The Welfare Consequences of Urban Traffic Regulations

(Preliminary, please do not cite without permission)

Isis Durrmeyer^{*} Nicolás Martínez [†]

May 11, 2022

Abstract

We develop a novel structural model to represent individual transportation decisions and the equilibrium road traffic levels and speeds inside a city. The model has two main advantages relative to the existing frameworks. First, it is a micro-founded equilibrium model with a high level of heterogeneity. The model accounts for individual heterogeneity in access to different transportation modes, values of travel time, and schedule constraints. Furthermore, our model considers heterogeneous road congestion technologies across different areas. The second advantage is that all the model parameters are estimated using multiple publicly available data. We apply our model to the Paris metropolitan area to predict the road traffic equilibria under driving restrictions and road tolls and measure each policy's welfare consequences. Our results suggest that all the policies decrease individuals' utilities: the benefits of relaxing road congestion and improving car speeds do not offset the losses for individuals from switching to other transportation modes or off-peak hours. However, road tolls raise significant tax revenues that, when they are redistributed to individuals, generate positive total surplus changes. In addition, these policies reduce emissions of global and local pollutants. However, they represent only a small gain once converted into monetary terms using standard social values for these emissions.

JEL Classification: L9, R41, Q52

Keywords: structural model, policy evaluation, transportation, congestion, distributional effects, air pollution

 $[\]label{eq:cond} \ensuremath{^*\text{Toulouse School of Economics, Universit\acute{e} Toulouse 1 Capitole and CEPR. E-mail: isis.durrmeyer@tse-fr.eu$

[†]Toulouse School of Economics, Université Toulouse 1 Capitole. E-mail: nicolas.martinez@tse-fr.eu

1 Introduction

Road traffic reduction has been crucial in large metropolitan areas because of the multiple negative externalities cars generate. For instance INRIX estimates an annual aggregate cost of congestion of 87 billion dollars for the U.S.¹ Pollution levels and air quality are also tightly related to the number of cars on the road.

However changes in road traffic level are difficult to predict because the road traffic level is the consequence of an equilibrium in which individuals make their transportation decisions independently. However these individual decisions affect everyone since the car speeds and individual trip durations ultimately depend on the traffic level. Predicting individual reactions to a change in their transportation environment is challenging since it requires knowing how the road traffic equilibrium is modified after individuals make their transportation decisions. Observational studies that measure the direct impact of a change in the transportation environment are limited by only being able to compare two equilibria, failing to separately identify the individual reactions from the equilibrium adjustments.

We define as transportation environment, all the factors that affect individual transportation decisions and are exogenous to the individuals. A transportation environment includes the trip characteristics such as the origin, destination and itinerary, the characteristics of the different transportation modes available to the individual, as well as possible traffic policies that would be in place. Car travel times are determined by the driving speeds, which are given by the traffic level equilibrium.

In order to analyze and predict the effects of changes in the transportation environment, we develop a novel framework to analyze individual responses to changes in their transportation environment in equilibrium. We build a new structural model to represent equilibrium traffic conditions in a metropolitan area. This model exhibits an important dimension of heterogeneity, at the individual and geographical levels. The first part of the model represents the choice of a transportation mode and a departure time by individuals with heterogeneous but fixed travel patterns (origin and destination). Since individuals have distinct travel patterns, different available transportation modes, and schedule flexibility, they are likely to react differently to a change in the transportation environment. Our model considers the different transportation mode alternatives to be imperfect substitutes. We also acknowledge possible schedule constraints, implying that commuting at different times are not perfectly substitutable. More precisely, we rely on a discrete nested logit choice model, which contains heterogeneity in choice sets, sensitivities to trip duration, and preferences for the different departure periods. The second part of the model represents the road congestion technologies; which model how driving speeds react to changes in the number of individuals using cars and the intensive margin of action (number of kilometers driven). Our model

¹Source: https://www.cnbc.com/2019/02/11/americas-87-billion-traffic-jam-ranks-boston-and-dc-as-worst-in-us.html.

takes into account spatial heterogeneity by allowing the road technology to be different across regions of the metropolitan area.

The model has the advantage of being transparent, tractable and can be estimated with combinations of data that are typically publicly available for many metropolitan areas. We also provide a methodology to verify whether the model parameters are such that the equilibrium is unique. This model differs from the existing ones in three key aspects. First, the model represent equilibrium transportation decisions for the entire metropolitan area rather than focusing on a specific road. Second, it accounts for the existence of different types of roads which have possibly different congestion technologies. Third, all the model parameters are estimated and we obtain the joint distribution of preferences, trip distance, characteristics, and transportation mode choice set, which is important to analyze the effect of changes in the transportation environment at the individual level. To allow for such individual heterogeneity, we consider some factors to be fixed. In particular we consider individual locations and destinations to be fixed and do not allow them to change in response to a change in the transportation environment. We also consider that the transportation modes available to an individual are fixed. Finally, we focus on trips whose motives are either work or study since they are unavoidable. We thus consider individuals have to take their trips and do not model the number of taken trips.

We apply our model of transportation decisions and congestion to Paris metropolitan area (Îlede-France region) and combine data from different sources to estimate the model parameters. We rely on a survey from 2010, where respondents provided detailed information about all the trips done the day before the interview. We use these data to estimate the transportation mode choice model. The survey does not provide trip durations using the non-chosen alternative transportation modes or alternative departure times. To overcome this issue, we supplement the survey with data on expected travel times using Google Directions and TomTom APIs for public transport and private vehicles during peak and non-peak hours. We estimate the congestion technologies using high-frequency data on traffic density and speed from road sensors at the day and hour level for 1,371 sensors, covering the highways, the ring roads, and the city center.

We the use our structural model to predict the effects of policy instruments that reduce road traffic in the Paris metropolitan area. More specifically we compare the effects of road tolls to simple driving restrictions. The advantage of driving restrictions is the simplicity of implementation, only requiring compliance controls. Paris and the surrounding region have used alternate traffic restrictions five times since 1997. The longest alternate-day travel scheme lasted four days from December 6th to 9th 2016. Since 2017, the emergency plan was triggered six times because of pollution peaks. The local authorities relied on driving restrictions based on car vintage and fuel type to control these air pollution episodes. Yet, road tolls have been introduced in many European cities. For instance, Stockholm and London have put in place systems of congestion charges, restricting access to the city center during peak hours of weekdays to those who pay a fee. Price mechanisms have the advantage of sorting individuals according to the benefits they get from driving: those who stop driving are those with good transportation alternatives to driving, limiting the welfare costs of traffic regulations. In comparison, driving restrictions affect all individuals identically. In addition, price instruments generate tax revenue that can be redistributed to individuals, mitigating the surplus losses.

We specify area-specific speed-density functions to represent the congestion technologies and estimate them for five heterogeneous driving areas: highways that connect the far suburb to the city center, the non-highway roads in the distant suburb, the far suburb, close suburb, the ring roads (*periphériques*) and the city center (*Paris intra-muros*). To better capture the heterogeneity between areas, we estimate the speed-density curves in a flexible way using Bernstein polynomials. We find large differences across areas, indicating the importance of infrastructure heterogeneity at determining driving speeds. We estimate downward-slopping non-linear speed-density relationships that reflect severe capacity constraints and are consistent with the fundamental traffic diagram.

For the demand side, we use a nested logit discrete choice model that captures both the transportation mode and the departure time choices of the individuals. Our estimation results for the transportation choice model reveal significant heterogeneity in sensitivities to travel times and schedule inflexibility. We use the estimated parameters of individuals' preferences and road congestion technology to simulate the effects of hypothetical transportation policies, endogeneizing traffic density, and its impact on trip durations.

First, we compare the effects of three simple policies: a driving restriction, a fixed toll, and a variable toll. We focus on comparing different transportation policies that imply identical traffic reduction and compare the policies' effects on individuals and emissions. We are able to measure the impact of the policies on global pollutant emissions (carbon emissions) and local pollutant emissions (nitrogen oxide, particulate matter and hydrocarbon emissions). The policies are restricted to peak hours, so we consider individuals are free to drive at non-peak hours. We first analyze the aggregate effects of different policy stringency levels. We find that all the regulations are costly for individuals, as speed gains cannot compensate the cost coming from the policies. However, the price instruments generate tax revenue that can improve the total surplus when fully redistributed to individuals. Moderate values of the tolls have a positive net impact on individuals if the revenue is redistributed. Without accounting for the tax revenue, driving restrictions hurt individuals less, as they force everyone to contribute to the traffic reduction. On the contrary tolls need to be large enough to induce enough traffic reduction. The variable toll is more efficient than the uniform toll since it targets long distance commuters, which exert the largest congestion externality. Next, we fix the policies' objectives in terms of traffic reduction and compare in more details tolls and driving restrictions at a specific stringency level that reduce traffic at peak hour by 40%. We observe that the variable tolls generate the largest inequalities across individuals.

Finally, we investigate whether the losses for individuals can be reduced by using alternative policy instruments or by combining the policies with other interventions. More specifically, we study car

vintage or fuel-based driving restrictions, driving licenses allocated through an auction like exist in some cities. These instruments do not perform better than the simple ones. In particular, the vintage and fuel-based restrictions improve the emission reductions through a better targeting of polluting cars. However, these policies have larger distributional effects since they put the cost of the traffic reduction on only a fraction of drivers. We also measure the potential gains from using differentiated tolls and non-linear variable tolls. Finally, we evaluate the role of access to public transport and public transport efficiency and cost to the surplus losses. We find that connecting the 25% of the population which currently does not have access to public transport is the best to reduce the policy surplus losses. As a robustness exercise, we perform the same analysis for stricter policies that are applied all day. We find that these policies hurt more individuals but are much more efficient at reducing emissions. The ranking across the different policy instruments is the same as for peak hours only policies.

This paper relates to the literature measuring the impacts of driving restrictions. Davis (2008) provides reduced form evidence that driving restrictions in Mexico City were unable to curb pollution in the long run because it induced individuals to buy a second car to bypass the regulation. Moreover, Gallego et al. (2013) provide evidence from additional cities and show that driving restrictions can successfully reduce pollution in the short run, but in the long run the policy impact is negative. Our analysis is very different both in terms of method and focus. We estimate a structural model that credibly represents transportation decision of individuals and use it to evaluate the effects of hypothetical policy. We focus on the short-run effects of transportation policies and we do not allow individuals to buy a second car or a car with different specifications in response to the policy. This is consistent with policies that are temporary and used to deal with a pollution peak episode. We also assume the total number of trips is unaffected by transportation policies, which is consistent with our restriction on work and study-related trips. Medium and long-term distributional consequences of transportation policies are also important to evaluate as shown by Tsivanidis (2018), Herzog (2021), and Barwick et al. (2021). While these approaches incorporate the long term effects of new transportation infrastructure or urban traffic policies on individuals' location decisions across a metropolitan area, they are not able to account for as much level of heterogeneity as we do.

The model that represents individuals' transportation decisions is similar to Lucinda et al. (2017) who identify demographic characteristics that affect individuals' preferences for the different transportation modes. We build a more realistic model that allows the alternatives available to vary across individuals and introduce imperfect substitution between peak and non-peak hour alternatives. The most significant novelty of our model is that trip durations endogenously change in response to policy interventions. Batarce and Ivaldi (2014) develop a model with endogenous congestion based on the number of trips made by individuals, where driving at any point in time generates congestion the whole day. Instead, our equilibrium congestion level varies according to the period, with only individuals driving during that period exerting congestion. Our approach is

closer to Basso and Silva (2014) who also allow peak and non-peak departure time substitution. However, they model a single road and limit the decision to driving or taking the bus. Instead, we allow for a richer set of transportation modes and focus on citywide congestion, with area-specific heterogeneity. Tarduno (2022) studies the role of route substitution in the efficiency of cordon pricing mechanisms to alleviate congestion. He estimates a structural model of route choice and departure time for the Bay area. While we abstract from route choice and have a more limited approach to departure time, his model focuses on individual driving and does not consider the important role of substitution towards other transportation modes.

Arnott et al. (1990) and Arnott et al. (1993) started the literature on bottleneck congestion models and the role of congestion pricing. Van Den Berg and Verhoef (2011) and Hall (2019) focus on the distributional effects of pricing policies in such models and show that the policies can be welfare improving in the absence of redistribution. Kreindler (2020) uses experimental data to estimate a model of departure time choice. His results show that while the bottleneck model allows for a detailed representation of the dynamics of congestion in a specific road, it fits poorly the congestion patterns of a whole city. Since our paper focuses on the effects of citywide policies, we consider alternative modeling assumptions for congestion that better fit our setup.

Our estimation of the road congestion technology builds on Couture et al. (2018) who estimate demand and supply for travel times and define speed as the equilibrium outcome. While we rely on the same variation to identify the road technology, we rely on a structural model that specifies individuals' utility. Our model allows us to evaluate more extensive sets of hypothetical policies and compute welfare costs in monetary terms. Other papers have focused on estimating the relationship between congestion and traffic density by relying on exogenous shocks on traffic density. For instance, Yang et al. (2020) leverage the driving restriction implemented in Beijing for identification. Relatedly, Anderson (2014) relies on a 35-days strike by public transport employees in 2003 to estimate the substitution between private vehicles and public transit in Los Angeles. His estimates suggest a considerable impact of the strike on traffic because public transit users are also those driving on the most congested roads. His study points out two essential features. First, it is crucial to account for road heterogeneity at the city level; we do this by estimating speed-density relationships for four different areas of Paris. Second, it is important to take into consideration the heterogeneity in individual trips' origins and destinations.

Previous literature on traffic restrictions in Paris also relies on a model with endogenous congestion. De Palma and Lindsey (2006) focus on the effects of road pricing using a bottleneck model simulation over the road network. Kilani et al. (2014) abstract from the road network and bottleneck dynamics but differentiate across different public transport modes. Our model builds on the insights from previous approaches but allows for a larger degree of heterogeneity at the individual level. Even though we consider area-wide congestion levels, when we compute individuals' travel times, we rely on the exact route the individual is taking. This implies that individuals departing from the same location but with different destinations face different congestion levels. We also allow for

heterogeneity in individuals' valuation of travel time by leveraging a large set of demographics variables included in the survey. Unlike previous literature that typically relies on calibrated parameters, we estimate all the parameters in our model, providing a transparent methodology that applies to other setups.

De Palma et al. (2017) analyze the role of public transport congestion on individuals' departure times. However, their application focuses on a specific urban train line. Our study estimates individuals' disutility of public transport overcrowding, and our overcrowding measure accounts for heterogeneity across the different urban railway transit lines. Bou Sleiman (2021) uses difference-in-differences to estimate the traffic impact of route and mode choice changes linked to a natural experiment, the closing of a riverbank road in Paris. To keep our model tractable, we keep route choices fixed in our counterfactuals. While this assumption limits the capacity of our model to study highly targeted policies, such as tolling a specific road, we can still analyze a large set of relevant policies.

2 A structural model of transportation decisions and traffic conditions

We develop an equilibrium model representing individuals' choice of a transportation mode and departure time. The model considers congestion levels on the roads as endogenous, depending directly on the number of drivers at each time period. To represent the relationship between speed and traffic density, we model the road congestion technologies for the different areas of the metropolitan area. Finally, we compute the equilibrium level of traffic.

2.1 Transportation mode choice model

First, we introduce the structural model representing how individuals decide which transportation mode to use as well as their departure time. We consider that the origin and destination of the trips are fixed and exogenous. We do not allow for an outside option, as we model the choice of individuals facing non-avoidable trips. We make the simplification that individuals choose between T time periods denoted by 1, ..., T. Our model is a nested discrete choice model, and we follow the standard distributional assumptions from the literature (see, Train, 2009). The nests are the different transportation modes. We assume individuals make a sequential decision: first, they choose a transportation mode, and then at which time they leave. The utility function of an individual n associated with transportation mode j, and departure time t is assumed to be linear in the mode and departure period characteristics:

$$u_{njt} = \beta'_n X_{njt} + \epsilon_{njt}$$

where X_{njt} are the observed characteristics of the transportation mode j for the period t, and include its monetary cost and duration. β_n is a vector of coefficients of preferences for these variables for consumer n, and ϵ_{njt} is a random idiosyncratic term that is assumed to be extreme value distributed. This assumption implies that the different modes and departure periods are imperfect substitutes. We allow for correlation between these idiosyncratic terms across different periods so that:

$$\epsilon_{njt} = \zeta_{nj} + \sigma \tilde{\epsilon}_{njt},$$

where σ represents the degree of independence between the preference shocks across different periods for the same transportation mode and is a parameter to estimate. When $\sigma = 1$, the preference shocks for different periods are independent while if $\sigma = 0$, the different time periods are perfect substitutes. Each individual has a choice set \mathcal{J}_n , which comprises all the transportation modes to which she has access. Each individual chooses the combination of alternative j^* and time t^* that maximizes their utility:

$$\{j^*, t^*\} = \underset{j \in \mathcal{J}_n, t \in \{1, \dots, T\}}{\arg \max} u_{njt}$$

We assume that ϵ_{njt} are identically and independently distributed across individuals and follow a type one extreme value distribution. The probability that individual n chooses the transportation mode j and departure time t is:

$$s_{njt} = \frac{\exp\left(\frac{\beta'_n X_{njt}}{\sigma}\right)}{D_{nj}^{1-\sigma} \sum_{j' \in \mathcal{J}_n} D_{nj'}^{\sigma}},$$

where $D_{nj'} = \sum_{t=1}^{T} \exp\left(\frac{\beta'_n X_{nj't}}{\sigma}\right)$. We further parameterize the individual heterogeneity in preferences and assume that β_n are functions of observable demographics characteristics Z_n :

$$\beta_n = \bar{\beta} + Z_n \Gamma$$

We estimate $(\bar{\beta}, \Gamma, \sigma)$ using the method of maximum likelihood; we want to find the values of the parameters that rationalize the best the observed choices given the theoretical probabilities. The log-likelihood function that we maximize is:

$$LL(\bar{\beta},\Gamma,\sigma) = \sum_{n=1}^{N} \sum_{j \in \mathcal{J}_n, t \in \{1,\dots,T\}} \omega_n d_{njt} \times \log\left(s_{njt}(\bar{\beta},\Gamma,\sigma)\right),$$

where d_{njt} is an indicator equal to one when the transportation mode j and time t are chosen and zero otherwise. ω_n corresponds to the weight of individual n. The identification of the parameters comes from observing a cross-section of individuals who have different choice sets, different trip characteristics, and different demographic characteristics. To evaluate welfare effects of transportation policies, we rely on consumer surplus. The consumer surplus of a trip of a given individual n is defined as the expected utility of the choice that maximizes the utility and is:

$$CS_n = \frac{1}{|\alpha_n|} \log \left(\sum_{j \in \mathcal{J}_n} \exp(\sigma \log(D_{nj})) \right) + C_n,$$

where C_n represents a constant utility term that cannot be identified and α_n is the parameter of sensitivity to the trip cost which converts the utility into monetary terms. The consumer surplus is not identified, but the variation of consumer surplus eliminates the constant C_n and thus is identified and given by:

$$\Delta CS_n = \frac{1}{|\alpha_n|} \left[\ln \left(\sum_{j \in \mathcal{J}_n^1} \exp(\sigma \log(D_{nj}^1)) \right) - \ln \left(\sum_{j \in \mathcal{J}_n^0} \exp(\sigma \log(D_{nj}^0)) \right) \right]$$

Where \mathcal{J}_n^1 and D_{nj}^1 represent respectively the choice set and the expected utility of transportation mode j under the counterfactual scenario, while \mathcal{J}_n^0 and D_{nj}^0 represent their initial values. We can further decompose the variation in consumer surplus into a partial policy effect which measures the policy at constant final speeds and an equilibrium speed effect measured at the initial situation, without policy. To make the expression clearer, we make apparent the dependence between the driving speed and the utilities associated with the transportation mode $D_{nj}^0(\mathbf{v_0})$ and $D_{nj}^1(\mathbf{v_1})$ where $\mathbf{v_0}$ and $\mathbf{v_1}$ represent the initial and final vectors of driving speeds. The decomposition is given by:

$$\Delta CS_n = \frac{1}{|\alpha_n|} \left[\underbrace{\log\left(\sum_{j \in \mathcal{J}_n^1} \exp\left(\sigma \log\left(D_{nj}^1(\mathbf{v_1})\right)\right)\right) - \log\left(\sum_{j \in \mathcal{J}_n^0} \exp\left(\sigma \log\left(D_{nj}^0(\mathbf{v_1})\right)\right)\right)}_{\text{policy effect at constant speed}} + \underbrace{\log\left(\sum_{j \in \mathcal{J}_n^0} \exp\left(\sigma \log\left(D_{nj}^0(\mathbf{v_1})\right)\right)\right) - \log\left(\sum_{j \in \mathcal{J}_n^0} \exp\left(\sigma \log\left(D_{nj}^0(\mathbf{v_0})\right)\right)\right)}_{\text{Equilibrium speed effect}}\right)$$

In our data, we observe a sample of M individuals representing the entire population in the Paris area using the survey weights. By aggregating the transportation mode decisions of all individuals, we obtain the total number of individuals using transportation mode j at period t as:

$$N_{jt} = \sum_{n=1}^{M} \omega_n s_{njt}$$

,

where ω_n is the weight of individual *n* in the sample.

2.2 Road traffic conditions and congestion technology

We model congestion technology at the local level, splitting the city in a = 1, ..., A mutually exclusive areas. The driving speed in an area depends on the road technology and the traffic density in that area only. Following the transportation literature, we base our congestion technology model on the Fundamental Traffic Diagram (see Small et al., 2007). We model speed as a weakly decreasing function of the traffic density, represented by the occupancy rate that traffic monitoring stations typically record. The occupancy rate is defined as the percentage of a fixed time period during which the traffic sensor detects a car. Geroliminis and Daganzo (2008) empirically show the existence of a FTD at the city level, called a Macroscopic Fundamental Diagram. Other applications include Yang et al. (2020), Couture et al. (2018), and Anderson and Davis (2020). We follow their approach but allow for heterogeneity within the city by considering area-specific congestion technologies. We also rely on minimal functional form assumptions.

Congestion levels can be different throughout the day, but we consider the road technology is fixed. The road technology represents all the elements that determine the speed at which individuals can drive for a given traffic density. It represents the type of road (high speed road or city road), the presence of traffic lights and intersections and the number of entry or exits that may force drivers to slow down. Formally, we define the speed in area a at time t, v_t^a to be a function of the occupancy rate τ_t^a :

$$v_t^a = f^a(\tau_t^a) + \nu_t^a$$

 f^a represents the congestion technology, which indicates how the speed decreases when the number of cars increases in the area. We make minimal function form assumptions on f^a by relying on basis polynomial. More specifically, we use Bernstein basis polynomials. We assume we observe a sample of independent realizations of traffic densities and speeds. We estimate the following equation:

$$v_t^a = \sum_{l=0}^{L} c_l^a B_l(\tau_t^a) + \nu_t^a.$$

The coefficients c_l^a are the parameters of interest and B_l are the Bernstein basis polynomials of degree L which expressions are:

$$B_l(\mathcal{T}) = \binom{L}{l} \mathcal{T}^l (1 - \mathcal{T})^{L-l}$$

Under the assumption that the occupancy rate τ_t^a is independent on the speed shock ν_t^a , we can estimate the parameters of the Bernstein polynomial c_l^a using standard linear methods. We can easily impose the relevant shape restrictions (positivity of the speed function and and a weakly decreasing function) through linear conditions on the parameters c_l^a :

$$\begin{array}{l} c_l^a \geq 0 \quad \forall l \\ c_{l+1}^a \leq c_l^a \quad \forall l=0,...,L-1 \end{array}$$

The congestion technology model determines speeds from the traffic occupancy rates. We thus need a mapping between the number of drivers predicted by the transportation mode choice model and the occupancy rate on the road. Because the traffic level depends on how many individuals choose to drive but also how long they drive we consider a mapping between the total number of kilometers driven and the occupancy rate. More specifically, we assume an affine relation between the occupancy rate and the total number of kilometers driven at the area level. Because we do not model the entire traffic source, we let the possibility that a fraction of the road traffic is irreducible (trucks, delivery cars, buses, etc...). Formally, we consider:

$$\tau^a_t = \phi^a_t \times K^a_t + \gamma^a_t$$

Where ϕ_t^a is a scale parameter that maps the number of kilometers driven to the surface of the area, and γ_t^a represents the irreducible traffic.

2.3 Equilibrium of the model

In this model, an equilibrium is defined by individual probabilities to drive at each period and speeds for each area and period. We substitute the individual probabilities in the speed function to express the equilibrium in terms of speeds only and get the following system of non-linear equations:

$$v_t^a = f^a \left(\phi^a \sum_{n=1}^I w_n \cdot s_{nt} K_n(\mathbf{v}) \cdot \mathbb{1}\{n \text{ crosses area } a\} + \gamma^a \right)$$

There is no general result that guarantees that the system of non-linear equations always have a unique solution. However, there are two special cases where the speed equilibrium is unique. The first special case is when there is only one period and one area, so we have just a non-linear equation to solve. Because the speed function f^a is monotonically decreasing, we are sure that if a speed equilibrium exists, it is unique. The second special case is when we have one area and multiple periods. The proofs relies on the fact that Jacobian of the system of equation has positive terms on the diagonal and negative terms off-diagonal. This is the consequence of two key features of our model: the speed function is decreasing and the different departure periods are substitutes. We provide the proofs of uniqueness under these two special cases in Appendix B.

Even though there is no proof of the uniqueness of the equilibrium, we provide a method to check if the system of equations in speeds has a unique solution given our set of estimated parameters. The approach consists of defining the function:

$$g_t^a(\mathbf{v},\kappa) = \kappa v_t^a + (1-\kappa) f^a(\mathbf{v})$$

and check whether there exists $\kappa \in [0, 1[$ such that g is a contraction. Recall that a function g is a contraction if it is K-Lipschitz, with K < 1, implying that $\forall \mathbf{v} \in [\underline{\mathbf{v}}, \overline{\mathbf{v}}]$:

$$||g(\mathbf{v}') - g(\mathbf{v})|| \le K ||\mathbf{v}' - \mathbf{v}||.$$

We use the supremum norm, so the Lipschitz coefficient K is given by:

$$\max_{a \in 1, \dots, A} \max_{t \in 1, \dots, T} \max_{a' \in 1, \dots, A} \max_{t' \in 1, \dots, T} \max_{\mathbf{v} \in [\underline{\mathbf{v}}, \overline{\mathbf{v}}]} \left| \frac{\partial g_t^a(\mathbf{v}, \kappa)}{\partial v_{t'}^{a'}} \right|$$

If we can find $\kappa \in [0, 1]$ such that K < 1, the function $\mathbf{g}(.)$ is a contraction and if the iteration process converges, it converges to the unique solution of the system. If there are a set of κ such that the function $\mathbf{g}(.)$ is K-Lipschitz with K < 1, then we select the value of κ that implies the lowest K to ensure the maximum speed of convergence. Therefore, we solve for:

$$\min_{\kappa \in [0,1[} \max_{j \in A \times T} \max_{\mathbf{k} \in A \times T} \max_{\mathbf{v} \in [\underline{\mathbf{v}}, \overline{\mathbf{v}}]} \left| \frac{\partial g_j(\mathbf{v}, \kappa)}{\partial v_k} \right|$$

We provide in Appendix B.2 the analytical formula for the Jacobian of $\mathbf{g}(.)$. and some results on how the Lipschitz coefficient varies with the tuning parameter κ and the policy environment.

3 Specification and estimation of the transportation choice model

We estimate the transportation model parameters using a combination of two main datasets. We rely on survey data on individuals' commuting patterns in the Paris area ("Enquête Globale Transport 2010", EGT hereafter).² It is combined with a second self-constructed dataset on expected trip durations and itineraries for cars and public transport from TomTom and Google Maps. We also leverage several useful datasets. First, we use proprietary car sales and characteristics data to estimate cars' fuel consumption and emissions of pollutants. Second, we obtain data on public transport passenger flows, train capacities and public transport schedules to estimate the overcrowding levels in the urban railway transit. Finally, we gather information about the past prices of public transport tickets and subscriptions.

3.1 Data, sample selection and choice set characteristics

The EGT data contain information about the departure and arrival times and precise locations, the transportation mode used for every trip, and the trip's motives. To be consistent with the assumption that the trips are non-avoidable, we restrict the sample to work and study-related trips. A trip's transportation modes choice set might depend on previous trips during the day, as an individual taking the car to work will probably take it to go back home. To eliminate possible

 $^{^2 \}mathrm{Enquête}$ Globale Transport (EGT) - 2010, DRIEA, ADISP. We provide more details about the survey in Appendix A.1.

correlation in the decisions, we restrict our sample to the day's first trip and trips in the morning widely defined (between 4:00 a.m. and 2:59 p.m).

In addition to the residence location, the survey records household and individual-level sociodemographic characteristics such as age, socio-professional activity, household size, income class, and housing characteristics. Since individuals within a household have different demographic characteristics and make different trips, we choose the individual as the observation unit rather than the household. Furthermore, we assume that individuals' decisions are independent within families, implying that more than one individual can use the car when the household owns a vehicle. This assumption ignores potential cost savingsand detours associated with carpooling. However, modeling such joint decisions would add too much complexity to the transportation mode choice model. Finally, survey recording errors, such as missing departure or arrival locations, force us to drop some observations, resulting in a final sample of 12,975 individual trips. With the sample weights, we have transportation decisions representing 4,034,801 million individuals, which corresponds to approximately one-third of the total population of the Paris metropolitan area (11.85 million inhabitants in 2011).

The survey only provides as household income information a bracket, with the full income distribution divided in ten brackets. A relevant income measurement should take into account the number of individuals in the household composition. We thus construct a proxy of households' wealth by leveraging the precise information we have about the house size and location. More precisely, we estimate the real estate value per square meter in the household's neighborhood using estate transaction data from the French tax authority. We then multiply this estimate by the apartment or house size. The real estate prices are estimated using all transactions in the municipality and the neighboring municipalities and weighting each match by the inverse of the distance to the household's location. Finally, we divide our estimated household wealth by the number of consumption units in the household. As defined by the French Statistical Institute (INSEE), the first individual in a household represents one consumption units. We provide more details about the data and the estimates in Appendix A.4. We find a correlation coefficient of 0.51 between the income proxy to the midpoints of the income bracket provided by the EGT data.

The survey contains helpful information to define individuals' choice set: ownership of a car, two-wheel motorized vehicle and bicycle, as well as some characteristics of the vehicles. Thus, we know precisely which transportation modes are available to each individual. We consider five transportation mode alternatives: bicycle, public transport, two-wheel motorized vehicles, walk, and car. If the household does not own a car or a motorized two-wheel vehicle, the alternative is considered non-available for the individual. If walking or biking takes more than 2.5 hours, we define those alternatives as non-available. For some trips, Google Maps cannot find a public transport itinerary, so we consider the option as non-available (this occurs for 33% of the trips).

The survey also records whether individuals own a public transport pass, the type of subscription and if they are partly or entirely reimbursed for their trip. Using this information, the vehicle characteristics, and some additional assumptions, we estimate the cost of each trip for each available transportation mode. The estimates are both trip and individual-specific and account for differences in the public transportation subscriptions and vehicle characteristics in addition to the trip distance. Consequently, the driving and public transport costs vary with the mode, the trip's distance, and individual characteristics. We provide the precise methodology and related assumptions made to estimate the costs of trips in Appendix A.2.

Since the survey data only provide trip durations for chosen transportation modes, we rely on additional data to compute travel times for non-taken alternatives. Trip durations are necessary to estimate the sensitivity to the values of travel time for all individuals, which constitute the key model parameters. Google Maps and TomTom Directions APIs (Application Programming Interfaces) provide directions and expected travel times associated with an origin and destination pair of GPS coordinates at a specified departure time. Google Maps predicts trip durations and provides car and public transport itineraries. Google Maps is widely used in most countries, and due to its accuracy, many individuals rely on it to obtain their trip itinerary. The data from its related API services have been used previously in the transportation literature (see Kreindler, 2016, Hanna et al., 2017, Akbar and Duranton, 2017 and Akbar et al., 2018). The TomTom directions API is fairly similar to the one from Google Maps, with similar reliability, but has the inconvenient of only being available for cars and not for public transport. For this reason, we use Google Maps data to estimate trip durations using public transport alternatives, while we rely on TomTom data for expected durations by car.³ In the model, we consider two time periods: peak and non-peak hours. We use the predicted car trip durations at 8.30 a.m. for peak-hour and 10.30 a.m. for non-peak hour. We specify the trip query to a future date and thus the predicted car trip durations are not subject to the current traffic conditions and their idiosyncrasies. We did the TomTom queries after the COVID crisis started but during a period without any restrictions. We may be worried that TomTom modified its prediction algorithm because of COVID, so we compare TomTom's expected trip durations with Google Maps predictions that we obtained before. We find that they predict very similar durations (see Appendix A.3 for more details on the queries and the comparison between the periods before and after COVID). We use the predicted car duration at non-peak hours for two-wheel motorized vehicles, assuming they can bypass traffic.

We also estimate trip distances and trip durations by walk and bicycle. We first obtain the walking trip distances from the GPS coordinates using Open Street Maps, which finds the fastest itinerary given the street network. Then, we compute the average speed for individuals who bike and walk using their declared trip durations and the distance. We find a walking speed of 5.52 km/hr and 10.23 km/hr for biking. Finally, we check that the speeds are not sensitive to rounding errors in

³We use TomTom data because the service is free while Google Direction API is not.

the reported trip durations by excluding observations reporting a multiple of 10 minutes.

Table 1 provides summary statistics of the main variables used. We can observe some heterogeneity in the number of alternatives across the population. While some individuals have access to only one transport mode, they still need to make a departure time choice and thus are useful for our estimation. Taken alternatives have, on average, shorter travel times and a cheaper cost than non-taken ones. However, the difference in cost between taken and non-taken transportation modes is smaller than the difference in duration, supporting the idea that duration is a key factor for choosing a transportation mode. The large standard deviation for the trip durations and distances show large differences between the characteristics of the taken trips.

		<i>J</i>			
Variable	Mean	Median	Std. dev.	Min	Max
No. modes available	2.77	3	0.9	1	5
Trip distance (km)	12.92	8.1	13.94	0.06	140.45
Car vintage	2004	2005	5.2	1969	2011
Duration for chosen alternatives	34.75	30	22.62	0.87	205
Duration for non-chosen alternatives	41.54	31.95	34	0.57	279.18
Cost: chosen alternatives	0.92	0.65	1.15	0	11.88
Cost: non-chosen alternatives	1.02	0.73	1.21	0	14.24

Table 1:	Summary	statistics.
----------	---------	-------------

Notes: Duration in minutes, cost in \in . All statistics are weighted using the survey weights.

Table 2 provides a comparison between trips' duration and cost by transportation mode. Taking the car is the fastest option and is available to the largest fraction of individuals, but the low initial shares suggest that the large monetary cost dissuades many from choosing it. Interestingly, public transport is not, on average, the fastest nor the cheapest alternative, yet 45% of the individuals choose it in the sample.

	*	*	-		
Variable	Mean	Median	Std. dev.	Min	Max
Duration					
Bike	52.12	39.65	40.66	0.57	149.98
Public transport	46.93	41.68	26.9	4.1	279.18
Motorbike	17.71	15.58	12.36	0.72	93.82
Walk	64.83	54.52	38.57	1.54	149.95
Car, peak	26.55	21.02	20.65	0.87	123.73
Car, non-peak	20.42	16.55	15.98	0.73	126.15
Cost					
Bike	0.64	0	0.82	0	1.7
Public transport	1.25	1.24	1.27	0	10.55
Motorbike	1.21	0.72	1.39	0	13.72
Car	1.17	0.76	1.25	0	14.24
	Mode a	availability		\mathbf{Sh}	ares
Bike	6	69.94	_	2	.08
Public transport, peak	7	71.54		30).33
Public transport, non-peak	7	71.54		14	4.52
Motorbike	12.85			2	.08
Walk	4	12.72		1	5.8
Car, peak	7	79.75		22	2.88
Car, non-peak	7	79.75		12	2.31

Table 2: Summary statistics by transportation mode

Note: Duration expressed in minutes, cost in \in . Mode availability and initial shares in %. All statistics are weighted.

To allow for a potential disutility of overcrowding in the public railway transport, we construct a line and period-specific measure of overcrowding for all the subway and urban railway transit lines. We use data from 2015 on passenger flows, the oldest year available, combined with data on hourly public transit schedules and estimates of the train capacity for each line. We measure the overcrowding level by the ratio between the hourly number of passengers divided by the line's capacity. This variable reflects the time heterogeneity in train passenger flow, schedules, and capacities. From the urban railway line-level overcrowding estimates, we obtain individual overcrowding levels by weighting the line-level overcrowding measures by the percentage of the trip duration spent in the line. We give all the details on the data used and the construction of the overcrowding variable in Appendix A.5. On average we find that the public railway transport is at 89% capacity at non-peak hours and reaches 143% at peak hour.

Table 3 shows the distribution of trip characteristics across the different demographic categories. For instance, we see that the youngest category (less than 18 years old) have, on average, trip distances three to four times lower than the rest of the population. The explanation is the proximity of schools to their residence. Interestingly, we do not observe large differences in car usage for the first three quintiles of the income proxy distribution but do observe a considerably lower usage for the top two quintiles. However, we see that the public transport's share increases with the income proxy, as well as the choice of traveling at peak hour. While the share of white-collar individuals using their car is very similar to those of workers using the same transport mode, the share of white-collar individuals using it at peak hour is much larger. We take this as evidence of differences in schedule flexibility across socio-professional categories.

		Mean			Mode shares		
	Freq.	Distance	Duration	Car	Peak/	\mathbf{PT}	Peak/
					car chosen		PT chosen
Age							
$Age \le 18$	25.4	4.21	18.9	21.3	93.7	31.8	81.7
Age \in]18- 25]	10.8	14.8	38.9	20.6	57.4	68.9	62.5
Age \in]25- 35]	17.8	15.5	35	36.7	57.6	52.7	66.6
Age \in]35- 45]	19.4	17	35.6	45.7	57.9	43.2	62.1
Age \in]45- 60]	24	16.5	34.9	45.3	63.5	43.9	65.9
Age > 60	2.56	13.4	30	49.9	56.1	39.7	64.6
Income proxy							
Wealth $\leq 110,000$	20	14.6	32.8	36.7	57	41.3	58.3
Wealth \in]110,000-152,000]	20	14.9	33.2	39.5	65.1	42.7	64.8
Wealth \in]152,000-205,000]	20	13.4	31.8	37.2	65.3	44	69.7
Wealth \in]205,000-283,000]	20	12	30.2	33.7	68.4	47.7	71.1
Wealth $> 283,000$	20	10.1	28.5	28.8	70.8	48.6	72.7
Socio-professional catego	ory						
Independent	3.58	17	30.5	59.5	51.8	24.4	61
White collar	33.3	15.9	35.2	39.3	70.5	49.6	74.5
Blue collar	30.6	16.4	35.3	44.5	50.2	45.2	54.4
Student \leq high school	25.6	4.09	18.8	21.6	93.3	31.5	82.8
Student - higher education	6.86	15.1	41.9	11.7	68.1	80.4	58.8
Household size							
Couple and/or children	84	13.1	31.2	35.6	65.9	42.9	68.6
Single	16	12.7	32.2	32.9	60.2	55.1	63.8
Average		13	31.3	35.2	65	44.8	67.6

Table 3: Summary statistics by demographic groups.

Note: Duration expressed in minutes, distance in km, wealth per household consumption units in \in . Frequencies and mode shares are in %. Distance and duration are those of the chosen transportation modes. All statistics are weighted using the survey weights.

3.2 Model specification and estimation results

We specify the utility as a linear function of the following variables: transportation mode-specific intercepts, the trip's monetary cost, and the logarithm of the trip duration. In these types of discrete choice models, only differences in utilities are identified so we consider walking as the baseline option and normalize the intercept to 0. Note that the total utility of walking is not normalized 0 because the utility contains the trip duration. We use the logarithm of the trip duration since it implies that the marginal valuation of travel time decreases with the trip duration: spending one additional minute hurts more someone taking a short trip than another one commuting for a longer time. We allow the sensitivity to the trip duration to vary with age (we rely on six age classes) and income (we rely on the quintiles of the wealth proxy). The trip duration might be differently valued if individuals have to make a physical effort or if they can use their time to read or listen to the radio. We, thus, add an interaction between a dummy for walk and bicycle and the trip duration in the utility function. While we allow for heterogeneity in the sensitivity to the trip duration, we do not allow heterogeneity in the cost sensitivity. We choose to put all the heterogeneity on the sensitivity to the duration for two reasons. First, the trip durations are rather precisely estimated, while the trip costs estimates rely on several assumptions and imputations. Second, the cost display much less variation across modes and individuals than trip durations.

Individuals may have different schedule flexibility, and it may be very costly for some individuals to leave at non-peak hour. Workers in some specific economic sectors and young individuals that go to their study place may have less flexibility regarding their departure time for work. Therefore, we interact the socio-professional category with a dummy for non-peak hour. Families typically try to coordinate on departure times and may be less flexible, so we also have an interaction between an indicator for families and non-peak hour.

Additionally, we add some controls to represent the characteristics of public transport. First, we use the overcrowding level in the metro and trains, a dummy if the public transport is composed by railway modes only and the number of layovers. Following the literature on public transport congestion, such as De Palma et al. (2017) and Haywood and Koning (2015), we assume the utility is linear in the level of overcrowding in the public transit.

The estimated coefficients, presented in Table 4, have the expected signs. Individuals value negatively the cost of the trip as well as its duration. The interactions between the trip duration and individual characteristics reveal that the sensitivity to trip duration is more heterogeneous across age than across income categories. Older individuals are more sensitive to the trip's duration. But the sensitivities to the trip duration are not so different for individuals between 25 and 60 years old. The sensitivity to the trip's duration depends non-linearly on income: individuals in the 4th income quintiles are the most sensitive to the trip duration. However, the coefficient is not significantly different from those of lower-income categories. The only significant income

heterogeneity coefficient is the highest income quantile. It reveals that the top-income individuals are less sensitive to trip durations than lower-income categories.

The transportation mode dummies reveal significant differences in the stand-alone preferences for the different transportation alternatives. We also find that driving a car or taking public transport at peak hour is preferred to using these modes at non-peak hour. The value of the nest parameter is high (0.79), indicating that leaving at peak and non-peak hours are subject to independent shocks and thus constitute imperfect substitutes. Yet, the coefficient is lower than one, implying a small correlation between the preference shocks within transportation modes. The value of the nest parameter also confirms the relevance of the nest to represent the substitution patterns between the different transportation modes and departure times.

The interactions between the non-peak hour dummy and the socio-professional categories reveal that white collars and young students are less flexible than the other categories. The most flexible individuals are the blue collars, students in higher education, and independent workers. We also find that families dislike traveling outside peak hours, indicating lower schedule flexibility. Finally, the public transport controls have the expected signs: the number of layovers and the overcrowding reduce the utility of public transport. The railway dummy is positive but not significant.

Coefficients	Mean	S.E.
Log(duration)	-0.56**	0.09
$Log(duration) \times wealth \in q2$	-0.05	0.08
$Log(duration) \times wealth \in q3$	-0.01	0.08
$Log(duration) \times wealth \in q4$	-0.11	0.08
$Log(duration) \times wealth \in q5$	0.15^{\dagger}	0.09
$Log(duration) \times age \in [18-25]$	-0.4**	0.1
$Log(duration) \times age \in [25-35]$	-1.59^{**}	0.09
$Log(duration) \times age \in [35-45]$	-1.7**	0.08
$Log(duration) \times age \in [45-60]$	-1.45^{**}	0.08
$Log(duration) \times age > 60$	-2.03**	0.2
$Log(duration) \times effort$	-1.66^{**}	0.06
Cost	-0.41**	0.02
Bicycle	-3.48**	0.08
Public transport, peak	-4.2**	0.21
Public transport, non-peak	-4.49**	0.23
Motorized 2-wheel	-7.35^{**}	0.21
Car, peak	-6.22^{**}	0.2
Car, non-peak	-6.81^{**}	0.22
Non-peak hour \times white collar	-0.57**	0.09
Non-peak hour \times blue collar	0.16^{\dagger}	0.08
Non-peak hour \times school \leq high school	-0.98^{**}	0.12
Non-peak hour \times school > high school	0.01	0.1
Non-peak hour \times family	-0.08^{\dagger}	0.04
No. layovers in public transport	-0.35**	0.04
Railway only	0.05	0.06
Public transport overcrowding	-0.06**	0.02
σ	0.79**	0.06
No. observations		12,975
Log-likelihood		-13,624

Table 4: Estimation results for the utility parameters

Notes: Walking is the baseline alternative. The reference categories are individuals with age < 18, the first wealth quintile and independent workers. Duration measured in minutes. Cost in \in .

3.3 Values of travel time and substitution across modes and periods

With these parameters, we compute individuals' values of travel time (or opportunity cost of time) in \in /hr. The value of travel time represents how much an individual should receive (in \in) to compensate for the decrease in utility related to an increase in travel time by an hour. Formally the value of travel time for individual *n* is:

$$\operatorname{VOT}_n = \frac{\beta_n^{\operatorname{duration}}}{\alpha} \times \frac{1}{\operatorname{duration}_n},$$

where α is the sensitivity to the trip monetary cost. Because we specify the utility as a function of the logarithm of the duration, the marginal utility of duration depends on the trip duration. We,

therefore, calculate the values of travel times associated with the chosen alternatives and departure periods.

Table 5 presents detailed information about the distribution of the value of time across individuals. We obtain an average value of time of $\in 15.9$, higher than the median value, which reaches only $\in 10.3$. We observe substantial heterogeneity in how individuals value their time in transport, reflected by the extreme minimum and maximum opportunity costs of time ($\in 0.44$ and $\in 389$). However, the maximal value is an outlier since the 99th percentile of the distribution is much more reasonable ($81 \in /hr$). The mean value of travel time is in line with a recent study by Buchholz et al. (2020) that estimates an average value of time of \$13.5 on a sample of cab riders in Prague. Given an average hourly wage of $\in 23.9$ in 2012 in the Paris metropolitan area⁴, our results are also consistent with the literature such as Parry and Small (2009) who estimates the mean valuation of travel time to be close to half the average hourly wage in London and Washington D.C. We find a value of travel time consistent with Kilani et al. (2012), who estimate an average of $\in 17$ for a working father. However, our estimate is for the whole population and not for a specific individual profile with a higher opportunity cost than average. We perform several robustness checks to our demand specification. Appendix E.1 presents average parameters and the distribution of VOT for each one of the robustness specifications.

We present the distribution of the value of travel time by age and income in Figure 1 using second-order local polynomials. Who dislikes the most to spend time commuting? The very young individuals are associated with large values of the opportunity cost of time, but the value of time decreases rapidly with age, reaching its minimum for 20 years old. After that age, the value of time increases, except for a flat part between 30 and 50 years old. The graph also indicates that the most senior individuals have the highest valuations. One of the reasons for this shape is that the estimated value of time for young children indirectly represents those of their parents, who have both to commute to work and drop off their children at school. The value of travel time rapidly decreases between 0 and 18, as older kids can commute on their own to school. The steep increase afterward reflects the role of professional commitments that raise travel time value. We also see heterogeneity in the value of travel time across income categories, but the heterogeneity is much less pronounced. Poor individuals have the lowest valuations of time on average, but the opportunity cost of time increases rapidly with income. The opportunity cost of time displays an inverted bell shape. From a wealth of €360,000, the opportunity cost increases.

⁴Source: https://www.lemonde.fr/les-decodeurs/article/2016/11/28/en-ile-de-france-le-salaire -horaire-depasse-de-41-celui-des-autres-regions_5039717_4355770.html.

Table 5: Distribution of the values of travel time.

Min	p1%	Mean	Median	p99%	Max	
0.44	1.34	15.9	10.3	81.1	389	
Notes: in \in/hr .						

Figure 1: Value of travel time and individual characteristics.



Note: Income is in $\in 100,000$. We windsorize income values greater than 10.

To better understand the substitution across departure periods for driving, we compute the own and cross elasticities to duration. We also calculate the own and cross elasticities to the car trip cost to put the values in perspective. The elasticities represent the change in the probability to drive at period t when the driving duration or cost at period t' increases by 1%:

$$\mathcal{E}_{n,t,t'}^{duration} = \frac{\partial \log s_{nt}}{\partial \log \operatorname{duration}_{nt'}} \qquad \mathcal{E}_{n,t,t'}^{cost} = \frac{\partial \log s_{nt}}{\partial \log \operatorname{cost}_{nt'}}$$

where t and t' can be peak hour or non-peak hour. Table 6 contains the main statistics of the distributions of the elasticities and we provide a graph with the distributions of the own elasticities in Figure 14 of Appendix C.

Individuals are less elastic to changes in duration or cost at peak hour than at non-peak hour, indicating greater substitution from driving at non-peak hour to other options. The cross elasticities of driving at peak hour to the trip duration at non-peak hour are very low (0.45 on average), indicating very low substitution. The elasticity of driving at non-peak hour to the duration at peak hour is slightly higher, reflecting the preference for peak hour. The direct elasticities to driving

costs are very low, indicating that the monetary trip cost is not the most critical barrier to driving. The elasticities are slightly lower for non-peak hour than peak, reflecting the preference for peak hour.

The direct elasticities are very heterogeneous across individuals. For instance, the own duration elasticity for peak hour is between -3 and -0.06. The heterogeneity partly reflects the differences in trips distance which is highly correlated to duration and cost. It is also the consequence of the heterogeneity in preferences and access to efficient alternatives to cars. The cost elasticities are even more dispersed, as they are between -6.8 and 0. We have elasticities of zero because some individuals are fully reimbursed for the driving costs by their companies, and some use electric vehicles that we assume have zero cost.

Table 6: Driving duration and driving cost elasticities						
	Min	p1%	Mean	Median	p99%	Max
Own elasticities						
$\mathcal{E}_{ ext{peak, peak}}^{ ext{duration}}$	-3	-2.51	-1.28	-1.29	-0.23	-0.06
$\mathcal{E}_{ ext{non-peak, non-peak}}^{ ext{duration}}$	-3.1	-2.66	-1.58	-1.78	-0.49	-0.43
$\mathcal{E}_{ ext{peak,peak}}^{ ext{cost}}$	-6.37	-2.23	-0.4	-0.22	-0.02	0
$\mathcal{E}_{\mathrm{non-peak, non-peak}}^{\mathrm{cost}}$	-6.78	-2.53	-0.48	-0.29	-0.02	0
Cross elasticities						
${\cal E}_{ m peak,\ non-peak}^{ m duration}$	0.06	0.07	0.45	0.4	1.21	1.65
$\mathcal{E}_{\mathrm{non-peak, peak}}^{\mathrm{duration}}$	0.17	0.3	0.87	0.78	1.88	2.09
$\mathcal{E}_{ ext{peak, non-peak}}^{ ext{cost}}$	0	0.0009	0.13	0.08	0.73	2.42
$\mathcal{E}_{ m non-peak,\ peak}^{ m cost}$	0	0.006	0.21	0.13	1.23	3.88

Notes: Values are in %.

4 Estimation of the road traffic congestion technologies

4.1 Data and sample selection

We split the Paris area into five zones: the city center, the ring roads, the close suburb, the far suburb, and the main highways that connect the city center to the suburb. We estimate the congestion technology for three areas: city center, ring roads, and highways. Because of data limitations, we cannot estimate the congestion technology for the suburbs. However, we still allow for adjustment of the equilibrium speeds in the close suburb and far suburbs by assuming the same congestion technology as in the city center.

To estimate road congestion technologies, we rely on hourly data on traffic conditions from remote

sensors for 2016 and 2017. The data come from two different sources. The highway traffic data come from the regional road maintenance agency (DRIF). Traffic data for Paris and for the ring roads come from the city of Paris.⁵ We provide the road coverage of the sensors in Figure 15 in Appendix C.

The traffic sensors typically contain three variables: the car flow (in number of cars per/hr), the traffic density (in units of cars/km), the occupancy rate, the percentage of time during which the sensor detects a car (in %/hr), and the speed (in km/hr). The car flow, density, and speed are related through the fundamental equation of traffic flow:

flow = speed \times traffic density.

The traffic density and occupancy rate are related through:

traffic density = $\frac{\text{occupancy rate}}{\text{mean effective car length}} \times \text{no. lanes}$

For traffic observations from the highways, we restrict the sample to sensors that record traffic going in the direction of the city center. We drop outliers in speed (below 0 or greater than the maximal highway speed limit, 130 km/hr) and occupancy rates (below 0% and above 60%). An occupancy rate of 60% represents extreme traffic conditions: the traffic monitoring institute in Paris defines traffic as pre-saturated from 15% and saturated from 30%. Using the fundamental relationship between traffic flow, occupancy rates, and speed, we detect inconsistent observations. More specifically, we compute the implied average car length using the flow-speed relationship above:

mean effective car length =
$$\frac{\text{occupancy rate} \times \text{speed} \times \text{no. lanes}}{\text{flow}}$$
.

We then drop observations with a car length lower than 3.6 meters (the length of a city car like Renault Twingo) and 18.75 meters (the size of a heavy truck). On the initial sample of 8.9 million observations, we keep 6.196 million of them. These observations come from 654 traffic monitoring stations and constitute an unbalanced panel.

The data on the city center traffic (*Paris intra-muros*) and the ring roads (*periphériques*) contains sensors' measurements of the traffic flows and the occupancy rates. Sensors cannot measure speed accurately because of traffic lights and multiple intersections. We, therefore estimate the speed using the formula provided provided below:

speed =
$$\frac{\text{flow}}{\frac{\text{occupancy rate}}{\text{mean effective car length}} \times \text{no. lanes}}$$
.

⁵Source for Paris data: https://opendata.paris.fr/explore/dataset/comptages-routiers-permanents -historique/information/.

Instead of relying on an assumption for the effective car length as often done in the literature (e.g., Geroliminis and Daganzo (2008) or Loder et al., 2017 for their study of Zurich), we rely on the highway data to predict the average car lengths in Paris. It has been documented by Jia et al. (2001) that the traffic composition varies over time, making the uniform car length assumption inappropriate. We, therefore, rely on a prediction model for the car length that we estimate on the highway data. Then, we use this model to predict hourly car length in the city center and the ring. Our model specifies the mean effective car length as a function of the distance to the city center and day-of-the-week interacted with hour fixed effects. Because the relationship between the car length and the distance to the city center may not be constant across locations, we rely on a piecewise linear specification with six intervals. This prediction model is estimated using 4.9 million observations from highway data, for which we observe the GPS coordinates of the measurement stations and obtain an \mathbb{R}^2 of 0.17. To predict the car length, we set the distance to the city center to 0. We then get expected car lengths that are hour- and day-of-the-week-specific. Our predictions are very realistic since they are all between 5.5 and 6.7 meters with an average of 5.9 meters. We do not directly observe the number of lanes in the city center traffic data, so we rely on additional data from Open Street Map. Finally, we exclude outliers in occupancy rate and estimated speed following the same criteria we used for the highway data. Our final sample contains 1.9 million hourly observations for the ring roads; 118 stations record these observations. We have 599 measurement stations recording 8 million observations for the city center.

Table 7 provides summary statistics for the speeds and occupancy rates in the different areas. We observe significant heterogeneity across areas, suggesting that our partition is relevant. Speeds at peak and non-peak hours are significantly different, supporting our differentiation across periods. In addition to providing evidence of heterogeneity across areas, the traffic speeds and occupancy rates are highly variable; we leverage these variations to estimate flexible road congestion technologies.

		Peak	Non-peak		All samp		
	Area	Mean	Mean	Mean	Median	Std. dev.	
Speed	Highway	44.9	67	86	91	25.6	
	City center	22.4	31.7	38.1	31.8	24.9	
	Ring roads	30.4	57.9	56.3	61.1	21.3	
Occupancy	Highway	23.4	14.1	7.92	5.84	7.67	
	City center	17.2	6.17	6.04	3.75	7.52	
	Ring roads	33.5	15.3	15.4	12.3	11.6	

Table 7: Road traffic conditions by area.

Note: Speed is in km/hr and occupancy rate is in %. Average at peak and non-peak hours are computed over workdays, excluding school and public holidays and weighted by the average traffic flow of the remote sensor. Peak hour is between 8:00 and 8:59 a.m. and non-peak hour is between 6:00 and 6:59 a.m.

4.2 Estimates of the congestion technologies

We use the observations on speed and occupancy rate to estimate the congestion technologies in each area. We approximate f^a by Bernstein polynomials of degree seven and impose the constrain that the functions are weakly decreasing.

The identification of the relationship between traffic density and speed comes from variation in traffic conditions in the data. First, we expect different hours to have heterogeneous traffic levels: traffic should be heavier in the morning and evening. Second, traffic conditions are presumably variable across weekdays, weekends, public holidays, and school holidays. In addition, we observe data for different roads within an area that may also have various traffic conditions. Finally, we have several measurement stations along the streets, generating variations in traffic conditions.

Table 8 represents the number of observations used to estimate the congestion functions for each area and the fit of the models, measured by the R^2 . The three congestion technologies have good fits with R^2 between 0.21 for the city center and 0.69 for the ring roads. The lower R^2 in the city center reflects more idiosyncrasies in traffic speed: traffic lights and intersections generate heterogeneous traffic flows, implying heterogeneous speeds. As Figure 2 shows, the estimated values of the maximal speeds are very much in line with the maximum speed limits in each area. We estimate it to be 99.8 km/hr for highways compared to speed limits that vary between 90 and 130 km/hr depending on the road and location. The speed limit is typically 70 km/hr on the ring roads, very close to our estimate of 71.6 km/hr. Lastly, the speed limit is usually between 30 and 50 km/hr in the city center, but individuals may not always comply with the speed limit. Therefore, our estimate of 53.8 km/hr shows high consistency again.

Figure 2 suggests that the congestion technologies are very different across the city areas. On highways, the maximum speed is higher than in the other zones, but it decreases more rapidly with the traffic density. The technology becomes inferior to ring roads from an occupancy rate of 19.6%, which may be because there are many interchanges and entries and exits in highways, decreasing the speed quickly when the road becomes congested. The rapid decrease in the rate on the highways implies a critical difference in speeds between peak and non-peak hours. While the average speed is 85.4 km/hr at non-peak hour, it is only 65.2 km/hr at peak hour. The congestion curve for ring roads remains flat initially: the slope remains higher than -0.3 until an occupancy rate of 4.2%. In contrast, the speed in the city center displays a convex relation, with speeds decreasing faster for lower occupancy rates than for larger ones.

Area	Number of observations	\mathbb{R}^2
Highway	$6,\!195,\!874$	0.65
City center	$8,\!013,\!979$	0.21
Ring roads	$1,\!907,\!088$	0.69

Table 8: Fit of the congestion technology by area



Figure 2: Estimated congestion technologies and initial traffic conditions. Note: Initial conditions are average speeds at peak and non-peak hour from TomTom predicted durations.

4.3 Robustness checks

We might be worried that our estimates are biased if the same unobserved shocks affect the occupancy rate and speed. We do two sensitivity analyses to check that our estimates are robust to adding more controls to decrease the propensity to have common unobserved shocks. First, we estimate the congestion technologies that are allowed to vary with time and across measurement stations. Second, we exclude observations with extreme weather. We discuss the underlying assumptions behind these robustness checks and their results in Section C.2 in Appendix C. We find that our benchmark estimated congestion technologies are not systematically higher or lower

when we estimate different specifications or use restricted samples, suggesting that we do not have biases from the endogeneity of the occupancy rate.

4.4 Mapping between individual decisions and initial traffic conditions

Our transportation mode choice model predicts the individual probabilities to drive at peak and non-peak hours given car trip durations. Car trip durations for individual n at period t are functions of the speeds in different areas:

duration_{nt} =
$$\sum_{a \in \mathcal{A}_n} \frac{K_n^a}{v_t^a} \times \varepsilon_{nt}$$

 \mathcal{A}_n denotes the areas in which individual n is driving. ε_{nt} represents the individual- and periodspecific speed shock. They capture individual-specific unobserved trip characteristics that make individual average speed lower or higher than the average. We assume these shocks are exogenous to the traffic conditions and hold them constant when introducing traffic-reducing policies. We take the logarithm of the previous equation and estimate the inverse of the initial area-specific speeds at peak and non-peak hours, $\frac{1}{v_{\text{peak}}^a}$ and $\frac{1}{v_{\text{non-peak}}^a}$ and the individual speed shocks from TomTom data. We use the non-linear least squares to estimate:

$$\log(\operatorname{duration}_{nt}) = \log\left(\sum_{a \in \mathcal{A}_n} \frac{K_n^a}{v_t^a}\right) + \tilde{\varepsilon}_{nt}$$

The estimated speeds are provided in Table 9, each period is estimated separately. Speeds are always higher at peak hours than non-peak hours. They are also much higher on the highway, then on the ring roads and far suburb and finally in the city center and the close suburb. These estimated speeds also reveal that the difference in speeds between peak and non-peak hours is very similar for the city center and the close suburb, which is consistent with our assumption that these two areas share the same congestion technology.

	Peak h	nour	Non-pea	k hour
Area	Average speed	Std. Errors	Average speed	Std. Errors
Highway	64.99	0.85	85.83	0.98
City center	13.79	0.14	18.25	0.16
Ring roads	27.67	0.46	46.52	0.72
Close suburb	15.89	0.1	20.15	0.11
Far suburb	25.47	0.13	29.21	0.13

Table 9: Estimated average speeds from TomTom data

Notes: Speeds are in km/hr. Standard-errors are computed using the delta-method.

We then calibrate the scale parameter and the irreducible traffic by matching the predicted number of kilometers driven at each period and the initial average occupancy rates. We do not directly observe the initial occupancy rates, but we obtain them from the initial speeds estimates above and by inverting the congestion technology:

$$\tau_t^a = \left(f^a\right)^{-1} \left(v_t^a\right).$$

 τ_t^a is unique since f^a is decreasing in speed.

We can only identify two parameters per area because we observe individual choices and speeds for only two periods. We therefore impose that $\phi_t^a = \phi^a$ and $\gamma_t^a = \gamma^a$. Instead, we could impose some restrictions across areas and let the parameters vary with the period. But given that the areas have very different sizes and may be subject to different levels of irreducible traffic, it does not seem to be a more realistic assumption.

Using these initial occupancy rates and the predicted number of individuals in each area given the initial speeds, we solve for ϕ^a and γ^a such that:

$$\tau_t^a = \phi^a \times K_t^a(\mathbf{v}) + \gamma^a$$

We have two linear equations and two unknowns to find a unique pair of parameters for each area. We further impose that the shares of irreducible traffic are between 0 and 60% and use the constrained least-squares method that minimizes the sum of square deviations from the equations above. The calibrated parameters are presented in Table 10. We obtain a rather large irreducible share of traffic in the close suburbs (21.6%) and the city center (13.3%) while we estimate no irreducible traffic for highways or the ring roads.

Area	Scale parameter	Irreducible traffic	% irreducible traffic
	(ϕ^a)	(γ^a)	$(\gamma^a/\tau^a_{\rm peak} \times 100)$
Highway	0.024	0	0
City center	0.373	3.18	13.3
Ring roads	0.372	0	0
Close suburb	0.115	4.37	21.6
Far suburb	0.011	6	57.3

Table 10: Calibrated parameters of the mapping between occupancy rates and number of cars.

Notes: The share of irreducible traffic is in % of peak hour occupancy rates.

We evaluate the model's fit by comparing the aggregate frequency of each transportation mode with the predicted using our model. The observed and predicted shares are provided in Table 11 below: the difference comes only from the windsorizing of a few outliers in terms of individual speed shocks. The trimming has virtually no impact on the predicted shares. In terms of speed, Table 12 shows the comparison between the average speeds from traffic data and those predicted from the model equilibrium. Our estimates are optimistic for the highways at peak hour but pessimistic for ring roads and the city center while correctly predicting the speeds ranking between areas for both periods.

Table 11: Shares of transportation modes observed and predicted by the model

	Observed	Predicted	Predicted
		TomTom durations	full model
Bicycle	2.1	2.1	2.09
Pub. transport, peak	30.32	30.28	30.28
Pub. transport, non-peak	14.52	14.56	14.57
2 wheels	2.08	2.08	2.08
Walking	15.8	15.8	15.8
Car, peak	22.88	22.85	22.91
Car, non-peak	12.3	12.33	12.27

Notes: Shares are in %.

	Peak hour		No	Non-peak hour		
Area	Traffic	Eq. speeds	Traf	fic Eq. speeds		
Highway	44.9	65.2	67	85.4		
City center	22.4	13.7	31.'	7 18.3		
Ring roads	30.4	28.8	57.9	9 44.2		
Close suburb		15.8		20.2		
Far suburb		25.5		29.2		

Table 12: Comparison between the estimated equilibrium speeds and average speeds from traffic data

Notes: Speeds are in km/hr.

4.5 Check of the equilibrium uniqueness

Applying the method developed in Section 2.3, we check that given the estimated model parameters, the model is such that it has a unique equilibrium. We compute the Lipschitz coefficients for values of the algorithm tuning parameter κ between 0 and 0.99. We solve for these coefficients at the initial speeds and at the equilibrium speeds under the different benchmark policy environments that we consider in section 6. Basically, we compute:

$$\max_{a \in A} \max_{t \in T} \max_{a' \in A} \max_{t' \in T} \left| \frac{\partial g_t^a(\mathbf{v}^*, \kappa)}{\partial v_{t'}^{a'}} \right|,$$

where \mathbf{v}^* denotes the vector of equilibrium speeds. Figure 3 shows several interesting insights. First, we do not have a contraction when κ is too low. The algorithm starts to be a contraction from $\kappa = 0.53$ without any policy in place. We find that the policies do not disrupt the convergence of the algorithm. We observe that the no-policy environment is the one that requires the highest κ to converge.



Figure 3: Lipschitz coefficients at the equilibrium speed.

We now check for which values of the tuning parameters our algorithm is a contraction in the entire set of possible speed values in their interval $[\underline{\mathbf{v}}, \overline{\mathbf{v}}]$. This time we calculate:

$$\max_{a \in 1, \dots, A} \max_{t \in 1, \dots, T} \max_{a' \in 1, \dots, A} \max_{t' \in 1, \dots, T} \max_{\mathbf{v} \in [\mathbf{v}, \overline{\mathbf{v}}]} \left| \frac{\partial g_t^a(\mathbf{v}, \kappa)}{\partial v_{t'}^{a'}} \right|$$

Figure 4 shows that there exist values of κ lower than one such that the function is a contraction on the entire space of the speeds. As before, we find that the no-policy environment is associated with the highest κ . The lowest value of κ that ensures contraction under the different policy environments and in the entire space of possible speed values is thus 0.65. We use this value to solve the equilibrium speeds in the counterfactual simulations.



Figure 4: Lipschitz coefficients in the space of possible speeds.

(a) I eak nour poncies. (b) An day p

We show in Appendix B.3 that the number of iterations and the computation time increase exponentially with the tuning parameter κ . But the value we choose, $\kappa = 0.65$, multiplies the convergence time by less than 2.9 relative to $\kappa = 0.5$, which is associated with the lowest convergence time.

5 Results from the equilibrium model

5.1 Value of driving and costs of congestion

First, we measure the value of driving at peak hour at the initial equilibrium speeds. We find that if the individuals were not able to take their cars at peak hour, it would generate a welfare loss of 2.56 million euros, which corresponds to 79 cents per potential driver. Then, we measure the value of driving at peak hours under maximal speeds (we use the maximal speeds given the irreducible portions of traffic for each area). We find that it corresponds to a welfare increase of 6.12 million euros, corresponding to 1.90 euros per potential driver. The value of driving at peak hour at maximal speeds is 2.4 times the value of driving at initial speeds, highlighting the importance of congestion levels when looking at individuals' welfare. The congestion generates a total increase in travel time of 396 thousands hours. This corresponds to an average increase in travel time of 7.4 minutes per driver; A considerable variation given that the average trip duration is 31.3 minutes.

We compute the marginal costs of congestion for each area at peak and non-peak hours, defined as the total surplus losses associated with having one additional kilometer or driver in the area. We also compute the marginal cost of an average car, by adding the average number of kilometers driven by an individual in that area. Both measurements change the driving speed marginally, which increases trip durations for everyone. We calculate the total marginal changes in individual trip durations into monetary terms, and then we sum the individual surplus changes.

We obtain area specific marginal costs for an additional kilometer between 2 and 53 cents as Table 13 shows. The marginal cost for one additional kilometer driven is the lowest on the highways and in the far suburb, reflecting that these areas are wide and one additional kilometer driven does not change much the traffic density. When we consider the marginal costs associated with an average driver, the costs are higher as they are between 34 cents and $\in 2.1$. The initial marginal costs are much higher at peak hours than at non-peak hours. The marginal costs depend on two key model parameters: first, the congestion function slopes are higher at peak hour than non-peak hour, and second, the trip durations are higher at peak hour than non-peak hour, thus the individual marginal cost of congestion per kilometer is 4.5 times larger at peak hour than at non-peak hour. Finally, we also calculate an average marginal cost of congestion by considering adding an individual driving with the average itinerary (i.e., distance traveled in each area) of those individuals owning a car. The average costs are $\in 2.97$ at peak hour and $\in 1.87$ at non-peak hour.

	One kilometer		O	One average driver		
	(in cents)			(in €)		
Area	Peak	Non-peak	Pe	ak	Non-peak	
Highways	7.6	1.7	1.1	14	0.42	
City center	53	42	1.5	58	1.26	
Ring roads	45	25	2.0)9	1.19	
Close suburb	45	36	1.2	29	1.04	
Far suburb	8.2	4.3	0.5	58	0.34	
Average			2.9	97	1.87	

Table 13: Estimated initial marginal costs of congestion by area.

Notes: Average marginal cost for an average driver in the whole metropolitan area.

Comparing our results with the literature is not a straightforward task. In most studies the cost of the congestion externality is considered as the difference between the socially optimum cost of congestion and the private cost of congestion (see Small et al., 2007). This relation is in turn estimated using aggregate supply and demand curves and thus do not consider the marginal cost of an extra individual driving but the cost of an additional kilometer driven.

We think that the most appropriate way to compare our results to the ones in the literature is to use our marginal cost of congestion per individual. Doing such a comparison puts our congestion values close to those from Akbar and Duranton (2017) and Yang et al. (2020).

5.2 Importance of endogenous speeds

To study the importance of accounting for speed adjustments to predict policy effects accurately, we compare the outcomes predicted from our model against those obtained under constant speeds. The comparison illustrates the importance of considering the policies' effects on congestion levels and driving rates, as they will influence substitution patterns between modes and departure periods. We thus analyze the predictions for our three principal policies: (1) driving restrictions that ban a fraction of the car from driving at peak hour, (2) a uniform toll at peak hour, and (3) a variable toll at peak hour that is linear in the total trip distance. We consider a range of policy stringency levels, such that the strongest stringency level achieves the same traffic at peak hour under all three policies. The traffic density is represented by the total number of kilometers driven.

Figure 5 shows the predictions of the fraction of individuals who drive under each policy. All scenarios and stringency levels show the same direction of the biases due to the usage of exogenous speeds: an overestimation of the number of individuals substituting away from using their cars at peak hour and an underestimation of those who choose to drive at non-peak hour. The equilibrium speed effects dampen the incentives to stop driving because the speeds at peak hour improve due to the regulations, and at the same time, the speeds at non-peak hour decrease.

Figure 5: Predicted car shares as function of the policy stringency level - peak-hour policies





Figure 6: Speeds under the different policies and stringency levels.

We compare the average speed in each area at peak and non-peak hours, associated to the different policies, for all policy stringency levels. As Figure 6 above shows, the speeds at peak hours increase monotonically with the policy stringency level for the three types of policies. In contrast, speeds at non-peak hours decrease monotonically, which is the consequence of important shifts towards driving at non-peak hours. The speeds change the most on the highways while in the city center there are much less speed improvement at peak hour. We can also observe that at non-peak hours the speed in the city center is almost constant. This reflects the fact that individuals driving in the city center have better alternatives to cars while those who drive on the highways and the ring roads are more likely to substitute towards driving at non-peak hours.

To further highlight the implications of endogenous speeds on policy analysis, we compare the tax revenues predicted under constant speeds with the predictions from our model. The results, in Figure 7 below, suggest that not accounting for the changes in speeds underestimates significantly the number of individuals paying the toll and the tax revenue. Moreover, the magnitude of the bias increases with the policy stringency levels, reflecting the prominent role of speed adjustments.


Figure 7: Predicted tax revenues under tolls at peak hour.

6 Quantifying the welfare costs of regulations

Using the transportation choice model and estimated road technologies, we study the welfare costs of driving restrictions and road tolls which constitute standard transportation policies. We focus on policies applicable at peak hour only so that driving at non-peak hour is never constrained. There are two reasons why we choose to focus on peak-hour policy. As shown earlier, the congestion problem is most severe at peak hour. The high concentration of cars is also associated with high pollution levels and deteriorated air quality. These two factors make the regulation at peak hour the most important. The second reason we choose to focus on peak hour policy is technical: in our sample, we observe individuals without alternatives to driving their cars. Under a restriction applicable all day, these individuals would not comply with the regulations, so we would need to assume the amount of the fine, the expected probability of being caught, and perhaps a disutility to breaking the law. We make such assumptions and provide an analysis of all-day policies in Appendix D.

Our model predicts individual probabilities of choosing each available transportation mode and the driving speeds at peak and non-peak hours in the five different areas of the region. Thus, we do not predict counterfactual choices. Instead, we use the predicted probabilities to compute the expected number of individuals choosing each transportation mode and departure period. Our measurement of welfare costs corresponds to one trip per person for the entire population of commuters in the Paris area. There are, on average, 253 working days annually and two commuting trips per day, so

we should roughly multiply the costs by 500 to convert them into annual welfare costs.

6.1 Analysis of different policy levels

3.5

2.

0.5

20 40 60 80 Reduction in distance driven at peak hours (in %)

(a) Surplus losses.

Total surplus loss

To properly compare the three policies across stringency levels, we present the stringency level of each policy as the reduction, in percentage, of kilometers driven at peak hour. First, we focus on the aggregate welfare implications of the different policies. Figure 8 presents the associated surplus losses, tax revenues, and welfare losses associated to the different policy stringency levels. Since the driving restriction affects all individuals in the same way (same reduction in the probability of driving), it is not surprising to see an almost linear relation between its stringency level and its surplus loss. Since the uniform toll charges everyone a same amount, independently of their trip duration, we find a concave relation between its stringency level and surplus losses. This relation comes in contrast to the variable toll, which by targeting high distance drivers can achieve higher traffic reductions with considerably lower surplus losses. Panel (b) shows how the uniform toll is much better at generating revenue than the variable toll for almost all stringency levels. However, redistributing this larger tax revenue is not enough for the uniform toll to have a lower welfare cost than the variable, which remains the most efficient policy at all stringency levels. Interestingly, as panel (c) shows, after tax revenue redistribution pricing mechanisms have a positive welfare impact for low stringency levels; stringency levels lower than 35% for the uniform toll and up to 50% for the variable toll.



60 80 ven at peak hours (in %) 20 40 60 80 Reduction in distance driven at peak hours (in %)

(c) Welfare losses (surplus losses

minus tax revenues).

100

20 40 Reduction in distance dri

Figure 8: Change in individual surplus, tax revenue and welfare

We now focus on the distributional implications of the policies. Figure 9 presents the surplus losses using distributional weights, The range of losses (difference between maximum and minimum surplus losses), and the difference in surplus variation between the top and bottom 10% of the income distribution. Using weights based on the individuals' income level does not same to considerably

(b) Tax revenues.

affect the ranking or size of the aggregate surplus losses, as seen in panel (a). However, panel (b) highlights the main disadvantage of the variable toll, its distributional impact. While the uniform toll and the driving restriction generate almost the same dispersion, the variable toll dispersion of surplus losses is twice as large. This result highlights how policymakers' objectives might affect the optimal policy, as the variable toll has an inversed ranking depending on the criteria used.



Figure 9: Change in distributional outcomes.

We compute an average cost of regulation associated with reducing emissions of local pollutants, in \in /ton of emissions avoided. Since we have multiple pollutants, we construct an equivalent NO_X emissions measure that includes emissions of NO_X, PM, and HC. Each pollutant is weighted by its society cost relative to the cost of NO_X . The values we use for the society costs come from a report of the DG MOVE (2014).⁶ We compute the cost of regulations by dividing the welfare cost (including the redistribution of the full tax revenue) by the change in the total emissions. Figure 10 presents the costs associated to the different policy stringency levels. First, we can see that the three policies have different cost of regulation functions. Under a uniform toll, the function is almost linear, while under a variable toll, the cost function is concave. In sharp contrast, the cost function decreases under driving restrictions, reflecting the fact that the welfare losses increase at a slower pace than the decrease in emissions. The function is almost linearly decreasing, but the costs are very high for all stringency levels. On the other hand, the average costs of reducing emissions have more variation under the two types of tolls. Price instruments can have negative costs for low stringency levels, implying a net benefit. These are policies with a uniform toll below $\in 2.7$ or a variable toll of $\in 0.15$ /km. We can find a more extensive set of policies with a net positive impact on individuals under the uniform toll because since, as seen in Figure 7 above, the uniform toll can raise a larger tax revenue than the variable toll. Another interesting specificity of driving

⁶The values are: 13,052 €/ton for NO_X, 1,695 €/ton for HC and 211,795 €/ton for PM.

restrictions is that the average cost of regulation is always lower for peak-hour policies than for all-day policies. In contrast, for both price instruments, we find that the ranking is reversed. For most of the policy stringency levels, the tolls applicable all day are associated with lower regulation costs.





(a) Change in Equivalent NO_X (b) Change in CO_2 emissions. (c) Cost of the regulation. emissions.

6.2 Comparison across policies at a benchmark level

To analyze the effects of the previous policy instruments in more detail, we set a benchmark policy stringency level. We calibrate the stringency level of the three policies to achieve the same traffic reduction at peak hour, represented by the total distance driven. We start by implementing a driving restriction that deprives individuals of driving with a probability of one half. This represents a ban based on whether the last digit of the license plate is odd or even. Such policy triggers a reduction of the number of kilometers driven by 34% at peak hour. We then calibrate the uniform and variable tolls to reach the same traffic reduction in terms of kilometers driven. An alternative benchmark would be to match the total distance driven at peak and non-peak hours across policies. We obtain slightly lower calibrated tolls to achieve such matching. The policy parameters are summarized in Table 14 below.

The uniform tolls would imply a total cost of $\in 5.4$ per day, close to the London congestion charge before 2007. The London toll implemented in 2003 was initially £5/day and increased to £8 in 2005, £10 in 2011, £11.5 in 2011, and £15 in 2020. The variable toll is $\in 0.10$ /km, implying an average cost of $\in 1.43$, lower than the uniform toll. This is the consequence of the variable toll being able to target individuals with long trips and thus achieving quickly a reduction in the total distance driven at non-peak hour. The maximum toll cost reaches $\in 14.1$ for the longest distance.

Outcome matched	Driving restriction	Uniform toll	Variable toll
Distance, peak hours	50%	2.69	0.1
Distance, peak and non-peak hours	50%	2.08	0.09

Table 14: Policy parameters for the benchmark stringency level

Note: Uniform tolls in \in and variable tolls in \in /km.

Modal shift

The variation in aggregate shares of transportation modes reflects the modal shifts of individuals. The results, reported in Table 15, indicate an important inter-temporal modal shift towards driving at non-peak hours. The share of individuals driving at peak hours drops under all policies but the magnitudes are different. While driving restrictions and uniform tolls decrease the number of car users by 4.6 to 5.8 percentage points, the variable toll only decreases it by 1.93 percentage point. This is because the variable toll discourages individuals with long distances to drive and keep the number of drivers relatively high. For the same reason, we also observe important differences for the modal shifts across policies: driving restrictions and the uniform toll increase the fraction of individuals who walk and take the public transport more than the variable toll.

Additionally, there are relatively large increases for walking and public transport, which are the main transportation alternatives chosen. Surprisingly, the modal shift towards public transport is larger under a variable toll than under the driving restriction or the uniform toll. These differences between the two tolls can be explained by the uniform toll discouraging short trips, which can be easily done by walking or using a bicycle. In contrast, the variable toll prevents individuals with long trip distances, which cannot be done by foot or using a bicycle. Which explains why the modal shifts are more oriented towards public transport and driving at non-peak hours than towards the other modes.

	Initial	Driving	Uniform	Variable
		restriction	toll	\mathbf{toll}
Bicycle	2.09	2.3	2.34	2.09
Pub. transport, peak	30.3	32.2	32.7	31.4
Motorbike	2.08	2.44	2.49	2.31
Walking	15.8	17.1	17.4	15.8
Car, peak	22.9	14.6	14	20
Car, non-peak	12.3	15.9	15.4	13.3
Pub. transport, non-peak	14.6	15.4	15.6	15.1
Total car share	35.2	30.6	29.4	33.2
Total pub. transport share	44.8	47.5	48.3	46.6

Table 15: Predicted shares of the transportation modes under the different policies.

Note: Shares are expressed in %.

Welfare impacts

We now turn to the policies' impacts on consumer surplus, which results are reported in Table 16. The uniform toll generates the largest total consumer surplus losses among the three policies but it raises the largest tax revenue, $\in 1.52$ million. The second most costly policy in terms of individual surplus is the driving restriction which costs $\in 1.11$ million and generates no tax revenue. Finally, the variable toll is the least costly for individual surplus and still generate substantial tax revenue ($\in 0.87$ million). If we abstract from the cost of public funds and redistribute the tax revenue entirely, the variable toll outperforms the uniform toll by a small amount ($\in 0.16$ million). The driving restrictions are much more costly than the tolls. The uniform toll generates tax revenue that covers 102% of the surplus losses while the variable revenue covers 128% of the total welfare costs for individuals. As long as the shadow cost of public funds is below 75%, the uniform toll outperforms the driving restriction policy, which is very likely to be the case in reality. The variable toll is preferred to the driving restriction, even when there is no recycling of the tax revenue.

To better understand the role of the changes in traffic levels induced by the different policies, we decompose the total variation in consumer surplus into two terms: (1) the variation in consumer surplus due to the policies keeping the speeds constant and (2) the changed in surplus induced by the change in equilibrium speeds. The first outcome measures the cost of the policy independent of the changes in driving speeds. The second outcomes measures the gains from speed improvement at peak hour and losses from reduced speeds at non-peak hours. As expected, the reduction in congestion levels always creates a welfare gain. These gains are enough to cover only between 13% and 34% of the consumer surplus losses linked to the policies, indicating that congestion improvements are not enough to offset the cost of traffic policies. The largest gain from speed

occurs under the variable toll since we provide incentives to stop driving to individuals with long distances that have a high impact on traffic.

When looking at the share of individuals with surplus gains, we find that no one is better-off under a uniform toll. However, a very small fraction of individuals is better off under driving restrictions (around 3,000 individuals out of a total of 4.03 million). In contrast, a significant fraction of individuals improve their welfare under a variable toll (18.9% of the population). These are individuals with short trips and high valuations of time who continue to drive at peak hours despite the toll. They are better-off because their gains from higher speeds compensate for the costs of paying the tolls. Among the better-off individuals, the average toll is ≤ 0.29 versus an average toll of ≤ 1.44 for individuals holding a car.

Across policies, there is a constant share of individuals with no change in their surplus. These are individuals without a car, which are not affected by the regulations. This result does not hold if the quality of public transport is downgraded from the modal shifts. In appendix E.2, we show the robustness of the results to an increase in the overcrowding in the public transport, finding very small changes in substitution patterns even after a 15% or 30% increase in overcrowding.

To assess the differences between policies, a key dimension is the distribution of the welfare losses across individuals. The uniform toll is the regulation associated with the smallest interval between the maximum and minimum welfare losses. While the driving restriction has a fairly similar range, the median welfare loss indicates fewer consumers have larger losses under driving restrictions. Meanwhile, the variable toll has the highest dispersion between the maximal and minimal welfare losses. This is consistent with the presence of winners and suggests a larger heterogeneity of the policy's impact across the population. In terms of average welfare cost, we note that the variable toll is associated with the lowest loss, at 17 cents, while the two other policies decrease, on average, the individual welfare by 28 to 37 cents.

	Driving	Uniform	Variable
	$\operatorname{restriction}$	toll	toll
Total $\Delta CS (M \in)$	-1.11	-1.49	-0.69
ΔCS , constant speed	-1.28	-1.69	-0.93
ΔCS from speed	0.164	0.199	0.24
Total $\Delta wCS (M \in)$	-1.16	-1.55	-0.752
Tax revenue	0	1.52	0.871
$\Delta W = \Delta CS + tax revenue$	-1.11	0.027	0.181
$\% \Delta CS = 0$	20.3	20.3	20.3
$\% \Delta CS > 0$	0.075	0	18.9
$\% \Delta CS < 0$	79.7	79.7	60.8
$Min \ \Delta CS$	-2.09	-2.04	-4.07
Mean ΔCS	-0.276	-0.37	-0.171
Median ΔCS	-0.163	-0.255	-0.031
$Max \ \Delta CS$	0.037	0	0.765

Table 16: Consumer surplus variation under the three main policies.

Note: ΔCS are in \in /trip.

Impacts on individual trip durations

We complement the welfare analysis by looking at how the policies impact the expected travel times. The results are presented in Table 17 below. Under the three policies, some individuals reduce their expected trip durations. The variable toll has the largest share of individuals with reduced travel times, with 56% of the individuals versus 28% and 29% under the two other policies. However, for many individuals the gain in terms of travel time does not compensate for the cost of the toll, as we obtain larger shares of individuals with surplus losses.

The average increase in expected travel time is between 0.03 and 0.42 minutes implying an increase in the total travel time between 2,320 and 24,900 hours. The variable toll offers slightly lower speed improvements than the two other policies. We analyze in more details the change in the different speeds below. Lastly, we can observe that the distribution of changes in travel time under all policies is skewed towards larger trips. The maximum time reductions are always lower than the maximum increases in travel time. The skewness is however less pronounced under the driving restriction and uniform toll than under the variable toll.

The value of the predicted speeds in each area and for each policy are provided in Table 18 below. We see that the policies have different impacts on equilibrium speeds depending on the area. On the highways, the speed improves the most at peak hour under the variable toll and is the most reduced at non-peak hours. This is because individuals that drive on the highways have typically long distance commutes and thus are are discouraged from driving at peak hour. But since they do not have good transportation alternatives, they end up driving at non-peak hour. Which in turn decreases the speed at that period.

The uniform toll is the policy that improves the most the speeds at peak hour in the city center, the ring roads and the close suburb. In contrast, in the far suburb, the speed is the highest under the driving restrictions. This indicates that many of the individuals who drive through the far suburbs prefer to drive under the uniform toll because they do not have efficient alternatives. Indeed, the public transport does not cover very well the far suburb.

The speeds at non-peak hours decrease in all areas but they do not reach the initial levels at peak hours. This reflects the imperfect substitution between driving at peak and non-peak hours.

	Driving	Uniform	Variable
	restriction	toll	toll
$\% \Delta duration = 0$	20.3	20.3	20.3
$\% \Delta duration > 0$	50.4	51.9	23.7
$\% \Delta duration < 0$	29.3	27.9	56
Min Δ duration	-10.4	-13	-10.9
Mean Δ duration	0.371	0.423	0.034
Median Δ duration	0.011	0.052	-0.044
Max Δ duration	24.4	26.9	40.5
Total Δ duration (in 1,000 hrs)	24.9	28.5	2.32
Average speed, peak (km/hr)	34.1	34.8	33.1
Average speed non-peak (km/hr)	32.4	32.9	32.9

Table 17: Trip duration variation under alternative policies.

Note: Δ durations are in minutes, except "Total Δ duration".

	Area	Initial	Driving	Uniform	Variable
			restriction	toll	toll
Speed	Highways, peak	65.2	82.7	82.3	87
	City center, peak	13.7	17.2	18.6	15.8
	Ring roads, peak	28.8	41.7	45.7	43.7
	Close suburb, peak	15.8	19.6	20.4	17.9
	Far suburb, peak	25.5	28.4	28.2	27.6
	Highways, non-peak	85.4	75.9	78.4	74.1
	City center, non-peak	18.3	17.6	17.6	18
	Ring roads, non-peak	44.2	41.6	42	41.8
	Close suburb, non-peak	20.2	18.8	18.9	19.6
	Far suburb, non-peak	29.2	27.1	27.5	27.7
Occupancy	Highway, peak	15.5	10.3	10.4	8.78
	City center, peak	23.9	18.3	16.5	20.3
	Ring roads, peak	30.5	23.8	21.9	22.8
	Close suburb, peak	20.3	15.4	14.6	17.4
	Far suburb, peak	10.5	8.67	8.79	9.13
	Highways, non-peak	9.38	12.4	11.7	13
	City center, non-peak	16.8	17.8	17.7	17.3
	Ring roads, non-peak	22.6	23.9	23.7	23.8
	Close suburb, non-peak	14.8	16.3	16.1	15.4
	Far suburb, non-peak	8.23	9.43	9.2	9.05

Table 18: Predicted speeds under the different po

Note: Speeds are expressed in km/hr and occupancy rates in %.

Winners and losers

To complement the previous analysis, we investigate the distribution of the welfare costs of the different policies across individuals. Table 19 presents the average change by age category, income category (using our proxy of wealth based on the expected value of the housing), socio-professional activity, and family size. We provide here the average change in individual surplus without any potential redistribution of the tax revenue

First, we can see that the ranking of the policies is identical for all subgroups of individuals. Individuals prefer the variable toll and then driving restrictions. Note that this ranking could be reverted if the tax revenue is redistributed to individuals. The category of individuals below 18 and, consequently, the individuals in education below high school is the one with the lowest cost of the variable toll. This is explained by the short distance of their trips; which implies larger gains from the improved speed at peak hours, partially compensating the toll cost. In our transportation mode choice model, we specify the utility as a function of the logarithm of the trip duration, implying that individuals with short trips have higher marginal utilities of an additional minute in their commute. Under the driving restriction, the category of individuals between 18 and 25 is the one with the lowest average cost of the policy, reflecting the low car availability of this subgroup of individuals.

The most affected individuals are those between 35 and 60 years old, those in the lower two quintiles of the income distribution, and those with a family. These profiles are consistent across all policies. Among the employed individuals, the white collars are the most affected by the policies, while the blue collars are the least affected. Our model suggests that the blue collars have the highest utility of driving outside peak hours, indicating higher flexibility in their schedules.

	Driving	Uniform	Variable
	restriction	toll	toll
Age < 18	-0.228	-0.303	-0.057
Age \in [18- 25]	-0.195	-0.269	-0.148
Age $\in [25-35[$	-0.272	-0.366	-0.184
Age \in [35- 45]	-0.322	-0.43	-0.237
Age \in [45- 60]	-0.327	-0.437	-0.238
$Age \ge 60$	-0.306	-0.403	-0.178
Wealth $\leq 110,000$	-0.303	-0.4	-0.207
Wealth \in]110,000-152,000]	-0.32	-0.423	-0.213
Wealth \in]152,000-205,000]	-0.299	-0.397	-0.188
Wealth \in]205,000-283,000]	-0.252	-0.341	-0.141
Wealth $\geq 283,000$	-0.208	-0.288	-0.106
Independent	-0.298	-0.419	-0.241
White collar	-0.331	-0.427	-0.223
Blue collar	-0.283	-0.394	-0.213
Education \leq high school	-0.233	-0.309	-0.059
Education $>$ high school	-0.128	-0.185	-0.113
Family	-0.29	-0.388	-0.18
Single	-0.205	-0.276	-0.126
Average	-0.276	-0.370	-0.171
			-

Table 19: Average surplus variation by demographic group.

Note: in \in /trip.

Impact on emissions

Next, we focus on the environmental impacts of the policies by studying the reduction in the emissions of different pollutants. Appendix A.6 outlines how we estimate emissions as a function of the car characteristics in the sample. The results displayed in Table 20 below show the differences between policy instruments' efficacy at reducing emissions.

The three policies are calibrated to have the same effect on the total number of kilometers driven at peak hour. Yet, we see that they do not have the exact same effects on emissions. The main reason is that the substitution towards driving at non-peak hour is different. And as seen above, the variable toll is the regulation that generates the most substitution towards driving at non-peak hours. A second factor contribute to obtaining different impacts on emissions: the different policies discourage different individuals from driving which are associated with different car characteristics.

As expected, the uniform toll is the most efficient policy to reduce emissions of all pollutants. This is because the policy discourages shorter trips for which the gains in travel time do not compensate the high uniform price of driving. Individuals with short trips are also usually individuals with good alternatives to cars. The driving restriction is, on the other hand the least efficient at reducing emissions.

As a robustness check, we also estimate car emissions that are non-linear functions of speed using Copert emission factors. The emission factors are pollutant-specific and vary with the fuel type, Euro norm, and car segment. We provide the results using these alternative estimates of car emissions in Table 32 in Appendix C. We obtain slightly lower emission decreases due to the policies, but the heterogeneity patterns across pollutants and policies remain. We also decompose the effect of the policies into the part due to speed changes and the part due to the modal shifts. The shares of emission reduction due to the speed improvements are shown in Table 33. Overall, the increase in speeds is responsible for a small share of the total decrease in emissions of most pollutants. Nonetheless, under some scenarios it represents around 20% of the change for HC and 10% for NO_X.

	Driving	Uniform	Variable
	restriction	toll	toll
ΔCO_2	-300	-383	-330
ΔCO	-0.638	-0.821	-0.688
ΔNO_X	-0.379	-0.481	-0.432
ΔHC	-0.104	-0.134	-0.112
ΔPM	-0.052	-0.068	-0.057
$\Delta Eq. NO_X$	-1.24	-1.59	-1.37
$\%\Delta CO_2$	-9.61	-12.3	-10.6
$\%\Delta CO$	-9.8	-12.6	-10.6
$\%\Delta NO_X$	-9.22	-11.7	-10.5
$\%\Delta HC$	-9.81	-12.7	-10.6
$\%\Delta PM$	-9.68	-12.5	-10.6
$\%\Delta Eq. NO_X$	-9.53	-12.2	-10.5

Table 20: Changes in total emissions under the different policies.

Note: Δ emissions are in tons.

Average costs of regulations

Lastly, we combine our estimates of the emission reductions and the welfare losses to compute an average cost of avoiding one ton of polluting emissions. More specifically, we calculate costs of reducing carbon emissions and emissions of local pollutants separately. For the latest, we rely on our measure of equivalent $-NO_X$ emissions presented above. Since the traffic regulations are typically used to prevent pollution peaks caused by emissions of local pollutants, we believe that it is relevant to measure the costs of the regulations per ton of equivalent NO_X emissions avoided. We provide the costs when the objective is reducing CO_2 emissions as a comparison benchmark.

The results, presented in Table 21, show that the costs of regulations are much lower for a ton of CO_2 emissions than for a ton of equivalent NO_X emissions; they are around 240 times lower. The report from DG MOVE (2014) suggests that the value of a ton of carbon is 326 times lower than the one for a ton of NO_X , suggesting our emissions costs are closer to their objective value for NO_X than for CO_2 emissions.

Ignoring tax revenue redistribution, we find that the more cost-efficient policy is the variable toll. Between the three main policies, the most costly is the uniform toll. This can be explained by the variable toll discouraging longer car trips, implying greater emissions reductions. However, when redistributing the tax revenue, the driving restriction becomes the costliest policy, and the tolls have negative costs. The most cost-efficient policy is the uniform toll under full redistribution of the tax revenue.

		Driving	Uniform	Variable
		restriction	toll	toll
w/o redistribution	ΔCO_2	3,720	3,893	2,094
	$\Delta EqNO_X$	897,728	$935,\!863$	$502,\!519$
	$\Delta EqNO_X$, Copert	$605,\!655$	$654,\!286$	$359,\!188$
with redistribution	ΔCO_2	3,720	-72	-550
	$\Delta EqNO_X$	897,728	-17,192	$-131,\!897$
	$\Delta EqNO_X$, Copert	$605,\!655$	-12,019	-94,277

Table 21: Costs of regulation for the different policies, with and without redistribution of the tax revenue.

Note: in \in /ton.

Marginal costs of congestion

Table 22 presents the marginal costs of congestion for each area and period under the different policies. As before, we compute the marginal costs associated to adding one average driver in each area. The traffic regulations reduce the marginal costs of congestion across all areas at peak hours under driving restriction and the uniform toll. However, we observe an increase in the marginal costs of congestion in the city center and the close suburb at peak hour under the variable toll. The raise in the costs are lower than 8% and we can observe the opposite pattern at non-peak hours where the costs decrease in the same two areas and the far suburb. This is because the individuals with short trips have more incentives to drive under the variable toll and individuals with short trips usually drive in the city center or the close suburb. Interestingly, the increases in the marginal costs of congestion at non-peak hours are fairly small. Thus, in overall, the average cost of congestion is reduced and not only exchanged between periods. The highways and the far suburb are the most benefited area in all policies, but it is also the most impacted one at non-peak hour.

Area	Initial	Driving	Uniform	Variable
		restriction	toll	toll
Peak hours				
Highways	1.14	0.49	0.51	0.41
City center	1.58	1.35	1.3	1.7
Ring roads	2.09	1.29	1.09	1.3
Close suburb	1.29	1.07	1.04	1.39
Far suburb	0.58	0.39	0.4	0.58
Non-peak hours				
Highways	0.42	0.75	0.66	0.76
City center	1.26	1.3	1.3	1.19
Ring roads	1.19	1.31	1.29	1.22
Close suburb	1.04	1.12	1.12	0.98
Far suburb	0.34	0.51	0.49	0.39

Table 22: Marginal costs of congestion under the different policies.

Note: Costs associated to adding an average driver distance, in \in .

7 Mitigating the surplus losses

7.1 Alternative policy instruments

In addition to the three standard policies we studied above, we analyze other regulations that exist in several cities. In particular, we focus on vintage based driving restrictions and driving permit auctions.

Vintage-based driving restrictions

The first is a vintage-based driving restriction forbidding driving cars older than a particular vintage. This resembles the low emission zones that have been particularly common in Germany. Under our implementation of the standard driving restriction, we restrict everyone but only 50% of the time. This implies that the loss from not being able to drive a fraction of the time is partly compensated by improved speeds when individuals can drive. In contrast, with the vintage-based driving restriction, we never allow old cars to drive so their drivers do not have any benefits from the policy. On the other hand, the individuals who drive newer cars have only benefits through the speed improvements. Thus, we can expect the vintage-based driving restriction to have larger distributional effects than the standard one, as the burden is put on a fraction of the population. Yet, this policy might be more cost-efficient than the usual restriction by targeting the old and more

polluting cars. We compare policies by calibrating the vintage to reach the same traffic reduction in terms of kilometers driven at peak hour as the regular driving restriction.

Since the car vintage is a discrete variable, we cannot exactly match the expected number of kilometers at peak hour with the vintage parameter only. To overcome this issue, we use the strictest vintage and assume that individuals are subject to the policy with a probability lower than 1. Then, we calibrate this probability to exactly match the traffic reduction across policies. This is a technical assumption that can be interpreted as the frequency at which the policy is implemented. The calibrated parameters is presented in Table 23 below. The vintage parameter is 2006, above the average car vintage of 2004. 50.8% of the population is affected by the vintage policy.

 Table 23:
 Parameters for the vintage-based driving restrictions

Policy type	Outcome matched	Vintage	Frequency
Peak-hour	traffic, peak hour	2006	0.90

As Table 24 suggests, we find that the vintage restriction is hurting a smaller fraction of individuals (50.8%) and benefit a larger fraction of the population (28.9%) than the standard driving restriction. Overall, the vintage-based policy is slightly more costly, generating 1.19 million euros of losses versus 1.11 for the standard driving restriction.

However, on average, individual surplus decreases by 30 cents under the vintage-based driving restriction, indicating that the gains from the unaffected individuals do not compensate for the losses of the affected ones. Surprisingly, the value of the gains from speed are identical under the two policies, even if the vintage-based policy increases total travel time by 29,000 hours while the increase due to the standard driving restriction is of 24,900 hours. Under this sophisticated instrument, the distribution of changes in expected travel times is skewed, with a large fraction of individuals increasing their travel time by large amounts, implying more inequalities.

From the emissions reduction perspective, the vintage-based driving restriction improves results significantly. This is particularly striking for the equivalent NO_X emissions that decrease 55% more. Depending on the policy maker target, the higher welfare cost will balance with the higher emission reduction through the implied average cost of regulation. We find that the vintage-based restriction is preferable to a standard restriction, implying a cost that is $\leq 125,000$ lower.

	Standard	Vintage-based
	restriction	restriction
$\% \Delta CS > 0$	0.075	28.9
$\% \Delta CS < 0$	79.7	50.8
$\operatorname{Min}\Delta\mathrm{CS}$	-2.09	-3.91
Mean ΔCS	-0.276	-0.296
Max ΔCS	0.037	1.3
Total $\Delta CS (M \in)$	-1.11	-1.19
ΔCS , constant speed	-1.28	-1.36
ΔCS from speed	0.164	0.164
Total $\Delta wCS (M \in)$	-1.16	-1.37
$\Delta \operatorname{CO}_2$ (ton)	-300	-363
$\Delta \text{ eqNO}_{X} (\text{ton})$	-1.24	-1.92
Implied cost $eqNO_X$	897,728	621,314

Table 24: Welfare costs of vintage-based driving restrictions.

Auctioned driving license

The second instrument is a quota of driving permits allocated through an auction. For example, Shanghai has this type of system to assign a fixed number of vehicle licenses. Finally, the last policy we analyze introduces a quota of driving licenses allocated through an auction. We consider a simple auction format: a uniform second-price auction. This implies that individuals bid their true valuation of the license. Individual valuations are equal to the difference between the expected utilities with and without the right to drive. The equilibrium price is the highest rejected bid associated with a fixed number of licenses. Individuals are assumed to perfectly anticipate the changes in speeds. Thus, the willingness to pay includes the gain in utility from better speeds at peak hour. We use an iterative algorithm to solve for the license price together with the equilibrium speeds. For some values of the quota of driving licenses, our algorithm cannot find a stable equilibrium. This is particularly the case when we consider stringent policies (i.e. with a low number of licenses) where the speed gains are large. We thus select the quota of driving license that implies the closest outcome to the one obtained in our main policies. We then re-calibrate the uniform toll to match the traffic reduction across policies, measured by the number of kilometers driven at peak hours. We analyze here policies that trigger a reduction in traffic at peak hours of 25.2%.

This policy seems more suited to a comparison with the uniform toll, as they both put a price on the right to drive. However, there is an important difference from the perspective of individuals: under the toll, individuals make their decisions to drive and pay the toll after they receive their preference shocks for the transportation modes and departure times. While under the auction, individuals have to submit their bid for the license before receiving their preference shocks. Thus, the bids reflect the willingness to pay for the license, given by the difference between the expected utility with and without the license.

Policy type	Parameter	Value
License	Quota of licenses	30.5%
	License price	1.34
Uniform toll	toll	1.95

Table 25: Parameters for the driving license and equivalent uniform toll policies

Note: License price and toll in \in /trip.

For the license auction, we see stronger substitution patterns towards other transportation modes rather than towards driving at non-peak hours. This is due to the nature of the auction: once an individual has the license, the cost of driving is the same across periods, so she chooses her preferred time. This policy reduces total traffic more, as the share of individuals driving becomes 28.9%, 1.9 percentage points below the value under driving restrictions. As a direct consequence of the larger substitution towards other transportation modes, the impact on emissions is larger. However, the toll remains better, as the individuals with long trips are typically those who get the license and continue to drive.

	Unif. toll	License
$\% \Delta CS > 0$	0	0.133
$\% \Delta CS < 0$	79.7	79.6
$\operatorname{Min}\Delta\mathrm{CS}$	-1.54	-1.31
Mean ΔCS	-0.281	-0.472
Max ΔCS	0	0.394
Total $\Delta CS (M \in)$	-1.14	-1.9
ΔCS , constant speed	-1.37	-2.31
ΔCS from speed	0.23	0.405
Total $\Delta wCS (M \in)$	-1.18	-1.95
Tax revenue	1.28	1.31
Δ welfare	0.142	-0.592
$\Delta \operatorname{CO}_2$ (ton)	-285	-440
$\Delta \text{ eqNO}_{X}$ (ton)	-1.19	-1.88
Implied cost local polluta	nts (€/ton]	NO _X)
Without redistribution	956,790	$1,\!015,\!204$
With redistribution	$-119,\!543$	$315{,}530$

Notes: ΔCS in \in . Total ΔCS in million \in , computed using survey weights.

Since individuals who get the license pay for it regardless of how often they decide to drive at peak hours, the policy generates a high welfare cost. Individuals can no longer react to some bad realizations of their preference shocks for driving. The license regulation generates $\in 1.9$ million loss for consumer surplus, against 1.14 million under the uniform toll. The quota of driving licenses and the toll generate comparable tax revenue of around $\in 1.3$ million. If the entire auction revenue is redistributed, the quota of driving licenses generates a net loss of $\in 0.6$ million.

7.2 Differentiated tolls

[This section is preliminary.]

We now investigate what are the consequences of applying differentiated tolls. First, we consider tolls that depend on the areas where individuals drive. Second, we analyze combinations of fixed and variable tolls.

Area-specific tolls

We consider a toll that takes two different values: one for the city center and ring roads and one for the highways, the close and the far suburb. This policy instrument is similar to the cordon pricing mechanism, which defines tolls based on the distance to the city center. We determine all the toll combinations that imply the same objective traffic at peak hours and find the best toll combination for each objective outcome. When individuals drive through the two toll zones, we assume they only pay the highest toll.⁷

The left graph in Figure 12 shows all the combinations of tolls that achieve the same objective traffic at peak hour. The right figure provides the change in total welfare (measured by the sum of consumer surplus and tax revenue) as function of the value of the toll in the suburb. First, we see that the uniform toll is not too far from the welfare-maximizing toll combination. By slightly reducing the value of the toll in the suburb and increasing the value of the toll inside the city center, we can multiply the total welfare surplus by 3.2. The welfare is maximized for the combination of $\in 2.03$ in the suburb and $\in 3.51$ in the city center. We also provide the tolls that maximize other outcomes and we see that the best toll combination for welfare is very close to the best one for tax revenue. In contrast, the best toll combination for consumer surplus (both with and without redistributive weights) consists of having zero toll in the suburb and the high value of $\in 7.01$ inside the city center.





Combination of fixed and variable tolls

We now consider another way to define tolls, given by a fixed and a variable part. We allow for negative values of the two components of the toll. Interestingly, we find that the best toll for consumer surplus is has the lowest fixed part and the highest negative part. This is because

⁷We could alternatively consider individuals driving through the two toll areas pay the sum of the tolls, but this situation would not nest the benchmark uniform toll and thus makes the comparison less straightforward.

individuals with short trip have high valuation of the travel times so their gains more than compensates the losses for individuals with long distance that end up paying high toll values. However this combination of tolls is not efficient at raising tax revenue and this is why the best combination for consumer surplus assuming the full tax revenue is redistributed is a toll with a smaller variable part than our benchmark variable toll (7 cents instead of 10 cents) and a small fixed part of $\notin 0.72$. Note that we could reach the maximal tax revenue with the large fixed value of $\notin 3.2$ and a negative variable part (-0.2 cents/km). This combination of tolls would be nevertheless welfare-decreasing as Figure 12 (b) suggests.



Figure 12: Combination of fixed and variable tolls and their welfare costs.

7.3 Improving public transport

Public transport is a crucial part of cities' infrastructure and its quality and density has a large impact on individuals transportation decisions. While our paper focuses on traffic policies, we can still give useful insights on how infrastructure improvements such as better coverage or quality of public transport mitigate the consumer surplus losses.

While public transport is the most used transportation mode in our data, 28.5% of the individuals do not have access to it. In the first scenario, we allow them to use an hypothetical public transport service with the median characteristics: an speed of 12.2 km/hr, a price of $1.52 \in$, and congestion levels of 1.67 at peak and 0.80 at non-peak hours. A commonly considered improvement for public transport corresponds to capacity increases at the station or train levels. In our second scenario, we interpret this type of investments as improvements to the quality of the public transport service via a decrease in the trip's duration of 40% for all trips. Finally, in the third scenario we look at

the impact of making public transport free, a policy that has been implemented during pollution peaks in Paris.

Table 26 shows the shares of transportation modes under the different scenarios. Extending public transport coverage is by far the most effective policy at decreasing car usage thanks to a large substitution towards public transport. Nonetheless, at non-peak hour, the coverage increase does not improve public transport usage, while the highest increase in usage comes from the quality improvement. A possible explanation for these results would be a lack of scheduling flexibility among those benefiting from the coverage extension. A full subsidy of the service is less effective than an increase in quality at nudging individuals towards public transport usage, showing the importance of infrastructure investments, rather than subsidies, at improving public transport usage. While a full subsidy or a decrease in durations of 40% could be seen as extreme cases of possible policies, these results give useful insights on how to rank the impact of different types of policies.

Welfare changes are showed in Table 27. An increase in public transport coverage reduces considerably consumer surplus losses, with a decreases between 35% and 40% across policies. Quality improvements via a 40% reduction in duration lead to smaller reductions of consumer surplus losses, of at most 15% of the loss under no improvement. Finally, fully subsidizing public transport has a very small effect at mitigating the costs of the policies. These results indicate that a coverage extension is the most useful improvement in order to mitigate traffic policies. In the subsidy scenario, the decrease in revenue from a a larger usage of public transport is bigger than the decrease in consumer surplus variation due to the policy implementation. Thus, making public transport free while implementing another traffic regulation can create a larger decrease in welfare than when only applying the restriction or pricing mechanism.

	Dr	iving r	estrict	ion		Unifor	rm toll			Variał	ole toll	
Mode	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Bicycle	2.4	1.7	1.9	2.2	2.4	1.7	1.9	2.2	 2.2	1.7	1.7	2
Pub. transport, peak	32.1	48.1	37.9	35.1	32.4	48.7	38.3	35.4	32.9	48.6	38.6	35.9
Motorbike	2.5	1.7	1.8	2.1	2.5	1.7	1.8	2.1	2.6	1.8	1.9	2.2
Walking	17.6	12	15.5	16.5	17.8	12.2	15.7	16.8	15.9	11.6	13.9	14.9
Car, peak	14	10.4	11.6	12.8	14	9.6	11.5	12.6	14	9.7	11.7	12.8
Car, non-peak	16.1	10.9	13.3	14.6	15.5	10.8	12.7	14	16.5	11.1	13.7	14.9
Pub. transport, non-peak	15.3	15.1	18.1	16.8	15.5	15.2	18.2	16.9	15.8	15.5	18.5	17.2
Total car share	30.1	21.3	24.9	27.4	29.5	20.4	24.2	26.6	 30.5	20.8	25.4	27.8
Total pub. transport share	47.4	63.2	56	51.9	47.9	63.9	56.5	52.4	48.7	64.1	57.1	53.1

Table 26: Shares of transportation modes under public transport improvements

Notes: In %. (1): Benchmark, no improvement. (2): Coverage improvement. (3): Duration improvement. (4): Free usage of public transport.

	D	riving r	estricti	on		Unifor	rm toll			Varial	ole toll	
Mode	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Total $\Delta CS (M \in)$	-1.27	-0.77	-1.15	-1.2	-1.55	-1.02	-1.38	-1.46	-1.64	-1.02	-1.39	-1.48
Tax revenue	0	0	0	0	1.5	1.06	1.26	1.39	1.1	0.77	0.79	0.91
Δ welfare	-1.27	-0.77	-1.15	-1.2	-0.01	0.04	-0.13	-0.07	-0.58	-0.24	-0.6	-0.57
$\% \Delta CS = 0$	20.3	20.3	20.3	20.3	20.3	20.3	20.3	20.3	20.3	20.3	20.3	20.3
$\% \Delta CS > 0$	0.5	0.57	0	0.02	0	0.01	0	0	14.3	14.4	12.6	13.6
$\% \Delta \mathrm{CS} < 0$	79.2	79.2	79.7	79.7	79.7	79.7	79.7	79.7	65.4	65.4	67.2	66.1
Min ΔCS	-2.49	-2.25	-2.51	-2.49	-2.18	-1.97	-2.18	-2.18	-5.3	-4.23	-5.31	-5.29
Mean ΔCS	-0.32	-0.19	-0.29	-0.3	-0.38	-0.25	-0.34	-0.36	-0.41	-0.25	-0.34	-0.37
Median ΔCS	-0.18	-0.15	-0.13	-0.16	-0.26	-0.21	-0.19	-0.23	-0.14	-0.11	-0.1	-0.13
Max ΔCS	0.05	0.05	0	0.02	0	0.01	0	0	1.5	1.43	1.15	1.15

Table 27: Consumer surplus variation under public transport improvements

Notes: ΔCS are in \in /trip. (1): Benchmark, no improvement. (2): Coverage improvement. (3): Duration improvement. (4): Free usage of public transport.

8 Conclusion

Using the data obtained from a detailed survey, Google maps, and TomTom, we estimate a nested logit specification to model the transportation decisions of individuals in the Île-de-France region for their daily trip to work or study place. The estimated parameters confirm the importance of trip duration for individuals' decisions and reveal serious schedule inflexibility making it challenging to discourage individuals from driving. We combine this model with a flexible reduced-form congestion model that predicts how road speeds vary when the number of drivers decreases in the different parts of the region. We simulate the effects of different transportation policies, measure their welfare costs at the individual level, and estimate their impacts on emissions. We find that all the regulations are costly for individuals. Still, simple driving restrictions are not as bad as we expected since, by definition, it forces everyone to contribute to traffic reduction. It generates fewer welfare losses than tolls on aggregate. However, driving restrictions do not raise revenue, unlike the tolls. When the tax revenue is redistributed, the ranking of the policies is inverted.

References

- Akbar, P., Couture, V., Duranton, G., and Storeygard, A. (2018). Mobility and congestion in urban india. Working Paper 25218, National Bureau of Economic Research.
- Akbar, P. and Duranton, G. (2017). Measuring the cost of congestion in highly congested city: Bogota. CAF Working paper.
- Anderson, M. L. (2014). Subways, strikes, and slowdowns: The impacts of public transit on traffic congestion. *American Economic Review*, 104(9):2763–96.
- Anderson, M. L. and Davis, L. W. (2020). An empirical test of hypercongestion in highway bottlenecks. Journal of Public Economics, 187:104197.
- Arnott, R., De Palma, A., and Lindsey, R. (1990). Economics of a bottleneck. Journal of Urban Economics, 27(1):111–130.
- Arnott, R., Palma, A. D., and Lindsey, R. (1993). A structural model of peak-period congestion: A traffic bottleneck with elastic demand. *The American Economic Review*, 83(1):161–179.
- Barwick, P. J., Li, S., Waxman, A. R., Wu, J., and Xia, T. (2021). Efficiency and equity impacts of urban transportation policies with equilibrium sorting. Working Paper 29012, National Bureau of Economic Research.
- Basso, L. J. and Silva, H. E. (2014). Efficiency and substitutability of transit subsidies and other urban transport policies. American Economic Journal: Economic Policy, 6(4):1–33.
- Batarce, M. and Ivaldi, M. (2014). Urban travel demand model with endogenous congestion. *Transportation Research* Part A: Policy and Practice, 59:331 – 345.
- Bou Sleiman, L. (2021). Are car-free centers detrimental to the periphery? Evidence from the pedestrianization of the Parisian riverbank. Technical Report 2021-03, Center for Research in Economics and Statistics.
- Buchholz, N., Doval, L., Kastl, J., Matějka, F., and Salz, T. (2020). The value of time: Evidence from auctioned cab rides. Technical report, National Bureau of Economic Research.
- Couture, V., Duranton, G., and Turner, M. A. (2018). Speed. Review of Economics and Statistics, 100(4):725-739.
- Davis, L. W. (2008). The effect of driving restrictions on air quality in mexico city. *Journal of Political Economy*, 116(1):38–81.
- De Palma, A. and Lindsey, R. (2006). Modelling and evaluation of road pricing in paris. *Transport Policy*, 13(2):115–126. Modelling of Urban Road Pricing and its Implementation.
- De Palma, A., Lindsey, R., and Monchambert, G. (2017). The economics of crowding in rail transit. *Journal of Urban Economics*, 101:106–122.
- DG MOVE (2014). Update of the handbook on external costs of transport. Report for the European Commission.
- Gale, D. and Nikaido, H. (1965). The jacobian matrix and global univalence of mappings. *Mathematische Annalen*, 159(2):81–93.

- Gallego, F., Montero, J.-P., and Salas, C. (2013). The effect of transport policies on car use: Evidence from latin american cities. *Journal of Public Economics*, 107:47 62.
- Geroliminis, N. and Daganzo, C. F. (2008). Existence of urban-scale macroscopic fundamental diagrams: Some experimental findings. *Transportation Research Part B: Methodological*, 42(9):759 770.
- Hall, J. D. (2019). Can tolling help everyone? estimating the aggregate and distributional consequences of congestion pricing. *Journal of the European Economic Association, forthcoming.*
- Hanna, R., Kreindler, G., and Olken, B. A. (2017). Citywide effects of high-occupancy vehicle restrictions: Evidence from "three-in-one" in jakarta. Science, 357(6346):89–93.
- Haywood, L. and Koning, M. (2015). The distribution of crowding costs in public transport: New evidence from paris. Transportation Research Part A: Policy and Practice, 77:182–201.
- Herzog, I. (2021). The city-wide effects of tolling downtown drivers: Evidence from london's congestion charge. Working papers.
- Jia, Z., Chen, C., Coifman, B., and Varaiya, P. (2001). The pems algorithms for accurate, real-time estimates of g-factors and speeds from single-loop detectors. In *ITSC 2001. 2001 IEEE Intelligent Transportation Systems*. *Proceedings (Cat. No. 01TH8585)*, pages 536–541. IEEE.
- Kilani, M., Proost, S., and van der Loo, S. (2012). Tarification des transport individuels et collectifs a paris. rapport final de recherche predit. In *Technical Report*.
- Kilani, M., Proost, S., and van der Loo, S. (2014). Road pricing and public transport pricing reform in paris: Complements or substitutes? *Economics of Transportation*, 3(2):175–187. Special Issue in Honor of Herbert Mohring.
- Kreindler, G. (2016). Driving delhi? behavioural responses to driving restrictions. Working paper.
- Kreindler, G. E. (2020). The welfare effect of road congestion pricing: Experimental evidence and equilibrium implications. *Unpublished paper*.
- Loder, A., Ambühl, L., Menendez, M., and Axhausen, K. W. (2017). Empirics of multi-modal traffic networks-using the 3d macroscopic fundamental diagram. *Transportation Research Part C: Emerging Technologies*, 82:88–101.
- Lucinda, C. R., Moita, R. M. S., Meyer, L. G., and Ledo, B. A. (2017). The economics of sub-optimal policies for traffic congestion. *Journal of Transport Economics and Policy*, 51(4):225 – 248.
- Parry, I. and Small, K. (2009). Should urban transit subsidies be reduced. American Economic Review, 99(3):700-724.
- Small, K. A., Verhoef, E. T., and Lindsey, R. (2007). The economics of urban transportation. Routledge.
- Tarduno, M. (2022). For whom the bridge tolls: Congestion, air pollution, and second-best road pricing. Working papers.
- Train, K. E. (2009). Discrete choice methods with simulation. Cambridge university press.
- Tsivanidis, N. (2018). The aggregate and distributional effects of urban transit infrastructure: Evidence from bogota's transmilenio. Working paper.

- Van Den Berg, V. and Verhoef, E. T. (2011). Winning or losing from dynamic bottleneck congestion pricing?: The distributional effects of road pricing with heterogeneity in values of time and schedule delay. *Journal of Public Economics*, 95(7):983–992.
- Yang, J., Purevjav, A.-O., and Li, S. (2020). The marginal cost of traffic congestion and road pricing: Evidence from a natural experiment in beijing. *American Economic Journal: Economic Policy*, 12(1):418–53.

A Additional information on data and sample construction

A.1 Detailed information about EGT

The EGT survey was conducted from 2009 to 2011 between October and May, excluding school holidays. Due to a low answer rate during the initial collection period (2009 to 2010), the second wave of data collection took place from 2010 to 2011. Regarding the survey sampling, the region (Île-de-France) was divided into 112 sectors. In each sector, between 400 and 500 individuals were interviewed to maintain representativeness at the sector level. Instead of relying on a trip diary where surveyed individuals self-report their trips, the EGT hinged on pollsters visiting households and recording the information of the trips performed the previous day using a computer application.

A.2 Cost estimation

The survey does not report the cost incurred by the individuals when taken a trip. Thus, we proceed to estimate the cost both for taken and non-taken alternatives. For all transportation modes, we neglect the fixed expenses. We assume that individuals do not consider yearly, monthly, or one-time-only fixed costs related to the transportation alternatives unless that cost is a service subscription. Thus, expenses related to the transportation alternatives like car purchase, insurance, taxes, or administration fees, are not included in the cost computation. The rationale is that we focus on the short-term reaction to policy introduction, implying that individuals cannot avoid these fixed costs.

Walking is always free, while cycling is free only for households that own bicycles. If the individual has a bike-sharing subscription, the cost of the subscription is divided by twice its duration (in days), which is equivalent to assuming two bike trips per day. In other cases, we consider biking has the cost of a single bike-sharing ticket (≤ 1.7).

A.2.1 Public transport

Public transport in Paris comprises a network of subway, buses, tramways, and suburban trains. This network is operated by two companies: "RATP" mainly covers the public transport inside Paris and close suburbs, while "SNCF" operates trains that connect Paris to the suburban areas.⁸ During the period of our data, the RATP pricing system was based on pricing zones, from 1 (inside

⁸Observations for which part of the trip used either Vogueo (ships) or any non-urban train are dropped. We keep the school buses.

Paris) to 5 (far suburb), of the trip's origin and destination.⁹ We use the prices stated in the price guide of RATP for July 2011.¹⁰

Public transport in Paris comprises a network of subway, buses, tramways, and suburban trains. This network is operated by two companies: "RATP" mainly covers the public transport inside Paris and close suburbs, while "SNCF" operates trains that connect Paris to the suburban areas.¹¹ During the period of our data, the RATP pricing system was based on pricing zones, from 1 (inside Paris) to 5 (far suburb), of the trip's origin and destination.¹² We use the prices stated in the price guide of RATP for July 2011.¹³

In contrast, the ticket price using the SNCF network depends on the exact stations of origin and destination rather than simply their zones. Since there is no exhaustive data with the prices for all combinations of origin and destination stations, estimate the train ticket price. We rely on a sample of ticket prices for 36 origin and destination pairs and estimate the train ticket price as a cubic function of the distance between stations using ordinary least squares. This regression has a good fit with an \mathbb{R}^2 of 0.82. We use this function to predict prices of the public transport trips that include taking an SNCF suburb train.

For individuals with a public transport subscription, we estimate a trip's average cost by dividing the daily price of the subscription by two, which is the average number of trips taken in a day conditional on using public transport. For some individuals, we have missing data about the coverage of their subscriptions. We assume individuals pay the regular ticket price as if they did not have a subscription. For those who take a public transport trip outside of their subscription coverage or do not have a subscription, we also assume a cost equal to the price of an individual ticket. Seven individuals stated to have used the service without paying (fraud); we assume a zero cost for their public transport trip. The survey includes information on whether individuals can buy subsidized tickets and access reduced-price subscriptions. We take this information into account when computing their cost of using public transport.

 $^{^{9}}$ Before 2011, the region was split into 8 pricing zones. In 2011, zones 6 to 8 were eliminated and included in zone 5. We follow the zoning system of 2011.

¹⁰Source: "Guide tarifaire", Juillet 2011, https://www.slideshare.net/quoimaligne/guide-tarifaire-ratp -sncf-ile-de-france-2011.

¹¹Observations for which part of the trip used either Vogueo (ships) or any non-urban train are dropped. We keep the school buses.

¹²Before 2011, the region was split into eight pricing zones. In 2011, zones 6 to 8 were eliminated and included in zone 5. We follow the zoning system of 2011.

¹³Source: "Guide tarifaire", Juillet 2011, https://www.slideshare.net/quoimaligne/guide-tarifaire-ratp -sncf-ile-de-france-2011.

A.2.2 Cars and motorbikes

We estimate the cost of using a car or a motorbike by combining the trip distance from the route provided by TomTom, estimates of the fuel consumption of the households' vehicles, and average fuel prices in 2011 from the National Survey Institute ("Insee").¹⁴ When the household has multiple vehicles, we assume the trip uses the most fuel-efficient one. We also assume each individual pays the total cost of the trip, regardless of the number of passengers.

The survey data does not contain vehicles' fuel consumption, so we predict it from the main car characteristics (vintage, fuel type, and fiscal horsepower). We rely on detailed data on car purchases that contains information on CO_2 emissions and car characteristics. We specify car's CO_2 emissions as a linear function of a time trend and the fiscal horsepower. The functions are fuel-specific. More details on the car purchase data and the regressions are provided below in Section A.6.

For two-wheel motorized vehicles, we assign the average fuel consumption by the number of cylinders provided by the French Energy Agency "ADEME".¹⁵

A.3 Map direction APIs

The survey maps the Paris region into a grid with 1,489,347 squares to locate individuals' trips, each one measuring 100 square meters. Instead of recording each trip's exact departure and arrival GPS coordinates, the survey records the starting and finishing squares in the grid. Thus, we use the GPS coordinates of the centroids of the grid squares. This approach limits any trip geocoding inaccuracy to a maximum of approximately 70 meters.

The public transport queries were done on June 2^{nd} , 2019, setting all trips to take place on Tuesday, June 4^{th} 2019, with a departure time at 9:30 a.m. The car queries were done in April 2021, setting the trips to take place the Thursday 16^{th} of September 2021 at 8:30 a.m. for peak hours and 6:30 a.m. for non-peak hours.

TomTom queries for future dates use historical trip data and not the live conditions. However, to solve concerns regarding the impact of Covid on the reliability of these predictions, we compare our TomTom queries at peak hours (8:30) with old Google maps queries from August 2019. We find that the average difference between the two data sources implies that TomTom queries predict trip durations that are 7.5% larger than the older ones from Google maps, with a median difference of 7.6%. The results suggest that the Covid crisis did not significantly modify the prediction algorithm

¹⁴See https://www.prix-carburants.developpement-durable.gouv.fr/petrole/se_cons_fr.htm.

¹⁵Source: https://www.statistiques.developpement-durable.gouv.fr/les-deux-roues-motorises-au -1er-janvier-2012.

of TomTom. Furthermore, they show the similarities between the two sources and thus TomTom's reliability.

A.4 Estimating real estate prices

We proxy the households' incomes using the expected price of their housing. To estimate each household's real estate price, we rely on a database containing all real estate transactions in France for 2014-2015 ("Demande de Valeurs Foncières - DFV+") from the French tax authority and CEREMA.¹⁶ We use this period since it is the earliest for which these data are publicly available.

We match each household to all the estate transactions from the municipality of residence and the neighboring municipalities. For Paris, we use the "arrondissement" level.¹⁷ We limit the sample of matches to apartments and houses' sales, properties sold in one transaction and exclude partial sales of the property. We drop transactions with a price of zero or above five million euros and keep only the properties with a built surface is between 15 and 500 square meters. We compute each transaction's price per square meter by dividing the transaction price by the property's built surface. We trim the sample and drop the top 0.5 % and the bottom five percent of the distribution of square meter prices for each municipality. Next, we estimate estate prices for our sample by taking a weighted average, using the inverse of the distance between the household and the match as weight. We finally multiply the average estate price by the households' apartment, or house surface area reported in the survey. As seen in table 28, the average real estate price that we find for each area is close to the area-specific price from official French data, supporting the credibility of our estimates.

 Table 28:
 Comparison real estate prices

	Paris	Close suburbs	Far suburbs
Average from matched data	8,030	4,789	3,175
Average from official data	8,074	4,338	2,998

Note: Sample average computed using survey weights. Average from official data obtained by averaging quarterly average price per square meter for the year 2014. Source official data https://basebien.com/PNSPublic/DocPublic/ Historiquedesprixdesappartementspardep.pdf

A.5 Estimation of the public transport overcrowding

To have an individual level measurement of overcrowding in public railway transport, we first compute a line-level measure of congestion. To do so, rely on data provided by the agencies in

¹⁶Source: https://www.data.gouv.fr/fr/datasets/dvf-open-data/#_.

¹⁷Paris is split in 20 "arrondissements".

charge of public transportation in the Paris region (SNCF and RATP) on the number of passengers at the metro or train station level. We rely on data for 2015, which is the oldest data available, and consider only the urban railway network, where overcrowding is the most problematic. These data only record the validations from passengers that use an electronic metro card. Unfortunately, there is no exhaustive data on traffic in public transport that accounts for passengers using tickets. Estimates suggest that the electronic validations represent two-thirds of the traffic for 2016.¹⁸ We believe that during morning peak and non-peak hours, the share is even higher than 67%.

The data are composed of two separate datasets. The first one contains daily entry flows of passengers at the railway station level. The second dataset contains "hourly profiles" at the station level: the distribution of validations (in %) across hours for different periods (business days outside holidays, business days during school holidays, and weekends). By combining these two datasets, we can obtain a daily estimate of the number of passengers in each metro station for a regular business day. We exclude weekends, school holidays, public holidays, and two dates with a relatively low total number of entries (less than a million versus an overall daily average of 7.5 million according to the official figures of the RATP). We interpreted this low number of passengers as indicating the occurrence of a strike. In the end, we average 172 days to compute the daily numbers of entries per railway station. We use the average profiles for the business days outside holidays for the second semester of 2015.¹⁹ We define our two periods as between 7 and 8:59 a.m. for peak and 6:00-6.59 a.m. for non-peak hours to be consistent with the expected driving durations.

We observe the number of passengers entering a station, but there are multiple railway lines in a station. We estimate the average number of passengers by line using the total annual traffic divided by the total railway traffic from official figures. More specifically, we calculate the percentage of passengers that use each line per station and then sum across stations of the line to get the total number of passengers per period.

We use the 2015 General Traffic Feed Specification (GTFS) schedules from RATP and SNCF that provide frequencies of trains at the station level.²⁰ We use schedules of September 2015 to be consistent with the station passenger flow data and compute the average number of trains per hour for peak and non-peak hours for each station and urban transit line. We average across stations to get the expected number of trains for each railway line. Additionally, we gather information about the passenger capacity of the train models used on each line.²¹ The passenger capacity represents the number of passengers a train can carry, assuming four passengers per square meter. We compute the total railway line capacity at peak and non-peak hours by multiplying the train

¹⁸See https://www.iledefrance-mobilites.fr/usages-et-usagers-des-titres-de-transport.

¹⁹We noticed some problems with the data from the first semester 2015: the percentages did not sum to 100% for 20 stations and thus preferred not to use that semester of data.

²⁰Source: https://transitfeeds.com/l/162-paris-france.

²¹We rely on Wikipedia and old reports containing information about the fleet of trains owned by RATP and SNCF.

capacity by the number of trains per hour. Finally, the overcrowding level c_{lt} for line l, at time period t is:

$$c_{lt} = \frac{p_{lt}}{2 \times t c_{lt}},$$

where p_{lt} is the hourly number of passengers in line l at time period t, tc_{lt} is the line total passenger capacity per hour. We multiply the total line capacity by two because we have two directions, and we use the total number of passengers going in both directions. Then from the overcrowding at the railway line level, we obtain an estimate of individuals' subjective overcrowding levels, which averages the different lines through:

$$C_{nt} = \sum_{l} c_{lt} \times \omega_{nl},$$

where c_{lt} is the estimated line l overcrowding level and ω_{nl} is the percentage of the trip duration spent in line l. The ω_{nl} are such that $\sum_{l} \omega_{nl} = 1$, $\forall n$. Note that the portion of the trip done by foot or using bus, tram or another transportation network for which we do not have data are excluded. ω_{nl} are obtained from the Google maps detailed public transport itineraries.

For 46 stations out of 537, we estimate more passengers at non-peak than at peak hours. These stations are all in the far suburb except one. The exception is Bercy station on the suburb train network. Across stations and lines, the number of passengers is multiplied by 3.2 between non-peak and peak hours, and the median ratio is 2.34. But the train and metro frequency increase a lot between peak and non-peak hours. Finally, we provide in Table 29 the estimates of the overcrowding levels in the different metro and train lines. On average, we estimate the overcrowding to be 0.89 at non-peak hour and 1.43 at peak hour. But these averages hide important heterogeneity across lines that provide key variation for estimating the sensitivity to overcrowding in public transport.

One caveat of our overcrowding measures is that is they use validations by electronic passes only, so we underestimate the traffic. However, we exploit the variation between peak and non-peak hours traffic across metro lines. As long as traffic is homogeneously underestimated over the network and periods, omitting a portion of the traffic is not a major problem.

	Met	ro	ç	Suburb trains				
Line	Peak	Non-peak	Line	Peak	Non-peak			
1	0.43	0.72	А	2.35	4.37			
2	0.5	0.77	В	0.55	1.07			
3	1.04	1.29	С	0.58	1.16			
3B	0.18	0.36	D	1.24	1.75			
4	0.6	1.11	Ε	0.78	1.37			
5	1.3	1.84	Η	0.4	0.71			
6	0.72	1.01	J	0.63	1.12			
7	1.5	1.71	Κ	1.39	2.22			
7B	0.19	0.41	L	0.61	1.36			
8	0.86	1.12	Ν	0.48	1.01			
9	0.81	1.07	Р	1.91	3.82			
10	0.59	1.13	R	0.93	1.3			
11	0.98	1.39	U	0.88	2.1			
12	1.07	1.44						
13	1.62	1.93						
14	0.56	0.95						

Table 29: Estimates of the overcrowding level in the railway.

A.6 Car emissions

A.6.1 Copert emissions factors

Our goal is to have emissions estimates that depend on the distance driven and the speed. To do so, we rely on Copert emissions factors for cars published in the *EMEP/EEA air pollutant emission inventory guidebook - 2009*. This report provides an emission function that links a car's emissions of local pollutants with its speed. The function parameters are specific to each pollutant, fuel type, emission standard, and car segment (in four categories: mini, small, medium, and large).

The EGT data does not directly provide the car's segment, so we predict it from the fiscal horsepower. We use a logit regression that links the segment to the fiscal horsepower and estimates it using data from car characteristics and their corresponding segment. This regression estimates the horsepower thresholds that classify the cars into segments from their horsepower.

Additionally, we assign cars to an emission standard from their vintage: cars with a vintage below 2000 are under Euro 2 standards. From 2000 to 2005 they are under Euro 3 standards. From 2005 onward, they are subject to Euro 4 standards. Since both the Copert emissions data and the survey include information on fuel types, we can directly match survey cars to the correct set of factors by

fuel type. We assume electric vehicles do not emit pollutants.

A.6.2 Alternative emissions' estimation

The Copert emission factors are only available for local pollutants, so we use another method to estimate CO_2 emissions. We also use this method to predict emissions of local pollutants as a robustness check. We estimate a model that predicts the average car emissions of pollutants from its fuel type, fiscal horsepower, a linear time trend, and the emission standards applicable at that time. To estimate this prediction model, we use car registration data in Île-de-France from 2003 to 2018.²² We complement these data with local pollutant emissions data by car model from the UK Vehicle Certification Agency.²³

We first match the French car registration data to the UK data on local pollutant emissions. This allows us to weigh each car model in the UK emissions data by their sales in the Île-de-France region. The two datasets do not contain the same car characteristics. We, therefore, rely on the following algorithm to match the French car sales to the UK emissions:

- 1. We compute the total car sales by year, brand, model name, fuel type, and CO_2 emissions in the French data.
- 2. We merge them with the UK emissions data by year, brand, model name and fuel type. Since there are several versions for the same combination in the UK data, we select the closest neighbor based on cylinder capacity and CO_2 emissions. We use the following norm to compute the distance to every potential matches:

distance =
$$\sqrt{\left(\frac{\mathrm{CO}_{2,FR} - \mathrm{CO}_{2,UK}}{\mathrm{CO}_{2,FR}}\right)^2 + \left(\frac{\mathrm{Cylinder}_{FR} - \mathrm{Cylinder}_{UK}}{\mathrm{Cylinder}_{FR}}\right)^2}$$

- 3. To ensure the matching accuracy, we drop observations for which either the percentage difference in CO_2 emissions or cylinder capacity between the two matches is larger than 10%.
- 4. For each pollutant, we drop car models for which the emission levels are above the corresponding Euro standard limit. We also drop cars whose emissions level is lower than a tenth of the Euro norm value, as we assume they correspond to reporting errors.

For hybrid cars and other fuel types (liquefied petroleum or natural gas), we observe that the top models in France are not available in the UK emissions data. Thus, we match the data on

²²These come from proprietary data obtained from the French Car Manufacturers syndicate "CCFA" (for 2003-2008) and from AAAData (for 2009-2018).

²³Source: https://carfueldata.vehicle-certification-agency.gov.uk/downloads/archive.aspx..

registrations to another dataset that provides car emissions data from 2012 to 2015. The data come from the French environment agency ("ADEME").²⁴ We did not use it for the conventional fuel cars because the sample period is limited to the latest emission standards. We follow the same procedure as described above, but we allow for more discrepancy between the potential matches. Specifically, we drop observations only if the percentage difference in CO_2 emissions or cylinder capacity between the two matches is greater than 30%. We also rely on these ADEME data to obtain estimates of PM emissions for gasoline cars since the UK data only provide PM emissions for diesel cars.

Once we have a final sample of car models and the corresponding sales and emissions, we estimate a prediction model. We specify the emissions level of a specific pollutant as a linear function of the fiscal horsepower, a linear time trend, and dummies for the years of changes in the emissions standard of this particular pollutant. We allow the parameters to be different by fuel type. For PM emissions, we regress the logarithm of the emission levels on car characteristics because PM data are less reliable and have more outliers. All regressions are weighted by the number of car sales.

	Observa	ations	\mathbf{R}^2	
Pollutant	Gasoline	Diesel	Gasoline	Diesel
$\overline{\mathrm{CO}_2}$	6,407	7,962	0.84	0.83
NO_{X}	$5,\!414$	6,759	0.31	0.78
HC	$6,\!138$	6,701	0.51	0.14
CO	13,68	86	0.36	3
РМ	3.57	0	0.85	5

Table 30: Fit of the emissions regressions

Notes: Pooled regressions used for the PM and CO emissions estimation.

For hybrid cars and other fuels, given the small matched sample size (91 observations), we use the weighted average emission levels by fuel type. Emissions for electric cars are set to zero for all pollutants. For some individuals in EGT data, the car vintage and horsepower are missing; in such cases, we attribute the average vintage or horsepower values in the EGT sample for the specific fuel type.

²⁴Source: https://www.data.gouv.fr/fr/datasets/emissions-de-co2-et-de-polluants-des-vehicules-commercialises-en-france/.

	Gase	oline	Diesel			
Pollutant	Observed	Predicted	Observed	Predicted		
$\overline{\mathrm{CO}_2}$	203	197	171	168		
CO	651	569	227	219		
NOX	100	62.1	371	346		
HC	100	88.2	114	55.6		
PM	0.44	1.59	3.26	23.1		

Table 31: Comparison between observed and predicted emissions

Note: Values computed for the sample of cars used to predict the emissions levels on the EGT sample. All the values are for the year 2003 (the average car vintage in the EGT sample) except for PM. For PM emissions, we use the earliest year in the sample with available data: 2012 for gasoline and 2005 for diesel.

B Additional results on the uniqueness of the model equilibrium

We provide here the proofs of uniqueness of equilibrium under special cases of our model and investigate the numerical properties of our algorithm for the general model.

B.1 Uniqueness under two special cases

Model with one period, one area. We consider here the case on only one endogenous speed in the model. The equilibrium speed v is given by:

$$g(v) := v - f(\phi \sum_{n=1}^{I} w_n s_n(v) K_n + \gamma) = 0$$

If this non-linear equation admits a solution, the solution is unique if and only if |g'(v)| > 0 $\forall v \in [\underline{v}, \overline{v}]$. The derivative is:

$$g'(v) = 1 - \underbrace{f'(\phi \sum_{n=1}^{I} w_n s_n(v) K_n + \gamma)}_{\leq 0} \cdot \underbrace{\phi}_{>0} \cdot \sum_{n=1}^{I} w_n K_n \underbrace{\frac{\partial s_n(v)}{\partial v}}_{\geq 0},$$

which is always positive for $v \in [\underline{v}, \overline{v}]$ as long as the speed function is weakly decreasing in the occupancy rate and the probability to drive increases with the speed.

Model with multiple periods, one area.

Now, we consider a model with a single area but multiple time periods which are substitutes for
individuals. In this setting we have a system of T non-linear equations, $\mathbf{g}(\mathbf{v}) = 0$, where:

$$g_t(\mathbf{v}) := v_t - f(\phi_t \sum_{n=1}^I w_n K_n s_{nt}(\mathbf{v}) + \gamma_t$$

We want to show that the Jacobian of the system is a Leontieff matrix, i.e. the diagonal terms are positive and the off-diagonal terms are non-positive. First we compute the diagonal terms, which resemble the previous derivative and is always greater than 1:

$$\frac{\partial g_t}{\partial v_t} = 1 - \underbrace{f'(\phi_t \sum_{n=1}^{I} w_n K_n s_{nt}(\mathbf{v}) + \gamma_t)}_{\leq 0} \cdot \underbrace{\phi_t}_{>0} \cdot \sum_{n=1}^{I} \underbrace{w_n K_n \frac{\partial s_{nt}(\mathbf{v})}{\partial v_t}}_{\geq 0}.$$

Then we compute the off-diagonal terms, which are always negative due to the substitutability between the different time periods.

$$\frac{\partial g_t}{\partial v_{t'}} = -\underbrace{f'(\phi_t \sum_{n=1}^{I} w_n K_n s_{nt}(\mathbf{v}) + \gamma_t)}_{\leq 0} \cdot \underbrace{\phi_t}_{>0} \cdot \sum_{n=1}^{I} \underbrace{w_n K_n \frac{\partial s_{nt}(\mathbf{v})}{\partial v_{t'}}}_{\leq 0}$$

The Jacobian of $\mathbf{g}(\mathbf{v})$ is thus a Leontieff matrix and by Theorem 5 from Gale and Nikaido (1965) it is a P-matrix. We can then apply the main theorem of Gale and Nikaido (1965) (Theorem 1) that states that if the Jacobian of a system of non-linear equations is a P-matrix, the system has a unique solution in its bounded support.

B.2 Uniqueness for the general model

We provide here the analytical formula for the Jacobian of the contraction defined as:

$$g_t^a(\mathbf{v},\kappa) = \kappa . v_t^a + (1-\kappa) . f^a(\mathbf{v}).$$

We can separate the Jacobian in three types of derivatives: $\frac{\partial g_t^a(\mathbf{v},\kappa)}{\partial v_t^a}, \frac{\partial g_t^a(\mathbf{v},\kappa)}{\partial v_t^{a'}}, \frac{\partial g_t^a(\mathbf{v},\kappa)}{\partial v_{t'}^{a'}}.$

$$\frac{\partial g_{t}^{a}(\mathbf{v},\kappa)}{\partial v_{t}^{a}} = \kappa + (1-\kappa) \underbrace{\int_{i=1}^{a'} \left(\phi^{a} \sum_{n=1}^{I} w_{n} K_{n}^{a} s_{nt}(\mathbf{v}) + \gamma^{a} \right)}_{\leq 0} \cdot \underbrace{\phi^{a}}_{>0} \cdot \sum_{i=1}^{I} w_{n} K_{n}^{a} \frac{\partial s_{nt}(\mathbf{v})}{\partial v_{t}^{a}}}_{\geq 0}$$

$$\frac{\partial g_{t}^{a}(\mathbf{v},\kappa)}{\partial v_{t}^{a'}} = (1-\kappa) \underbrace{\int_{i=1}^{a'} (\phi^{a} \sum_{n=1}^{I} w_{n} K_{n}^{a} s_{nt}(\mathbf{v}) + \gamma^{a})}_{\leq 0} \cdot \underbrace{\phi^{a}}_{>0} \cdot \sum_{i=1}^{I} w_{n} K_{n} \frac{\partial s_{nt}(\mathbf{v})}{\partial v_{t}^{a'}} \mathbb{1}\{n \text{ crosses } a'\}}_{\geq 0}$$

$$\frac{\partial g_{t}^{a}(\mathbf{v},\kappa)}{\partial v_{t'}^{a'}} = (1-\kappa) \underbrace{\int_{i=1}^{a'} (\phi^{a} \sum_{n=1}^{I} w_{n} K_{n}^{a} s_{nt}(\mathbf{v}) + \gamma^{a})}_{\leq 0} \cdot \underbrace{\gamma^{a}}_{>0} \cdot \sum_{i=1}^{I} w_{n} K_{n}^{a} \frac{\partial s_{nt}(\mathbf{v})}{\partial v_{t'}^{a'}} \mathbb{1}\{n \text{ crosses } a'\}}_{\leq 0}$$

The signs of the derivatives of the probabilities are obtained from the analytical formulas

$$\frac{\partial s_{nt}(\mathbf{v})}{\partial v_{t'}^{a'}} = \underbrace{\frac{K_n^{a'}}{(v_{t'}^{a'})^2} \times \varepsilon_{nt'} \times 60}_{\geq 0} \times \underbrace{\frac{\beta_n^{\text{duration}}}{\underbrace{\text{duration}_n}}_{\leq 0}}_{\leq 0} \times \underbrace{s_{nt}\left((1-\sigma)\frac{s_{n,t'}}{\sum_{\tilde{t}=1}^T s_{n\tilde{t}}} + \sigma s_{nt'}\right) \mathbb{1}\{\text{n crosses a'}\}}_{\geq 0} \text{ if } t \neq t'$$

$$\frac{\partial s_{nt}(\mathbf{v})}{\partial v_t^{a'}} = -\underbrace{\frac{K_n^{a'}}{(v_t^{a'})^2} \times \varepsilon_{nt} \times 60}_{\geq 0} \times \underbrace{\frac{\beta_n^{\text{duration}}}{\underbrace{\text{duration}_n}}_{\leq 0}}_{\leq 0} \times \underbrace{s_{nt}\left(1-(1-\sigma)\frac{s_{n,t}}{\sum_{\tilde{t}=1}^T s_{n\tilde{t}}} - \sigma s_{nt}\right) \mathbb{1}\{\text{n crosses a'}\}}_{\geq 0}$$

B.3 Additional results of the algorithm

We show additional numerical results about the convergence by plotting the average number of iterations needed to converge for the possible values of κ between 0.45 and 0.95. More specifically, we draw 10 different initial speed values from a uniform distribution over $[\underline{\mathbf{v}}, \overline{\mathbf{v}}]$ and solve for the speed equilibrium with different values for the tuning parameter. As expected, the number of iterations and the time increases exponentially from $\kappa = 0.45$ onward. Furthermore, we always converged to the same equilibrium speeds regardless of the policy environment. This shows that the choice of setting $\kappa = 0.5$ is efficient in terms of speed of convergence.





(a) Number of iterations. (b) Average convergence time.

C Additional tables and figures

C.1 Additional results for the transportation choice model



Figure 14: Own and cross duration elasticities for car

C.2 Additional results for the congestion technology estimates

For highways, we use 6.2 million observations consisting of hourly data from 654 roadway traffic measuring stations for the years 2016 and 2017. The stations are located on 10 highways (A1, A10, A12, A14, A15, A3, A4, A6a, A6b, A86) and 3 national roads (N104, N118, N315).



Figure 15: Traffic sensors' road coverage

Robustness analysis. We perform a sensitivity analysis to check whether the estimates of the congestion technologies are biased due to potential endogeneity issues. The concern is that there might be some speed shocks that also affect traffic density. Such shocks may be for example weather conditions that may change the incentives to drive (e.g. rain) as well as the speed (e.g. the rain may reduce visibility so drivers slow down). To check whether our estimated congestion technologies are subject to such issue, we use hourly data on weather conditions in Paris to define an extreme weather event and drop theses observations. We define an extreme weather event as satisfying at least one of the following characteristics: temperature below the 5% quantile (1.39° C) or above the 95% quantile (25.28° C), a rain amount greater than the 95% quantile (0.46 mm/hr), snow or a wind speed greater than the 95% quantile (20.5 km/hr). We find that the estimates are robust to excluding up to 39.5% of the observations where am extreme weather condition was observed within a three-hours window (see Figure 16).



Figure 16: Robustness checks: elimination of extreme weather conditions.

Next, we may be worried that road or time unobserved heterogeneity cause an endogeneity problem for the estimation of the congestion technologies. To check this we allow congestion technologies to vary across stations or time. More specifically, we estimate congestion technologies by subgroup and then aggregate them to construct an average congestion technology. Since we approximate the speed-traffic density functions by Bernstein polynomials of order seven we can only consider subgroups with more than eight observations. The aggregate congestion technologies are displayed in Figure 17. We see that the inclusion of the fixed effects do not change by a lot the shapes of the congestion technologies.

Figure 17: Robustness checks: estimate technology by subgroups.



	Driving	Uniform	Variable
	restriction	toll	toll
ΔCO	-0.604	-0.996	-0.77
ΔNO_X	-1.05	-1.16	-1.3
ΔHC	-0.071	-0.102	-0.088
ΔPM	-0.048	-0.053	-0.06
$\Delta Eq. NO_X$	-1.84	-2.03	-2.28
$\%\Delta CO$	-12.4	-20.4	-15.8
$\%\Delta NO_X$	-11	-12.1	-13.6
$\%\Delta HC$	-15.2	-22.1	-19
$\%\Delta PM$	-10.2	-11.2	-12.7
% $\Delta Eq. NO_X$	-10.6	-11.7	-13.2

Table 32: Changes in emissions under the different policies using Copert emission factors.

Note: Δ emissions are in tons.

Table 33: Relative importance of the speed changes for the emission reductions.

	Driving	Uniform	Variable
	restriction	toll	toll
СО	12.6	10.1	8.98
NO_X	11.6	9.9	10.3
HC	19.7	18.8	20.5
PM	8.12	6.88	7.38
Eq. NO_X	10.2	8.66	9.09
01			

Note: in %.

D Results for all day policies

[This section is preliminary.]

We look at the impacts of the stricter policies applicable during the whole day. We calibrate the policy parameters to match the total traffic reduction, both at peak and non-peak hours. Since 4.06% of the individuals in our sample do not have an alternative to cars, we allow them to be non-compliant to the regulation. In exchange, they pay a fee of $\in 68$. This value is inspired by the fine for breaking the driving restriction rule applied to Paris in 2016. The value is also a proxy for the cost of using a taxi instead of using individuals' cars. We only allow those without alternatives

to be non-compliant in the main analysis. However, we check the robustness of our predictions to adding an endogenous compliance behavior. For a fine of $\in 68$, we find that the probability of being non-compliant is very small. These all-day policies imply a larger reduction of total traffic. While we reach a total share of individuals driving between 27.5% to 30.5% under the peak policies, we can decrease the share of individuals driving to 21.8% under these strict policies. As a consequence, we observe larger effects of the policies in terms of welfare loss and emissions reductions.

Modal shifts

Under the mild policies, we predict a large modal shift towards driving at non-peak hours. However, under strict policies, we observe large modal shifts towards other transportation modes. Bicycles and two-wheel vehicles see the largest relative increase: for instance, the share of individuals biking doubles under the driving restriction. The variable toll implies a significantly larger substitution to public transport than the other policies; We predict that 55.6% of the individuals use public transport. Since the number of individuals taking public transport may affect the congestion level in the railway transit and create disutility to users, we evaluate two robustness scenarios in which we raise the public transportation overcrowding level by 10% and 30%. The results show that the shares are barely affected. Given these results, we estimate the percentage of individuals using public transport to commute to be overestimated by at most 0.75% when we do not change overcrowding levels.²⁵

Welfare impacts

The total welfare losses from all-day policies are considerably higher than from peak-only policies. These results highlight the impact of closing the inter-temporal substitution channel that allowed individuals to keep driving. These strict policies decrease the total consumer surplus by \in 3.6 to \in 12.6 million, depending on the type of instrument used.

There are also large differences in the total revenue raised by the different policies. The highest tax revenue is obtained under the vintage-based driving restriction. The revenue is generated by the 4% of the population that do not have an alternative to cars. Most of them have cars older than 2005, and thus they pay the fine. But unlike the standard driving restriction, the fin has to be paid in any case and not only 50% of the time. The standard price instruments also generate large revenue as the tolls reach higher values than under mild policies . Yet, even under full redistribution of the tax revenue to the individuals, all the policies are welfare decreasing and are more costly than the mild policies. As for the mild policies, the uniform toll is close to be a neutral policy under redistribution.. The ranking of the five policy instruments is the same for peak-only and all-day policies.

²⁵Results are available upon request.

Trip durations

All policies imply a reduction in the share of individuals with a decrease in commuting time relative to the mild policies. These policies reach very similar speeds at peak hours but generate higher speeds at non-peak hours. Individuals who want to drive have to pay regardless of the hour and thus prefer to drive at peak hours. Consequently, the traffic does not move from peakhour towards non-peak hours. Individuals unaffected by the driving restriction do not substitute massively towards non-peak hours to take advantage of the faster speeds. This can be explained by the large disutility of driving at non-peak hours.

Under mild policies, the average change in travel time was almost constant across regulations. However, under all-day policies, the differences appear to be more significant. The expected trip duration increases between 3 minutes (under the license auction) to 4.6 minutes (under driving restrictions). The changes can be explained by the individuals with long-distance trips, who value the driving license the most and are the ones paying for it. This, in turn, implies that most individuals substituting to another transport mode have shorter trips and do not incur large travel time increases.

Emission reduction

The same ranking between policies holds both for all-day policies and the peak-only policies in terms of emission reductions. However, there is an important difference in the magnitudes: the strict policies lead to greater reductions in emissions for all pollutants. The reductions more than double between peak-only policies and all-day ones.

The variable toll is again the policy associated with the greatest emissions reductions. The reductions are between 51.2% and 56.8% depending on the pollutant. Unexpectedly, but consistent with the results for the peak policies, the vintage restriction is inefficient at curbing pollution even though it allows only the most modern cars to stay on the road.

Average cost of regulation

We analyze the welfare costs of the strict policies with their impact on emissions by computing the average cost of reducing emissions by one ton. Overall, we find fairly similar costs to those of the mild policies for the three price instruments (fixed and variable tolls and the license). The ranking across these three policies in terms of cost-efficiency is also the same for peak and all-day policies. However, for the standard and vintage-based driving restrictions and without redistribution of the tax revenue, we find that costs are much higher for all-day policies than for peak-only ones. This is because individuals who do not have alternatives to cars have to pay a large fine. When the fine revenue is redistributed to individuals, we estimate significantly lower costs of regulations. The average cost of regulation is 30% lower when the vintage-based policy is implemented all day, while it is 42% lower when we consider standard driving restrictions.

E Robustness checks

E.1 Comparison to alternative demand models

[This section is preliminary.]

Weather controls

We use historical hourly data from OpenWeather for the city of Paris to control for a possible role of weather on individuals' preferences. We match the OpenWeather data to the exact departure hour and date provided in the survey. With this information at hand, we estimate the model with temperature quintile categories, rain categories, and snow categories dummies. The dummies are interacted with the bike and walking alternative constants, as these modes are the ones susceptible to be the most affected by weather shocks. Average coefficients can and VOT estimates can be seen on tables 34 and 35. From the results, we can that the inclusion of weather controls has no significant impact on the average duration and cost estimates, as well as almost no impact on the VOT distribution. Thus, weather conditions do not seem to play an important role in individuals' transportation decisions.

Unobserved heterogeneity

We consider three extensions of the benchmark specification that allow for unobserved heterogeneity in the sensitivity to duration, cost, and the non-peak hour dummy. For all cases we estimate a random coefficients version of the benchmark model where we assume that the unobserved heterogeneity follows a normal distribution. Looking at the parameter estimates from table 34, we find non-significant unobserved heterogeneity for the random coefficients model on duration and for the one of the non-peak dummy. These results indicate that the demographic interactions that are included in the model account for the large majority of the sensitivity heterogeneity of those variables. VOT estimates from table 35.

Models using duration

We compare our benchmark results with those of a model using duration instead of the logarithm of duration, as well as with a model using a third degree polynomial on duration. The VOT estimates from table 35 indicate how the valuation of travel time by individuals is highly dependent on the assumed functional form. Using the logarithm of duration allows for shorter trips to have a larger disutility from additional travel. As expected, forcing all individuals to have the same valuation of travel time regardless of how long is their trip (duration specification) considerable reduce the average valuation of travel time. We check an alternative specification that allows for a variable marginal disutility of time by estimating a model where individuals have preferences over a polynomial of degree three on duration. The polynomial assumption induces the existence of outliers in the VOT distribution but if we check the median VOT value we see it remains close to the average VOT from our benchmark model.

Alternative model specifications

We consider several alternative model specifications to check different main assumptions of our model. First we consider an alternative definition of non-peak hour driving durations. Instead of taking the average duration between the TomTom queries at 6:30 am and 9:30 am, we use only the 6:30 am durations. This change has very limited effects on the average coefficients and almost no effect on the distribution of VOT estimates. We evaluate an extension with three periods: early non-peak hour (departure earlier than 7:00 am), peak hour (departure between 7:00 am and 9:00 am), and late non-peak hour (9:30 am onwards). The distribution of VOT values indicates a lower value of travel time, with can be explained by a large increase in the sensitivity to the trip's cost. We also consider a model where we inverse the nest structure, i.e., we allow individuals to choose between peak and non-peak hour for all modes and allow for correlation in the utility shocks of all the alternatives of a same time period. The parameter estimates from table 34 show a negative nest parameter (σ), indicating that this nest structure is not appropriate to represent the data.

Finally, we consider two possible changes to our choice set definition. In first change we allow every individual to drive, even those that do not have access to car. For those without a car, we input the cot of a car with the average characteristics of those available in the sample. For the travel times, we query the TomTom service for those trips, as with the rest of the sample. As expected, adding a non-available alternative bias the results by making individuals less cost sensitive and thus increasing their value of travel time. The second change we study corresponds to changing the availability of bike to only individuals who either report having a bike or a bike-sharing pass. While the change seems to increase the upper bud of the VOT distribution, the median of the distribution remains close to our benchmark model.

Coefficients	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Log(duration)	-2.26^{**}	-2.26^{**}	-1.93^{**}	-2.08**	-1.92^{**}			-2.84^{**}	-2.77^{**}	-1.42^{**}	-2.14^{**}	-2.26**
	(0.07)	(0.07)	(0.24)	(0.18)	(0.18)			(0.07)	(0.06)	(0.24)	(0.06)	(0.07)
Duration						-0.75**	-1.5^{**}					
						(0.02)	(0.07)					
Cost	-0.51^{**}	-0.51^{**}	-0.41^{**}	-0.75^{**}	-0.41^{**}	-0.83**	-1.05^{**}	-0.67**	-0.76**	-0.35**	-0.4**	-0.47^{**}
	(0.02)	(0.02)	(0.05)	(0.11)	(0.05)	0.02	0.01	(0.02)	(0.02)	(0.03)	(0.02)	(0.02)
Unobserved heterogeneity			-0.16	0.71**	0.1							
			(0.2)	(0.09)	(0.24)							
Bicycle	-3.48**	-3.48**	-3.48**	-3.61**	-3.48**	-2.61**	0	-3.49**	-3.4**	0.94^{**}	-3.59**	-3.18**
	(0.08)	(0.08)	(0.39)	(0.37)	(0.39)	0.07	0	(0.08)	(0.08)	(0.3)	(0.08)	(0.08)
Public transport, peak	-4.89**	-4.84**	-4.91**	-5.27**	-4.89**	-0.81**	-0.49**	-4.85^{**}	-4.92**	0.01	-5.05**	-4.82**
	(0.2)	(0.22)	(0.54)	(0.54)	(0.54)	0.09	0.13	(0.2)	(0.2)	(0.02)	(0.19)	(0.2)
Public transport, non-peak	-5.52^{**}	-5.47^{**}	-5.55^{**}	-5.97^{**}	-5.52^{**}	-1.36^{**}	-2.96^{**}	-5.34^{**}	-5.74^{**}	0.53^{**}	-5.72^{**}	-5.45^{**}
	(0.4)	(0.44)	(1.22)	(1.22)	(1.22)	0.17	0.19	(0.4)	(0.4)	(0.06)	(0.39)	(0.4)
Motorized 2-wheel	-7.35**	-7.31^{**}	-7.38**	-7.91^{**}	-7.36**	-3.06**	-1.4**	-7.42^{**}	-7.3**	0.64^{\dagger}	-7.66**	-7.3**
	(0.23)	(0.24)	(0.54)	(0.51)	(0.54)	0.13	0.07	(0.22)	(0.21)	(0.36)	(0.23)	(0.23)
Car, peak	-6.22**	-6.17^{**}	-6.25^{**}	-6.75^{**}	-6.22**	-2.06**	-1.74^{**}	-6.19^{**}	-6.13^{**}	0.04	-6.8**	-6.16^{**}
	(0.21)	(0.23)	(0.85)	(0.85)	(0.86)	0.1	0.15	(0.21)	(0.21)	(0.21)	(0.21)	(0.21)
Car, non-peak	-7.29**	-7.24^{**}	-7.32**	-7.92^{**}	-7.29**	-2.92^{**}	-3.82**	-7**	-7.28**	-0.83**	-7.91^{**}	-7.22**
	(0.21)	(0.23)	(0.68)	(0.68)	(0.68)	0.1	0.12	(0.21)	(0.21)	(0.07)	(0.21)	(0.21)
σ	0.8**	0.8**	0.81**	0.89**	0.8**	0.7**	0.47^{**}	0.61**	0.56^{**}	-0.31**	0.85^{**}	0.79**
	(0.06)	(0.06)	(0.07)	(0.08)	(0.07)	0.06	0.06	(0.05)	(0.03)	(0.1)	(0.07)	(0.06)
No. observations	12,975	12,975	12,975	12,975	12,975	12,975	12,975	12,975	12,975	12,975	12,975	12,975
Log-likelihood	-13624	-13615	-13623	-13561	-13624	-13421	-13442	-13587	-15823	-14991	-14679	-13547

Table 34: Robustness checks - demand specification

Notes: Walking is the baseline alternative. The reference categories are individuals with age < 18, the first wealth quintile and independent workers. Duration measured in minutes. Cost in \in . Column (1) shows the benchmark model. Column (2) adds weather controls. Column (3) corresponds to the benchmark model with random coefficients in the log of duration. Column (4) corresponds to the benchmark model with random coefficients in the cost. Column (5) corresponds to the benchmark model with random coefficients in non-peak indicator. Column (6) model with duration (in minutes). Column (7) model with third degree polynomial of duration (in minutes). Column (8) benchmark model with alternative duration for non-peak hours. Column (9) Three period model. Column (10) model inversing the nest structure. Column (11) car availability for everyone. Column (12) alternative definition of bike availability. Standard-errors are computed using the delta-method.

	Min	Q1	Mean	Median	Q99	Max
log(Duration) models						
Benchmark	0.44	1.35	15.96	10.29	81.2	388.87
Benchmark + weather controls	0.44	1.34	15.98	10.29	81.34	390.99
Random coefficients models						
$\log(duration)$ - normal distribution	0.43	1.34	15.99	10.31	81.58	388.84
Cost - normal distribution	0.29	0.84	9.43	5.99	47.67	236.54
Dummy non-peak - normal distribution	0.44	1.35	15.96	10.28	81.24	388.77
Duration models						
Duration (in min)	-0.27	-0.27	5.39	5.83	13.73	15.74
Polynomial duration	-3.6	-2.95	55.22	17.71	483.79	2032.89
Alternative specification						
Alternative non-peak duration	0.44	1.31	15.86	10.22	81.51	387.15
3 period model	0.21	0.76	13.7	8.59	72.4	335.01
Inverse nest	0.89	2.7	21.8	15.61	110.39	587.11
Alt choice set - car	0.76	1.82	19.58	12.14	97.75	528.59
Alt choice set - bike	0.49	1.47	17.14	11.09	87.78	422.9

Table 35 Value of travel time (VOT) for the demand robustness specifications

Note: VOT in \in/hr

E.2 Sensitivity to public transport level of congestion

Table 36: Shares of transportation modes under different overcrowding levels

	No policy	Driving restriction			Uı	Uniform toll			Variable toll		
Mode	(1)	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)	
Bicycle	2.09	2.42	2.43	2.44	2.42	2.42	2.44	2.21	2.22	2.24	
Pub. Transport, peak	30.28	32.08	31.91	31.58	32.36	32.19	31.85	32.94	32.77	32.43	
2 wheels	2.07	2.45	2.46	2.49	2.47	2.48	2.5	2.6	2.62	2.64	
Walking	15.78	17.56	17.58	17.63	17.8	17.82	17.86	15.88	15.9	15.94	
Car, peak	23	14.01	14.04	14.1	14.01	14.04	14.11	14.01	14.03	14.09	
Car, non-peak	12.22	16.14	16.18	16.26	15.48	15.52	15.59	16.53	16.58	16.66	
Pub. Transport, non-peak	14.56	15.34	15.39	15.5	15.47	15.53	15.64	15.82	15.88	16	
Total car share	35.21	30.15	30.22	30.36	29.48	29.56	29.7	30.54	30.61	30.75	
Total PT share	44.84	47.42	47.3	47.08	47.83	47.72	47.49	48.76	48.65	48.44	

Notes: In km/hr. (1): No improvement. (2): 15% increase in overcrowding. (3): 30% increase in overcrowding

		No policy	Driving restriction			Un	Uniform toll			Variable toll		
	Area	(1)	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)	
Peak	Highways	65.7	84.3	84.2	84	83.5	83.4	83.2	93.5	93.4	93.4	
	City center	13.7	17.3	17.2	17.1	18	17.9	17.8	19.3	19.3	19.2	
	Ring roads	28.4	44.1	43.9	43.4	46.4	46.2	45.6	63	62.9	62.6	
	Close suburb	15.8	19.4	19.4	19.3	19.6	19.6	19.5	20.5	20.4	20.4	
Non-peak	Highways	84.8	75.7	75.6	75.3	77.7	77.6	77.3	70.4	70.2	69.9	
	City center	18.4	17.8	17.7	17.6	17.9	17.9	17.8	17.7	17.6	17.5	
	Ring roads	45	42.3	42	41.5	42.9	42.7	42.2	38.9	38.6	38.2	
	Close suburb	20.2	19	19	18.9	19.2	19.2	19.1	18.7	18.6	18.6	

Table 37: Speeds under different overcrowding levels

Notes: In km/hr. (1): No improvement. (2): 15% increase in overcrowding. (3): 30% increase in overcrowding