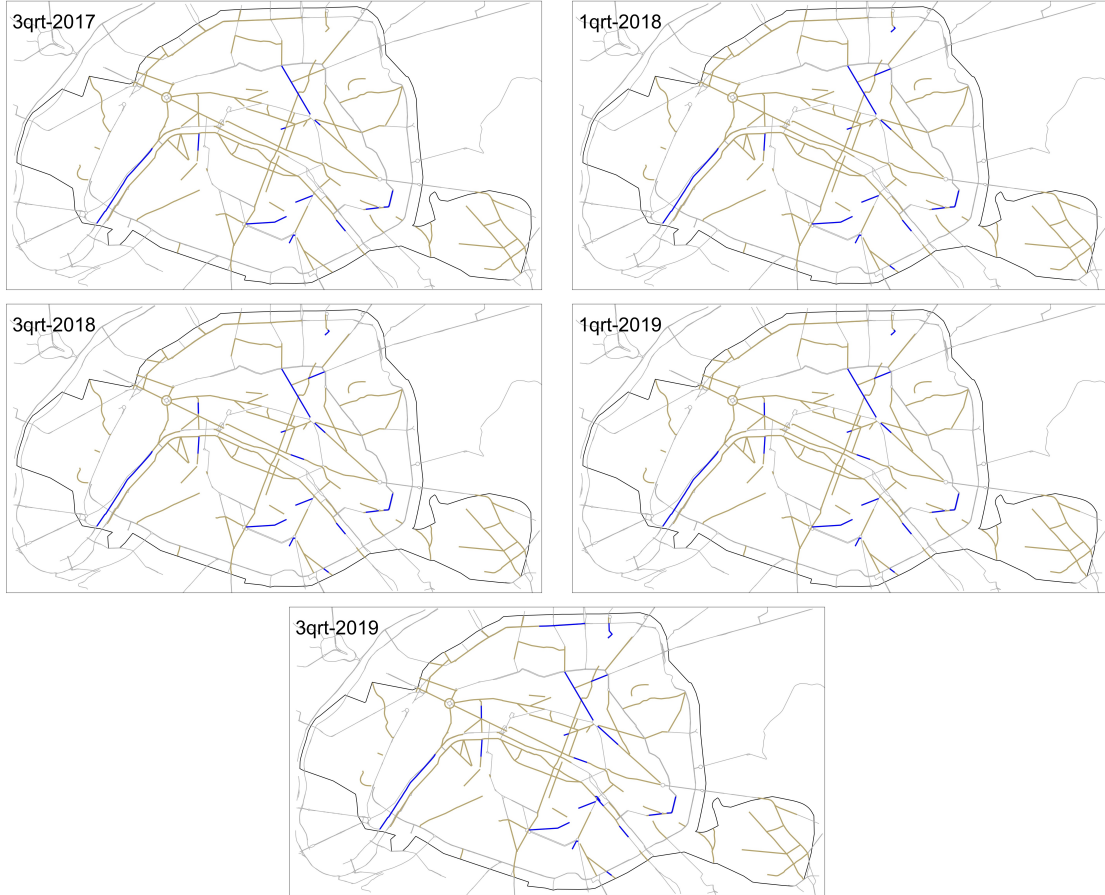
The construction of the network was coordinated at the municipal level, and it thus left little leeway to district mayors to steer the development towards their places of interest, thus relaxing potential concerns about the endogeneity of location and/or timing. The timing of development appeared rather 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 in terms of development status and timing across different parts of the city.

The analysis focuses on the municipality of Paris exclusively. The city of Paris is part of a broader metropolitan area, which according to OECD definitions comprises more than 1800 municipalities (OECD, 2023). The disproportionately high administrative fragmentation that characterises the Paris metropolitan area is the source of governance challenges and spillovers. For example, the set of car restrictions implemented by the municipality of Paris in 2016 triggered significant negative spillovers on inhabitants of other municipalities (Bou Sleiman, 2024). Compared to that policy, the one analysed in this paper is however significantly more local: the development of bike lanes is unlikely to have affected the travel choices related to the spending of inhabitants of municipalities other than Paris. This, to-

⁴The 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.

⁵See the press statement about the observatory launch here: https://parisenselle.fr/wp-content/uploads/2017/02/PeS_Observatoire-Plan-Velo_Presse_14022017.pdf.

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*.

gether to align the scope of the evaluation with the electoral base of the government that enacted the policy, justifies the decision to focus on the municipality of Paris exclusively.

3 Conceptual framework

We model the demand for non-tradable goods in a partial equilibrium setting where consumers must choose in what location to consume the non-tradable good, as well as the mode of transport they will use to travel to that location.⁶ Assume that there are $j \in J$ locations, each populated by R_j residents. A resident of location j maximises a 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:

⁶This partial equilibrium model is similar to [Gorback \(2020\)](#), but in our case, consumers must also choose the mode of transport they will use. This is especially important in our case, given that most consumption trips in Paris are not done by bike.

$$\max_{n_j^i, c_j, h_j} \left(\frac{h_j}{\beta} \right)^\beta \left(\frac{c_j}{\alpha} \right)^\alpha \left(\frac{n_j^i}{1 - \alpha - \beta} \right)^{1 - \alpha - \beta} \frac{z_{ij}}{d_{ij}} \quad \text{s.t.} \quad I_j = q^j h_j + c_j + p^i n_j^i \quad (1)$$

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 . β and α are the share of expenditure spent on housing and on the tradable good, z_{ij} is the idiosyncratic preference shock for a person residing in j and wishing to consume the non-tradable good in location i , and d_{ij} is the disutility of travelling from location j to location i .

The choice of where to purchase the non-tradable good depends on an idiosyncratic preference term, z_{ij} , which is Fréchet distributed, $F(z_{ij} = \exp(-E_i z_{ij}^{-\varepsilon}))$, with E_i being a destination-level amenity parameter and ε governing the substitutability between alternative spending destinations. Secondly, it depends on the disutility of travel from home j to the shopping destination i , d_{ij} .⁷

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 the expected value ρ . They can choose to 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 travel via transport mode $m \in \{\mathcal{M}_{Private} \cup \mathcal{M}_{Public}\}$ is given by $d_{ijm} = \exp(\kappa t_{ijm} - b_m + v_{ijm})$, where t_{ijm} is the bilateral cost, expressed in terms of time it takes to go from j to i via transport mode m , κ is the disutility of travel elasticity to transport costs, b_m is a mode-specific common preference shifter, and v_{ijm} is a mode-specific idiosyncratic preference shock.

Following [McFadden \(1974\)](#), preference shocks are drawn from a Generalised 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 $\lambda_{Private} = 1$ by construction, allows for correlation within the public transport modes nest, with the correlation increasing as $\lambda \rightarrow 0$.

⁷While 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. 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 ([Miyachi et al., 2021](#)).

Expected disutility from travelling before the realization of the idiosyncratic travel preference shocks is given by $d_{ij} = \exp(\kappa \bar{t}_{ij})$, where \bar{t}_{ij} is a weighted average of the bilateral travel cost across transport modes, with weights that depend, in part, on the share of commuters choosing a given transport mode. Specifically:

$$\begin{aligned}\bar{t}_{ij} &= -\frac{1}{\kappa} \ln [(1 - \rho) \exp(-\kappa t_{ij,0}) + \rho \exp(-\kappa t_{ij,1})] \\ t_{ij,0} &= -\frac{\lambda}{\kappa} \ln \left(\sum_{k \in \mathcal{M}_k} \exp \left(\frac{b_k - \kappa t_{ij,k}}{\lambda} \right) \right) \\ t_{ij,1} &= -\frac{1}{\kappa} \ln (\exp(b_{car} - \kappa t_{ij,car}) + \exp(-\kappa t_{ij,0}))\end{aligned}\tag{2}$$

Using Equation 2 into the shopping location choice problem, the probability of purchasing the non-tradable good in location i is:

$$Pr_{ij} = \frac{E_i \exp(-\nu \bar{t}_{ij})}{\sum_s E_s \exp(-\nu \bar{t}_{sj})}\tag{3}$$

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

The probability defined in Equation 3 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 3 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 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.

An expression for market access (from the firm point of view) is obtained by multiplying the expression in Equation 3 by location j 's total number of residents, R_j , and their income spent on the non-tradable good, $(1 - \alpha - \beta)I_j$:

$$MA_i = (1 - \alpha - \beta) \sum_j \frac{E_i \exp(-\nu \bar{t}_{ij})}{\sum_s E_s \exp(-\nu \bar{t}_{sj})} \times R_j \times I_j\tag{4}$$

Section 5 describes how the measure of market access defined in Equation 4 is employed for the evaluation of the economic impact of the new cycling infrastructure.

4 Data

4.1 Geographical unit of analysis and time-frame

Our geographical unit of analysis is 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 (Schivavina 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.⁸

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 the disruptions to public transport caused by the national strike in December 2019, and the start of the COVID-19 pandemic in March 2020. Both these events strongly impacted the mobility of Parisians and their spending behaviour (Boumie et al., 2023). The analysis is conducted at the quarter level.

4.2 Outcome variables

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, 2022; 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 1984.⁹ This

⁸Specifically, we construct weights based on the overlapping surface between polygons and grid cells.

⁹In 2020, *Groupement des Cartes Bancaires CB* had more than 100 members (including payment service providers, banks and e-money institutions).

dataset is exceptional in its coverage, allowing us to capture a significant proportion of all consumer expenditure in France. To appreciate the richness of the Groupement des Cartes Bancaires CB data, consider a few comparisons with national statistics provided for the full year 2019 by the National Institute of Statistics and Economic Studies (INSEE). GDP in France in 2019 was estimated as €2,427 billion, with €1,254 billion (52 percent of GDP) representing household consumption expenditure. Excluding fixed charges (rents, financial services, insurances) from household consumption expenditure, as these are typically paid by checks, direct debits and credit transfers, the remaining part of consumer expenditure amounts to €828 billion (34 percent of GDP). Comparing these figures with total CB card payments (€494 billion), the value of CB card payments represents 20 percent of French GDP, 39 percent of total household consumption expenditure, and finally 60 percent of total household consumption expenditure excluding fixed charges.

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 71 million cards in use in the CB system, and 1.8 million CB-affiliated French merchants ([Groupement des Cartes Bancaires, 2019](#)). Figure A2 displays the evolution of the quarterly nominal total value of transactions recorded on CB payment system during the period 2015-2019.¹⁰

CB data contain the merchant business identification number (SIRET code), thus allowing us to match it with the national business registry (SIRENE), containing information on 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 traditional non-tradable sectors: retail commerce, restaurants, accommodation services, travel agencies, personal services, bakeries, sports clubs, cinemas and theatres.¹¹ Our final sample comprises 67,230 unique SIRET codes, accounting for 61% of total card economic 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 the nominal total value added according to the national accounts. For the three largest non-tradable industries employed in this analysis (i.e., retail commerce, restaurants and accommodation), this ratio is well above 50%, which suggests that our dataset provides a good coverage of

¹⁰The growth of nominal total value is partly related to increasing card usage in retail payments and substitution from cash. Nonetheless, in one of our robustness checks, we show that increasing card usage and substitution from cash does not appear to threaten our analysis.

¹¹The 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).

total economic activity.¹²

Our main outcomes of interest are total revenues, the number of transactions (transactions' volume) and average revenue per transaction. We collapse them at the grid cell/quarter level by taking averages and we subsequently log-transform them.

Other outcome variables We look at a few additional outcome variables, which might also be affected by the development of the new network. First, we get information on business entry, by counting all newly created establishments at any given point in time and grid cell according to the national business registry. We replace the number of newly created establishments with a dummy taking value 1 if any entry takes place and 0 otherwise. This is because, given the highly geographically disaggregated nature of our analysis, the number of new establishments is close to zero for the majority of observations.

Next, we build an index of house prices using microdata on the universe of house transactions occurring in Paris from 2015 to 2019 (*demandes de valeurs foncières, DVF*). This data contains information on the sale price and several house characteristics. The data are geolocalised, allowing us to directly assign each transaction to their respective grid cell. After a simple cleaning procedure (Cailly et al. (2019)), we run a hedonic regression of the log of the sale price on house characteristics (the number of bedrooms, the number of rooms, the number of squared meters, and a set of dummies identifying different housing types), in addition to grid cell and time fixed effects. Subsequently, we build a house price index by adding the constant, the fixed effects and the residuals and aggregating the value at the grid cell-quarter level.

Finally, we build a measure of car traffic using the total number of cars transiting across a given grid cell. We use publicly available data from *données de comptage routier*, which collects information from 3,342 sensors distributed across the city of Paris. The data measure the number of cars flowing in front of the sensors at given times of the day. We use the information on traffic at 5PM in line with travel time data described in the next section and collapse the daily frequency to the quarterly one by taking quarter-specific averages. We construct the number of cars transiting through a given grid cell by computing the weighted average of the closest sensors and using as a weighting factor the relative distance.

4.3 Bilateral travel times

We obtain information on the development of the *Plan Vélo* from the *Observatoire du Plan Vélo de Paris* (see Section 2), which has maintained since July 2017 a geolocalised repository

¹²Further details about the representativeness of the sectors can be provided upon request.

keeping track of the daily development of the plan. We use this geolocalised information to update our transport model at a quarterly level and estimate new cycling travel times on the first day of February, May, August, and November of every year from 2015 to 2019.

To construct bilateral travel times by mode of transport, we combine the information on the *Plan Vélo* infrastructure with OpenStreetMap data,¹³ and General Transit Feed Specification (GTFS) files for the city of Paris for each quarter from 2015 to 2019, which are provided by the RATP group, Paris’s main public transit operator. We use OpenStreetMap and the data on the *Plan Vélo* to construct cycling, driving and walking networks for the city of Paris. Next, we rely on Dijkstra’s algorithm to find the minimum travel time between the centroids of every pair of grid cells.¹⁴ For public transit, we rely on Conveyal’s R5 Routing Engine to calculate minimum travel times by public transit between the centroids of every pair of grid cells. This routing engine allows for multimodal trips combining multiple forms of public transit as well as walking, thus providing a more realistic estimation of travel times by public transit.

With these travel time matrices by travel mode, we define the bilateral travel time for every origin i /destination j pair, transport mode m , and quarter t : $t_{ijm,t}$. Figure A3 shows the evolution over time of bilateral travel times by bike from the city hall of Paris (*Hôtel de Ville*) to all possible destinations. Travel times by bike decreased primarily in more peripheral districts, where commute takes on average longer and where cycling infrastructure was at the beginning of the sample period nearly absent.

4.4 Control variables

Table 1 summarises the different variables employed in the analysis. We assemble a rich dataset of time-varying, grid cell-level characteristics from multiple sources. Annual socioeconomic and demographic characteristics, such as total population, the population aged between 25 and 39 years old, foreign population, number of job seekers and working age population, are sourced from the *Institut national de la statistique et des études économiques* (INSEE). In an attempt to include only pre-determined variables as controls to avoid endogeneity, these variables enter our specification with a three-year lag period. INSEE data are at the IRIS geographical level (the equivalent of census tracts in France). We report variable x to the grid cell level as follows. We calculate for each IRIS unit k the share of surface overlapping with grid cell i , s_{ik} . Subsequently, we set $x_i = \sum_k s_{ik}x_k$, where k indexes all

¹³OpenStreetMap is a community-managed geographic database that includes, among many other things, detailed information on the road, pedestrian and cycling network in Paris.

¹⁴See Appendix Section B for more details on the construction of these networks.

Table 1: 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 (€/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*.

IRIS units overlapping with grid cell i .

5 Empirical strategy

We estimate the elasticity of local economic activity to market access through the following specification:

$$\ln(Y_{it}) = \alpha_i + \alpha_{dt} + \beta \ln(MA_{it}) + \gamma X_{it} + \delta LBLD_{it} + e_{it} \quad (5)$$

where Y_{it} can be 1) total revenues, 2) transactions' number, 3) average revenues per transaction, all calculated across merchants located in grid cell i at time t . Equation 5 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.

The measure of market access, MA_{it} , employed for the estimation follows the model described in Section 3. It leverages the modal choice problem in Section 3 and observed bilateral travel times described in Section 4. More specifically, observed bilateral travel times are combined with data on commuting flows from the 2018 Census (INSEE, 2018) to estimate the structural parameters of the modal choice problem. The estimated parameters are subsequently used to construct bilateral expected travel costs. Together with observed bilateral travel flows, these are needed to estimate the semi-elasticity of consumption-related

travel flows to travel costs, ν , which is finally used to construct an empirical counterpart of market access according to Equation 4, as follows:

$$\text{MA}_{it} = \sum_{ij} \frac{\exp(-\widehat{\nu} \bar{t}_{ij,t})}{\sum_s \exp(-\widehat{\nu} \bar{t}_{sj,t})} \text{Population}_j \times \text{Median income}_j. \quad (6)$$

where $\widehat{\nu}$ is estimated and equal to 0.01 (see Appendix C). 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 (fixed at 2015), as well as pre-plan travel times by public transport, car and walking while constructing $\bar{t}_{ij,t}$. The destination-specific amenity parameter E_i is excluded from the measure of market access but accounted for in the regressions through grid cell fixed effects.¹⁵


Figure 2 portrays the change in market access induced by the development of cycling infrastructure. First, market access rises in places situated in the proximity of a new bike lane if the latter connects with the rest of the network. For example, the construction of bike lanes in the southwest quadrant of the city did not trigger a sizeable expansion in market access since these lanes were poorly connected with other parts of the network. Second, market access in given locations rises only if the reduction in bilateral travel costs exceeds the average reduction experienced by other locations. This is a direct consequence of the conceptual framework, according to which consumers shop in only one location.

The potentially non-random placement of transport infrastructure is a challenge to the estimation of the elasticity of local economic activity to market access. A local government might push for the new infrastructure to cross a given neighbourhood to revitalise it. The endogeneity of transport infrastructure placement might also stem from the initiative of private interest groups, who might lobby the municipal or local government into developing or not developing new infrastructure in their neighbourhood for their 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, forthcoming). This identification strategy requires including a local bike lane density measure in our specification. By controlling for the intensity of local bike lane

¹⁵As bike travel costs go down two things happen: 1) the unweighted average travel cost to go from j to i drops, and 2) the share of consumers choosing to go shopping by bike increases, thus magnifying the impact of the reduction in bike travel costs on the bilateral disutility of travelling (which also drops). While the number of consumers going shopping from j to i by bike unambiguously rises, the number of consumers choosing to drive might decline if the negative impact on the car modal share outweighs the expected increase in the number of consumers living in j choosing to go shopping in i . The chosen setup thus allows to account for substitution across transport modes. However, it should be noted that the choice of using pre-plan car travel costs might lead to an underestimation of the just described substitution pattern, especially if the development of bike lanes occurred alongside the introduction of speed limits or the elimination of car lanes that led to an increase in car travel costs.

Figure 2: Bike-induced change in quarterly market access (2015q1-2019q4) (EUR)



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

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 allows 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 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 refer to as “projects”). In our favourite specification, $LBLD_{it}$ corresponds to the total length of a given bike lane or project crossing grid cell i as of time t . We take this as our favourite definition of local bike lane density since we deem it likely that development in a given grid cell is influenced by the 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 neighbouring 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 neighbours.

Finally, we run an instrumental variable version of Equation 5, without the local bike lane density control. Market access is instrumented with a market access proxy constructed using exclusively bilateral travel times with locations farther than one kilometre:

$$\text{MA}_{it}^{1km} = \sum_{ij|dist_{ij}>1km} \frac{\exp(-\widehat{\nu}_{ij,t})}{\sum_s \exp(-\widehat{\nu}_{sj,t})} \text{Population}_j \times \text{Median income}_j. \quad (7)$$

Alternatively, we also test an instrument version corresponding to market access excluding origin-destination pairs located along the same bike lane or project.

6 Results

6.1 Baseline results

Table 2 presents the estimates of β from Equation 5, divided into 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 across all specifications. In column 1, we exclude from Equation 5 the local bike lane density control. In columns 2 and 3, we include in turn our two measures of local bike lane density. The instrumental variable results are presented in columns 4 and 5 and are quantitatively similar to the OLS and similarly significant, thus suggesting that the potential bias caused by the endogeneity is LBLD is less of a concern. Conversely, we do not detect a statistically significant elasticity of average revenues per transaction in any of the specifications.

The magnitude of the coefficients oscillates between 4.10 and 5.17. To interpret them, we rely on the distribution of changes in our market access (MA) measure across time. Because our measure of MA allows for consumers to travel by different modes (driving, walking, cycling and public transit), and cycling is the least common mode choice, even relatively sizeable decreases in travel times by bike will lead to fairly modest increases in our MA measure. In grid cells where MA increased, market access went up by 0.93% on average 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 over the whole period.¹⁶ We provide a more detailed quantification of the impact of the infrastructure 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 A4, 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

¹⁶We multiply 0.0093 times 5.173 for the first one, and 0.0093 times 4.603 for the second one.

Table 2: Elasticity of local economic activity to market access: baseline evidence

Panel A:	Log total revenues				
	(1)	(2)	(3)	(4)	(5)
Log MA	4.107** (1.882)	5.173** (2.280)	4.500* (2.447)	4.406** (1.91)	4.125** (1.85)
Panel B:	Log transactions' volume				
Log MA	5.062** (2.041)	4.603** (2.160)	4.996** (2.298)	4.754** (2.05)	5.005** (2.00)
Panel C:	Log average revenues per transaction				
Log MA	-0.933 (1.202)	0.577 (1.404)	-0.482 (1.484)	-0.339 (1.24)	-0.859 (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	Neighbours	-	-
Instrument				Exclude 1km	Exclude project pairs
FS F-stat				1400.93	110718.35
Estimation		OLS			2SLS

Notes: coefficients from the estimation of Equation 5. Standard errors are clustered at the grid cell level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Groupement des Cartes Bancaires CB*.

further back in time are considered.

Because of 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 affect positively the revenues of local merchants also through an amenity channel by, for example, improving the appearance of the sidewalk or street in which they are located. These two channels are correlated but not entirely collinear. Let us consider the example of two businesses both located on streets where the city council decides to build spacious bike lanes. Both streets become cleaner and safer, as the city council will also likely implement speed restrictions for cars and car traffic will go down. As a result, more people might choose to go shopping or eat out in those two places by an “amenity” channel. However, if we assume that only the first of the bike lanes connects with the existing cycling network, then business volume should rise even more for the first business since this is not only a nicer place to shop, but it has also become easier to reach. In our setup, the local bike lane density control, $LBLD_{it}$, will act as a control for the amenity channel, thus allowing MA_{it} to identify the connectivity channel.

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, 2023).¹⁷ In column 2 of Table 3, we follow Chandra and Thompson (2000) and exclude central districts, specifically *arrondissements* 1 to 4. In column 3, we remove transport hubs. Using information on the public transport network (metro, tramway, and suburban trains),¹⁸ we define as transport hubs those grid cells located less than 500 meters away from stations featuring three or more public transport connections.¹⁹ All coefficients in Table 3 remain positive and statistically significant, thus eliminating centrality bias as a concern in our setting. Finally, in columns 4 and 5 we report the coefficients of a long difference version of Equation 5. Here the starting period is an average between 2015Q1 and 2015Q2 and the end period is an average between 2019Q3 and 2019Q4. Column 5 adds as a control the log of the distance of each grid cell from the city centre (Paris city hall). If central areas receive disproportionate increases in market access due to the intrinsic network structure of the problem and the bias increases with proximity to the city centre, this specification would allow to correct for it (Coşar et al., 2022). In line with the results of the other robustness checks, the coefficients of interest are marginally larger but not significantly different.

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 through the following regression:

$$\ln(Y_{it}) = \alpha_i + \alpha_{dt} + \sum_t \beta^t \Delta \ln(MA_{i,15-19}) \times \tau_t + \gamma X_{it} + \delta LBLD_{it} + e_{it} \quad (8)$$

¹⁷The centrality bias test proposed by Borusyak and Hull (2023) is not computationally feasible in our context due to the very granular scope of our analysis. Currently, it takes approximately 7 hours to calculate all travel times by bike for every quarter between 2015 and 2019. If followed Borusyak and Hull (2023) and calculated 999 counterfactual travel time matrices for every quarter between 2015 and 2019, this would take approximately 291 days of computation time.

¹⁸Data come from Île-de-France Mobilités website.

¹⁹The 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.

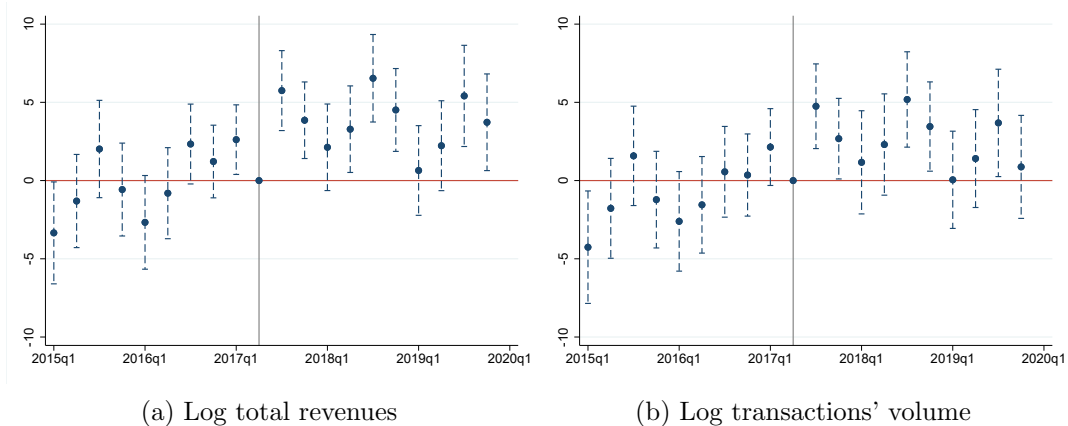
Table 3: Robustness tests: dealing with centrality bias

Panel A:	Total revenues					
	Log-log			Long difference		
	(1)	(2)	(3)	(4)	(5)	
MA	5.173** (2.280)	5.225** (2.353)	5.278** (2.284)	10.655*** (2.491)	11.027*** (2.506)	
Panel B:	Transactions' volume					
	MA	4.603** (2.160)	4.793** (2.230)	4.770** (2.158)	8.348*** (2.746)	8.633*** (2.763)
	N	27,097	25,197	26,817	1,352	1,352
Test	Baseline	Remove central districts	Remove transport hubs	Baseline	Log distance from city centre	

Notes: baseline estimation as in Table 2 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; long-difference specification: $\Delta \ln(Y_i) = \alpha_d + \Delta \beta \ln(MA_i) + \Delta \gamma X_i + \Delta \delta LBLD_i + \epsilon_i$ (col.4); long-difference specification: $\Delta \ln(Y_i) = \alpha_d + \Delta \beta \ln(MA_i) + \Delta \gamma X_i + \Delta \delta LBLD_i + \varphi \ln dist_i + \epsilon_i$ (col.5). Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Groupement des Cartes Bancaires CB*.

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 3, places that experienced greater 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.

Figure 3: Pre-trends analysis



Notes: estimated β^t from Equation 8 on the y-axis. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Groupement des Cartes Bancaires CB*.

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 i again to be developed at time $t = t_0$, $D_i^{t_0} = 0$ (and eventually developed), and regress the development date on the set of demographic variables employed as controls in the main specification:

$$\text{Development date}_i | (D_i^{t_0} = 0) = \alpha + \beta X_i^{t_0} + e_i \quad (9)$$

We run Equation 9 for different choices of $t = t_0$ and report the results of the estimation in Table A5. 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 of our sample period, during which 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 control variables 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 should have featured some bike lane development according to the original *Plan Vélo* (see Figure A7). Since these cells were selected to be part of the original plan according to the same logic, we can expect this subsample to be more homogeneous than the baseline one. We test this hypothesis through a balancing test (Table A7), 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 A6. Columns 1 and 3 reproduce the baseline results from column 2 of Table 2, 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. The average improvement in market access in this subsample (0.0125) implies an increase in total revenues and transactions' volume by 10.57% and 11.56%, respectively.

Household sorting — Evidence has shown that investment in infrastructure can impact household sorting [Tsivanidis \(forthcoming\)](#). If this is the case, it could be that part of the effect of an improvement in market access on business revenues is due to a different household composition rather than the development of bike lane infrastructure. For instance, the development of bike lanes might have triggered the inflow of young people, who tend to go out and consume more in bars/restaurants. To clean our estimates from the potential effect of changing local household composition, in our baseline specification we control for (a three-year lag of) local demographic characteristics. Here, we further investigate how neighbourhoods with different levels of market access changed during the 5 years of the analysis. We do so by running the pre-trend analysis as in Equation 8, using as an outcome the available demographic controls.²⁰ Results reported in Figure A6 show that places that experienced greater market access improvements by the end of the period did not experience any statistically significant variation in demographic composition. This evidence is consistent with household sorting being a process occurring in the medium-run horizon, while the time frame of our analysis ends three years after the beginning of the intervention.

Other potentially confounding factors — In Table A8, 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.²¹ Second, in columns 3 and 4, we test the robustness of our results to a law passed in 2015 that allowed businesses located in certain parts of the city to stay open on Sundays.²² 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 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 greater market access improvement is partly driven by an increase in card usage and substitution

²⁰We remove the controls $X_{i,t}$ for the purpose of this test. Due to data availability, this test is performed annually instead of quarterly.

²¹Specifically, we include the lagged share of revenues for each grid cell and time in the following non-tradable 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).

²²See the map of concerned places in Figure A8. Data come from APUR, Mairie de Paris and DRIEA IF/UD75.

away from cash payments. To do so, we check if market access increases the share of establishments featuring card activity in our database. First, we build a card usage intensity index by dividing the number of establishments present in the *Groupement des Cartes Bancaires CB* dataset over the number of establishments that should be active in that same place and quarter according to the national business registry. Then, we re-run Equation 5 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 A9 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 establishment characteristics. The characteristics on which we run the clustering algorithm are 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.²³ Next, we run a *k-means* clustering algorithm (Bock, 2007) for different values of the number of clusters, k , and we select $k = 5$ through an elbow test as shown in Figure A9.²⁴

The characteristics of the five clusters in terms of the variables used for the clustering exercise are reported in Table A10. The degree of specialisation in stores that sell essential goods, 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 in fast food, restaurants or bars. Table A11 contains the estimated coefficients from a fully-interacted version of Equation 5. The evidence suggests a positive and statistically significant impact of a

²³We 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.

²⁴The 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 picking 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.

market access improvement on economic activity for grid cells specialised in 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 helps improve establishments' salience, particularly for younger, smaller and thus less known businesses. Thanks to the new bike lanes, people commute more on the surface and through a transport mode that is easier to park than cars. Hence, they become more easily 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. An improvement in market access may encourage new firms to enter the market. However, our estimates in column 1 of Table A12 do not support this hypothesis. A potential explanation can be a rise in commercial rents. As market access improves in a given area, the rental rate a prospective 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 data on commercial rents to test this hypothesis.

Next, we investigate whether an improvement in market access has an impact on residential housing prices. We construct a price index by running a hedonic regression of the (log) of housing prices on individual properties' characteristics, in addition to grid cell and time fixed effects (see Section 4.2 for more details on the dataset used). Our estimated coefficient is not statistically significant (column 2 of Table A12). 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 kilometre 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 housing prices underscores the weaker capitalisation of cycling infrastructure improvements in housing 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.²⁵ We find a negative impact of market access on car traffic (column 3 of Table A12). Multiple mechanisms can explain this result. 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 discourage car usage. By favouring modal switching, the introduction of bike lanes helps attenuate the negative externalities associated with car congestion, with a positive impact likely extending beyond the areas directly interested by the development of bike lanes (Hall, 2021).

6.5 Distributional consequences of *Plan Vélo*

The new infrastructure entailed significant spatial reallocation of spending. The approach taken in this paper is indeed well-suited to analyse the distributional consequences of the new infrastructure, while it lacks the features necessary to analyse the absolute gains. Market access in a given place increases if bilateral travel costs to that place decline, on average, more than to other places in the city.²⁶ This means, essentially, that market access gains (and, thus, potential demand) in a given place can take place only at the expense of other locations. Our framework enables us to identify what are the places that have gained in relative terms and those that lost with 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 2), 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 2, 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 on average by lower income (Figure A10), owing primarily to the distribution of population. Hence, *Plan Vélo* redistributed 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 located respectively in the south-west and

²⁵see Section 4.2 for more details on the data.

²⁶See Section 3 for the formal definition of our market access measure.

north-west quadrants 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 parts of the city with the city centre, a larger fraction of residents choose 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 it easier to go shopping in peripheral ones. However, central districts in Paris are mostly shopping locations and have low population density, so the change in spending patterns is more likely to benefit businesses located in the city centre at the disadvantage of those located in peripheral districts.

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 different ways. It can reshape the geography of spending towards locations that become better accessible. Furthermore, by favouring a switch to active mobility, it can benefit certain local businesses by making them more salient, and 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 empirically evaluate the impact of the construction of a large-scale bike infrastructure network, the *Plan Vélo*, occurring between 2017 and 2019 in the city of Paris, on businesses operating in the non-tradable sector. We find robust evidence in favour of an increase in economic activity, as proxied by the total value and volume of card transactions in parts of the city subject to an increase in market access triggered by the development of the new infrastructure.

Our analysis of the distributional impact of the new infrastructure highlights how the *Plan Vélo* redistributed economic activity away from more peripheral/more densely populated districts towards more central/more scarcely populated ones. The analysis of the distributional consequences are useful for current policy-makers, especially given that the second edition of *Plan Vélo* 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.
- BERNARD, L. (2023): “The impact of segregated cycling lanes on road users,” *Available at SSRN 4353625*.
- 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 (2023): “Nonrandom exposure to exogenous shocks,” *Econometrica*, 91, 2155–2185.
- BOU SLEIMAN, L. (2024): “Are car-free centers detrimental to the periphery? Evidence from the pedestrianization of the Parisian riverbank,” Working paper.
- 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,” .
- 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.
- COŞAR, A. K., B. DEMIR, D. GHOSE, AND N. YOUNG (2022): “Road capacity, domestic trade and regional outcomes,” *Journal of Economic Geography*, 22, 901–929.
- 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., National Bureau of Economic Research.
- 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 (2023): “Social Benefits and Private Costs of Driving Restriction Policies: The Impact of Madrid Central on Congestion, Pollution, and Consumer Spending,” *Journal of the European Economic Association*, 21, 1227—1267.
- GARCIA-LÓPEZ, M. A., M. MAGAGNOLI, AND E. VILADECANS (2024): “People on bikes getting coffee. The impact of cycle lanes in cities,” *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,” *Manuscript, University of Pennsylvania*.
- GROUPEMENT DES CARTES BANCAIRES (2019): “Cartes Bancaires en chiffres,” <https://www.cartes-bancaires.com/a-propos/cb-en-chiffres/>.
- HALL, J. (2021): “Can Tolling Help Everyone? Estimating the Aggregate and Distributional Consequences of Congestion Pricing,” *Journal of the European Economic Association*, 19, 441—474.
- 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 (forthcoming): “Growth off the rails: Aggregate productivity growth in distorted economies,” *Journal of Political Economy*.
- IEA (2023): “Greenhouse gas emissions from Energy Data Explorer,” <https://www.iea.org/data-and-statistics/data-tools/greenhouse-gas-emissions-from-energy-data-explorer>.
- 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.
- LANDAIS, C., D. BOUNIE, Y. CAMARA, E. FIZE, J. W. GALBRAITH, C. LAVEST, T. PAZEM, AND B. SAVATIER (2020): “Consumption Dynamics in the COVID Crisis: Real Time Insights from French Transaction & Bank Data,” CEPR Discussion Papers 15474, C.E.P.R. Discussion Papers.
- MAIRIE DE PARIS (2015): “Paris Capital du Vélo 2020,” *report*.
- MALIGORNE, C. (2017): “À Paris, la voiture est déjà une espèce en voie de disparition,” <https://www.lefigaro.fr/economie/le-scan-eco/dessous-chiffres/2017/10/12/29006-20171012ARTFIG00166--paris-la-voiture-est-deja-une-espece-en-voie-d-e-disparition.php>, accessed: 2024-7-29.

- MAYER, T. AND C. TREVIEN (2017): “The impact of urban public transportation evidence from the Paris region,” *Journal of Urban Economics*, 102, 1–21.
- McFADDEN, D. (1974): “The measurement of urban travel demand,” *Journal of Public Economics*, 3, 303–328.
- MIYAUCHI, Y., K. NAKAJIMA, AND S. J. REDDING (2021): “The economics of spatial mobility: Theory and evidence using smartphone data,” Tech. rep., National Bureau of Economic Research.
- OECD (2023): “OECD Regional Database,” .
- 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. (2022): “Is online retail killing coffee shops? estimating the winners and losers of online retail using customer transaction microdata,” .
- 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*.
- TASSINARI, F. (2024): “Low emission zones and traffic congestion: evidence from Madrid Central,” *Transportation Research Part A: Policy and Practice*, 185, 104099.
- THORNE, V. (2022): “Cycling Towards Cleaner Cities? Evidence from New York City’s Bike Share Program,” Tech. rep.
- 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 Additional figures and tables

Table A1: 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 Appendix Section B.

Table A2: Initial travel length and changes in travel times by bike

District	Initial commute				Decrease in commute Bike
	Bike	Car	Public transit	Walking	
Most central districts					
1	Fast	Fast	Fast	Fast	Intermediate
2	Fast	Fast	Fast	Fast	Slow
3	Fast	Fast	Fast	Fast	Intermediate
4	Fast	Fast	Fast	Fast	Intermediate
Intermediate districts					
5	Intermediate	Intermediate	Intermediate	Intermediate	Slow
6	Fast	Fast	Intermediate	Fast	Slow
7	Intermediate	Intermediate	Intermediate	Intermediate	Intermediate
8	Intermediate	Intermediate	Intermediate	Intermediate	Intermediate
9	Intermediate	Intermediate	Fast	Intermediate	Slow
10	Intermediate	Intermediate	Intermediate	Intermediate	Intermediate
11	Intermediate	Intermediate	Intermediate	Intermediate	Fast
Less central districts					
12	Intermediate	Intermediate	Intermediate	Slow	Fast
13	Intermediate	Intermediate	Intermediate	Intermediate	Intermediate
14	Intermediate	Intermediate	Slow	Intermediate	Slow
15	Slow	Slow	Slow	Intermediate	Intermediate
16	Slow	Slow	Slow	Slow	Fast
17	Slow	Slow	Intermediate	Slow	Fast
18	Intermediate	Intermediate	Intermediate	Intermediate	Intermediate
19	Slow	Slow	Slow	Slow	Fast
20	Slow	Slow	Slow	Slow	Intermediate

Notes: Initial refers to 2015Q1. Bilateral travel times are averaged by district. “Fast” initial commute (decrease in commute) if initial commute (decrease in commute) lies in the bottom 25th percentile of the district-specific travel time (decrease in travel time) distribution; “intermediate” if it lies between the 25th and 75th percentile; “slow” if it lies in the top 25th percentile.

Table A3: Estimated modal choice parameters

Description	Parameter	Value
Disutility of travel 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 A4: 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 2 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 A5: Robustness tests: testing random development timing

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

Notes: the dependent variable is the date on 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 A6: Robustness tests: keeping only Plan Vélo subsample

	Log total revenues		Log transactions' volume	
	(1)	(2)	(3)	(4)
Log MA	5.173** (2.280)	8.129*** (2.953)	4.603** (2.160)	8.897*** (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 2 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 A7: 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		Difference	t-stat	p-value
	Mean	Std Dev	Mean	Std Dev			
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 A8: Robustness tests: miscellanea

Panel A:	Log total revenues			
	(1)	(2)	(3)	(4)
Log MA	5.173** (2.280)	4.264* (2.206)	6.814*** (2.575)	5.211** (2.280)
Panel B:	Log transactions' volume			
Log MA	4.603** (2.160)	4.140* (2.307)	5.970** (2.418)	4.652** (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 2 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 A9: Robustness tests: elasticity of 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 2 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 A10: 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 A11: Testing heterogeneous effects with respect to local merchant characteristics

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


Notes: baseline estimation as in Table 2 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 A12: Elasticity of other outcomes to market access

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


Notes: coefficients from the estimation of Equation 5, 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 A1: Total number of bike trips recorded in Paris over time



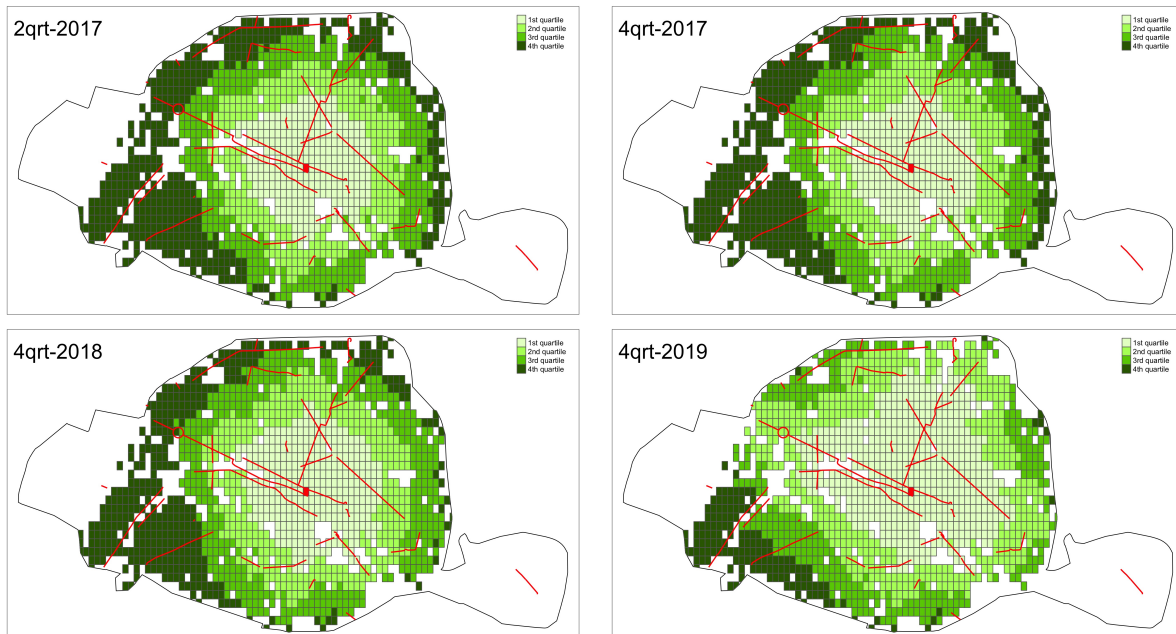
Notes: all bike trips are recorded by sensors distributed across the city. Source: *Comptage vélo - Données compteurs* dataset from <https://www.data.gouv.fr/en/datasets/comptage-velo-historique-donnees-compteurs-et-sites-de-comptage/>.

Figure A2: Total card transaction revenues taking place in Paris over time



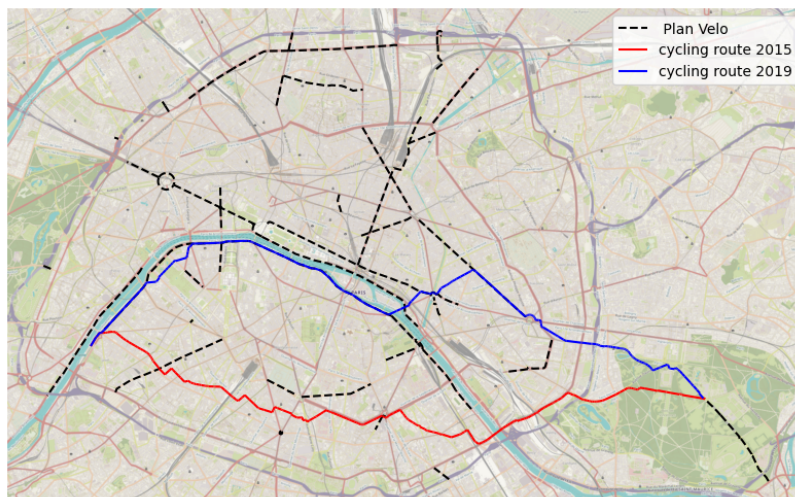
Source: *Groupement des Cartes Bancaires CB*.

Figure A3: 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 A4: Example of cycling route before and after *Plan Vélo*



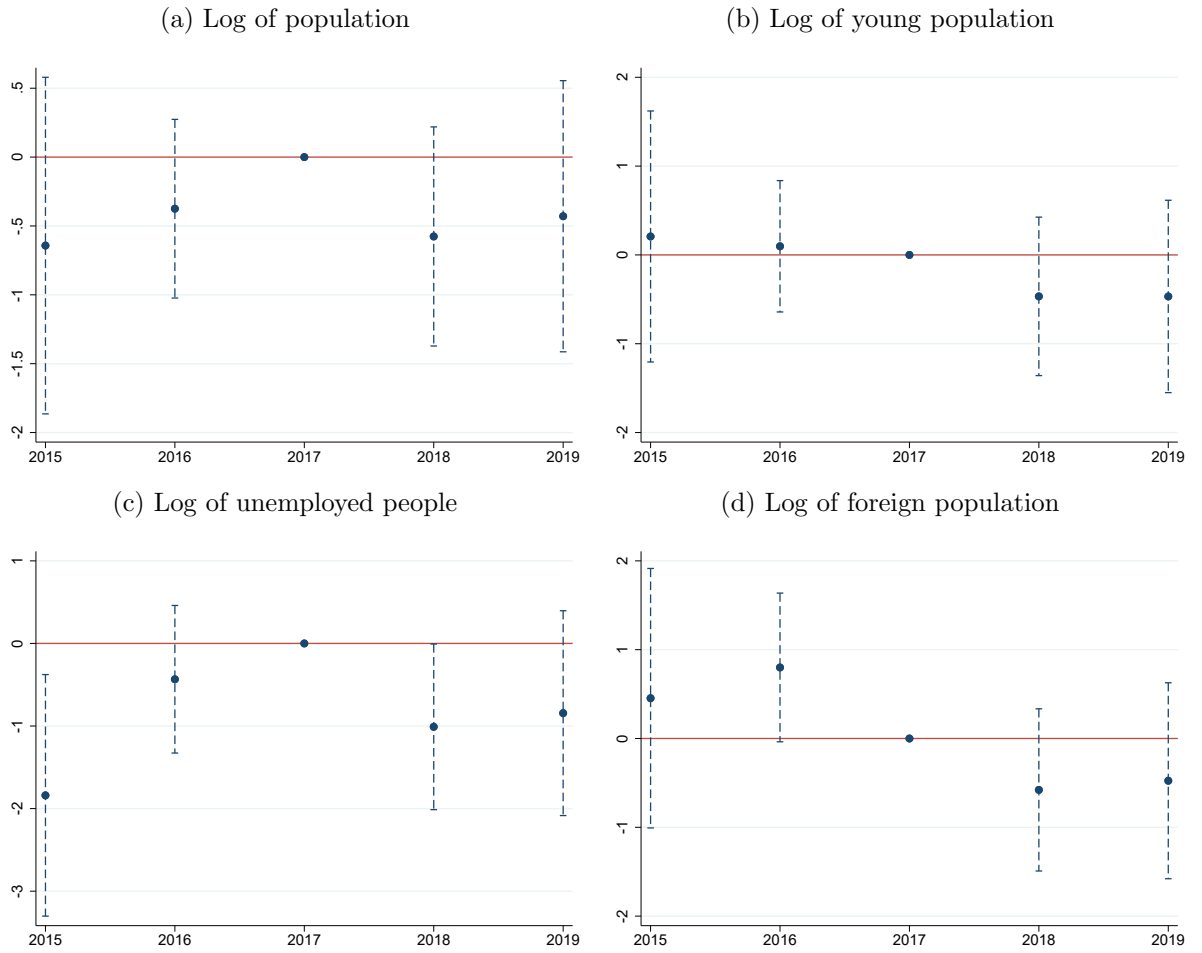
Notes: example of the least-cost path for cycling calculated by our routing algorithm for the first quarter of 2015 (before the *Plan Vélo*), in red, and for the last quarter of 2019 (after the *Plan Vélo*), in blue. The trip in 2015 was estimated to take 114 minutes, while the trip in 2019 would take 92 minutes, implying a 20% reduction in travel time. Back to Appendix Section B.

Figure A5: Bilateral travel times by bike: Google Maps vs own estimates



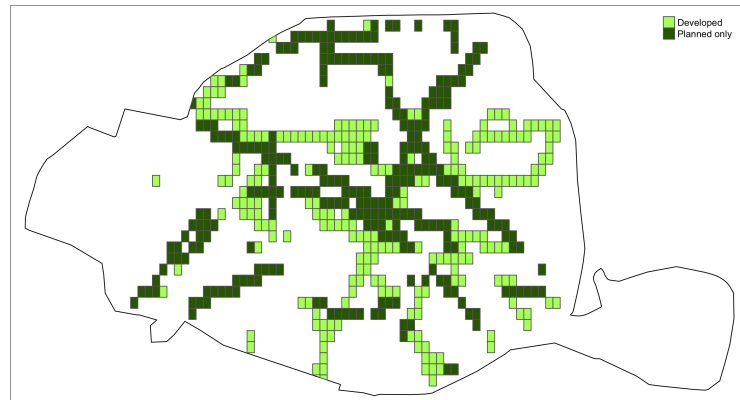
Notes: comparison between bilateral travel times estimated by Google Maps on June 2024 for cycling trips in Paris and the travel time estimates for those same origin-destination pairs for the fourth quarter of 2019 using our routing algorithm. 8,000 origin-destination pairs were chosen at random from the set of all trips from the centroid of one cell to the centroid of another cell in Paris. We then queried these trips on Google Maps and compared the estimated travel times to those produced by our routing algorithm for Q4, 2019. The red line represents a linear fit of the data, with a coefficient of 0.61 and an intercept of 1.2 minutes. The R-squared of this regression is 0.9. Back to Appendix Section [B](#).

Figure A6: Household sorting: changes in demographic characteristics



Notes: estimated β^t from Equation 8 on the y-axis. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Groupement des Cartes Bancaires CB*. Back to Section 6.2.

Figure A7: Grid cells crossed by 2015 Plan Vélo




Notes: blue grid cells are those that were 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 A8: 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 A9: 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 A10: Spatial distribution of purchasing power



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

B Calculating bilateral travel times by mode

OpenStreetMap is used to estimate transport-mode-specific bilateral travel times between the centroid of each pair of grid cells forming the geographical unit of our analysis. More specifically, 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) for each quarter from 2015 to 2019²⁷; iii) and General Transit Feed Specification (GTFS) files for the city of Paris for each quarter from 2015 to 2019, provided by the Paris main public transit operator, RATP group.²⁸

Bilateral travel time matrices (in minutes) for public transit are obtained through the `travel_time_matrix()` function of the *Rapid Realistic Routing with R5* package for R.²⁹ This package uses Conveyal’s R5 Routing Engine to calculate realistic travel times allowing it 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.

The construction of travel time matrices for the other transport modes (walking, cycling, and driving) follows a different, more flexible approach. We start by extracting the network of all *ways*³⁰ that is 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 speeds of the cycling network, we rely on information from OSM to classify each edge into six categories, summarised in Table A1. 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.³¹ The *Plan Vélo* provides this type of cycling infrastructure. As the new cycling infrastructure is built, information from the *Plan Vélo* is matched with the OSM network to transform the roads that overlap with the *Plan Vélo* into cycle tracks, irrespective of their status in OSM data. By doing this, we can match the timing of the new infrastructure more precisely, given that OSM is a crowd-sourced project that might suffer from a small lag in updating the true

²⁷Information corresponding to the first day of February, May, August, and November of every year from 2015 to 2019.

²⁸GTFS 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.

²⁹See [Pereira et al. \(2021\)](#) for more details on this package.

³⁰A way is defined as any linear feature of a map, such as a road, a sidewalk, a river, etc.

³¹These are equivalent to what we refer to as “bike lanes” when discussing the *Plan Vélo* infrastructure.

conditions on the ground. 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 lanes³² and residential roads) to 0.17 times the speed of cycle tracks for primary, trunk and motorway roads.³³

Having derived travel speeds for each edge in each type of network (driving, cycling and walking),³⁴ we apply Dijkstra’s algorithm to find the minimum travel time between each pair of centroids, by network.³⁵ We allow travel times to be asymmetric, namely the travel time from point A to point B can be different from the travel time from point B to point A, and repeat the calculation for every type of network (cycling, walking and driving) and quarter.

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 A4 illustrates how our routing algorithm captures the impact of the new bike infrastructure on the least-cost path for cycling. In this figure, we show the route calculated by our routing algorithm for the same trip taken by bike in the first quarter of 2015 (in red) and then for the last quarter of 2019 (in blue). We can see that after the *Plan Vélo* is implemented, the optimal path favours streets with bike lanes, resulting in a slightly longer trip (measured in distance), but that is 20% faster (going from 114 minutes in 2015 to 92 minutes in 2019).

In Figure A5, we compare the travel times by bike estimated by Google Maps on June 2024 for a set of 8,000 random origin-destination pairs to the travel times for these same trips estimated by our routing algorithm in the last quarter of 2019. We can see that our estimates are highly correlated with those produced by Google Maps. A linear regression of our estimates on the Google Maps travel times produces a coefficient of 0.61 (significant at a 1% level), and an intercept of 1.2 minutes. This suggests that our travel time estimates are largely proportional to those estimated by Google Maps.

³²Cycle 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.

³³These adjustments to the cycling speed follow a similar logic to Broach (2016), which provides generalised cost formulas for cycling, where the cost of cycling depends on the type of infrastructure and the density of traffic, among other factors.

³⁴An 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.

³⁵To 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. Technically, we 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.

C Derivation of the market access measure

The derivation of the market access measure is articulated in the following steps.

Step 1: conditional logit estimation — Observed bilateral travel costs are converted into disutilities (Equation 1), d_{ij} , through a modal choice conditional logit estimation (Section 3) based 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 transport mode chosen to commute to work. Information on the place of residence/work in the survey is available at the district level, for a total of 20 districts, or *arrondissements*. The 2230-by-2230 commuting cost matrices are then aggregated into 20-by-20 matrices to match 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.³⁶

The output of the conditional logit estimation is displayed in Table A3. We estimate a disutility of travel 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).³⁷ The smaller size of the municipality of Paris compared to Bogotá could explain the discrepancy: shorter distances reduce consumers’ responsiveness to differences in travel times across transport modes.³⁸ 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, implying a preference by consumers for walking between two destinations holding travel time constant, as opposed to cycling or driving. Conversely, the common preference shifter for public transport is positive, denoting a preference for public transport compared to walking, once again assuming identical travel times.

The estimated parameters are combined with information on the car ownership rate in the city of Paris $\rho = 0.37$ to obtain an estimate of expected bilateral travel costs \bar{t}_{ij} according to Equation 2.^{39,40}

³⁶Specifically, we tried aggregating travel costs using the median as well as bilateral population-weighted means.

³⁷The 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.

³⁸The modal shares are: 71% by public transport, 7% by bike, 16% by walking and 6% by car.

³⁹A 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 $-(1/\kappa) \ln(\exp(b_{car}) + 1)$.

⁴⁰Paris car ownership rate is taken from Maligorne (2017)

Step 2: estimating the semi-elasticity of consumption-related travel flows to travel costs — To build an empirical counterpart of the market access measure described in Equation 4, an estimate of ν , namely the semi-elasticity of travel flows with respect to bilateral travel costs, is needed. However, bilateral travel flows are not directly observed in our dataset. Hence, we developed an imputation procedure to calculate a proxy for them in 2019. Specifically, we observe daily transactions indexed by the merchant and card identifier. We do not know where the cardholder lives, nor do we have any demographic information on him/her. However, we can impute a “most likely” residence location based on each cardholder’s shopping history.

1. First, we retain transactions occurring in the city of Paris on weekends and on weekdays after 18h;
2. Second, 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. Next, we 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 . Bilateral travel flows, x_{ij} , are calculated by summing across transactions carried out by cardholders with imputed residence location j towards merchants with (known) business location i , and they are used to estimate the empirical counterpart of Equation 3:

$$\ln x_{ij} = \alpha_i + \alpha_j + \nu \bar{t}_{ij} + e_{ij} \quad (10)$$

In Equation 10, \bar{t}_{ij} are the bilateral travel costs obtained in the previous step, and α_i and α_j are, respectively, business and residence location fixed effects. The estimation is repeated for four different quarters of 2018. We consistently estimate $\hat{\nu} = 0.1$, similar to 0.07 in Ahlfeldt et al. (2015).

A market access measure is constructed in line with Equation 4 taking as inputs the estimate for $\hat{\nu}$ and $\bar{t}_{ij} \forall i, j$, and used in the main estimating equation of the paper (Equation 5).