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, David Bounie, Youssouf Camara

April 2022

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

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

[1]Banca d'Italia

[2]CY Cergy Paris Université

[3]i3, CNRS, Telecom Paris, Institut Polytechnique de Paris

April 1, 2022

Abstract

Cities are increasingly interested in developing cycling infrastructure. Yet, little is known about the (potentially heterogeneous) economic impact of such investments. We evaluate the economic impact of a large-scale cycling infrastructure investment in Paris, the *Plan Vélo*. Using geolocated data covering nearly the universe of French card transactions we estimate a positive and statistically significant elasticity of local revenues to bike market access. We find a larger elasticity in areas with smaller and younger establishments. The project increased the non-tradables consumption share of central/less densely populated neighborhoods 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: federica.daniele@bancaditalia.it. 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, La Banque Postale, la Caisse des Dépôts et Consignations, Telecom Paris and University of Paris 2 Panthéon-Assas. We are grateful for the insightful comments received by: Gabriel Ahlfeldt, Raffaello Bronzini, Ana Moreno-Maldonado, Clara Santamaria and Marta Santamaria.

1 Introduction

The use of the bicycle as a commuting vehicle is increasingly common in a large number of European cities. Local governments are ramping up their efforts to develop appropriate supporting infrastructures, such as the construction of large bike lane networks or the adoption of bike-sharing systems. While there is general consensus on the positive impact of developing bicycle infrastructure on pollution reduction, the impact on local economic activity is instead a relatively under-investigated subject. Yet, local shop-owners tend to share the conviction that the construction of bike lanes threatens their business (since it may make more difficult for car drivers to park their car), which often brings them to campaign against it.[1]

This article analyses how the development of a large-scale bicycle network during 2017-2019, the *Plan Vélo* in the city of Paris, affected local economic activity and reshaped the geography of consumption in the non-tradables sector. We leverage the staggered development of the infrastructure and the changes induced on bilateral transport costs to construct a time and space-varying bike market access metric. We then estimate the elasticity of local revenues to bike market access using geolocated data covering nearly the universe of French card transactions. Armed with this elasticity, we conduct a counterfactual exercise where we estimate the cumulative gain/loss for local businesses from the development of the new infrastructure.

Until very recently, the commonly accepted norm in empirical studies evaluating the impact of newly developed transport infrastructure was to take the simple existence in a given location of a new infrastructure as the treatment. These studies tended to ignore the network dimension of transport infrastructure benefits, namely that the local gains from a new transport infrastructure depend on the characteristics of the destinations becoming easier to access. In line with recent advances in the economic geography literature, we overcome this shortcoming by deriving the economic impact of the new transport infrastructure through a market access approach. Simply put, market access from a merchant viewpoint corresponds to how many customers can reach this merchant with relatively little effort/time via the available transport network. We measure market access from the viewpoint of cyclists: at any given point in time, we construct a pixelised cost surface of the city of Paris, where low-cost pixels correspond to pixels that are relatively easier to cross by bike. Next, we run a least cost path algorithm to obtain the full set of bilateral commuting costs for each

[1] Various case-study show that retailers systematically overestimate the size of biking customers. See the press article on Bloomberg CityLab "The Complete Business Case for Converting Street Parking Into Bike Lanes" from https://www.bloomberg.com.

origin/destination pair and quarter. Finally, combining these with the spatial distribution of income, we construct a bike market access measure for each merchant location and quarter, given essentially by a weighted sum of spending power across all consumer locations, with weights inversely proportional to the bilateral commuting cost by bike.

Using the bike market access metric, we estimate the elasticity of local revenues, as proxied by card transaction total revenues, to bike market access. The decision to develop a bike lane might be endogenous to economic activity in a given location. To bypass this potential identification threat, we let the elasticity be identified by the development of bike lanes in more distant parts of the network, which can be considered exogenous with respect to the economic conditions of a given location. We estimate an elasticity to bike market access equal to 0.45 for total revenues and 0.50 for transactions' volume. We furthermore find that the elasticity is weakly increasing in the time elapsed since variation in bike market access. Conversely, the elasticity on average revenues p/transaction is not statistically distinguishable from zero. We argue that the estimated elasticities identify primarily a connectivity channel, as opposed to an amenity one, which is being taken care of through the inclusion of the local bike lane density measure.

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

We investigate the existence of heterogeneous effects by dividing the city into clusters with similar merchant characteristics. We find that a positive and statistically significant elasticity of local economic activity to bike market access seems to be upheld by the local presence of younger merchants and/or by the specialisation into food-related industries that tend to be more youth-oriented or characterised by lower average end-of-the-meal bills, such as cafes, fast-food or bars.

The heterogeneity in the estimated elasticity allows us to further shed some light on the underlying mechanisms behind our findings. The key mechanism reflected in the construction of our market access metric and therefore directly tested in our specification is related to the change in the bilateral commuting costs. Thanks to the development of the new infrastructure certain merchant locations become more easily accessible by consumers, that therefore decide to shop there in larger numbers. This mechanism can be amplified by demand-side "footfall" externalities kicking in when consumers rely on active mobility as a transport mode for retail-related shopping trips (Koster et al., 2019). The positive and statistically significant elasticity of transactions' volume to bike market access in retail-specialised neighborhoods supports this amplification channel. Finally, while not being able to formally test it, our results highlight a third potential mechanism, related to the existence of information asymmetries affecting disproportionately younger merchants and to the greater salience subsequently favored by the new cycling infrastructure. The construction of a new cycling infrastructure can in fact favor an increase in the cycling modal share. The latter can in turn benefit local businesses beyond the already mentioned mechanisms if cycling (more than car ridership or public transport) helps commuters/consumers be better aware of the available consumption opportunities. The reason can be that they travel on surface or that they can more easily stop to visit a store located along their ride.

Next, we explore the elasticity of further outcomes to bike market access. While we do not detect a statistically significant positive elasticity of firm entry to bike market access, which would be consistent with a model of monopolistic competition and free entry, we estimate a positive and statistically significant elasticity of house prices. This last finding might help reconcile the first one of the absence of an increase in firm entry following a bike market access improvement, since an increase in the rental rate of commercial real estate might increase the entry cost of starting a new business. More specifically, we find a positive and significant elasticity for house prices equal to 0.063 that appears after a 3-quarter lag. Finally, with the aid of webscraped TripAdvisor data, we confirm the existence of a positive and statistically significant elasticity to bike market access of local economic activity also when the latter is proxied by the number of TripAdvisor reviews.

In the last part of the paper, we identify what are the locations that have gained and those that have lost in relative terms from the development of the new infrastructure. We find that the pre/post difference in total revenues generated exclusively by the variation in bike market access has been positive for 40% of units in our sample and negative for the remaining others. The uneven impact of the infrastructure can be explained by analysing the characteristics of the locations that gained from the development of the bike lane network. Locations where bike market access improved tend to be 1) centrally located, and 2) characterised on average by lower purchasing power (or income). Our analysis highlights that an important consequence of the development of Plan Vélo was to increase the non-tradables consumption share of central/less densely populated neighborhoods at the disadvantage of peripheral/more densely populated ones.

Relation with the literature This paper contributes to several strands of the economic literature. First, we relate to the body of literature that looks at the link between transport infrastructure and economic activity. The first wave in this strand of literature explores this link by looking at the impact of getting access to a new infrastructure on outcomes such as employment (Duranton and Turner, 2012), population (Baum-Snow, 2007; Gonzalez-Navarro and Turner, 2018) and property prices (Gibbons and Machin, 2005; Billings, 2011). A more recent set of papers has adopted a market access approach with respect to the exploration of this link. The aim of this approach is to account for the fact that gains from transport infrastructure investment depend on whether locations get connected to more or less attractive areas (Ahlfeldt et al., 2015; Heblich et al., 2020; Gorback, 2020; Tsivanidis, 2019). With this paper, we contribute to this literature by providing the first empirical assessment of the economic impact of a cycling infrastructure.

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

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

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

2 The *Plan Vélo*

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

In February 2017, an independent observatory, the *Observatoire du Plan Vélo de Paris*, was set up with the purpose of monitoring the advancement of the development of the Plan.[3] At that time, two years had passed since the start of the mandate, and only 4% of planned bike lanes had been effectively developed. The monitoring conducted by the observatory had to be accurate since it represented the main tool to credibly exert pressure on the local administration to speed up the works.

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

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

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

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

little leeway to district mayors to steer the development towards their areas 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 closeby the two main axes (North-South and West-East), or the bike lanes that had a longer total extension. We argue that these two features of the development process are amenable from the point of view of our identification strategy, which relies on variation of development status and timing across different parts of the city.

3 Conceptual framework

We use the partial equilibrium model laid out in Gorback (2020) as the guiding framework of our analysis. In this framework, non-tradable sector revenues in a given location depend on all pairwise bilateral commuting costs, and therefore on the structure of the transport network. The model is useful to the extent that it provides a rationale for the connection between local economic activity and the transport network.

Assume that there are $j \in J$ locations, each populated by R_j residents. A resident of location j maximises Cobb-Douglas utility by choosing how much to consume of a housing (h_j), tradable (c_j), and non-tradable (n_j^i) good, and in which location i to purchase the latter:

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

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

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

The choice on where to purchase the non-tradable good depends on 1) an idiosyncratic preference term (z_{ij}), and on 2) the bilateral commuting cost between origin and destination (d_{ij}). The idiosyncratic preference term is Fréchet distributed, $F(z_{ij} = e^{-E_i z_{ij}^{-\varepsilon}})$, where E_i is a destination-level amenity parameter and ε governs the substitutability between alternative consumption destinations. The utility semi-elasticity to commuting is τ: the greater this parameter, the more preferred are closeby shopping destinations.

Solving for the probability of purchasing the non-tradable good in location i, this is equal to:

$$\rho_{ij} = \frac{E_i \exp(-\tau\varepsilon d_{ij})}{\sum_s E_s \exp(-\tau\varepsilon d_{sj})} \tag{3}$$

The probability is increasing in the destination-specific amenity parameter, E_i, and decreasing in the bilateral commuting cost, d_{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 d_{ij} 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 neighborhood j only if $d_{i,j}$ drops more than $d_{i',j}$ with $i' \neq i$. Importantly, since consumers can shop in one and one only location, a generalised decline in d_{ij} does not entail an increase in aggregate expenditure, but rather a recomposition of expenditure towards locations that become easier to reach in relative terms.

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, $(1-\alpha-\beta)I_j$, we obtain a proxy of firm-level market access:

$$MA_i = (1-\alpha-\beta) \sum_j \frac{E_i \exp(-\tau\varepsilon d_{ij})}{\sum_s E_s \exp(-\tau\varepsilon d_{sj})} \times R_j \times I_j \tag{4}$$

In what follows, we rely on Equation 4 to guide the construction of a market access indicator in our setting.

4 Data

In this section we describe the various sources of data covering the area of the city of Paris that we have assembled for the purpose of the analysis presented in this paper.

4.1 Geographical unit of analysis and time-frame

Our geographical unit of analysis are equally-sized squared grid cells (used interchangeably with units 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 a grid of 2,230 units of 180×180 meters (Schiavina et al., 2019). Of these 2,230 units, we keep only those with less than 75% of green surface, thus dropping units falling within Paris' two urban forests. This leaves us with 1787 units. We further drop from the sample those units 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 1,418 final units. All our variables are computed at the grid cell level. In the case of geolocated data, it is straightforward to report them at the unit level. In the case of data available at the polygon level, we use weights to report all the information at the grid level.[4]

The observation period we consider goes from January 2015 until November 2019. We selected this time frame in order to have two years of pre-period before the development of Plan Vélo started in mid-2017. We end our analysis in November 2019 because of 1) the disruptions to public transport caused by the national strike in December 2019, 2) the later explosion of the Covid-related pandemic. Both these events strongly impacted the mobility of Parisians and their consumption behavior (Bounie et al., 2020). The frequency adopted is the quarterly one.

4.2 Card transaction data

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

Our dataset on card transactions comes from *Cartes Bancaires* (CB), a consortium including the near totality of French banks created in 1984.[5] We have information at the merchant-month level for the period ranging from 2015 to 2019. CB collects each month the value and volume of transactions made via CB cards, i.e., cards issued by banks part of the CB network. As of 2019, there were 77 million cards in use in the CB system, and 1.8 million CB-affiliated French merchants (Cartes Bancaires, 2019). Figure 3 displays the evolution of quarterly nominal total value of transactions recorded on CB payment system during the period 2015-2019. The growth of nominal total value is related also to increasing card usage in retail payments and substitution from cash. In one of our robustness checks, we show however that increasing card usage and substitution from cash does not appear to threaten our analysis, and the generalisation of our findings to all payments, via both card

[4]We construct weights based on the overlapping surface between the polygons and the grid cells.

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

and cash.

CB data contain the merchant business identification number (SIRET code) that allows us to match them to the French business registry (SIRENE). For each merchant, we thus have the date of creation, the sector of activity (NAF code) and the exact geographical location. For our analysis, we keep merchants located within the city of Paris and operating in the following sectors: retail commerce, restaurants, accommodation services, travel agencies, personal services, bakeries, sports clubs, cinemas and theaters.[6] Our final sample comprises 67,230 unique SIRET codes out of a total of 106,695 located in Paris (across all sectors), accounting for 61% of total card activity. We measure the coverage of our dataset by calculating for each industry the ratio between the nominal total value of transactions recorded in the (full) CB dataset and nominal total value added according to the national accounts. For the three largest non-tradables industries employed in this analysis (i.e., retail commerce, restaurants and accommodation), Table 1 shows that this ratio is well above 50%, which suggest that our dataset provides a good coverage of total economic activity.[7]

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

4.3 Measuring bike market access

We now turn to the measurement of our key independent variable, *bike* market access, i.e., the size of potential demand that can reach shopping destination i "sufficiently easily" by bike, for which we provide a theoretical rationale in Section 3 and a measurable expression in Equation 4.

The measurement of bike market access is articulated in several steps. The first step consists of gathering information on different road infrastructure types, including Plan Vélo. Starting from the latter, data on the development of the plan come from the *Observatoire du Plan Vélo de Paris* (see Section 2), which has maintained since July 2017 a geolocalised repository keeping track of the daily development of the plan. Further, we gather information on the road network, by type of road, through OpenStreetMap (see Figure 4).

As a second step, we divide the surface of the city of Paris into 500-by-600 pixels, each measuring approximately 20-by-20 meters. We overlay the transport infrastructure network at different points in time onto this finer grid and constructed quarterly cost rasters spanning

[6] The sectors' NAF codes are: 47, 56, 79, 55, 96, 10, 9312Z, 5914Z, 9001Z and 9004Z.

[7] Notice that we are only able to perform this test at the national level. A similar test conducted at the level of the city of Paris would probably show an even higher coverage, since card usage is more widespread in cities.

the entire city. Allen and Arkolakis (2014) provide one of the earliest implementations of numerical methods to get bilateral transport costs and study their impact on the spatial distribution of economic activity. We follow them in choosing a set of cost parameters - one for each available type of transport infrastructure - reflecting the time needed to travel across a given pixel. However, commuting costs depend on other factors beyond time, such as comfort, safety, cleanness, etc. Given our focus on cycling, we emphasise the comfort/safety component by setting the rank of different street types in terms of commuting costs equal to the inverse of their rank in terms of the degree of comfort experienced by a cyclist in using them. We normalise to 1 the cost for pixels crossed by Plan Vélo and assigned a value of 2 to pixels crossed by living/residential streets, 4 to pixels crossed by tertiary and secondary roads, and 6 to pixels crossed by primary roads.[8][9] An example of the cost raster is provided in Figure 5, where the left panel shows the road network (black lines) and Plan Vélo bike lanes (red lines) around the city hall, the middle and right panel show the cost raster pre- and after-development of Plan Vélo.

To ensure that our results survive to an alternative cost parameter configuration deduced from Google Maps routing data, we perform a robustness check in which we use a bike market access measure based on a more sophisticated calibration of the cost parameters (see Appendix B). Results using this alternative bike market access measure are similar to our baseline results (see Section 6.2).

As a third step, we use the cost rasters (describing the ease of moving across the city by bike at different points in time) to obtain the full set of bilateral commuting costs for each origin i/destination j pair and quarter, by means of the Fast Marching Method.[10] As the plan gets developed, bilateral commuting costs on average decline, and more so in the proximity of well-connected branches. Figure 6 shows the evolution over time of commuting costs from the city hall to all possible destinations.

As a fourth step, we adapt the formula in Equation 4 and build a measure of *bike* market access as follows:

[8] We also assign a value of 200 to pixels that are not crossed by any transport infrastructure or that belong to waterways.

[9] After obtaining the bilateral commuting costs by means of the Fast Marching Method algorithm, we exploit the degree of freedom earned with the initial normalisation, and rescaled the values in a way that the average commuting cost was approximately equal to one hour (expressed in minutes), the duration of the average commute in Paris. Essentially, our approach is therefore such that time pins down the order of magnitude, while comfort drives the ranking of different street types.

[10] The algorithm essentially delivers the minimum cost to go from origin i to destination j given the available transport network under the assumption of symmetric costs, $d_{ij} = d_{ji}$. See Allen and Arkolakis (2014) for details on this algorithm.

$$\mathrm{BMA}_{it} = \sum_j \frac{\exp(-\tau\varepsilon d_{ij,t})}{\sum_s \exp(-\tau\varepsilon d_{sj,t})} \mathrm{Population}_j \times \text{Median income}_j \tag{5}$$

where $\mathrm{Population}_j$ and Median income_j are pre-plan population and median income and $d_{ij,t}$ are the just obtained bilateral commuting costs.[11] Existing estimates have found the commuting (semi-)elasticity to be $\tau = 0.01$ (Ahlfeldt et al., 2015; Gorback, 2020). These elasticities are estimated based 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. Since consumption-related trips are at the centre of our analysis, we decided to run a reduced form estimation of a gravity model using a subset of our CB dataset, to properly calibrate the value for $\tau\varepsilon$. Our estimated values fall in the $(0.04, 0.06)$ range, depending on the quarter we use, and are thus very similar to existing estimates of τ, assuming an elasticity of substitution $\varepsilon = 5$ (Broda and Weinstein, 2006). For this reason, we set $\tau\varepsilon = 0.05$. See Appendix B for details on this estimation.

Figure 7 portrays the evolution of bike market access over time. A few considerations emerge from the inspection of this figure. First, market access in a given area situated in the proximity of a new bike lane rises if the latter connects with the rest of the network. For example, the construction of a bike lanes along the south-west Seine riverbank occurring between 2017q4 and 2018q4 did not trigger an expansion in bike market access due to their initial lack of connectedness with other parts of the network. Second, market access in given location increases only if the reduction in bilateral commuting costs exceeds the average reduction experienced by other locations. This is a direct consequence of the conceptual framework, according to which consumers shop in one and one location only. For example, the 15^{th} district located in the south-west quadrant of the city experienced a loss of bike market access during 2018q4 and 2019q4 in spite of the development of connected bike lanes. The loss in bike market access occurred because other areas - most notably the central districts - experienced a more sizeable reduction in bilateral commuting costs.

4.4 Control variables

We assemble a rich dataset of time-varying neighborhood characteristics from multiple sources. We get annual socioeconomic and demographic characteristics from the *Institut national de la statistique et des études economiques* (INSEE) such as total population, population aged between 25 and 39 years old, foreign population, number of job seekers and working age population. Given the lag with which these data are released, these control variables enter

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

our specification with a three-year lag period. These data are at the IRIS geographical level (the equivalent of census tracts in France). In order to report variable x to the grid cell level, we calculate for each IRIS unit i the share of surface overlapping with grid cell j, s_{ij}. Subsequently, we set $x_j = \sum_k s_{kj} x_k$, where k indexes all IRIS units overlapping with grid cell j. See Table 2 for the summary statistics of the different control variables employed in the analysis.

4.5 Other outcome variables

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

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

Finally, we webscraped data on restaurants registered on TripAdvisor and operating in Paris. After some data cleaning and address geolocation we ended up with an unbalanced panel of 13,102 merchants divided into: pizzerias (4.8%), cafes (4.7%), bars (8.8%) and restaurants not belonging to any of the other categories (81.7%). We use this alternative data source to further corroborate our findings on the increase in economic activity for establishments experiencing larger increases in bike market access following the development of Plan Vélo.

5 Empirical strategy

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

We address these identification concerns by letting our coefficient of interest be identified by variation in bike market access triggered by bike lane development in distant neighborhoods (Hornbeck and Rotemberg, 2021). This identification strategy requires to include in our specification a local bike lane density measure. By controlling for the intensity of local bike lane development, the bike market access elasticity is identified by variation in bike market access stemming from infrastructure development occurring further away. The identifying assumption is that development in distant neighborhoods is independent of local economic conditions measured in individual neighborhoods.

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

$$ln(Y_{it}) = \alpha_i + \alpha_{dt} + \beta ln(BMA_{it}) + \gamma X_{it} + \delta LBLD_{it} + e_{it} \tag{6}$$

where Y_{it} can be 1) total revenues, 2) transactions' number, 3) average revenues p/transaction, all calculated across merchants located in grid cell i at time t. Equation 6 also includes unit 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 unit and time level. Control variables are (log) population, ratio of foreigners, unemployment rate and (log) population aged between 25 and 39 years old.

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

6 Results

6.1 Baseline results

Table 3 presents the estimates of β from Equation 6, testing the robustness across different measures of local bike lane density. In column 1, we report the estimated coefficients without including any local bike lane density control. When we control for our preferred measure of local bike lane density in column 2, both coefficients stay statistically significant and remain in the ballpark range of 0.4-0.5. The results are qualitatively similar when we employ alternative measures of local bike lane density (column 3 and 4). Conversely, we do not detect a statistically significant elasticity with respect to average revenues p/transaction, regardless of the local bike lane density control included.

According to our preferred specification in column 2, the estimated coefficients indicate that a 1% improvement in bike market access, namely an increase in the number of consumers within reach *by bike* of a given shopping destination, leads to a 0.45% (0.50%) increase in total revenues (transactions' volume). Our estimated elasticities must be placed into context: the average improvement in bike market access (in absolute terms) is about 7 percentage points. It follows that the average improvement in bike market access triggered a slightly larger than 3 percentage points increase in total revenues. 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 bike market access on local economic activity materialises with a delay. In Table 4 we replace our baseline measure with its lagged values. When we do so, we find the elasticity of total revenues to be positive across all lags, but statistically significant only for the contemporaneous measure of bike market access and its third lag. Conversely, the elasticity of transactions' volume is statistically significant across all lags, and it grows in magnitude as lags of bike market access further back in time are considered.

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

6.2 Robustness

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

Centrality bias — We test the robustness of our results to the exclusion of highly-connected neighborhoods, which might mechanically benefit more from transport infrastructure development (Borusyak and Hull, 2020). In column 2 of Table 5, we follow Chandra and Thompson (2000) and exclude central districts, specifically districts 1 to 4. In column 3, we remove central areas, where centrality is defined from the standpoint of the public transport network. Using information on the public transport network (metro, tramway, and suburban trains)[12], we define as "central" those grid cells located less than 500 meters away from stations featuring three or more public transport connections.[13] All coefficients in Table 5 remain positive and statistically significant, thus eschewing centrality bias as a concern in our setting.

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

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

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

$$ln(Y_{it}) = \alpha_i + \alpha_{dt} + \sum_t \beta^t \Delta ln(BMA_{i,15-19}) \times \tau_t + \gamma X_{it} + \delta LBLD_{it} + e_{it} \quad (7)$$

where we regress our outcome variables on the (log) change in bike market access that occurred during 2015q1-2019q4. The (log) change in bike 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 8, areas that experienced higher bike market access improvements started featuring higher levels of economic activity only after development began, thus suggesting that non-random bike lane *placement* is not a major challenge for our strategy.

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

$$\text{Development date}_i | (D_i^{t_0} = 0) = \alpha + \beta X_i^{t_0} + e_i \quad (8)$$

We run Equation 8 for different choices of $t = t_0$ and report the results of the estimation in Table 6. During the first months of construction it appears that areas characterised by a higher population and percentage of job seekers received access to the bike lane network faster. However, towards the end our sample period, during which most of the construction occurred, none of the covariates is associated in a statistically significant way with the development date. We interpret the absence of systematic correlation between the covariates and the development date as supportive evidence of development timing mostly independent of local economic conditions.

Exploiting the unfinished Plan Vélo — In a further check, we restrict the sample to units that were planned to feature some bike lane development according to the original Plan Vélo (see Figure 9). We can expect this subsample to be more homogeneous than the baseline one, since all included units were initially considered to be "in need" of a bike lane and regardless of the fact that only some had been effectively developed by the end of 2019.

We test this by means of a balancing test (Table 8), which confirms that units that received some bike lane development did not differ in a statistically significant way from units 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. Results from re-running the baseline specification on this subsample are displayed in Table 7. We observe that the coefficients remain statistically significant and that they are slightly larger than in the baseline estimation.

Other potentially confounding factors — In Table 9, 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 local economic activity composition.[14] In column 3, we test the robustness of our findings to an alternative bike market access measure accounting for cycling infrastructure that was present in the city before the development of Plan Vélo.

First, we collect data on all Paris' old cycling network bike lanes from the city of Paris open data web portal. We keep only those bike lanes from which we know the development date and whenever this is anterior to the construction of Plan Vélo. Using this additional shapefile, we update the cost raster definition (see Section 4.3) and recalculate bike market access.[15] In column 4, we make sure that the development of a tramway line (T3b) during our period of study does not represent a confounding factor, by removing grid cells less than 500 meters away from it. Finally, in column 5, we test the robustness of our results to a law passed in 2015 that allowed businesses located in selected areas to stay opened on Sundays.[16] 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. However, we include in the baseline specification an interaction term between the grid cell-specific share of surface concerned by the law and time dummies, and make sure that our estimates do not change. Across all these tests, the elasticity of economic activity to bike market access remains positive and statistically significant.

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

[15] We assign to pre-Plan Vélo bike lanes a cost value of 2, the same value we assign to residential streets to capture the inferior quality of these bike lanes compared to Plan Vélo. The alternative bike market access measure employed for this test is obtained by re-running the FMM algorithm on this alternative cost raster definition. A disadvantage of this dataset is that we know the exact development date for less than 60% of total bike lane length featured within it, which is the main reason for not including it in our baseline estimation.

[16] See the map of concerned areas in Figure 10. Data come from APUR, Mairie de Paris and DRIEA IF/UD75.

Alternative transportation modes — Next, we tackle the possible interaction between bikes and other transportation modes. In particular, we notice that our coefficient may capture the reduced form effect of the impact on local revenues of an increase in bike market access net of the negative impact driven by a potential decline in market access through alternative transport modes. While the reduced form impact of a new infrastructure is what matters from a policy perspective, we seek to provide an estimate of the impact of the new infrastructure *channeled by the increase in bike market access only*. We do so by adding to our baseline regression two measures proxying for *car* or *public transport* market access.

In column 2 of Table 10, we proxy for the negative impact of a potential decline in car market access by controlling for the (log of the) volume of cars transiting through a given grid cell in a given quarter. 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.[17] Similarly, in column 3 of Table 10, we proxy for the negative impact potentially driven by substitution away from public transport by controlling for the total number of metro trips starting from stations located within 500 meters of a given grid cell in a given quarter.[18] We observe that the coefficients remain statistically significant in either case, and that they increase when controlling for potential substitution with public transport, in line with the greater substitutability of this transport mode with cycling. Conversely, the estimated coefficients do not change substantially when controlling for car traffic. One possible explanation has to do with the already low modal share of cars in the French capital and limited availability of street parking.[19]

Bike market access with calibrated cost parameters — Next, we make sure that our results are unchanged when adopting a cost parameter configuration likely to be closer to the true parameters. Specifically, we deploy an alternative measure of bike market access where the cost parameters of the cost raster have been calibrated by means of a route-discrepancy minimisation procedure between our least cost path routes and Google Maps routes. The calibration procedure selects a cost value of 2 for streets and 3.5 for boulevards and avenues, very much in line with the ones chosen for the baseline measure. The details of

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

[18] The data on metro ticket stamps comes from Île-de-France Mobilités website.

[19] See for example the results from the analysis of the 2018 Commuting Survey `https://www.omnil.fr/IMG/pdf/presentation_egt_v_publique_vf.pdf`.

this calibration are provided in Appendix B. The results from running the baseline estimation using this alternative bike market access definition are displayed in Table 11. The coefficients comfortably align with the baseline results displayed in Table 3. In spite of this, we prefer retaining as the baseline our simpler cost parameter configuration given the difficulty in showing the well-behavedness of the loss function when departing from a 2D minimisation problem and, especially, the difficulty in replicating Google Maps routing algorithm.

Test for card usage — Finally, we test if the increase in revenues in areas with higher bike market access improvement is partially driven by an increase in credit cards usage and substitution away from cash payments. To do so, we check if bike market access increases the share of establishments using credit cars. First, we build a card usage intensity index by dividing the number of merchants (or establishments) present in the *Cartes Bancaires* dataset over the number of establishments that should be active in that same area and quarter according to the business registry. Then, we run Equation 6 with this index on the left hand side. If the estimated coefficient were to be found positive and statistically significant, this would entail that a bike 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 12 reassures us though that this is not the case.

6.3 Heterogeneity

In this section, we test the existence of heterogeneity in the estimated impact of bike market access improvements on local economic activity. First, we build clusters of neighborhoods with similar local establishments characteristics. Even if non-tradable products or services are at the core of our analysis, neighborhoods will still feature some degree of specialisation into activities with certain characteristics. First, we select the characteristics on which we run the clustering algorithm: a set of (dummy-based) sub-industry indicators as of 2015 (supermarkets and malls, specialised food retail stores, specialised non-food retail stores, fast food restaurants/bars, restaurants, bars); a size dummy taking value 1 if in 2015 average merchant size in a given unit is greater than the median value; an age dummy taking value 1 if in 2015 average merchant age in a given unit is greater than the median value calculated across all units.[20] Subsequently, we run a *k-means* clustering algorithm (Bock, 2007) for

[20]We define a unit 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.

different values of the number of clusters, k, and we select $k = 5$ by means of an elbow test as shown in Figure 11.[21]

The characteristics of the five clusters in terms of the variables used for the clustering exercise are reported in Table 13. Industries that are more unevenly spatially distributed generate greater differences across clusters: for example, the degree of specialisation into stores that sell essential goods, such as specialised food retail stores, is fairly homogeneous across clusters, while neighborhoods tend to differ quite substantially in terms of their degree of specialisation into fast-food, restaurants or bars. Table 14 contains the estimated coefficients from a fully-interacted version of Equation 6. The evidence suggests a positive and statistically significant impact of a bike market access improvement for the units specialised into smaller/younger establishments or establishments operating in food-related industries that tend to be more youth-oriented or characterised by lower average end-of-the-meal bills, such as cafes, fast-food or bars. Clusters specialised in other industries, such as non-food/specialised retail industries, also display a positive and statistically significant impact of a bike market access increase but only on transactions' volume and only when establishments are on average younger (see the comparison between cluster 3 and 5).

These findings show that the gain from access to better infrastructure appears to be greater for smaller/younger merchants. A potential explanation is that the development of new infrastructure helps removing information asymmetries (i.e., consumers not being aware of the existence of certain businesses) that can be especially penalising for this kind of establishments. Further, the statistically significant elasticity in areas specialised in the retail sector is in line with Koster et al. (2019), who find that footfall demand externalities are an important driver of profitability in this sector. The development of cycling infrastructure is typically accompanied by supplementary interventions aimed at making streets safer for bikers and, by reflection, also for walkers, thus increasing the take-up of the new infrastructure (e.g., reduction in driving speed limits, removal of car parking to discourage car traffic, etc.). We do not unfortunately observe these interventions nor their timing, but the empirical evidence that we find is compatible with the idea that the implementation of Plan Vélo triggered a mobility shift towards all active mobility transport modes, both cycling and walking, that benefitted the retail sector by means of footfall externalities.

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

6.4 Other outcomes

Next, we inspect how a set of additional outcomes are related to our measure of bike market access. First, we build a dummy indicator that takes value one in the case of positive firm entry. While the sign of the coefficient goes in the right direction, we are unable to assign a statistically significant impact of bike market access on firm entry, regardless of whether we measure it with respect to current bike market access or its lagged value (see panel A of Table 15).

Second, we investigate whether an improvement in bike market access has an impact on house prices. We construct a price index by running a hedonic regression of (log) of house prices on individual properties characteristics and unit and time fixed effects (see Section 4 for more details on the dataset used). We find a positive, statistically significant and increasing impact starting from the second quarter following an improvement in bike market access, with an elasticity reaching 0.063 three quarters after (panel B of Table 15). The observed increase in house prices can help explaining the lack of statistically significant impact on firm entry: as bike market access improves in a given area, the rental rate a potential new business must pay in order to enter the market also rises, thus offsetting the benefit accruing from a market expansion.

Finally, we zoom into the impact of bike market access for the restaurant sector by assessing the impact on the number of reviews registered on TripAdvisor.[22] In order to do so, we estimate our baseline regression on either the total number of reviews, or the average review grade.[23] The results shown in Table 16 confirm the existence of a positive and statistically significant (semi-)elasticity of local economic activity to bike market access. In particular, we find that an increase by 1% of bike market access translates into two more reviews per individual establishment one quarter after (panel A). Conversely, we do not detect a statistically significant impact of a bike market access improvement on average review grade (panel B). Taken together, these results suggest that the new infrastructure increased the business volume for restaurants located more closeby, but it didn't trigger a re-composition towards better quality businesses, if one takes the review grade as a proxy for quality.

[22] We consider reviews written in French only, in order to minimise the impact of touristic flows on our estimates.

[23] In this regression we control for the age of individual establishments based on TripAdvisor records, i.e., for each establishment, we count how many quarters have passed since its first appearance in our TripAdvisor dataset. Adding this control is important since establishments tend to concentrate their efforts in accumulating numerous reviews when they are new on the platform in order to gain visibility. A main limitation of this dataset is that we are unable to determine with certainty when an establishment goes out of business.

6.5 Distributional consequences of Plan Vélo

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

Our framework enables us, however, to identify what are the locations that have gained and those that have lost, always in relative terms, from the development of the new infrastructure. We investigate the distributional consequences of the plan by proceeding in the following way. We multiply the difference between 2019q4 and pre-Plan Vélo bike market access by 0.45 (the baseline elasticity - see column 2 of Table 3), 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 units (+2 p.p. on average across "winners") and negative for 60% ($-1,2$ p.p. on average across "losers").

The uneven impact of the infrastructure can be explained by analysing the characteristics of the locations that gained from the development of the bike lane network. In panel A of Figure 8, we show how bike market access has changed during our period of analysis. The first thing to notice is that the locations where bike market access improved tend to be 1) centrally located, and 2) characterised on average by lower purchasing power (or income) (Figure 12).[24] Hence, Plan Vélo acted to redistribute demand for non-tradables away from more peripheral but densely populated neighborhoods towards more central and less populated ones.[25] As an example, before the development of Plan Vélo residents of the 15^{th} or the 17^{th} district (situated respectively in the south-west and north-west quadrant of the city) might have preferred to go shopping locally since it was for them costly to reach the city centre. After the development of new bike lanes connecting these two districts with the city centre, a larger fraction of residents chooses to take a bike ride and go shopping in the city centre instead. The same reasoning holds for residents of the central districts, who may now find easier to go shopping in peripheral neighbourhoods. However, central districts in Paris are mostly shopping locations and have low population density, so that the change in consumption habits is more likely to advantage merchants located in the city centre at the disadvantage

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

[25] The conclusion of the redistribution from the periphery towards the centre would have not changed under the original plan, as the comparison with panel B of Figure 8 makes clear.

of those located in peripheral neighborhoods.

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

7 Conclusion

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

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

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

References

AGARWAL, S., J. B. JENSEN, AND F. MONTE (2017): "Consumer mobility and the local structure of consumption industries," Tech. rep., National Bureau of Economic Research.

AHLFELDT, G. M., S. J. REDDING, D. M. STURM, AND N. WOLF (2015): "The economics of density: Evidence from the Berlin Wall," *Econometrica*, 83, 2127–2189.

ALLEN, T. AND C. ARKOLAKIS (2014): "Trade and the Topography of the Spatial Economy," *The Quarterly Journal of Economics*, 129, 1085–1140.

ALLEN, T., S. FUCHS, S. GANAPATI, A. GRAZIANO, R. MADERA, AND J. MONTORIOL-GARRIGA (2020): "Is Tourism good for Locals? Evidence from Barcelona," .

ATHEY, S., D. BLEI, R. DONNELLY, F. RUIZ, AND T. SCHMIDT (2018): "Estimating heterogeneous consumer preferences for restaurants and travel time using mobile location data," in *AEA Papers and Proceedings*, vol. 108, 64–67.

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

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

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

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

BOUNIE, D., Y. CAMARA, AND J. W. GALBRAITH (2020): "Consumers' Mobility, Expenditure and Online-Offline Substitution Response to COVID-19: Evidence from French Transaction Data," *available at SSRN 3588373*.

BRODA, C. AND D. E. WEINSTEIN (2006): "Globalization and the Gains From Variety," *The Quarterly Journal of Economics*, 121, 541–585.

CAILLY, C., J.-F. CÔTE, A. DAVID, J. FRIGGIT, S. GREGOIR, A. NOBRE, F. PROOST, S. SCHOFFIT, N. TAUZIN, AND H. THÉLOT (2019): "Les indices Notaires-Insee des prix des logements anciens Méthodologie v4," .

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Tsivanidis, N. (2019): "Evaluating the impact of urban transit infrastructure: Evidence from bogota's transmilenio," *mimeo.*

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.

Appendix

A Figures and Tables

Table 1: Total revenues in *Cartes Bancaires* and INSEE data in selected non-tradable industries

NAF Code	Name of industry	Revenues (CB)	Revenues (INSEE)	VAT (%)	Revenues + VAT (INSEE)	Ratio
47	Retail Commerce	276,301	472,733	10%	520,006	53.1%
56	Restaurants	44,633	75,636	10%	83,199	53.6%
79	Travel Agencies	9,067	12,682	10%	13,950	65.0%
55	Accommodation	15,355	29,223	10%	32,146	47.8%
96	Personal Services	6,835	15,726	10%	17,298	39.5%
1071+4724Z	Bakeries and pastry shops	4,384	24,489	10%	26,938	16.3%
9312Z	Sports clubs	428	2,651	20%	3,181	13.5%
5914Z	Cinemas	748	2,146	5.5%	2,264	33.0%
9001Z+9004Z	Theatre and shows	545	3,548	3.8%	3,683	14.8%

Notes: values in thousands. Source: *Cartes Bancaires*.

Table 2: Descriptive statistics

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

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

Table 3: Elasticity of local economic activity to bike market access: baseline evidence

Panel A:	Log total revenues			
	(1)	(2)	(3)	(4)
Log BMA	0.334*	0.449**	0.359	0.356*
	(0.180)	(0.207)	(0.225)	(0.209)
Panel B:	Log transactions' volume			
Log BMA	0.534**	0.497**	0.502**	0.393*
	(0.208)	(0.222)	(0.241)	(0.224)
Panel C:	Log average revenues p/transaction			
Log BMA	-0.190	-0.038	-0.132	-0.026
	(0.130)	(0.141)	(0.147)	(0.136)
N	27,617	27,617	28,297	27,617
Controls	X	X	X	X
Unit FE	X	X	X	X
District×Time FE	X	X	X	X
LBLD	None	Same project	Neighbors	Same project/ same unit

Notes: coefficients from the estimation of Equation 6. Standard errors are clustered at the unit level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Cartes Bancaires*. Back to Section 6.1.

Table 4: Elasticity of local economic activity to bike market access: lagged impact

Panel A:	Log total revenues			
	(1)	(2)	(3)	(4)
Log BMA	0.319*			
	(0.191)			
First lag log BMA		0.298		
		(0.208)		
Second lag log BMA			0.299	
			(0.213)	
Third lag log BMA				0.388*
				(0.221)
Panel B:	Log transactions' volume			
Log BMA	0.341*			
	(0.203)			
First lag log BMA		0.440*		
		(0.236)		
Second lag log BMA			0.469*	
			(0.253)	
Third lag log BMA				0.595**
				(0.262)
N	23,480	23,480	23,480	23,480

Notes: baseline estimation as in Table 3 column 2, estimating the elasticity to lagged bike market access. Standard errors are clustered at the unit level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Cartes Bancaires*. Back to Section 6.1.

Table 5: Robustness tests: dealing with endogenous bike lane development

Panel A:	Log total revenues		
	(1)	(2)	(3)
Log BMA	0.449**	0.454**	0.423*
	(0.207)	(0.222)	(0.224)
Panel B:	Log transactions' volume		
Log BMA	0.497**	0.517**	0.450*
	(0.222)	(0.238)	(0.241)
N	27,617	25,697	23,237
Test	Baseline	Remove central districts	Remove connected areas

Notes: baseline estimation as in Table 3 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 unit level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Cartes Bancaires*. Back to Section 6.2.

Table 6: Robustness tests: testing random development timing

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

Notes: the dependent variable is the date in which the units 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 7: Robustness tests: keeping only Plan Vélo subsample

	Log total revenues	Log transactions' volume
Log BMA	0.677**	0.974***
	(0.264)	(0.295)
N	9,480	9,480

Notes: baseline estimation as in Table 3 column 2 restricted to the subsample of grid cells intersected by the original Plan Vélo. Standard errors are clustered at the unit level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Cartes Bancaires*. Back to Section 6.2.

Table 8: Balancing test of local characteristics in the Plan Vélo subsample between grid cells where development had taken place by the end of 2019 and those where it did not

	Developed		Not developed				
	Mean	Std Dev	Mean	Std Dev	Difference	t-stat	p-value
BMA (in000s)	3874	1075	3892	1078	19	-0.20	0.84
Roads (m)	1153	346	1123	312	-30	1.01	0.31
Planned bike lanes (m)	205	117	170	90	-35	3.71	0.00
Population	1357	858	1375	755	17	-0.24	0.81
Foreigners (%)	16	5	15	4	-1	1.46	0.14
Jobseekers (%)	9	2	9	2	-0	0.91	0.36
Population 25-39	380	277	369	232	-11	0.48	0.63
Entrant firms (#)	0	0	0	1	0	-0.01	0.99
Car flow (#)	25747	23754	20641	21100	-5105	2.55	0.01
House price (p/m2)	8867	1568	8992	1610	124	-0.88	0.38
Total revenues (in000s)	1713	3238	2663	8258	949	-1.75	0.08
Transactions' volume	32112	46523	37834	53585	5722	-1.29	0.20
Avg. revenues p/transaction	70	122	70	71	0	-0.01	0.99
Merchants (#)	30	30	35	28	5	-1.94	0.05
N	271	.	237	.	.	.	.

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

Table 9: Robustness tests: miscellanea

Panel A:	Log total revenues				
	(1)	(2)	(3)	(4)	(5)
Log BMA	0.449**	0.348*	0.535**	0.518**	0.443**
	(0.207)	(0.203)	(0.236)	(0.227)	(0.208)
Panel B:	Log transactions' volume				
Log BMA	0.497**	0.453*	0.447*	0.527**	0.490**
	(0.222)	(0.238)	(0.233)	(0.237)	(0.222)
N	27,617	26,238	27,617	26,437	27,617
Test	Baseline	Sectoral shares	Including non-PV bike lanes	Remove closeby by new tram	Sunday Law trend

Notes: baseline estimation as in Table 3 column 2 (col.1); augmented to include lagged sectoral shares as controls (col.2); accounting for pre-existing bike lanes (col.3); excluding grid cells located within 500 meters from the itinerary of tramway T3b (col.4); augmented to include an interaction term between the unit-specific share of surface concerned by the 2015 "Sunday Law" and time dummies (col.5). Standard errors are clustered at the unit level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Cartes Bancaires*. Back to Section 6.2.

Table 10: Robustness tests: controlling for substitution with alternative transportation modes

Panel A:	Log total revenues		
	(1)	(2)	(3)
	0.449**	0.429**	0.528***
	(0.207)	(0.208)	(0.194)
Panel B:	Log transactions' volume		
Log BMA	0.497**	0.477**	0.503**
	(0.222)	(0.223)	(0.214)
N	27,617	27,467	25,887
Test	Baseline	Cars volume	Metro trips

Notes: baseline estimation as in Table 3 column 2 (col.1); augmented to include the log of car traffic as control (col.2);augmented to include the log of total metro trips (col.3). Standard errors are clustered at the unit level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Cartes Bancaires*. Back to Section 6.2.

Table 11: Robustness tests: alternative bike market access measure based on calibrated cost parameters

	Log total revenues	Log transactions' volume
Log BMA	0.488**	0.399*
	(0.211)	(0.234)
N	27,617	27,617

Notes: baseline estimation as in Table 3 column 2 using an alternative bike market access measure based on calibrated cost parameters. Standard errors are clustered at the unit level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Cartes Bancaires*. Back to Section 6.2.

Table 12: Robustness tests: elasticity of credit card usage intensity to bike market access

	(1)	(2)	(3)	(4)
Log BMA	0.031			
	(0.056)			
First lag log BMA		0.029		
		(0.076)		
Second lag log BMA			0.031	
			(0.086)	
Third lag log BMA				0.030
				(0.090)
N	23,341	23,341	23,341	23,341

Notes: baseline estimation as in Table 3 column 2 applied to the ratio between the number of establishments reporting transactions in the *Cartes Bancaires* 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 *Cartes Bancaires*. Back to Section 6.2.

Table 13: 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 units per each cluster specialised in 2015 in the corresponding activities. Col.7-8 contain the % of units in each cluster such that in 2015 average merchant size (col.7) or age (col.8) was greater than the median value. Source: *Cartes Bancaires*. Back to Section 6.3.

Table 14: Testing heterogeneous effects with respect to local merchant characteristics

	Log total revenues	Log transactions' volume
Log BMA×Small and new businesses	0.672**	0.500
	(0.305)	(0.305)
Log BMA×Spec. food stores/fast food/bars	0.637*	0.641
	(0.343)	(0.403)
Log BMA×Spec. retail + old businesses	-0.201	-0.218
	(0.492)	(0.453)
Log BMA×Retail + old businesses	0.709	0.445
	(0.450)	(0.381)
Log BMA×Spec. retail/restaurants	0.425	1.310*
+ new businesses	(0.452)	(0.696)
N	27,617	27,617

Notes: baseline estimation as in Table 3 column 2, testing heterogeneous effects through the inclusion of interaction terms between bike market access and cluster-specific dummies. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Cartes Bancaires*. Back to Section 6.3.

Table 15: Elasticity of other outcomes to bike market access

Panel A:	Entry			
	(1)	(2)	(3)	(4)
Log BMA	0.042			
	(0.082)			
First lag log BMA		0.049		
		(0.096)		
Second lag log BMA			0.067	
			(0.100)	
Third lag log BMA				-0.055
				(0.107)
N	23,480	23,480	23,480	23,480
Panel B:	Log house prices			
Log BMA	-0.004			
	(0.011)			
First lag log BMA		-0.021		
		(0.015)		
Second lag log BMA			0.028*	
			(0.017)	
Third lag log BMA				0.063***
				(0.020)
N	23,382	23,382	23,382	23,382

Notes: coefficients from the estimation of Equation 6, testing the elasticity of other outcomes to bike market access. Standard errors are clustered at the unit level. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Cartes Bancaires*. Back to Section 6.4.

Table 16: Elasticity of local economic activity to bike market access using TripAdvisor data

Panel A:	Number of reviews			
	(1)	(2)	(3)	(4)
Log BMA	1.147*			
	(0.622)			
First lag log BMA		1.996***		
		(0.762)		
Second lag log BMA			1.943**	
			(0.846)	
Third lag log BMA				1.848**
				(0.915)
N	134,273	134,273	134,273	134,273
Panel B:	Average review grade			
Log BMA	-0.029			
	(0.124)			
First lag log BMA		-0.139		
		(0.153)		
Second lag log BMA			-0.188	
			(0.169)	
Third lag log BMA				-0.011
				(0.184)
N	104,753	104,753	104,753	104,753

Notes: coefficients from the estimation of $Y_{nt} = \alpha_n + \alpha_{dt} + \beta BMA_{it} + \gamma X_{nt} + \delta LBLD_{it} + e_{nt}$, where Y_{nt} is either the number of reviews a given restaurant n receives during quarter t, or the average review grade associated to those reviews. Source: TridAdvisor, *Observatoire du Plan Vélo de Paris*, INSEE and *Cartes Bancaires*. Back to Section 6.4.

Figure 1: Development of Plan Vélo

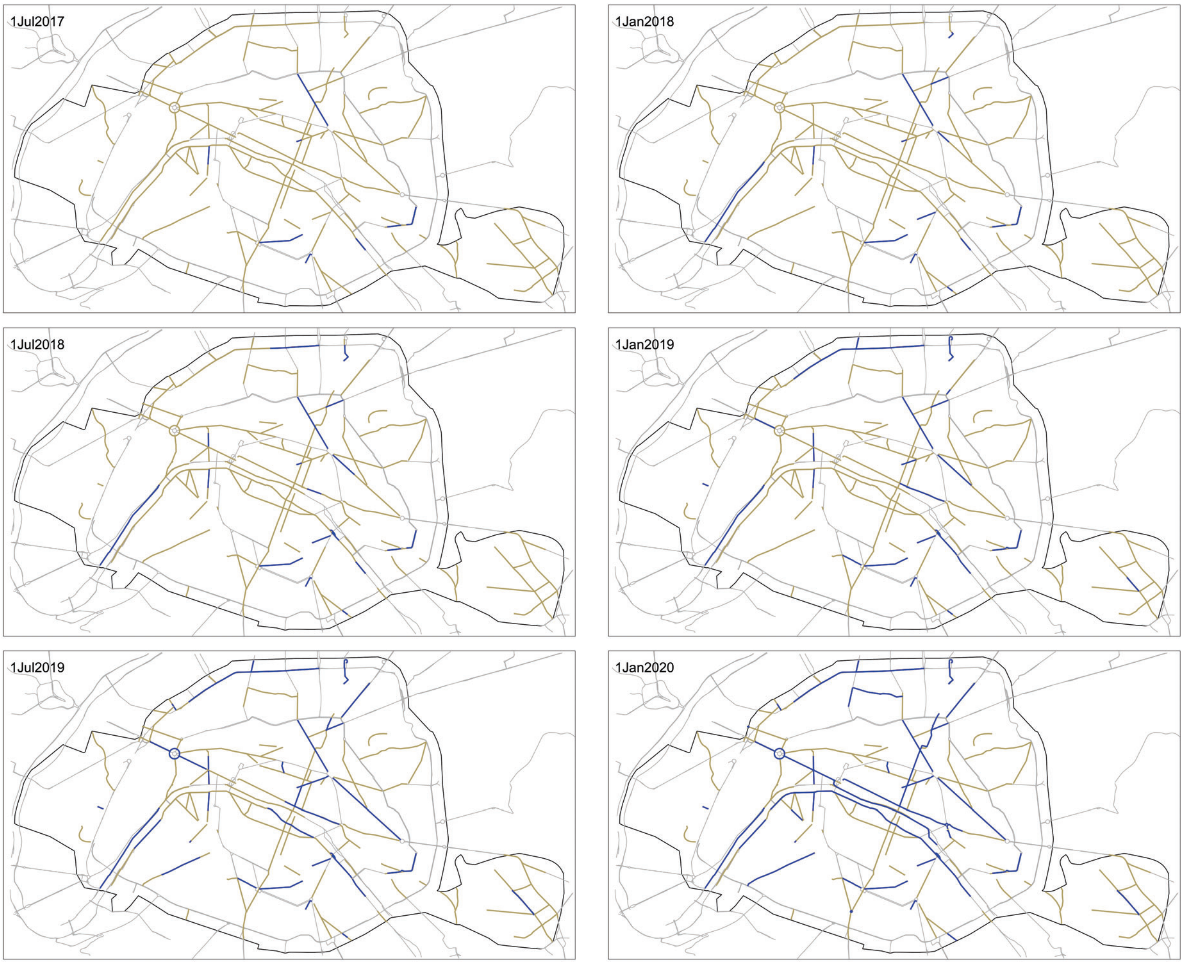

Notes: blue lines show the development of Plan Vélo; gold lines show the original Plan. Source: *Observatoire du Plan Vélo de Paris*.

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

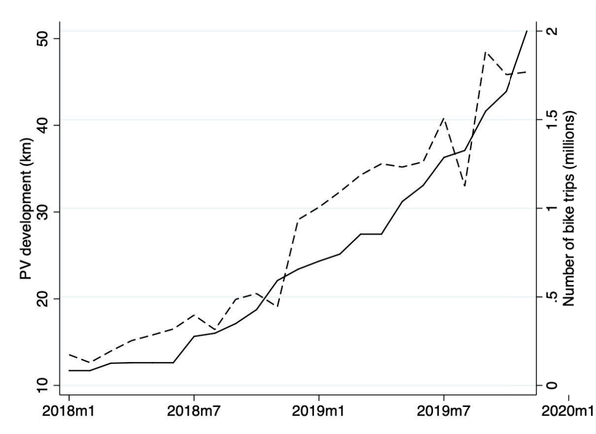

Notes: all bike trips are recorded by sensors distributed across the city. Source: *Comptage vélo - Données compteurs* dataset from https://www.data.gouv.fr/en/datasets/comptage-velo-historique-donnees-compteurs-et-sites-de-comptage/.

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

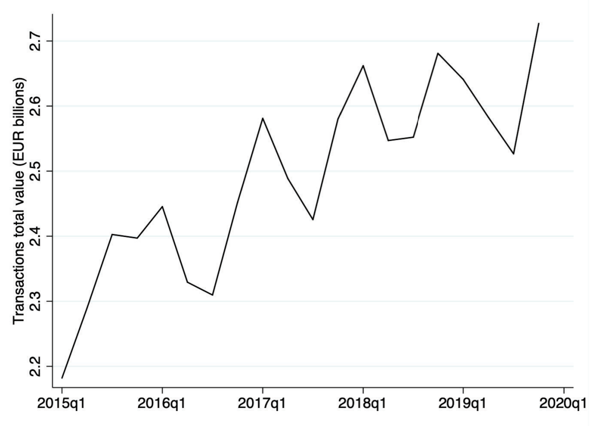

Source: *Groupement Cartes Bancaires.*

Figure 4: OpenStreetMap selected features for least cost path calculations

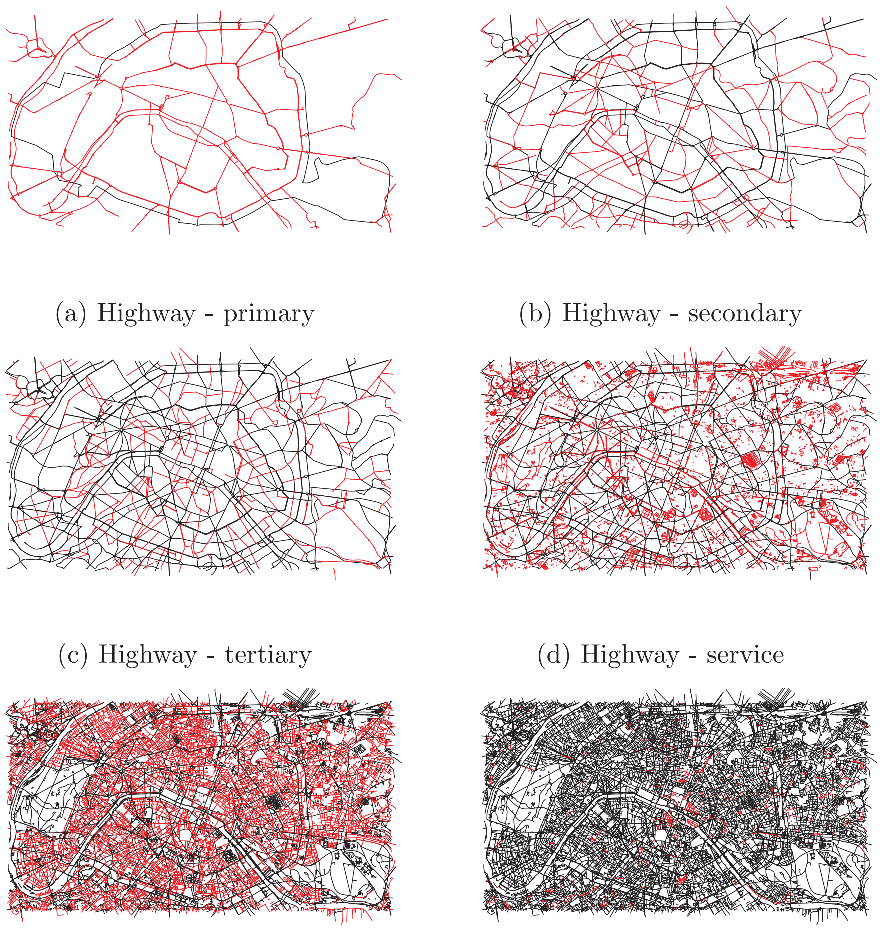

(a) Highway - primary

(b) Highway - secondary

(c) Highway - tertiary

(d) Highway - service

(e) Highway - residential

(f) Highway - living street

Source: OpenStreetMap. Back to Section 4.3.

Figure 5: Example of cost raster before and after the construction of Plan Vélo

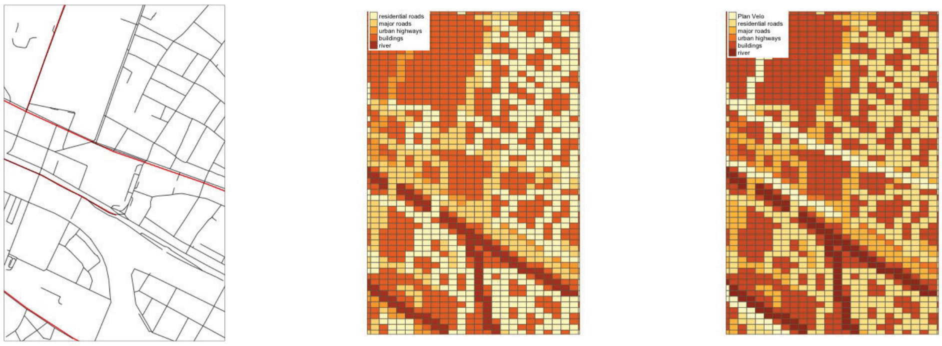

(a) Open Street Map (b) Cost raster (before PV) (c) Cost raster (after PV)

Notes: snapshots of the cost rasters centred around Hôtel de Ville. Back to Section 4.3.

Figure 6: Bilateral bike commuting costs at different points in time

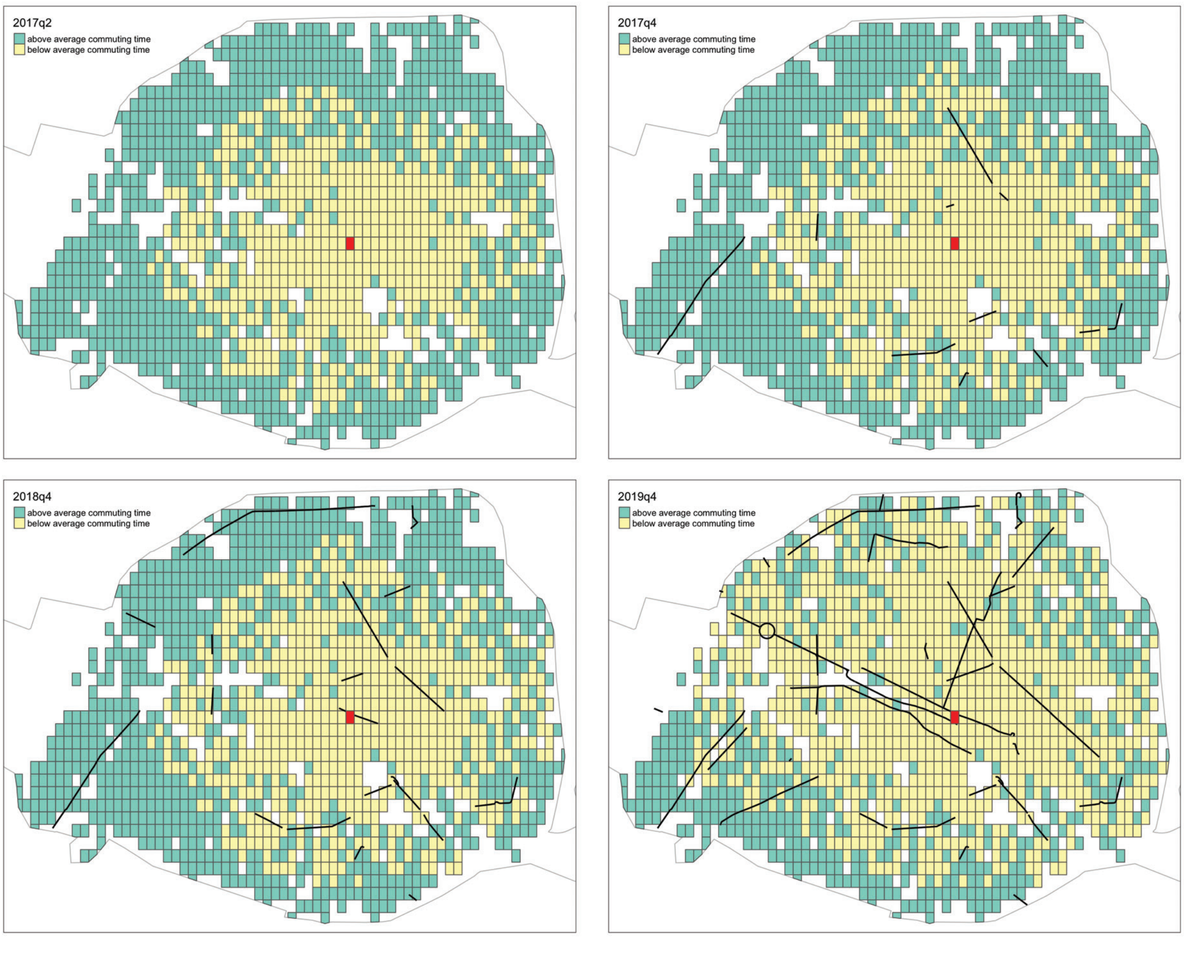

Notes: bilateral commuting costs over time from Hôtel de Ville (red dot) to other parts of the city. Overlaid black lines capture Plan Vélo development over time. Back to Section 4.3.

Figure 7: Quartiles of bike market access at different points in time

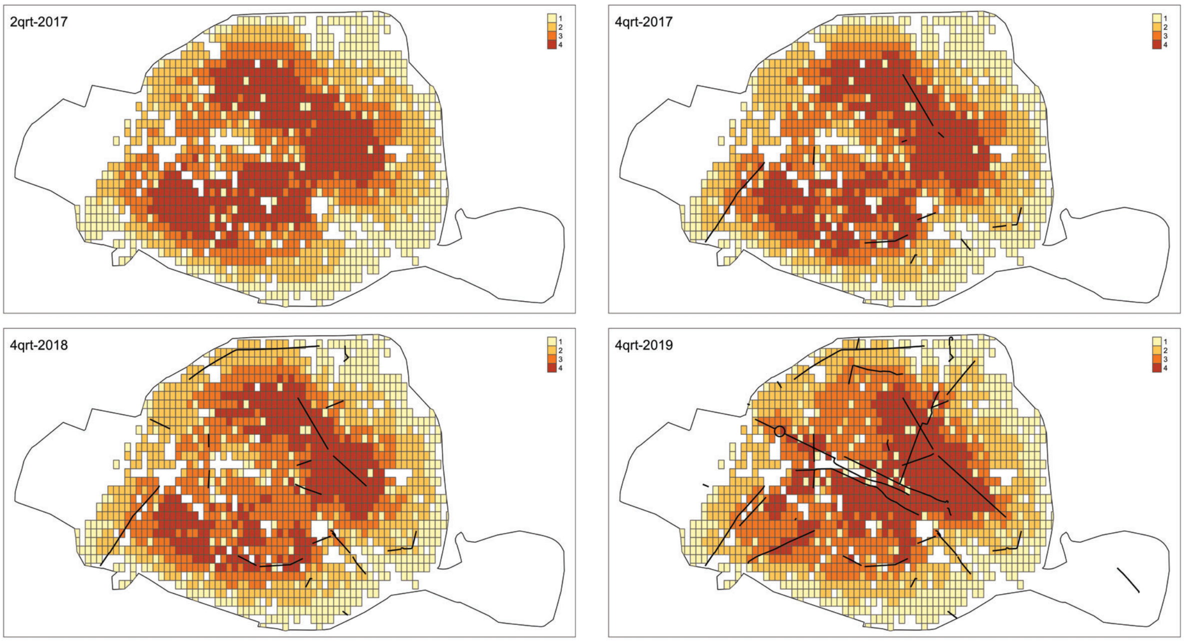

Notes: overlaid black lines capture Plan Vélo development over time. The quartiles refer to the cross-space and time market access distribution and are thus held constant. Back to Section 4.3.

Figure 8: Pre-trends analysis

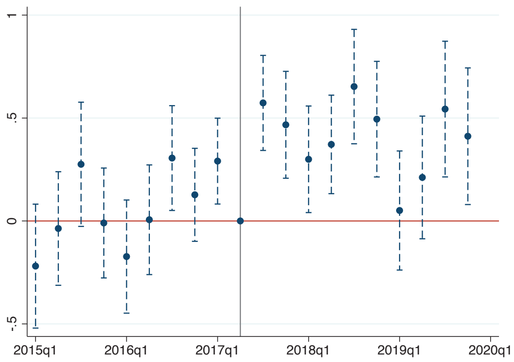

(a) Log total revenues

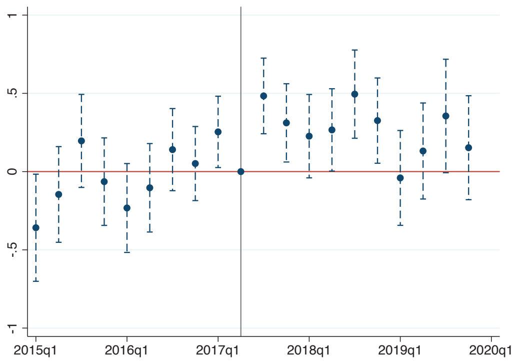

(b) Log transactions' volume

Notes: estimated β^t from Equation 7 on the y-axis. Source: *Observatoire du Plan Vélo de Paris*, INSEE and *Cartes Bancaires*. Back to Section 6.2.

Figure 9: Grid cells crossed by 2015 Plan Vélo

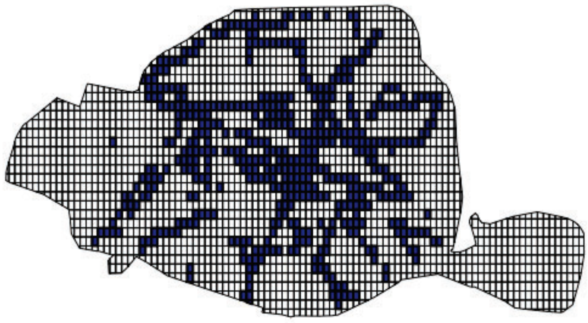

Notes: blue grid cells are those that were supposed to experience some bike lane development according to the Plan published in 2015. Source: *Observatoire du Plan Vélo de Paris*. Back to Section 6.2.

Figure 10: Areas concerned by the Sunday Law

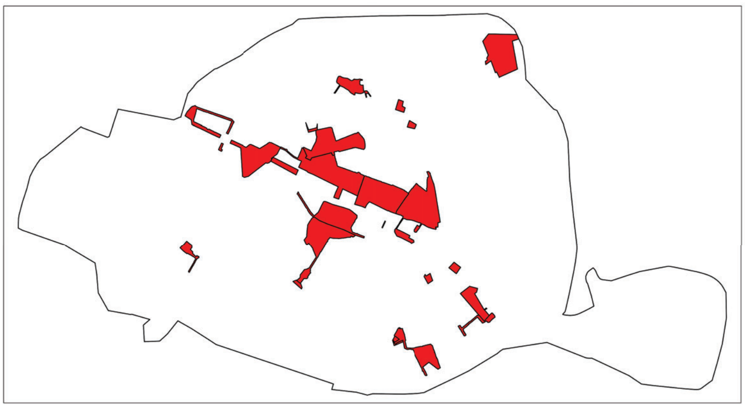

Notes: red areas include international tourism areas, tourism areas, commercial areas and train stations. Source: APUR, Mairie de Paris and DRIEA IF/UD75. Back to Section 6.2

Figure 11: Elbow test for the selection of the optimal number of clusters

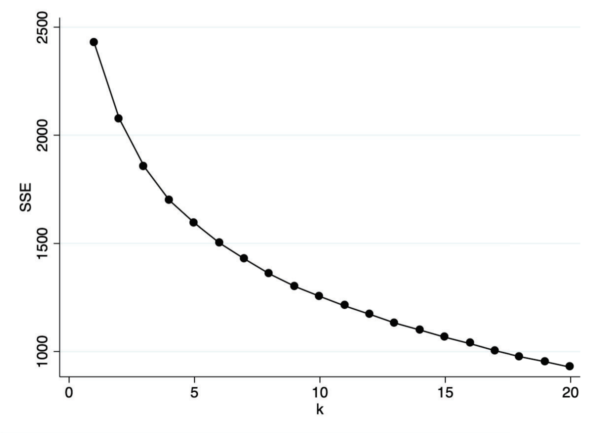

Notes: sum of squared errors on the vertical axis. Back to Section 6.3

Figure 12: Spatial distribution of purchasing power

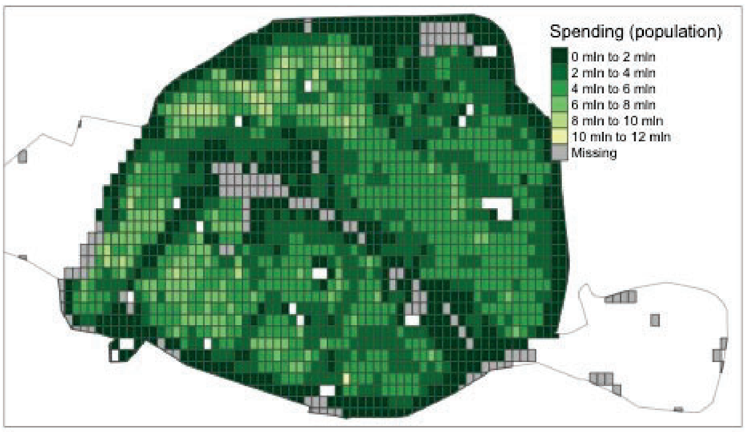

Notes: spending in a given grid cell j corresponds to Population$_j$ $\times$ Median income$_j$ in 2015. Source: INSEE. Back to Section 6.2.

Figure 13: Bike market access: before and after Plan Vélo development

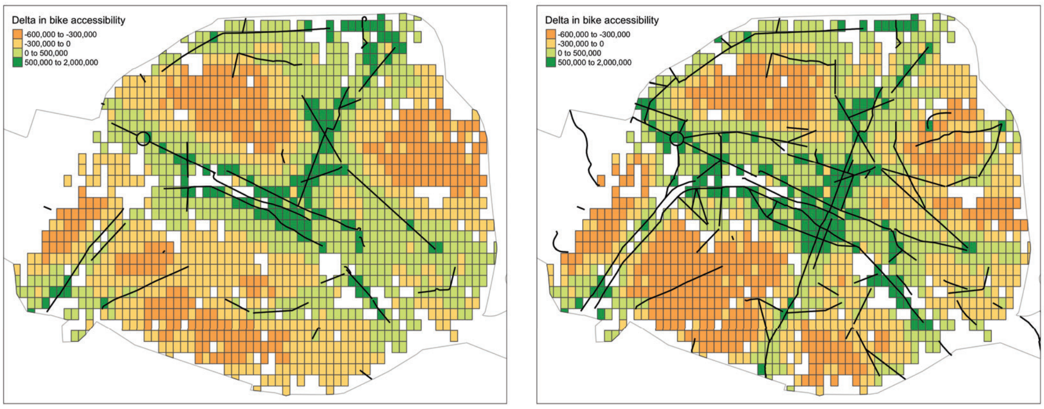

(a) Change in bike market access (2015q1-2019q4)

(b) Change in bike market access (2015q1-2019q4) - original Plan

Notes: bike market access change (in absolute terms) stemming from the actually developed portions of Plan Vélo (a), bike market access change (in absolute terms) that would have resulted from the development of Plan Vélo in its entirety (b). The bike lanes effectively developed as of 2019q4 are overlaid in panel (a), the original Plan is overlaid in panel (b). Back to Section 4.3.

B Bike market access measurement: further details

B.1 Calibrating commuting costs

In this section, we detail the calibration procedure followed to obtain the bike market access metric used in Table 11. First, we selected $N = 30$ itineraries within the city of Paris, making sure to include itineraries spanning the 20 districts and different distance ranges. We used Google Maps API to obtain the optimal route connecting the two endpoints by bike.

Second, we created a grid of $K = 64$ possible parameter combinations. Specifically, we cut the number of cost parameters to calibrate down from 3 to 2, since a 2-D minimisation problem allows us better to explore whether the loss function is concave or not. We then bundled together primary and secondary/tertiary roads, thus ending up with two parameters to calibrate, the cost parameter for roads, c_r, and the cost parameter for streets, c_s. In constructing the possible parameter combinations, we let each parameter to range between 1 (the normalised value we attribute to Plan Vélo bike lanes) and 5 and we do not impose $c_r > c_s$, thus letting in principle local streets potentially more costly to go through for a biker than larger and faster speed roads.[26]

Next, for each parameter combination $k \in K$, we construct the corresponding cost raster combining the transport infrastructure layers at our disposal and the corresponding cost parameters. Using the thus obtained cost raster, we get our optimal route applying the least cost path algorithm to each itinerary $i \in N$ and calculate the percentage of Google Maps optimal route that does not overlap with it. Setting this percentage as the residual for itinerary i and combination k, r_{ik}, the loss function that we seek to minimise is given by $loss_k = \frac{1}{N} \sum_i r_{ik}$.

The loss function is locally concave as Table 17 - binning the possible values for each cost parameter into three groups and displaying the average loss within each group relative to the lowest one - shows.

Table 17: Concavity of the loss function

	Roads (Low)	Roads (Middle)	Roads (High)
Streets (Low)	1.05	1.08	1.20
Streets (Middle)	1.01	1.00	1.20
Streets (High)	1.11	1.22	1.33

Notes: each cell contains the average loss for given set of parameter combinations rescaled by the lowest value.

The lowest value for the loss function is obtained for $c_s = 2$, which is identical to the

[26] We do not try values higher than 5 in our final calibration as we noticed that the loss function exceeded a given cutoff starts displaying a monotonically increasing behavior.

value we adopt for the construction of the baseline market access measure, and $c_r = 3.5$, which is slightly below the value for secondary/tertiary roads for the construction of the baseline market access measure. We then construct the calibrated cost raster combining the transport infrastructure layers at our disposal and the thus found cost parameters. Using the calibrated cost raster, we recalculate bilateral commuting costs and get a bike market access measure alternative to the one used as baseline, which we use as robustness.

B.2 Calibrating the commuting elasticity

Bilateral consumption flows are needed to calibrate the reduced form commuting elasticity $\tau\varepsilon$. While these are not directly observed in our dataset, for the last year in our sample we developed an imputation procedure to calculate a proxy for them. Specifically, for the last year in our sample, we observe daily transactions indexed by the merchant and card identifier. We do not know where the cardholder lives, nor we have any demographic information on him/her. However, we can, based on each cardholder shopping history, impute a "most likely" residence location. We do so by taking the following steps:

1. we retain transactions occurring in the city of Paris on weekends and on weekdays after 18h;
2. we further retain transactions that are usually carried out in the proximity of one's residence, which we identify as those transactions occurring in merchants identified by one of the following sectoral codes: 1071, 1072, 4724 (bakeries), 4773 (pharmacies), 4711B-D (supermarkets, minimarkets), 4721, 4722, 4723, 4725, 4729 (food stores);
3. we finally keep cardholders for which the number of observed transactions is $N \geq 9$.

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

$$\ln x_{ij} = \alpha_i + \alpha_j + \beta d_{ij} + e_{ij} \tag{A71}$$

where d_{ij} are bilateral commuting costs derived as described in Section 4, and α_i and α_j are, respectively, business and residence location fixed effects. Then, $\tau\varepsilon = -\widehat{\beta}$. We repeat this estimation for four different quarters of 2019, and obtain very comparable results regardless of the quarter choice, with $\widehat{\beta} \in (-0.06, -0.04)$.